This portfolio contains non-proprietary submissions received by the National Coordination Office (NCO) for the Networking and Information Technology Research and Development (NITRD) Program, in response to three public Requests for Input (RFI) to the National Cyber Leap Year (see 1 NCLY RFI-1, 2 NCLY RFI-2, and 3 NCLY RFI-3 documents in this portfolio). The submissions were used as input by the NITRD Senior Steering Group for Cybersecurity in the development of a conceptual framework to focus various cyber security R&D activities into a coherent model for the greatest impact. The framework focuses on identifying what could be changed (and why) to be able to play a different cyber "game" (as in "if you are playing a game you can't win, change the game!").

For further information or to send comments email: nco@nitrd.gov

Request for Input (RFI) – National Cyber Leap Year

Overview: This Request for Information (RFI) is issued under the Comprehensive National Cybersecurity Initiative (CNCI), established within Homeland Security Presidential Directive (HSPD) -23. The RFI was developed by the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity to invite participation in a National Cyber Leap Year whose goal is an integrated national approach to make cyberspace safe for the American way of life.

Background: We are a cyber nation. The U.S. information infrastructure – including telecommunications and computer networks and systems and the data that reside on them – is critical to virtually every aspect of modern life. This information infrastructure is increasingly vulnerable to exploitation, disruption, and destruction by a growing array of adversaries. The President's CNCI plan calls for *leap-ahead* research and technology to reduce vulnerabilities to asymmetric attack in cyberspace. Unlike many research agenda that aim for steady progress in the advancement of science, the leap-ahead effort seeks just a few revolutionary ideas with the potential to reshape the landscape. These *game-changing* technologies (or non-technical mechanisms that are made possible through technology), developed and deployed over the next decade, will fundamentally change the cyber game into one where the good guys have an advantage. Leap-ahead technologies are so-called because they enable us to leap over the obstacles preventing us from being where we want to be. These advances may require years of concerted research and development to be fully realized; good ideas often do. However, the intent is to *start now* and gain momentum as intermediate results emerge.

Objective: The National Cyber Leap Year has two main goals: (1) constructing a national research and technology agenda that both identifies the most promising ideas and describes the strategy that brings those ideas to fruition; and (2) jumpstarting game-changing, multi-disciplinary development efforts.

The Leap Year will run during fiscal year 2009, and will comprise two stages: *prospecting* and *focusing*.

Stage One, which we open with this formal announcement and describe in detail below, canvasses the cybersecurity community for ideas. Our aim is to hear from all those who wish to help.

The heart of Stage Two, which begins February 1, 2009, is a series of workshops to develop the best ideas from Stage One. As the year progresses, we will publish four types of findings: (1) *Game-changers*—descriptions of the paradigm-busters that technology will make possible; (2) *Technical Strategy*—as specifically as possible, the invention and/or research that needs to be done; (3) *Productization/Implementation*—how the capability will be packaged, delivered, and used, and by whom; and (4) *Recommendations*—prescriptions for success, to include funding, policies, authorities, tasking—whatever would smooth the way to realization of the game-changing capability.

***Deadline for Submission under this RFI*:** We anticipate multiple cycles of Stage One opportunities. The first Stage One cycle is covered by this RFI and will close **December 15, 2008**. Subsequent cycles will be announced by separate RFIs. All Stage One cycles are expected to be complete by April 15, 2009.

Stage One description

What we are looking for:

Contributors may submit up to 3 leap-ahead technology concepts. Multidisciplinary contributions from organizations with cybersecurity interests are especially encouraged.
Cognizant of the limits of conventional studies and reports, we have given substantial thought to what framework and methodology might render the community's best ideas understandable, compelling, and actionable to those who need to support them, fund them, and adopt them. Since our search is for game-changing concepts, we ask that submitters explain their ideas in terms of a game. Many ideas will fall into the following three categories. Ideas that:

Morph the gameboard (change the defensive terrain [permanently or adaptively] to make it harder for the attacker to maneuver and achieve his goals)

Example: non-persistent virtual machines – every time the enemy takes a hill, the hill goes away

Change the rules (lay the foundation for cyber civilization by changing network protocols and norms to favor our society's values)

Example: the no-call list – direct marketers have to "attack" on customer terms now

Raise the stakes (make the cost to play less advantageous to the attacker by raising risk, lowering value, etc.)

Example: charging for email – making the SPAMmer ante up means a lot more fish have to bite for SPAM to pay

Ideas that change the game in some other dimension are also welcome; just be sure to explain how.

Who can participate:

This RFI is open to all and we especially encourage public- and private-sector groups (e.g., universities, government laboratories, companies, non-profit groups, user groups) with cybersecurity interests to participate. Collaborative, multidisciplinary efforts are also highly encouraged. Participants in Stage One must be willing to participate in Stage Two should one of their ideas be selected. Participants must also be willing to have their ideas posted for discussion on a public website and/or included in our final report.

How we will use it:

The best ideas from Stage One will go on to Stage Two. Stage One submissions may be posted on our website for elaboration and improvement, as a key goal of the Leap Year is to engage diverse sectors (e.g., government, academia, commercial, international) in identifying multidimensional strategies and, where it makes sense, in rolling up their sleeves and starting to work. Submissions crafted with that larger community in mind will be the most compelling and influential.

Leap Year interim results and emerging guidance will be posted at: www.nitrd.gov/leapyear/.

Questions and submissions should be addressed to: leapyear@nitrd.gov.

In accordance with FAR 15.202(3), responses to this notice are not offers and cannot be accepted by the Government to form a binding contract. Responders are solely responsible for all expenses associated with responding to this RFI, including any subsequent requests for proposals.

All responses must be no more than two pages long (12 pt font, 1" margins) and in this form:

Who you are – name, credentials, group membership

Game-changing dimension – Board, Rules, Stakes?

Concept – What is the idea and why does it change the game?

Vision – Make us believe in your idea (What would the world look like if this were in place? How would people get it, use it? What makes you think this is possible? What needs to happen for this to become real? Which parts already exist; which parts need to be invented?)

Method – What process did you use to formulate and refine your concept? What assumptions or dependencies underlie your analysis?

Dream team – Who are the people you'd need to have on your team to make this real? If you just know disciplines that's OK. If you have names, explain what those people do. If your idea is selected for further consideration, we will do our best to bring these people together for a Stage Two workshop.

Responses must be submitted via www.nitrd.gov/leapyear/ or emailed to leapyear@nitrd.gov.

Responses to this RFI must be received by December 15, 2008 so that arrangements can be made for Stage Two activities beginning on or about February 1, 2009. Additional Stage One cycles, if any, will be announced by separate RFI with all Stage One activities expected to be complete by April 15, 2009.

Appendix A contains a sample submission and review considerations.

Appendix A

Sample submission

Who you are – quieteveningathome.org – We are a 501c3 group with 50,000 members dedicated to the preservation of the dinner hour as the core of American civilization.

Game-changing dimension – Change the rules

Concept – Telemarketers are using our resources and time to market their products. They can call and interrupt our dinners and use our own telephones to reach us. What if we changed the rules to "don't call us, we'll call you?"

Vision – The vision is a national do-not-call register. People should be able to go to donotcall.gov and register their phone number. It would be illegal for telemarketers who have not been given permission to call someone. If a telemarketer makes an illegal call, the recipient should be able to report them to a government agency and they should be fined. The technology to do this is easy, we are not sure about the laws and policies. Courts have ruled differently on this issue at different times. We think the political climate is friendly today for Federal legislation.

Method – We announced our search for ideas on our website and submissions were made there. We also publicized through restaurant and catering associations with whom we often partner, who offered interruption-free free meals for brainstorming sessions. Participation was not limited to members, but could not be anonymous, since it was our intention to follow up with submitters. The Board of Directors of QEAH enlisted the aid of Prandia University to work with the submitters of the best ideas to develop them into even better ideas. The Board ensured all the aspects described in the Leap Year RFI were addressed in our final submissions.

Dream team – Federal Trade Commission, Federal Communications Commission, constitutional lawyer, Telemarketers' Association, Consumers Union, Oracle or other database company.

Review considerations

Submissions will be reviewed by the NITRD Senior Steering Group for Cybersecurity using the following considerations:

Would it change the game?

How clear is the way forward?

What heights are the hurdles that may be found in the way forward?

Request for Input No. 2 (RFI-2) – National Cyber Leap Year

Overview: This Request for Information No. 2 (RFI-2) is the second issued under the Comprehensive National Cybersecurity Initiative (CNCI), established within Homeland Security Presidential Directive (HSPD) -23. RFI-2 was developed by the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity to invite participation in a National Cyber Leap Year whose goal is an integrated national approach to make cyberspace safe for the American way of life. Over 160 responses were submitted to the first RFI issued by the NITRD SSG (October 14, 2008), indicating a strong desire by the technical community to participate. RFI-2 expands the opportunities for participation by permitting submitters to designate parts of submissions as proprietary.

Background: We are a cyber nation. The U.S. information infrastructure – including telecommunications and computer networks and systems and the data that reside on them – is critical to virtually every aspect of modern life. This information infrastructure is increasingly vulnerable to exploitation, disruption, and destruction by a growing array of adversaries. The President's CNCI plan calls for *leap-ahead* research and technology to reduce vulnerabilities to asymmetric attack in cyberspace. Unlike many research agenda that aim for steady progress in the advancement of science, the leap-ahead effort seeks just a few revolutionary ideas with the potential to reshape the landscape. These *game-changing* technologies (or non-technical mechanisms that are made possible through technology), developed and deployed over the next decade, will fundamentally change the cyber game into one where the good guys have an advantage. Leap-ahead technologies are so-called because they enable us to leap over the obstacles preventing us from being where we want to be. These advances may require years of concerted research and development to be fully realized; good ideas often do. However, the intent is to *start now* and gain momentum as intermediate results emerge.

Objective: The National Cyber Leap Year has two main goals: (1) constructing a national research and technology agenda that both identifies the most promising ideas and describes the strategy that brings those ideas to fruition; and (2) jumpstarting game-changing, multi-disciplinary development efforts. The Leap Year will run during fiscal year 2009, and will comprise two stages: *prospecting* and *focusing*.

Stage One canvasses the cybersecurity community for ideas. Our aim is to hear from all those who wish to help.

The heart of Stage Two, which begins February 1, 2009, is a series of workshops to explore the best ideas from Stage One. As the year progresses, we will publish four types of findings: (1) *Game-changers*—descriptions of the paradigm-busters that technology will make possible; (2) *Technical Strategy*—as specifically as possible, the invention and/or research that needs to be done; (3) *Productization/Implementation*—how the capability will be packaged, delivered, and used, and by whom; and (4) *Recommendations*—prescriptions for success, to include funding, policies, authorities, tasking—whatever would smooth the way to realization of the game-changing capability.

***Deadline for Submission under this RFI-2*:** The second round of the Stage One cycle is covered by this RFI-2 and will close **February 20, 2009**. Subsequent cycles will be announced by separate RFIs. All Stage One cycles are expected to be complete by April 15, 2009.

Stage One description

What we are looking for:

Contributors may submit up to 3 leap-ahead technology concepts. Multidisciplinary contributions from organizations with cybersecurity interests are especially encouraged. Cognizant of the limits of conventional studies and reports, we have given substantial thought to what framework and methodology might render the community's best ideas understandable, compelling, and actionable to those who need to support them, fund them, and adopt them. Since our search is for game-changing concepts, we ask that submitters explain their ideas in terms of a game. Many ideas will fall into the following three categories. Ideas that:

Morph the gameboard (change the defensive terrain [permanently or adaptively] to make it harder for the attacker to maneuver and achieve his goals)

Example: non-persistent virtual machines – every time the enemy takes a hill, the hill goes away

Change the rules (lay the foundation for cyber civilization by changing network protocols and norms to favor our society's values)

Example: the no-call list – direct marketers have to "attack" on customer terms now

Raise the stakes (make the cost to play less advantageous to the attacker by raising risk, lowering value, etc.)

Example: charging for email – making the SPAMmer ante up means a lot more fish have to bite for SPAM to pay

Ideas that change the game in some other dimension are also welcome; just be sure to explain how. The rationale for why the idea is game-changing should be the central focus of each submission.

Who can participate:

This RFI-2 is open to all and we especially encourage public- and private-sector groups (e.g., universities, government laboratories, companies, non-profit groups, user groups) with cybersecurity interests to participate. Collaborative, multidisciplinary efforts are also highly encouraged. Participants in Stage One must be willing to participate in Stage Two should one of their ideas be selected. Excluding proprietary information, participants must also be willing to have their ideas posted for discussion on a public website and/or included in our final report.

How we will use it:

The best ideas from Stage One will go on to Stage Two. Non-proprietary elements of Stage One submissions may be posted on our website for elaboration and improvement, as a key goal of the Leap Year is to engage diverse sectors (e.g., government, academia, commercial, international) in identifying multidimensional strategies and, where it makes sense, in rolling up their sleeves and starting to work. Submissions crafted with that larger community in mind will be the most compelling and influential.

Leap Year interim results and emerging guidance will be posted at: www.nitrd.gov/leapyear/.

Questions and submissions should be addressed to: leapyear@nitrd.gov.

In accordance with FAR 15.202(3), responses to this notice are not offers and cannot be accepted by the Government to form a binding contract. Responders are solely responsible for all expenses associated with responding to this RFI-2, including any subsequent requests for proposals.

All responses must be no more than two pages long (12 pt font, 1" margins) and in this form:

RFI Name: RFI–2 – National Cyber Leap Year

Title of Concept

RFI Focus Area (Morph the gameboard, Change the rules, Raise the stakes)

Submitter's Contact Information – Name, Organization, Address, Telephone number, Email address

Summary of who you are – credentials, group membership

Concept – What is the idea? Explain why it would change the game. Introducing a good idea alone is not sufficient; you must explain how it changes the game.

Vision – Make us believe in your idea (What would the world look like if this were in place? How would people get it, use it? What makes you think this is possible? What needs to happen for this to become real? Which parts already exist; which parts need to be invented?)

Method – What process did you use to formulate and refine your concept? What assumptions or dependencies underlie your analysis?

Dream team – Who are the people you'd need to have on your team to make this real? If you just know disciplines that's okay. If you have names, explain what those people do. If your idea is selected for further consideration, we will do our best to bring these people together for a Stage Two workshop.

Labeling of Proprietary Information – Clearly label any part of the submission designated as proprietary. The proprietary information will be restricted to government use only. If the submission is selected for Stage Two, we will work with the submitter to determine exactly what information warrants proprietary protection and to establish appropriate controls for managing, protecting, and negotiating as appropriate the relevant intellectual property rights.

Responses must be submitted via www.nitrd.gov/leapyear/ or emailed to leapyear@nitrd.gov, and must be received by February 20, 2009. Additional Stage One cycles, if any, will be announced by separate RFI with all Stage One activities expected to be complete by April 15, 2009.

Appendix A contains a sample submission and review considerations.

Appendix A - Sample submission

Who you are – quieteveningathome.org – We are a 501c3 group with 50,000 members dedicated to the preservation of the dinner hour as the core of American civilization.

Game-changing dimension – Change the rules

Concept – Telemarketers are using our resources and time to market their products. They can call and interrupt our dinners and use our own telephones to reach us. What if we changed the rules to "don't call us, we'll call you?" Changing this rule changes the game to one where we decide which marketers to contact and when, returning control of the dinner hour to us.

Vision – The vision is a national do-not-call register. People should be able to go to donotcall.gov and register their phone number. It would be illegal for telemarketers who have not been given permission to call someone. If a telemarketer makes an illegal call, the recipient should be able to report them to a government agency and they should be fined. The technology to do this is easy, we are not sure about the laws and policies. Courts have ruled differently on this issue at different times. We think the political climate is friendly today for Federal legislation.

Method – We announced our search for ideas on our website and submissions were made there. We also publicized through restaurant and catering associations with whom we often partner, who offered interruption-free free meals for brainstorming sessions. Participation was not limited to members, but could not be anonymous, since it was our intention to follow up with submitters. The Board of Directors of QEAH enlisted the aid of Prandia University to work with the submitters of the best ideas to develop them into even better ideas. The Board ensured all the aspects described in the Leap Year RFI were addressed in our final submissions.

Dream team – Federal Trade Commission, Federal Communications Commission, constitutional lawyer, Telemarketers' Association, Consumers Union, Oracle or other database company.

Review considerations

Submissions will be reviewed by the NITRD Senior Steering Group for Cybersecurity using the following considerations:

Would it change the game?

How clear is the way forward?

What heights are the hurdles that may be found in the way forward?

Request for Input No. 3 (RFI-3) – National Cyber Leap Year

Overview: This Request for Input No. 3 (RFI-3) is the third issued under the Comprehensive National Cybersecurity Initiative (CNCI), established within Homeland Security Presidential Directive (HSPD) -23. RFI-3 was developed by the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity to invite participation in a National Cyber Leap Year whose goal is an integrated national approach to make cyberspace safe for the American way of life. Over 160 responses were submitted to the first RFI issued by the NITRD SSG (October 14, 2008), indicating a strong desire by the technical community to participate. RFI-2 (issued on December 30, 2008) expanded the opportunity for participation by permitting submitters to designate parts of submissions as proprietary. RFI-3 presents prospective cyber security categories derived from responses to RFI-1 for further consideration.

Background: We are a cyber nation. The U.S. information infrastructure – including telecommunications and computer networks and systems and the data that reside on them – is critical to virtually every aspect of modern life. This information infrastructure is increasingly vulnerable to exploitation, disruption, and destruction by a growing array of adversaries. The President's CNCI plan calls for *leap-ahead* research and technology to reduce vulnerabilities to asymmetric attack in cyberspace. Unlike many research agenda that aim for steady progress in the advancement of science, the leap-ahead effort seeks just a few revolutionary ideas with the potential to reshape the landscape. These *game-changing* technologies (or non-technical mechanisms that are made possible through technology), developed and deployed over the next decade, will fundamentally change the cyber game into one where the good guys have an advantage. Leap-ahead technologies are so-called because they enable us to leap over the obstacles preventing us from being where we want to be. These advances may require years of concerted research and development to be fully realized; good ideas often do. However, the intent is to *start now* and gain momentum as intermediate results emerge.

Objective: The National Cyber Leap Year has two main goals: (1) constructing a national research and technology agenda that both identifies the most promising ideas and describes the strategy that brings those ideas to fruition; and (2) jumpstarting game-changing, multi-disciplinary development efforts. The Leap Year will run during fiscal year 2009, and will comprise two stages: *prospecting* and *focusing*.

Stage One canvasses the cybersecurity community for ideas. Our aim is to hear from all those who wish to help.

The heart of Stage Two, which begins March 1, 2009, is a series of workshops to explore the best ideas from Stage One. As the year progresses, we will publish four types of findings: (1) *Game-changers*—descriptions of the paradigm-busters that technology will make possible; (2) *Technical Strategy*—as specifically as possible, the invention and/or research that needs to be done; (3) *Productization/Implementation*—how the capability will be packaged, delivered, and used, and by whom; and (4) *Recommendations*—prescriptions for success, to include funding,

policies, authorities, tasking—whatever would smooth the way to realization of the game-changing capability.

***Deadline for Submission under this RFI-3*:** The third, and final round of the Stage One cycle is covered by this RFI-3 and will close **April 15, 2009**.

Stage One description

What we are looking for:

Contributors may submit up to 3 leap-ahead technology concepts. Multidisciplinary contributions from organizations with cybersecurity interests are especially encouraged. Cognizant of the limits of conventional studies and reports, we have given substantial thought to what framework and methodology might render the community's best ideas understandable, compelling, and actionable to those who need to support them, fund them, and adopt them. Since our search is for game-changing concepts, we ask that submitters explain their ideas in terms of a game. Many ideas will fall into the following three categories. Ideas that:

Morph the gameboard (change the defensive terrain [permanently or adaptively] to make it harder for the attacker to maneuver and achieve his goals)

Example: non-persistent virtual machines – every time the enemy takes a hill, the hill goes away

Change the rules (lay the foundation for cyber civilization by changing network protocols and norms to favor our society's values)

Example: the no-call list – direct marketers have to "attack" on customer terms now

Raise the stakes (make the cost to play less advantageous to the attacker by raising risk, lowering value, etc.)

Example: charging for email – making the SPAMmer ante up means a lot more fish have to bite for SPAM to pay

Ideas that change the game in some other dimension are also welcome; just be sure to explain how. The rationale for why the idea is game-changing should be the central focus of each submission.

Submitters are encouraged to explore the following categories, which were derived by the NITRD SSG from the review of RFI-1 submissions. These categories encompass promising concepts identified by compelling submissions and may be fruitful themes for additional game-changing insights:

Attribution – Technologies and methods to establish that a particular entity (person, host, event) is the originator of an object (e.g. data) or the cause of an effect

Cyber Economics – Security decision-making frameworks that incorporate economic insights; understanding and altering economic value-chains to make cyber security exploits increasingly expensive for attackers

Disaster Recovery – Recovery in the event of a large-scale disruption of national cyber assets
Network Ecology – Incorporating end-to-end network management techniques to control access to and allocation of network resources; modeling of acceptable host and network activities
Policy-based Configuration/Implementation – Standards-based security policy definitions and enforcement frameworks; architectures and techniques for implementing fine-coarse access and permission controls
Randomization/Moving Target – Software diversity that randomizes code structure; virtualization techniques that hide, obscure, move, and alter; randomizing and obfuscating network resources, IP addresses, and the operating system; time-varying, crypto-based identities to identify services, hosts, interfaces, networks and users
Secure Data – Building provenance and access controls into the fabric of digital data
Software Assurance – Security-focused system assurance programming languages
Virtualization – Cloud-based virtual desktops for stateless thin clients; high-security hypervisors; least-authority execution via adaptive sandboxes

Submissions in areas outside these categories will also be considered.

Who can participate:

This RFI-3 is open to all and we especially encourage public- and private-sector groups (e.g., universities, government laboratories, companies, non-profit groups, user groups) with cybersecurity interests to participate. Collaborative, multidisciplinary efforts are also highly encouraged. Participants in Stage One must be willing to participate in Stage Two should one of their ideas be selected. Excluding proprietary information, participants must also be willing to have their ideas posted for discussion on a public website and/or included in our final report.

How we will use it:

The best ideas from Stage One will go on to Stage Two. Non-proprietary elements of Stage One submissions may be posted on our website for elaboration and improvement, as a key goal of the Leap Year is to engage diverse sectors (e.g., government, academia, commercial, international) in identifying multidimensional strategies and, where it makes sense, in rolling up their sleeves and starting to work. Submissions crafted with that larger community in mind will be the most compelling and influential.

Leap Year interim results and emerging guidance will be posted at: www.nitrd.gov/leapyear/.

Questions and submissions should be addressed to: leapyear@nitrd.gov.

In accordance with FAR 15.202(3), responses to this notice are not offers and cannot be accepted by the Government to form a binding contract. Responders are solely responsible for all expenses associated with responding to this RFI-3, including any subsequent requests for proposals.

All responses must be no more than two pages long (12 pt font, 1” margins) and in this form:

RFI Name: RFI–3 – National Cyber Leap Year

Title of Concept

RFI Focus Area (Morph the gameboard, Change the rules, Raise the stakes)

Submitter's Contact Information – Name, Organization, Address, Telephone number, Email address

Summary of who you are – credentials, group membership

Concept – What is the idea? Explain why it would change the game. Introducing a good idea alone is not sufficient; you must explain how it changes the game.

Vision – Make us believe in your idea (What would the world look like if this were in place? How would people get it, use it? What makes you think this is possible? What needs to happen for this to become real? Which parts already exist; which parts need to be invented?)

Method – What process did you use to formulate and refine your concept? What assumptions or dependencies underlie your analysis?

Dream team – Who are the people you'd need to have on your team to make this real? If you just know disciplines that's okay. If you have names, explain what those people do. If your idea is selected for further consideration, we will do our best to bring these people together for a Stage Two workshop.

Labeling of Proprietary Information – Clearly label any part of the submission designated as proprietary. The proprietary information will be restricted to government use only. If the submission is selected for Stage Two, we will work with the submitter to determine exactly what information warrants proprietary protection and to establish appropriate controls for managing, protecting, and negotiating as appropriate the relevant intellectual property rights.

Responses must be submitted via www.nitrd.gov/leapyear/ or emailed to leapyear@nitrd.gov, and must be received by April 15, 2009.

Appendix A contains a sample submission and review considerations.

Appendix A - Sample submission

Who you are – quieteveningathome.org – We are a 501c3 group with 50,000 members dedicated to the preservation of the dinner hour as the core of American civilization.

Game-changing dimension – Change the rules

Concept – Telemarketers are using our resources and time to market their products. They can call and interrupt our dinners and use our own telephones to reach us. What if we changed the rules to "don't call us, we'll call you?" Changing this rule changes the game to one where we decide which marketers to contact and when, returning control of the dinner hour to us.

Vision – The vision is a national do-not-call register. People should be able to go to donotcall.gov and register their phone number. It would be illegal for telemarketers who have not been given permission to call someone. If a telemarketer makes an illegal call, the recipient should be able to report them to a government agency and they should be fined. The technology to do this is easy, we are not sure about the laws and policies. Courts have ruled differently on this issue at different times. We think the political climate is friendly today for Federal legislation.

Method – We announced our search for ideas on our website and submissions were made there. We also publicized through restaurant and catering associations with whom we often partner, who offered interruption-free free meals for brainstorming sessions. Participation was not limited to members, but could not be anonymous, since it was our intention to follow up with submitters. The Board of Directors of QEAH enlisted the aid of Prandia University to work with the submitters of the best ideas to develop them into even better ideas. The Board ensured all the aspects described in the Leap Year RFI were addressed in our final submissions.

Dream team – Federal Trade Commission, Federal Communications Commission, constitutional lawyer, Telemarketers' Association, Consumers Union, Oracle or other database company.

Review considerations

Submissions will be reviewed by the NITRD Senior Steering Group for Cybersecurity using the following considerations:

Would it change the game?

How clear is the way forward?

What heights are the hurdles that may be found in the way forward?

Accenture, LLP
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Application Layer Information Assurance

Who We Are

Accenture, LLP is a global management consulting, technology services and outsourcing company. Combining unparalleled experience, comprehensive capabilities across all industries and business functions, and extensive research on the world's most successful companies, we work with clients to help them become high-performance businesses and governments.

Accenture has collaborated with Professor Barry Horowitz from the Department of Systems and Information Engineering at the University of Virginia (UVA) on this concept, based on a three-year relationship with Dr. Horowitz and a research consortium he chairs.

Game-Changing Dimension: Change the Rules

Concept

The growth in IT supply chain risks seriously reduce confidence in information assurance (IA) approaches that depend on the commodity IT components that are the basis for building information systems. Nonetheless, the economics and technology performance advantages of commodity IT components are sufficiently compelling that, except for very unusual circumstances, system builders are not able to avoid their use and satisfy system requirements (both economic and performance requirements). The risks of commodity IT components face additional problems from Service-Oriented Architectures (SOAs), whose information resources and services are more fluidly communicating within and outside of the enterprise.

We propose to focus on the unique attributes of SOAs that can be exploited to *reduce* IA risk. In particular, we propose building an Application Layer IA solution. This solution would enhance security through the use of SOA frameworks. It also would develop an ability to detect or avoid adversarial actions occurring at the application layer. For this solution, we would develop:

- Techniques for exploiting application layer knowledge about structure of objects and integration into services as basis for detecting data disruptions stemming from Trojan Horse-type malware.
- Techniques for exploiting information about historical resource use that relates to providing SOA services as a basis for recognizing unusual patterns of computing activity for recognizing adversarial activity (e.g., bulk data stealing)
- Techniques for deception built on dynamic control of SOA infrastructure
- Application control services for dynamic command and control of software layers and their deployment

For example, if we have some of the software objects that create a service replicated to run on multiple machines, and if the SOA system dynamically adjusts the selection of which machines to use for which functions, then adversaries could not be confident about which machine to target

within an integrated service at a given time. This technique is particularly adept at creating moving targets within cloud or virtualized architectures. Some agents could run as decoys, increasing the number of moving targets, while diffusing the risk to critical service disruption.

Vision

In addition to particular IA solutions, this effort would highlight new concepts for using application specific infrastructural information as an important contributor to IA. It would provide system designers with tools that offer the ability to develop customized adaptations of particular solutions. It would pose problems for adversaries trying to develop confident exploits. This class of solutions would complement a variety of solutions operating below the application layer. It would address specifically the risks of the actual IT applications as determined by the operational community, rather than the risks of the infrastructure supporting applications, as determined by the technical community.

Research has been done on design of mobile software agents that dynamically move to alternate machines when an intrusion is discovered. This research has been based on the premise that special operating systems and protocols to support mobilization are part of the solution. In the proposed research case, we deal with managed, regularized mobility and operation in an existing SOA environment. It would use application-specific risks as the basis for solution selection.

Method

Prior research at UVA in the field of overlay based networks has yielded an integrated software solution called Overlay-Based SOA (OBSOA). This solution could serve as the ground work for an infrastructure-hopping software layer. Experiments in this area could occur using the PlanetLab distributed computing environment available to university researchers. This laboratory environment includes more than 100 nodes to download application software and application-layer networking protocols. Distributed computing experiments can occur over the Internet.

UVA has developed a model to integrate simulated performance on some of the PlanetLab machines and actual service software objects on others. This will allow the conducting of large-scale system experiments to relate user service request rates to latencies caused by the OBSOA framework. Dr. Horowitz has combined UVA's innovative work on rapidly reconfigurable enterprise systems focusing on SOAs and distributed computing with his professional interest and experience in information assurance in generating this concept.

Dream Team

- **Cloud Platform Providers:** Citrix (Xen Server), Amazon Cloud, Microsoft Azure, VMWare, Sun Microsystems
- **Key Application Management Stakeholders:** DISA, NSA
- **Technology Research & Development:** UVA, Accenture, Los Alamos National Lab

Accenture, LLP
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Harvard Architecture Hardware and Capabilities-Based Operating System

Who We Are: Accenture LLP is a global management consulting, technology services and outsourcing company. Combining unparalleled experience, comprehensive capabilities across all industries and business functions, and extensive research on the world's most successful companies, Accenture collaborates with clients to help them become high-performance businesses and governments.

Game-Changing Dimension: Morph the Board

Concept: The hardware and software tools we use in modern information technology are unnecessarily dangerous. They are like power tools without safety shields. Unless programmers, operators, and architects remain constantly vigilant, mistakes can, and do, occur. These mistakes compromise information assurance. We can provide our IT users and professionals with safer tools, designed to assist, rather than resist, them in the development and operation of safe information processing environments.

A large number of the vulnerabilities in modern software are a direct result of the underlying data processing architecture of the computers on which they run. The current fundamental computer architecture is referred to as the "Von Neumann architecture", named for Dr John Von Neumann, a researcher in cellular automata and Turing machines. In the Von Neumann architecture, a computer's instructions (its "program") and the data on which these instructions operate are mixed together in storage (memory). The central processing unit (CPU) of a Von Neumann architecture computer cannot differentiate its instructions from its data and will "execute" data as if it were a program. It is all too easy for attackers to upload malicious programs in the guise of "data" and then trick the computer into executing this data, compromising the computer. These attacks, that turn the computer's own architecture against it, are commonly referred to as "buffer overflows" and "stack smashing exploits". It is extremely difficult for computer programmers, regardless of the language, to write programs that are not vulnerable to this type of attack. It is like building a fireplace out of wood. There should be no surprise when the house frequently burns down.

The Von Neumann architecture was developed to support Turing Machines, or cellular automata type programming. However, today we do not use Turing Machines as the underlying model for our software. Consequently, there is no need for the support of the Von Neumann processing architecture. A much "safer" alternative is commonly referred to as the "Harvard architecture". In a Harvard architecture computer, the instructions are stored in a physically separate part of the computer, isolated from the data they operate on. It is not possible to trick a Harvard architecture – based computer into executing a malicious application disguised as data. Because modern programs do not depend on executing data, most well-written programs will run on a Harvard architecture computer with little or no modification. By transparently changing the underlying hardware architecture, computers can become immune to the most common form of malicious exploitation.

A similar, nearly transparent change in the underlying architecture of our operating systems could yield similar gains in information assurance. The most common operating system in use today was

designed to facilitate ease-of-use and an enhanced consumer multimedia experience. The information assurance problems associated with it should come as no surprise. The UNIX and Linux-derived operating systems were designed with security in mind, but for an academic environment, not the modern world of multi-level, compartmented information assurance. On the other hand, several notable commercially successful secure operating systems have been developed, including Honeywell's Multics, Tymshare's KeyKOS, and IBM's System/38-AS/400-iSeries5. These operating systems, designed with real world information assurance situations in mind, use a common underlying security mechanism called "capabilities". Capabilities are strong cryptographic keys that define exactly what "permissions" a program or user can invoke. Capabilities ensure the integrity of the operating environment; strong cryptography ensures the integrity of the capabilities. While a native capabilities-based operating system offers the most security, a capabilities-based "hypervisor" can be added to existing operating systems, without modification, to significantly enhance their security and reliability. As with the Harvard architecture hardware, we need to give our developers, operators, and users safe tools that help, rather than hinder their ability to do their jobs.

Vision: We recommend that the Federal government take the lead and set the example for private industry in the use of intrinsically safe computer hardware and software, starting with the Harvard architecture powering our servers and capability-based operating systems and hypervisors protecting our information. By phasing out the acquisition of hardware and software that do not meet these architectures, the Federal government can help create the demand for these types of products. Fundamentally, humans are the source of information assurance problems. People make mistakes in designs, in configurations, and in operations. We need hardware and software that make it difficult for humans to make mistakes

Method: Too often, game morphing ideas require "big bang" types of changes. Fortunately, there is a migration path to the implementation of safer, more secure computing hardware and operating system software. Most modern microprocessors support an optional Harvard architecture-like feature called "Data Execution Prevention (DEP)". Typically, DEP is not activated because it prevents the execution of a small number of applications that have not been coded to current best practices. Remediation of these applications is generally easily accomplished, yielding a more stable and secure product that will run in the safer environment of DEP. Mandating the activation of DEP on all Federal computers, followed by the mandatory acquisition of only DEP-compatible applications will be a highly effective first step on the path to full Harvard architecture computers. Likewise, the Federal government can sponsor the development of capability-based hypervisors, to be run with the common operating systems in use today, as an effective first step toward the creation or adoption of a native capabilities-based operation system.

Dream Team: Norm Hardy, Agoric Corporation; Dr. Jonathan Shapiro, Johns Hopkins University; Dr. Frank Soltis, IBM Corporation; Intel Corporation; Advanced Micro Devices Corporation; the Secure Enterprise Networks Consortium (SENC), including Sun Microsystems.

Accenture, LLP
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Accenture Collaborative Innovation Solution (ACIS)

Who We Are

Accenture, LLP is a global management consulting, technology services and outsourcing company. Combining unparalleled experience, comprehensive capabilities across all industries and business functions, and extensive research on the world's most successful companies, we work with clients to help them become high-performance businesses and governments.

Game-Changing Dimension: Change the Rules

Concept

The concept we present here is a new way of collecting and evaluating ideas that will change the rules for addressing this nation's cyber security challenges. RFIs such as this one require individual geniuses to develop an idea independently and then present it for review to a panel of experts. This process is fundamentally limiting: it does not encourage geniuses to work collaboratively. Further, it requires substantial time to review all the separately submitted ideas.

The Accenture Collaborative Innovation Solution (ACIS) is a new application of existing technologies. It allows thinkers and creators to not only work collaboratively when answering difficult questions. It also allows them to subject their ideas to review *before* submitting to decision-makers, thus expediting review. The game-changing potential of such our process is that early peer reviews inspire new thinking—sooner.

The process for participating in ACIS is simple and straightforward. Users create an account to log into a website created for a particular campaign. Logged-in users gain access to available questions in which they may participate. Selecting question gives users access to ideas already submitted. A visual graph guides users, which can come to resemble a vine or tree. Users can read through existing ideas. They can rate them. Then they can submit their own, new and original thinking, add to existing idea, or both. Visually, the vine grows with user participation.

ACIS offers further options. Users can share their ideas with colleagues. They flag them for administrator review. They can attach documents to their ideas. They can select highlight favorite ideas—and much more.

At the end of a campaign, ACIS organizes all ideas into an executive summary that allows stakeholders to review the entire vine of ideas and associated metrics (readership, ratings, etc.).

ACIS's architecture is a Software-as-a-Service (SaaS) application built on an open-source, wiki platform with an Adobe Flex RIA front-end. ACIS capabilities facilitate the idea-creating process: leader boards, permission structure, activity feeds, email alerts, RSS feeds, and more.

Accenture, LLP
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Vision

The vision for this capability is clear and compelling:

- Engage the best thinkers nationwide from public, private, and academic sectors
- Facilitate their collaboration
- Focus their best thinking collaboratively on this nation's most difficult questions

Because our ACIS application is SaaS, users can gain access easily: go online, create an account, log in, and begin.

More than 10,000 users in more than 40 countries at more than 50 sites are using ACIS as part of their process to help Accenture clients answer their most difficult questions. Since a stated, key goal of the National Cyber Leap Year "to engage diverse sectors" and since our ACIS already exists, the NSF can begin immediately with ACIS—e.g., to facilitate workshops in Stage 2 of this RFI and across other RFI's.

Method

The methodology we used to develop ACIS rests on our relationships with many of the world's largest corporations. We recognize that the innovation process model where individual thinkers create game changing ideas is suboptimal if not flat-out unrealistic. By creating ACIS, we are helping our clients harness existing knowledge and experience within their workforces and customers to help solve their most difficult questions.

In fact, we used ACIS to generate ideas and collaborate across the Secure Enterprise Networks Consortium (SEN-C) in advance of this and accompanying responses to the NSF's National Cyber Leap Year RFI

The chief assumption underlying our concept is that participants want to work collaboratively. Our experience shows a resounding "yes!" they do—with the proviso of clear communications regarding stakeholders, timelines, and incentives. We have learned that successful incentive programs recognize not only final ideas but also participants.

Dream Team

Our dream team of participants includes leaders in IT security, computer security, economics, critical infrastructure, national security, and law enforcement, as well as private industry experts and consortiums such as our own Secure Enterprise Networks Consortium (SEN-C).

Accenture would be happy to open the site for our National Cyber Leap Year collaboration to the NSF and the Office of Science and Technology Policy (OSTP) for evaluation and demonstration.

Who We Are: We are the Center for Cyberspace Research (CCR) at the Air Force Institute of Technology (AFIT) located on Wright-Patterson Air Force Base in the Dayton, Ohio area. AFIT is a graduate school of engineering and management serving both military and civilian students. The AFIT faculty has expertise across the broad spectrum of technical and non-technical domains that comprise cyberspace. *POC: Michael R. Grimaila, PhD, CISM, CISSP, NSA IAM/IEM, Associate Professor, Air Force Institute of Technology (AFIT); Tel: (937) 255-3636 Ext. 4800;*

Game-changing dimension: Change the board and change the rules.

Concept: Virtually all modern organizations have embedded information systems and networking technologies into their core business processes as a means to increase operational efficiency, improve decision making quality, reduce delays, and/or maximize profit. Unfortunately, this dependence can place the organization's mission at risk when the loss or degradation of the confidentiality, integrity, availability, non-repudiation, or authenticity of a critical information resource or flow occurs. Organizations that are effective at mitigating cyber related mission risk typically have well documented, timely, and relevant risk assessment processes in place. An accurate, well documented risk assessment is THE foundational element necessary to conduct meaningful risk management so that resources can be protected at a level commensurate with their value. Unfortunately, many organizations fail to implement formal risk assessment processes because it is labor intensive, time consuming, and must periodically be revisited. This is further complicated by dynamic business objectives, relationships, processes, and structure. What if we changed the game board by providing organizations a means to more easily conduct, maintain, and document their risk assessment?

Vision: We believe that a combination of technology and structured organizational processes will enable organizations to more easily 1) identify and document the cyber resources they utilize in support of their business objectives, and 2) to understand the value that these resources contribute in fulfilling their organizational mission. Collectively, this will provide the capability to automate the identification and documentation of information dependencies; securely document the valuation of information asset dependencies by information consuming organizations; track and prioritize information as it travels from source to sink through infrastructure elements; and allow information providers the capability to deterministically identify all those who depend upon their information resources. Our solution will exploit automation where possible; provide a more accurate, timely, and relevant mission impact estimation following an information breach; and provide an accountability chain which can be used to enforce compliance with organizational policy and risk management objectives.

Method: We have identified several barriers which impede the ability to implement and maintain risk assessment information in dynamic organizations. We propose to investigate alternative methods, models, processes, and technologies that overcome the identified limitations in order to enable <u>any</u> organizations to implement and maintain a cyber risk assessment. Our work draws upon research activities in both the technical (e.g., system architecture, information architecture, event correlation) and non-technical (e.g., mission/process/task representation, decision making, valuation, visualization) domains to achieve the research objectives. The initial research thrusts include <u>Information Asset Identification</u> and <u>Information Asset Valuation</u>.

Information Asset Identification will examine the process by which information assets are identified and explore alternative automated methods which may be used to dynamically document their existence. The identification of an information asset consists of two discrete areas: identification and definition of a true and discrete information asset; and development of the ontology of information that supports the core missions of a given organization. Research is required to formally define the core asset on which cyber-dependent missions rely and to develop the correct ontology of the mission-supporting critical information. The working hypothesis for this thrust is that implementing new methods, processes, and technologies that take advantage of automation will significantly reduce the burden of manually updating information asset inventory documentation in dynamic environments. Information Asset Valuation will examine the process by which information assessment valuation is conducted and explore methods for capturing and refining over time information valuation assessments by decision makers. Understanding the value of information within an organization is a highly complex task since the value drivers are often intangible and difficult to quantify. We will explore how information assets support decisions and investigate methods for formally representing information valuation as a function of how, where, and when it is used in support of a mission. The working hypothesis for this thrust is that the value of information is dependent on the assessor's frame of reference, how the information is used in decision making, the pervasiveness of the information, and when the information is needed in support of command decision making to support mission objectives.

Dream team: Joint Task Force – Global Network Operations (JTF-GNO); Air Force Research Laboratory (AFRL); Air Force Information Operations Center (AFIOC); Carnegie Mellon University (CMU); Texas A&M University (TAMU); MITRE; Secure Decisions - Applied Decisions, Inc.

References:

Anderson, E., Choobineh, J., and Grimaila, M.R., "An Enterprise Level Security Requirements Specification Model," Proceedings of the 38th Annual Hawaii International Conference (HICSS 2005), Jan. 2005, pp. 186-196.

Grimaila, M.R. and Fortson, L.W., "Towards an Information Asset-Based Defensive Cyber Damage Assessment Process," Proceedings of the 2007 IEEE Computational Intelligence for Security and Defense Applications (CISDA 2007); Honolulu, HI; April 1-5, 2007, pp. 206-212.

Grimaila, M.R., Mills, R.F., and Fortson, L.W., "An Automated Information Asset Tracking Methodology to Enable Timely Cyber Incident Mission Impact Assessment," Proceedings of the 2008 International Command and Control Research and Technology Symposium (ICCRTS 2008), Bellevue, WA, 17-19 June 2008.

Grimaila, M.R., "Improving the Cyber Incident Mission Impact Assessment Process," Cyber Security and Information Intelligence Research Workshop (CSIIRW 2008), Oak Ridge National Laboratory, Oak Ridge, TN, May 12-14, 2008.

Hellesen, D., Grimaila, M.R., Fortson, L.W., and Mills, R.F., "Information Asset Value Quantification," Proceedings of the 2008 International Conference on Information Warfare and Security (ICIW 2008), Peter Kiewit Institute, University of Nebraska Omaha, 24-25 April 2008.

Zero-Security-Defect Virtual Machine Monitors

Who you are – Lee Badger, a Computer Scientist in the Computer Security Division at the National Institute of Standards and Technology. I am a computer security researcher with 20 years of experience in computer security, focusing on operating system security, security models, and software assurance. I recently completed a 6-year tour at DARPA where I funded and managed multiple programs in software assurance and self-defending systems.

Game-changing dimension – Change the board.

Concept – Virtualization is a set of techniques that allow a single physical computer to "pretend" to be multiple computers at the same time. Virtualization saves money and adds flexibility because virtual computers can easily be created and moved to adapt to computing demands: virtualization is expanding rapidly on the Internet.

Marketing literature depicts virtualization software as safe and secure, but testing shows that implementation of virtualization layers is complex, and often insecure. Pervasive deployment of such vulnerable layers could undermine the Internet. What if we can guarantee core security properties of virtualization layers, and turn virtualization into a security asset rather than a security liability?

Vision – A software component implementing virtualization is called a Virtual Machine Monitor (VMM). In a nutshell, the vision is to demonstrate a practical VMM that has no exploitable security defects. This is admittedly an extremely ambitious goal but it is also game-changing as the RFI requests: such VMM technology could add security pervasively throughout the U.S. information infrastructure.

Why it's hard: VMM software is extremely complex, and increasing complexity dramatically increases the likelihood of flaws that can be exploited. The complexity comes from the need to provide the illusions of many low-level system structures, such as virtual memory, asynchronous interrupts, and attached devices (e.g., displays). Moreover, VMM software must exhibit very high performance, which strongly biases implementation towards fast but unsafe languages such as C. For example, Xen, the most popular open source VMM, is comprised of over 910,000 source lines of code including its bundled utilities, and over 90% of the system is programmed in C. This single system represents, approximately, two and a half centuries of programmer effort.[1] Analyzing a system this complex is impossible without highly automated code and system analysis tools.

Why it's possible: Three factors combine to give confidence that VMMs can be realistically secured. **First**, software analysis tools have recently made dramatic progress in analyzing large (100,000s line) programs by: 1) discovering program invariants (e.g., dynamically using learning), 2) checking models of systems (e.g., validating temporal safety properties), 3) analyzing whole programs (because memory is available), 4) implementing system-specific analyses (e.g., for high-value components such as the Linux kernel), and 5) using symbolic execution for

[1] These numbers were obtained using David Wheeler's SLOCCount tool.

checking logical properties for all possible executions of a given path through the code (e.g., that division by zero does not happen). These techniques offer the hope of high scrutiny (and repeated scrutiny) of complex VMM software. **Second**, VMMs appear amenable to clear and formal specification of security properties (e.g., virtual machine isolation, prevention of virtual machine escape, prevention of malicious VMMs). Concise security obligations will reduce analysis burdens. **Third**, VMMs are extremely complex as discussed above, but they are still an *order of magnitude* less complex than operating systems; in combination with ongoing tool improvements and concise security requirements, they represent an opportunity to reach for assurance levels that could not be achieved in the past with operating systems.

How we'd show it: We will show results in three ways. **First**, we will show them constructively, by documenting the validation techniques that have been performed and the logic showing the absence of exploitable flaws. **Second**, we will use 3'rd party evaluation, by subjecting the VMM technology to adversarial red team evaluation. **Third**, we will subject our VMM technology to public scrutiny, by making the implementation public.

Method – By managing several DARPA software programs I became aware of the really impressive recent results from the software analysis community. A short study, unrelated to this RFI, on virtualization security at NIST then revealed the obvious connection and opportunity. The research could be conducted independently based on open source, available components, without any apparent dependencies or assumptions.

Dream team – The following organizations have the expertise needed. Names of researchers will follow if this idea is encouraged and if they agree to participate. *Disclaimer:* no-one has been asked and some may choose not to participate.

SRI	Security; formal analysis; system development.
Stanford	Automated software analysis.
Kestrel Institute	Formal verification; specifications.
University of MD	Protocol verification; software assurance.
University of Texas Austin	Formal verification; verification tools.
Cornel	Security property specification; system design.
Berkeley	Automated software analysis.
MIT	Dynamic software analysis and monitoring.
SRA	Red team testing.
anonymous	Testing.
Microsoft	Virtualization security.

BAE SYSTEMS

White Paper "Shorten the OODA Loop" Response to NITRD Cyber Leap Year RFI

Who We Are: BAE Systems, with significant history of delivering innovations in security and networking technologies for public and private sectors.

Game Changing Dimension: Morph the Gameboard. Attackers count on an open window of risk from when a vulnerability is published until system security administrators close it. This concept changes the game by accelerating the game clock (the defensive OPS tempo) based on risk that can be predicted and simulated prior to an actual attack.

Concept: Attacks *will* happen; systems must be able to *quickly* adapt and react to the risks, before an attacker can exploit the vulnerabilities. By shortening the timeframe to complete a cycle of the Observe-Orient-Decide-Act (OODA) loop, adversaries will have much less time to exploit published vulnerabilities, and success of attacks will be greatly reduced.

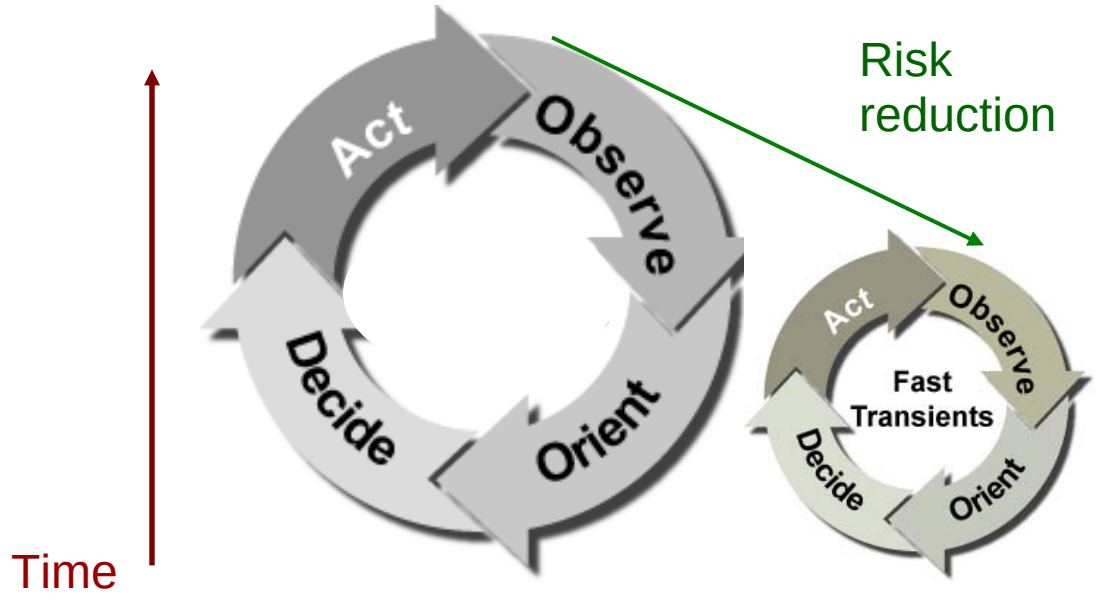

The OODA loop can be used to visualize the response and recovery from a computer systems attack. This project proposes to utilize improved processes aided by automated tools and visualizations to shorten the time required for each step in this cycle.

Observe: instead of waiting for an attack to occur, this step is shortened by Predictive Risk Analysis techniques to study and visualize the impact of an unprotected vulnerability being exploited by a Threat.

Orient: The impact of such an attack is studied for potential mitigations, and simulating the potential outcomes.

Decide: The best-fit mitigation is chosen using an automated Analysis of Alternatives decision process, based on risk mitigation and feasibility for the specific system and operation supported.

Act: Utilizing advanced implementation testing, installation, and recovery techniques, the mitigation is rolled out to all effected systems.

Vision: Existing "patch and pray" solutions are too slow and ineffective against determined threats in today's environment. Techniques such as Predictive Risk Analysis - assessing the risk inherent in systems and networks can analyze vulnerabilities for their risk potential and recommend effective safeguards to mitigate them. For this to be effective, the time required must be accelerated to study a published vulnerability for its risk impact and recommend a mitigation for the specific system under analysis.

When fully developed, the solution will have the capability to predict and simulate the risk of attack, the damage potential, and propose mitigation options within an automated environment, potentially closing the risk gap from days to minutes.

Method: Elements of this solution exist today in COTS products, open source and vendor databases, and research projects. The development of this solution, if funded, would integrate these disparate capabilities into a comprehensive solution which:

+ Utilizes existing mathematical models, such as those created by AFIWC or CMU to calculate risk
+ Automates the Analysis of Alternatives process to find the best combination of safeguards to reduce risk to acceptable levels
+ Applies these models as a solution to the "hard problem" of Cyber Security Metrics.
+ Utilizes advanced visualizations of Risk in an easy to understand GUI as a decision support aid.
+ Maximizes the use of existing COTS capabilities and standards-based databases
+ Integrates the risk model and safeguard analysis to existing design and engineering tools.

Dream Team: In addition to the BAE Systems team: an NSA Center of Excellence (university) with R&D in the area of IA Architecture, a university or small company with Information Visualization R&D capability, a COTS vendor of Risk Analysis tools, a COTS vendor of real-time Vulnerability, Safeguard and Threat reporting products, interested US Government agencies, and an Industry Enterprise Architecture Consortium.

Who we are – University and Industry researchers and engineers with a track record of significant world class innovations in security and networking technologies for over 2 decades, working together with Stakeholder Communities to create Secure LANs Across AMerica (SLAM).

Game-changing dimension of SLAM – Recent events such as the Russian DDOS attack against Georgia's cyber infrastructure, perpetrated concurrently with a kinetic attack, demonstrate that large-scale cyber forces can be marshaled easily and quickly to launch a coordinated attack that can target and reach most every part of the Internet irrespective of geography. Most experts agree that the core infrastructure of the Internet is vulnerable to such large-scale debilitating attacks. In response, there are many plans and ongoing efforts sponsored by various government and private IA R&D programs to transform the cyber infrastructure to defend against strategic damage and to make the Internet resistant to attack. *But what if we do not succeed?* What if the Internet were degraded or entirely disabled for a significant period of time? No one really knows the consequences. No one can possibly simulate such a large-scale event or create a cyber war game to measure the effects of such damage. It is imperative that we plan for catastrophe and devise feasible and effective solutions to reconstitute the cyber infrastructure as quickly as possible. We are proposing an important element of a National Cyber Recovery Plan.

Concept – We propose a feasible solution to rapidly deploy a backup communication capability sufficient to permit emergency personnel to re-constitute the Internet after significant damage from a "Cyber Katrina" event, limit its downtime and recover to a "clean" state. Most of the large Commercial Internet and military Infrastructure will likely collapse, but the small and diverse LANs Across America will survive. We propose to leverage existing capabilities of ham radio operators, SATCOM communication devices and new commercial technologies such as WiMAX to organize and constitute a CONUS-only terrestrial Internet infrastructure, a nation-wide network created by the density of common WiFi-enabled laptops, mobile devices (such as cellphones with multiple radios), and other transportation-borne devices. The essential capability and use of this ad-hoc Peer-to-Peer radio-enabled network is to provide emergency communication to provide the means for emergency personnel to "reboot" the Internet and other high-bandwidth networks if disaster strikes. In times of national emergency, we can shift our most critical cyber activities to LANs across America, made up of diverse and heterogeneous micro-networks and use that fall-back environment to bootstrap, regenerate and recover the rest of the Internet in a timely manner.

Vision – Several experiments demonstrate the feasibility of this concept even today without next-generation WiMAX technology. From a suburban NJ household, at least 4 neighboring WLANs are accessible (some secured). In an urban area in Arlington VA, 9-15 WLANs were accessible from an 8th floor apartment (most secured). 21 WLANs are accessible from an office on the campus of Columbia University, some secure and about half of them residential. Large-scale surveys have indicated that in general, only about 1/3 of WLANs are protected, though that is probably changing. One important factor in our favor: the limited range of WiFi makes it likely that in suburban areas, at least, most signals within range are operated by a neighbor one would likely know at least casually. The fundamental necessary components of a systemic WiFi enabled infrastructure exists; we need to enable SLAM via new software we shall develop and education to solicit ordinary citizen participation to download and install and have ready on board their own machines the capability to self-generate SLAM at times of need. Metropolitan-wide WLAN mesh networks have been demonstrated elsewhere; we propose to expand the concept across many metropolitan areas, and to make the system self-managing and self-configurable.

Method –The team will solicit broad participation to attack the problem of creating the SLAM network which requires solutions to a wide range of engineering, business and policy issues such as:

- Technology Challenges
  - Organic Ad-hoc MANs (bridging solutions to construct an efficient WAN capability using heterogeneous technologies: WiFi, WiMAX, HAM radios, SATCOM, mobile nodes, etc.).
  - Assured connections while using decentralized control (e.g., routing, scheduling, and power control) over large scale wireless networks.
  - Mesh networking solutions to develop cooperation between self-organizing WLANs.
  - Rapid deployment of trusted wide-area overlays.
- Engineering and Systems Challenges – maintaining long-range connectivity by using cognitive radio networks, delay-tolerant networking, Intel's Rural Connectivity Platform (RCP) and vehicular WiFi.
- Standards Challenges - we may need to standardize secured modes analogous to the Cold War era Emergency Broadcast System or CONELRAD. There are many challenges in getting priority QoS working properly across a large-scale, hetergeneous ad hoc network.
- Business Challenges – regulation and/or incentives will be required to enable private devices to participate in an emergency ad-hoc network.
- Stakeholder Communities Cultural Challenges
  - Amateur radio, WiFi (IEEE 802.11), WiMAX (IEEE 802.16), Service Providers, IETF, ground vehicles, airlines, ATC, FCC, FHA, FAA, NSF, Marine, Satellite, etc. all play a role in SLAM.
  - Policies and configuration of Internet routers to adopt SLAM. For the SLAM to provide either additional capacity to the Internet or just failover, we will need to have the equivalent of peering points where the traffic can hop back into wired infrastructure when it is reconstituted.
- International Cultural and Policy Challenges (there are many socio/political issues that are more challenging than the technical issues)
  - Canada is within close enough proximity to provide Alaskan connectivity, for example.

Dream team –The SLAM team includes a large systems integrator and academic researchers working in partnership with expertise in networking, security and software engineering. We have the capability to provide concrete technical information describing the required core algorithms, protocols and software implementations necessary to realize SLAM. We have demonstrable systems engineering experience ranging from Mobile Ad Hoc Networks, robust distributed data transmission mechanisms for self-deploying networks, security solutions for anti-DDoS "push back" and overlay techniques that can be proactively deployed atop rapidly varying network infrastructures. We will augment this expertise with the stakeholder communities including: Commercial Internet Service Providers, Network Device Vendors, Computing Device Vendors and Operating System Vendors, Standards Bodies and Policy Experts throughout the US Government. Our team is ready to provide substantial technical content to vigorously accelerate the SLAM workshop process.

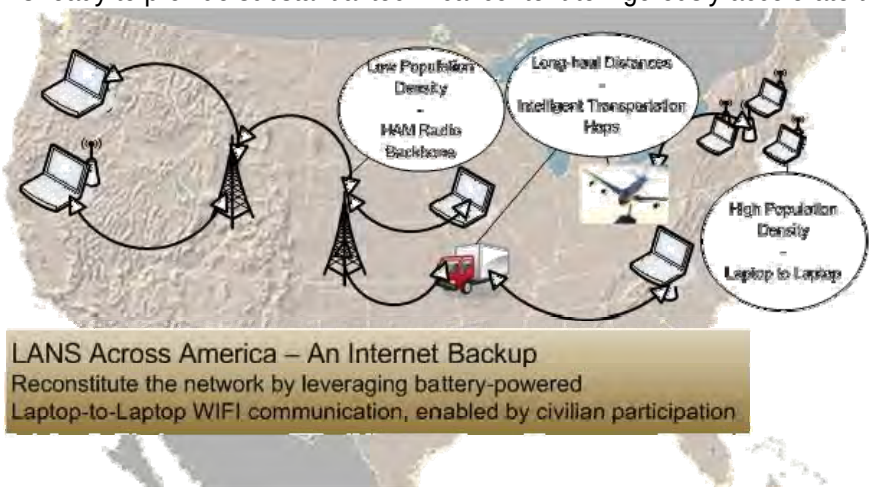

LANS Across America – An Internet Backup
Reconstitute the network by leveraging battery-powered Laptop-to-Laptop WIFI communication, enabled by civilian participation

BAE SYSTEMS

White Paper "Virtual Gameboards" Response to NITRD Cyber Leap Year RFI

Who We Are: BAE Systems, with significant history of delivering innovations in security and networking technologies for public and private sectors.

Game Changing Dimension: Morph the Gameboard.

Concept: When an adversary successfully attacks a system the "game board" changes. This would be accomplished in one of two (configurable) ways: 1) the game board (virtual system) shuts down automatically, or 2) a new "board" (virtual system) is created and the attacked system used as a Honeypot to deceive the attacker with false information. Attacks *will* happen; systems must be able to dynamically adapt and autonomically react to the attacks. This changes the game by making the real game disappear upon attack.

Vision: When people use computer applications, the applications and servers would be protected by providing security layers that ensure continued operational integrity by removing an attacked system from normal operation and continuing with a new, clean application or service. The base technology for this exists today in hardware and software techniques used for virtualization. The technology needs to expand to preserve application/system state information and re-create a new instance with the "good" state restored and "appropriate" connections transferred without risk of also restoring the compromised data and connections. Techniques will be needed to expand upon rapid, lightweight system and application virtualization techniques. Technology for communicating between virtual systems will make it possible to keep a compromised instance active but with dis-information as a Honeypot to make an attacker think they are continuing to achieve their goal(s).

Method: Two approaches or configurations can be developed utilizing lightweight virtualization techniques for client, server, and web-services systems. Both would utilize rapid, lightweight virtualization and host-based intrusion detection mechanisms. Virtualized systems and applications would frequently save all state and connection information (checkpoints) for re-instantiation. When a system is attacked, the attacked (virtual) system recognizes the attack and creates a new instance with the last known good state and connections. Depending upon configuration (or implementation), the attacked system then either exits, or continues to run with the attackers' connection(s) still active and all other data (state, database, file system, etc) provided through a connection to a "disinformation application" – thereby acting as a Honeypot to keep an attacker actively engaged.

Dream Team: In addition to the BAE Systems team: an NSA Center of Excellence (university) with R&D in the area of secure virtualized computing; virtualization vendors or open source providers such as VMware or XenSource; and interested Government agencies such as the NSA.

Global Immune Response to Network Intrusion

RFI Focus Area
Morph the Game Board

Submitter
Daniel Wyschogrod
BAE Systems, Advanced Information Technologies
6 New England Executive Park
Burlington, MA 01803

Summary of Who You Are
During his four years at BAE Systems, Daniel Wyschogrod has been the Principal Investigator on two projects involving Automatic Signature Generation (ASG), a technique to leverage anomaly detection results from network monitoring into arbitrarily low false alarm rate string signatures against repeated attacks. One project was sponsored by the Department of Homeland Security while the other was funded from BAE internal research funds. Mr. Wyschogrod also worked on intrusion detection for both the DQW and DCA projects sponsored by DARPA. Prior to coming to BAE, Mr. Wyschogrod was the Director of Content Inspection Software for SafeNet Inc. He was also on the staff of MIT Lincoln Laboratory for ten years where he worked on both network intrusion detection and image processing.

Concept
Currently network security is a reactive game. When a new worm, virus or botnet capture exploit is first released, it can take hours or days for analysts to identify the new attack and develop a "vaccine". By this time, many thousands of hosts may already be infected with the new virus, and modern viruses are increasingly intelligent about interfering with a user's ability to identify or heal an infected computer, further complicating the problem.

What if the "vaccine" to the new virus could be generated in just a few minutes and propagated across the internet faster than the worm? With the system constantly inoculating itself to new threats, attackers will be left to react to a playing field that now favors the defenders.

The effect is much like the human body's response to a new virus. An antigen detected anywhere in the body will trigger the amplifying response of creating antibodies throughout the entire body. Similarly, a new worm detected anywhere on the network will trigger the creation of signatures that spread throughout the entire network. A single detection leads to a cascade resulting in complete inoculation.

Vision
The vision is a network of distributed sensors and centralized actuators that will automatically detect and inoculate against threats. Network-based and host-based

network sensors will be distributed all across the internet – at the core routers and switches, at enterprise gateways, on individual hosts, etc. When a new attack is detected by any sensor, new signatures are automatically generated. These signatures are then distributed to in-line signature matchers placed at critical points in the internet – e.g. core routers at the ISP's – that will block the attack's spread far from vulnerable hosts.

There are already a number of sensor technologies – anomaly-based sensors, protocol-adherence-based sensors, honeypots – that can be used to detect the initial attack, and ISP's already have hardware in place that can inspect packets on the wire at line rates.

The missing link is the automatic signature extraction step. BAE Systems has developed such a technology that:

a) operates on a single instance of an attack,
b) has an arbitrarily low false alarm rate,
c) produces simple string-based signatures compatible with common pattern-matching hardware,
d) produces signatures that can detect re-occurrences of polymorphic attacks from the single original instance

The major gap is the integration with the existing infrastructure. Technically, there is a lot of integration work to be done to allow all the different component systems interoperate securely. Politically, the major hurdle is gaining acceptance of such a large infrastructure change. Both the Government and the major Internet Service Providers would need to be a part of the entire process. We believe that the technology is ready and that the political climate is open to such an infrastructure change.

Method

We examined the effect of recent worms and large scale attacks, such as the Conficker worm, and brainstormed what defensive changes would need to happen to limit the impact of – or even prevent – future attacks like Conficker. One of the ideas that emerged was to extrapolate an existing BAE Systems technology, our automatic signature generation capability, into an internet-wide defensive net. We then discussed at length what such a system might look like on a large scale and what the major obstacles and milestones would be for such a project. After several rounds of refinement, we developed the idea presented here.

Dream Team

Core infrastructure (Cisco, Juniper)
Internet Service Providers (Verizon, AT&T, Comcast)
Security product vendors (Symantec, McAfee)
Line-rate pattern matching hardware vendors (cPacket)

Credentialed In-Stack Protection

RFI Focus Area
Change the Rules

Submitter
Jeffrey Opper
BAE Systems, Advanced Information Technologies
6 New England Executive Park
Burlington, MA 01803

Summary of Who You Are
Mr. Opper is the Network Security Section Leader of BAE Systems Advanced Information Technologies. Mr. Opper is the Chief Architect for the Principles for Intrinsically-Assured Network Operation (PIANO) system being built as part of the DARPA IAMANET program. His prior work at BAE Systems included the development of algorithms for deep packet inspection and malware signature generation and the development algorithms for globally scalable cryptographic key distribution systems. Prior to BAE Systems, he was a Principal Engineer at the MITRE Corporation.

Concept
As computers and networking further pervade our critical infrastructure, it becomes increasingly vital to ensure that malicious code is not executed on critical networks. Although there are current methods for verifying that an executable is from a trusted source before running it, there is no means for detecting that a process has not been subverted after execution begins.

What if processes were able to continually check themselves – and be checked by other parties – to ensure that they had not been subverted? With processes intrinsically secured against subversion, attackers will be unable to gain control of critical processes and systems.

Vision
We envision an integrated capability in all computing systems to enforce periodic authentication of process credentials. These credentials, which are inextricably bound both to the executable image and registered with trusted hardware on the computing platform. Particularly for critical infrastructure networks, we envision a self-monitoring mesh of nodes that both observe themselves but also all other nodes in the network. In this way, an attacker would need to dynamically subvert the entire network simultaneously in order to remain undetected. Additionally, the network would be able to rapidly heal itself in the case of an attack.

Method

Techniques to secure a running system via executable and compiler modifications already exist with products such as StackGuard and ProPolice. Additionally, credentialing as a means to determine trust of an executable or node is common practice in network security. We have built upon these ideas by requiring each process to augment the stack with credentials that are verified by the kernel using trusted hardware. Subversion of the process via techniques such as NOP sleds and return-to-libc attacks obliterate the stack credentials and cause subsequent integrity checks to fail.

The most difficult challenge is in integrating the large number of players that need to contribute for the system to work. Securing a system in this manner begins with creating secure executables, which requires that all software, from the operating system to the web server and all additional software, be compiled and distributed in secure form. This would require the cooperation of a large number of software developers and systems integrators.

Dream Team
Hardware and operating system vendors (Microsoft, Apple, Dell, HP, Linux)
Processor manufacturers (Intel, AMD)
Compiler developers (GCC team)
National Security Agency
University partners (cryptographic algorithm development and validation)

From: Brad Barkett

Sent: Monday, March 09, 2009 3:21 PM
To: Leapyear
Subject: LeapYearRFI-3
Regarding: http://www.nitrd.gov/leapyear/

Here are some basic ideas off the top of my head:
1. Encourage the use of two factor authentication (such as passwords plus RSA security tokens) and make single factor passwords obsolete on both internal federal networks and websites.

2. Make posting username and password information to http/80 (rather than https/443) on any federal systems which are responsible for password management illegal.

3. Make unsecured WEP networks illegal.

4. Migrate all government client browsing to Firefox with NoScript cross domain script blocking plugin, to stop drive-by infections.

5. Push for the State Department to define regulations for what constitutes "cooperative" security responsibility on the part of sovereign states and impose sanctions upon countries that do not adequately prosecute individuals and maintain accountability for hosting and perpetration of computer crimes against American citizens and organizations.

6. Consider implementing a national perimeter firewall system like the one in use in China, without the media censorship, for the express purpose of protecting America's networks from malware and known blacklist websites.

7. Defeat the looming threat of Cross Site Request Forgery by moving all critical goverment webservers which involve logins and persistent state management to software which appends cryptographic nonces to URLs.

8. Define a federal standard for secure VPN and telecommuting practice requiring the certification of any home machine which will be connecting remotely to federal systems.

9. Consider deploying fake federal honeypot systems to track and profile potentially state sponsored computer intrusions from rival nations.

10. Require schools to teach basic computer security concepts as a general education requisite in high schools and colleges!

Thanks,
Brad Barkett

--
--
Bradley A. Barkett

BDNA Corporation
339 North Bernardo Avenue
Mountain View, CA 94043
T 650.625.9530
F 650.625.9533
www.bdnacorp.com

NSF – RFI National Cyber Leap Year

Who we are -- BDNA Corporation, A US Corporation based in Silicon Valley. Clark Campbell, Director of Public Sector

Game-changing dimension – Board – Lift the "Fog of War" in any environment

Concept –BDNA has developed a game-changing capability to discover all IP-enabled devices on a network without requiring any knowledge of the environment. Lifting this "Fog of War" in an environment and overcoming the data deficit in these areas is revolutionary. This is done from a centralized server without any operational impact to the environment (network or assets) and without any distribution of software or agents on any devices. Detailed information about all types of hardware and software is discovered and presented in an easy to read format for reporting and external use.

Organizations using this BDNA technology can support their initiatives for security verification, configuration verification, and for improving their information assurance posture. Armed with this visibility, organizations can also support other initiatives around compliance, finance, procurement, IT operations, and IT strategy.

Vision – If an organization knew exactly what IT assets were in an environment, there would be few unknowns about the organization's cyber footprint. With this transparency, you could easily determine the current state of exposure and risk to any number of security vulnerabilities.

If you could get a detailed inventory of all your hardware and software in a matter of hours from an implementation that took less than a day, what would you do with it? If you could easily ask questions of a repository that had all your hardware and software available for review, what would they be? Would you ask how many wireless access points you had in your environment? Would you ask how many unsupported operating systems you had on your network? Would you ask how many machines had no virus protection software? Would you ask which machines had their encryption software disabled? Would you ask which systems had utilities with known vulnerabilities installed? Would you ask questions you were taught to stop asking a long time ago because you received no credible answer back?

This concept is currently available from BDNA and is continually enhanced. BDNA provides a discovery catalog that is constantly updated with new information about the market and assets. The reporting capabilities are robust and often improved. For instance, we recently added an entire series of reports revolving around virtualization technology (which we discover with the same detail as physical assets).

Method – BDNA has developed a core discovery engine that uses an external catalog of fingerprints to forensically determine what a device is and how to gather all the attributes of the various devices. This allows the engine to discover any type of asset, from servers and workstations to IP phones and networking equipment using a centralized scanning server

BDNA Corporation
339 North Bernardo Avenue
Mountain View, CA 94043
T 650.625.9530
F 650.625.9533
www.bdnacorp.com

(multiple servers are possible for scalability) without the distribution of any software or agents on the discovered devices. The fingerprints also allow BDNA to gather detailed information about all software loaded and processes running on any asset. The catalog is constantly updated with changes to the market and can be extended for custom assets and software. BDNA maintains the catalog with information such as the current vendor, product, version, warranty, support, and energy information about IT assets.

The engine is designed to be non-invasive to the network and the discovered devices. It uses a targeted, heuristic set of inquiries to focus discovery on individual assets. This also allows the engine to operate very quickly to inventory large networks of assets in a matter of hours. Because of its lightweight nature, it can be run during core working hours as often as desired.

Reporting and analytic capabilities provide over a hundred base reports to immediately derive answers from the discovered data. The interface is easy to use and allows for custom reporting. Analytics also allow comparison over time. It is extremely easy to compare the state of any environment to any prior situation. BDNA has made the information available to external consumers through programmatic interfaces.

The engine operates at three levels discovery that provide insight into the devices. The first level of discovery simply takes an IP range, from a single IP address up to any range of addresses, and discovers all asset and their attributes in that range. Some assets will only provide detailed information to authorized users. BDNA has developed a second level of discover that has the ability to pass read-only, non-administrative (not root user or domain administrator) users to the device to gather detailed attributes. A third level of discovery provides the ability to gather detailed information from server applications such as databases and ERP applications. Provided a read-only credential to the server application, we can inventory key aspects of these applications.

Dream Team – What is our dream team?

While the BDNA technology is real today, there are many tangential areas where the capability can be extended. Ranging from monitoring enterprise standards and compliance to increased industry content to broader data collection, BDNA has many immediate and future benefits. To determine the best course of development, BDNA would recommend bringing together the BDNA founders and engineering staff with key individuals in organizations that manage the strategic, tactical and financial aspects of information technology. Specifically, BDNA would like to meet with CIOs, CFOs and CTOs who are responsible for aligning the IT resources with the organizational mission. In this forum, BDNA could determine the most pressing issues to those constituencies and accelerate development of the solutions with the fastest ROI.

For more information on this game-changing technology, please visit www.bdna.com

RFI Name: RFI-2 - National Cyber Leap Year

Title of Concept: Federal Desktop Core Configuration White List

RFI Focus Area: Morph the Game board / Change the rules

Contact Information: Brian Seaberg / DoD – Federal Liaison
Bit9, Inc.
266 2nd Avenue
Waltham, MA
(813) 373-1214

Who We Are: Bit9, Inc. is a funded technology start-up organization that is focused on changing the paradigm for windows desktop security and control. Born out of a prestigious NIST grant, Bit9 is poised to redefine the historical methodology for preventing malicious software. Bit9 has built a hardened agent that deploys to the desktop endpoint preventing unauthorized software from executing. Bit9 has received McAfee ePO certification which is the underlying platform for the DoD's HBSS program. Bit9 also supports the largest knowledgebase of windows program executables (the Bit9 Global Software Registry) which serves to provide trust ratings on executable code for our award winning agent technology. This knowledge combined with our agent control capability allows entities to establish and maintain a secure and trusted windows software environment.

Title of Concept: Federal Desktop Core Configuration White List

Concept: Today's cyber security standards are based on a game of "cat and mouse". Malicious software developers build new and more adaptive techniques ("new mice") to penetrate the environments while anti-virus vendors continually deploy signature files and behavioral tools to try and catch the new mouse. With laser focused and sophisticated adversaries, a tipping point has been reached where the ability of the adversary has surpassed the resources of the defender. In essence, there are more mice than the cat can catch, and some of the mice are larger than the cat.

The Federal Desktop Core Configuration White List will morph the game board and change the rules for our adversaries. The current approach for blocking malicious software requires computing environments to analyze and determine the threat from millions of unknown files, and instantaneous determine those that are safe. In today's computing environments, there are a finite number of approved applications and an infinite number of unapproved and potentially harmful applications. The Federal Desktop Core Configuration White List approach will focus on identifying and permitting the "good" software vs. catching the "bad" software. Imagine a

secured government facility. Would the security team be provided a list of the six billion (6,000,000,000) people who are not allowed to enter? Or would the security team be given a list of the five hundred (500) people that are allowed to enter? Further, would everyone be allowed to enter until they did something bad, or would each person not on the list be reviewed and approved prior to entry? The Federal Desktop Core Configuration White List would allow cyber security teams to treat the desktop operating environment like a secured government facility.

Vision: The Federal Desktop Core Configuration White List would be a cross organizational deployment solution that provides an agent on the desktop to prohibit unapproved software from threatening the computing environment. The agent would reference a library of approved software that included generally approved and agency specific approved software. With the Federal Desktop Core Configuration White List established, any attempt to load unapproved software or launch malicious code would be stopped and a notification sent to the designated security team that supports the environment.

This vision is in immediate reach. Bit9 has developed the underlying agent technology and approval methodology. Additionally, Bit9 has built a sophisticated repository that provides identification and knowledge to partition between trusted and malicious software. By gathering constituents from the key information assurance organizations, a comprehensive Federal and DoD specific White List could be determined allowing organizations to efficiently implement a broadly defined but totally secure computing environment.

Under the new paradigm, the control point would change from reactive (updating AV signature files to address the changing threats) to proactive. A trusted and secure computing baseline established by the Federal Desktop Core Configuration White List would eliminate the introduction and propagation of viruses, worms, Trojans and other malicious software. Adversarial approaches would be muted because any malicious software introductions would not launch within the environment.

Method: The concept was originally driven by a NIST grant. It has been further refined through collaboration with companies focused on the HBSS standards and point specific implementations within the DoD and other Federal agencies. The current spike in virus and malicious software activity has magnified the need to change the paradigm from filtering to blocking and approving.

Dream Team: Cyber security and information assurance representatives from NIST, ISAP / SCAP, OSD, DHS, NSA, ESSG and Microsoft.

National Cyber Leap Year RFI Submission

Who you are: Eric Fleischman, member of the Boeing Phantom Works' Digital Communications and Network Technology organization, an applied research group within The Boeing Company.

Game-changing dimension: There are multiple: *Change the rules* by architecting the Internet Protocol (IP) to enable strong authentication of network packets. *Raise the stakes* by including inherent denial-of-service resistance in the core protocol mechanisms, requiring attackers to devote more resources to attack. *Morph the game board* by allowing hosts to selectively reveal their location and identity and to securely use middleboxes as proxies.

Concept: Extend the IP protocol stack to disambiguate the current semantic overload of IP addresses being both (host) identifiers and (network topology) locators. Use cryptography to secure the binding between the identifier and possibly time-varying locator identities.

Vision: Imagine the security and privacy implications if telephony behaved like the Internet by permitting rogue individuals to impersonate the caller-id phone number of someone else or redirect other people's calls to a location where an eavesdropper could lurk or all calls would be dropped. A single design flaw in the IP protocol provides this attack surface. The Internet architecture relies on IP addresses to identify hosts but when a networked entity receives a packet, it currently can't reliably prove that the sender of that packet legitimately is the source IP address identified in the packet header. Consequently, IP networks suffer from a variety of impersonation attacks at the network layer and above, including network penetration, denial-of-service, phishing, spam, and routing reset attacks.

In our vision, hosts and middleboxes are capable of strongly authenticating the originator of packets. It is computationally infeasible to impersonate another host's identity. A host could prove that it is talking to another host that it knows by name, and discard or deprioritize packets from unknown entities. Network administrators could prove that packets on any given network segment are authorized to be present, and block access to unauthorized hosts. The architecture permits hosts to delegate authority to other entities to act on their behalf. This would, for instance, allow a host that wants to hide its current location to use network proxies to forward its traffic. Denial-of-service attacks could also be deflected to network proxies.

Method: The above vision is currently realizable as an architectural extension to the Internet. The Internet Engineering Task Force (IETF; see http://www.ietf.org/) has recognized since the ROAD Group's work (early 1990s) that the semantic overloading of IP addresses was a fundamental weakness of the IP protocol family. The overloading undermines routing aggregation and scaling impacting routing performance and mobility. It causes the "IP Identity Problem" that undermines Internet security by enabling IP identities to be spoofed and thereby undermine existing authentication and authorization systems. It impacts application session coherence enabling session hijacking.

The Host Identity Protocol (HIP; see http://www.ietf.org/html.charters/hip-charter.html) has been proposed as a strategic mechanism to architect a separation between the end-point identifier (i.e.,

host identifier) from locator (i.e., where the node is currently located within the network topology) identities. HIP uses public keys to serve as endpoint identifiers at the IP security protocol layer (IPsec; see RFC 4301) and above, and uses best-current-practice techniques for avoiding denial-of-service vulnerabilities. The IETF HIP working group has created an experimental protocol variant (RFC 5201) and is currently talking about creating a standards track version. The working group is seeking to create an incrementally deployable extension for current IP networks. Our organization has been one of the leaders in the IETF's HIP work, creating one of the HIP reference implementations. More importantly:

1) We co-founded and co-lead the creation of the Secure Mobile Architecture (SMA; see http://www.opengroup.org/products/publications/catalog/e041.htm) within The Open Group. SMA leverages HIP as its foundation for mobile security. We deployed SMA within Boeing's own factories in order to secure our Supervisory Control and Data Acquisition (SCADA) machine control systems. SCADA security limitations represent a significant United States' cyber vulnerability impacting many core industries including the electric (T&D, fossil, hydro, nuclear), oil/gas (e.g., processing, pipelines), water (e.g., dams), chemical, railroads, as well as numerous manufacturing industries, including aircraft. Our SMA implementation is the most viable technique (that we are aware of) to secure IP-networked SCADA-devices today.
2) For many years our organization has been doing applied research for the Office of Naval Research on mobile ad hoc networking (MANET) technologies (e.g., for first responder (police, ambulance, fire) and tactical military networks). This work has included demonstrations using HIP as a basis for coherently maintaining sessions within highly mobile and rapidly changing MANET networking environments. Our work demonstrated that HIP is an effective technique for minimizing the impact of mobility and multihoming in military contexts.

Deploying HIP represents an opportunity to correct a fundamental weakness of the IP protocol and, by so doing, to solve multiple problems with a single solution. A goal is for HIP to become widely available as a protocol stack within commercial operating systems. Until then, it can be deployed incrementally without disrupting a host's existing protocol stack or requiring applications or routers to change. Despite this, deploying HIP or other architectural extensions is challenging because of multiple stakeholders having competing economic interests and priorities. This challenge is not limited to HIP; very few new network-level services have actually deployed since the 1980s. We believe that the major hurdle to HIP (or any other enhancement) is in creating and motivating a deployment strategy. We therefore recommend that the National Cyber Leap Year focus on the problem of how to deploy secure architectural extensions to the Internet, and we offer HIP as an excellent starting point. We suggest the approach include a sponsored Open Source version of HIP and, if The Open Group permits, SMA as well. We also recommend the construction of a large-scale test-bed (e.g., USC maintains an emulab-based test-bed called DETER as a test-bed for network attacks and defenses) including periodic events to verify interoperability between implementations. We suggest that it become integrated with other Federal cyber activities, perhaps including the DARPA Military Network Protocol BAA. We urge the DHS to actively encourage viable SCADA security (e.g., SMA) deployments.

Dream Team: Department of Homeland Security (DHS), IETF HIP Working Group, The Open Group, USC, Boeing.

Title of Concept: Deploying secure identity-based Internet protocols

RFI Focus Area: Attribution

Change the rules:	Use cryptographic identities to allow **attribution** of every data packet or flow, and end-to-end standards-based **policy enforcement**.
Morph the gameboard:	Allow hosts to **selectively reveal** their location and identity, and to securely use middleboxes as proxies.
Raise the stakes:	Provide denial-of-service resistance in the core network protocol mechanisms, requiring attackers to devote more resources to attack

Submitter's Contact Information:

Thomas Henderson, Eric Fleischman, Steven Russert, and Steven C. Venema
Networked Systems Technology (NST) organization within The Boeing Company

Summary of who we are: Our NST organization is an applied R&D organization within Boeing's Engineering & Information Technology group that focuses upon network technologies and security for complex large scale systems. Boeing is a Fortune 100 aerospace company.

Concept: Extend the IP protocol stack to disambiguate the current semantic overload of IP addresses used as both identifiers and locators. Use cryptography to secure the binding between identifiers and possibly time-varying locators (network addresses).

Vision: Imagine the security and privacy implications if telephony behaved like the Internet by permitting rogue individuals to impersonate the caller-id phone number of someone else or redirect other people's calls to a location where an eavesdropper could lurk or all calls would be dropped. A single design flaw in the IP protocol provides this attack surface. The Internet architecture relies on IP addresses to identify hosts but when a network protocol stack receives a packet, it currently cannot reliably prove that the sender of that packet is located at the source IP address identified in the packet header. Consequently, IP networks are vulnerable to a variety of impersonation attacks at the network layer and above, including network penetration, denial-of-service, phishing, spam, and routing reset attacks.

In our vision, hosts and middleboxes are capable of strongly authenticating the originator of packets. It is computationally infeasible to impersonate another host's identity. A host could prove that it is talking to another host that it (securely) knows by name, and discard or deprioritize packets from unknown entities. Network administrators could prove that packets on any given network segment are authorized to be present, and block access to unauthorized hosts. The architecture permits hosts to delegate authority to other entities to act on their behalf. This would, for instance, allow a host that wants to hide its current location to use network proxies to forward its traffic. Denial-of-service attacks could also be deflected to network proxies.

Method: Many in the Internet Engineering Task Force (IETF; see http://www.ietf.org/) have recognized since the early 1990s that the semantic overloading of IP addresses was a fundamental weakness of the IP protocol family. The overloading undermines routing aggregation and scaling which impacts Internet routing performance and mobility. It also causes

the "IP Identity Problem" that undermines Internet security (particularly in mobile environments) by enabling IP identities to be spoofed and thereby undermining existing authentication and authorization systems. It also impacts application session coherence enabling session hijacking. A theoretical solution of separating the identifier function (i.e., IP address signifying an end point identity) from the locator function (i.e., IP address identifying where the node was located within the routing (network) topology) is popular with many researchers. Unfortunately, the IETF has not modified the IP architecture to systematically implement this insight but rather has made local *ad hoc* changes to a very few specific protocols (e.g., Mobile IPv6).

It would be simplistic and unfair to claim that the IETF acted carelessly in this regard; the Internet is a tremendous commercial success, and the security needs must be balanced with operational and deployment realities. The present-day vulnerabilities in the Internet are not necessarily due to lack of available mechanisms, but instead we must learn from deployment history and understand why some security solutions succeed (e.g. SSL/TLS) and others fail to deploy completely or in a timely fashion (e.g. end-to-end IPsec, DNSSEC, router authentication, browser patching, patches to DNS to prevent the Kaminsky attack, etc.).

In this environment, the Host Identity Protocol (HIP; see http://www.ietf.org/html.charters/hip-charter.html) has been proposed as a strategic mechanism to architect a separation of end-point identifier from locator functionalities, with a goal of devising an incrementally deployable solution based on the current TCP/IP architecture. HIP uses public keys to serve as endpoint identifiers at the IPsec layer and above. Deploying HIP represents an opportunity to correct a fundamental weakness of the IP protocol family and, by so doing, to solve multiple problems with a single solution. HIP provides applications with a unique nodal cryptographic identity that is closely coupled with IP security (IPsec; see RFC 4301). It provides a secure identity that existing authentication and authorization systems can leverage. It provides session protection services to thwart session hijacking and other risks. It also provides a needed framework for improving IP security, IP routing, and IP mobility. It can be deployed incrementally without disrupting a host's existing internet stack or requiring applications or Internet routers to change. And it can be implemented in proxies for legacy or embedded hosts that cannot upgrade.

HIP specifications are poised to move to the standards track in the IETF. However, deploying any change to the Internet's core protocols is frustratingly slow due to incremental deployment and incentive issues (see, for instance, IPv6 transition), more work is needed on deploying HIP or related security frameworks. In particular, use cases must be clarified, deployment paths must be carefully calculated, operational concerns (not always present in the research environment) must be addressed, and software must be exceptionally robust for HIP to transition. We suggest that the National Cyber Leap Year organize future workshops around the problem of how to jumpstart large scale architectural experimentation or deployments such as embodied by HIP.

Dream Team: Department of Homeland Security, IETF HIP Working Group, Boeing. Our organization (within The Boeing Company) has been one of the leaders in the IETF's HIP work, including creating one of the HIP reference implementations. Our HIP work for the Office of Naval Research demonstrated that HIP is an effective mobility and multihoming solution.

Labeling of Proprietary Information: To the best of our knowledge, HIP is not encumbered by intellectual property claims, and we are not submitting proprietary information.

Title of Concept: Securing the national SCADA infrastructure

RFI Focus Area:

Morph the gameboard: Inherently **secure data** communications between all SCADA devices

Change the rules: Using crypto-identities allows **attribution** of every data packet and end-to-end standards-based **policy enforcement**.

Submitter's Contact Information:

Eric Fleischman, Steve Russert, and Steven C. Venema
Networked Systems Technology (NST) organization within The Boeing Company

Summary of who we are:

Our NST organization is an applied R&D organization within Boeing's Engineering & Information Technology group that focuses upon network technologies and security for complex large scale systems. Boeing is a Fortune 100 aerospace company.

Background:

Industries such as telecommunications, oil and gas refining, electrical production and distribution, transportation, and manufacturing use Supervisory Control and Data Acquisition (SCADA) systems to monitor and control processes through the collection and analysis of real-time data. SCADA systems have historically relied upon physical isolation for their security. Attackers required physical access in order to compromise or disrupt system operations. However, SCADA systems continue to become increasingly interconnected by shared or otherwise insecure communications links (e.g., public Internet and wireless communication links). Physically isolated enclaves have become a thing of the past creating the very real risk of compromise or disruption by remote adversaries.

Concept:

Creating a "virtual enclave" capability that is compatible with existing SCADA infrastructure would securely restore the isolation of past systems while still allowing communications between physically separate clusters of existing SCADA components over modern shared network infrastructures. We call this concept, "SCADAnet".

SCADAnet uses cryptographic identities and encryption on every packet to intrinsically secure data flowing between SCADA systems to automatically provide attribution of both source and destination. This creates a network ecology that utilizes standards-based security policy definition and enforcement leveraging cryptographic identities. SCADAnet provides virtual enclaves by creating secure, isolated OSI Layer 2/3 overlay networks. Current SCADA systems continue to operate "as is" while a secured infrastructure transparently shields them from integrity loss and electronic attacks.

Vision:

Securing our vast existing SCADA infrastructure, most of which is privately owned, is a daunting and impractical task unless it is accomplished by an inexpensive solution that does not require any change to existing SCADA operations or equipment. SCADAnet leverages

commercial off-the-shelf technologies and standard protocols. By deploying SCADAnet to transparently secure existing systems, the United States can leap ahead in the cyber game by significantly reducing our national infrastructure's vulnerability to remote network attack. This is a currently available technology that cleanly dovetails with possible longer-term modifications to SCADA infrastructure.

The SCADAnet capability is scalable and leverages policy-based configuration and implementation in order to minimize deployment and operational costs while maximizing the security of the entire SCADA system.

Method:

Boeing has been actively engaged during the past several years in making this vision become a standards based, commercially viable reality in order to satisfy our own internal SCADA security requirements. This effort has had three primary components:

1. **Pilot implementation:** We have implemented a prototype SCADAnet capability within Boeing which has been published as Open Source software. Our initial capability has been deployed within critical elements of our 777 aircraft manufacturing line for the past 14 months. This implementation is based on emerging protocols from the IETF and the Trusted Computing Group public standards organizations.
2. **Active participation in related public standards organizations:** In order to accelerate the availability of multiple interoperable SCADAnet products, we are actively engaged in related efforts in the "Host Identity Protocol" or "HIP" working group of IETF, the Security Forum of The Open Group (TOG), Trusted Network Connect (TNC) working group of the Trusted Computing Group, and the International Automation Society (ISA) ISA100 Wireless Security group for control systems.
3. **Commercialization of SCADAnet capability:** We are working with an existing supplier of SCADA security appliances to have them augment their product with this SCADAnet capability. They are using our released open source code as an implementation baseline. We hope to see other suppliers offer compatible SCADAnet products in the future.

The Boeing prototype SCADAnet implementation consists of small, relatively inexpensive embedded computers, called "SCADAnet end boxes" that are placed in front of each cluster of SCADA devices in order to mediate access to standard IT network infrastructure and create overlay networks for the SCADA devices. These "end boxes" utilize standard network services such as DHCP, DNS, and a new intra-device secure coordination capability out of TNC called "Metadata Access Points" (MAP) to access the Enterprise network using either Ethernet or IEEE 802.11 wireless interfaces and include self-configuration capabilities. A standard SIM chip (similar to what is found in GSM phones today) that is embedded in each end box serves as the secure cryptographic identity store used to authenticate each SCADAnet endpoint and to enforce connectivity policy for the SCADAnet virtual enclaves

Dream Team:

Boeing to expand and maintain over time the SCADAnet open source codebase; DHS to encourage and support the deployment of SCADAnet; IETF, TCG, TOG and ISA members to help create standards that provide SCADAnet interoperability; NIST and/or NSA to verify the security algorithms used in SCADAnet; and Byres Security, Inc. and other SCADA security appliance suppliers to natively implement SCADAnet within their product lines..

Information in Response to National Cyber Leap Year RFI No. 2

Aaron Burstein
UC Berkeley School of Information
20 February 2009

Title: An International Auditing Framework to Counter Cyberespionage

Game-changing dimension: Change the rules: nations would agree to open themselves to certain verification measures to control the increasingly serious problem of cyber-based espionage.

Contact information: Please direct correspondence to Aaron Burstein

Who I am: I am a research fellow at the University of California, Berkeley School of Information, with funding from two NSF Science and Technology Centers and the Institute for Information Infrastructure Protection (I3P).[1] I am a lawyer (J.D., UC Berkeley, 2004) and have written extensively about cybersecurity, privacy, and intellectual property law and policy. Prior to beginning my research fellowship, I was an attorney in the U.S. Department of Justice.

Concept: Cyberespionage—the use of computers and networks by nations to appropriate secret or economically sensitive information from other nations—is a critical threat to U.S. cybersecurity. Though limiting the damage from cyberespionage presents tough technical challenges, there are also daunting legal and political challenges. Among them are the lack of internationally shared norms concerning the proper scope of state-sponsored information collection over the Internet, and a paucity of mechanisms and forums for addressing shaping and enforcing these norms. Solving the technical problems would not address these other challenges; indeed, it might exacerbate them.

Accordingly, I offer an idea to help to define and propagate international norms to dis-courage cyberespionage. The idea is to draft an international agreement that would allow any nation that signs to review the others signatories' funding for scientific and technolog-ical research and development. In effect, nations that agree to this framework would be able to audit the others' books on scientific research and development, allowing a better identification of who is working on what, and how much they are spending on it. Ideally, participating nations would provide the names of organizations that receive funding, the amount of funding, and contracts between the government and funding recipients.

The motivation behind this idea is not to improve technical capacities to associate cyber attacks with specific network hosts, but rather to attack two of the mechanisms that facil-itate cyberespionage: state sponsorship for private sector groups that engage in their own espionage activities, and the passing of state-collected information to private-sector firms to use in the actual production of goods and services. The United States, of course, has laws that prohibit computer abuse, trade secret theft, and disclosures of classified information; and U.S. intelligence officials have publicly disavowed the practice of passing intelligence to the private sector to use to its advantage. But these rules are obviously limited in reach.

[1] http://www.ischool.berkeley.edu/.

It is effectively impossible to prosecute foreign nationals who are found abroad, and the restraint shown by U.S. intelligence agencies does little to affect foreign agencies' practices. Current treaties do not address these problems. They merely set standards for signatories' national laws and terms for cooperation in civil and criminal law enforcement, leaving the development of cyberespionage doctrines and capacities beyond all examination. An additional difficulty is that much of the information targeted in the United States is stored in an incredibly diffuse, privately owned environment that is subject to no coherent monitoring or control. It is difficult to conceive of changes in coercive U.S. or international law that would ameliorate any of these problems.

Thus, I believe that a more cooperative and international approach is in order. This auditing approach would "change the game" in two ways. First, it would remove some of the cover of deniability that surrounds most cyberespionage. Audits would allow governments to present evidence that they are not funding groups to take an economic interest in cyberespionage, which could be helpful in assessing currently observed attacks. Second, by creating a forum for routine, non-public interactions between governments, audits would provide the basis for gradual shaping of norms. This measure of transparency would allow governments to tout good practices while deterring them from funding efforts that other nations find objectionable.

Vision: The appeal of this idea is threefold. First, it is technology-neutral. Since the emphasis of this idea is on examining funding and organizational structure, the framework would not become outmoded after changes in technology. Second, it would not require direct exposure of government information to the public. Though the idea could be extended to include public reports, direct exchanges of information among government officials are really the key to its success. Audits could be implemented with one national government at a time, and could proceed incrementally within those bilateral relationships. Third, the idea depends on information that is not generally considered to be within the scope of trade secrecy or individual informational privacy. My idea depends upon building trust and exchanging information with other nations to gain a better sense of the threats to U.S. interests. This is in contrast to other approaches that may involve monitoring communications or coordinating activities with a highly decentralized private sector that is often chary of government involvement.

Method: This idea originated in the context of research I have been conducting as part of an I3P project that examines legal and economic influences on the business rationale for cybersecurity. Other researchers involved with that project have emphasized that the lack of internationally shared norms concerning information collection is a major obstacle to improving cybersecurity. Conversations with my colleagues, as well as I3P Executive Director Charles Palmer, led me to focus on an idea that would help define norms for state-sponsored information collection, as well as a structure for enforcing them.

Dream team: Representatives from: National Security Council; DHS Office of Cybersecurity and Communications and/or Office of Infrastructure Protection; FBI Cyber Division; Office of the U.S. Trade Representative; international law and international relations experts.

KC Claffy response to http://edocket.access.gpo.gov/2008/E8-24257.htm

Who we are: Cooperative Association for Internet Data Analysis, http://www.caida.org/ .
A collaborative undertaking among organizations in the commercial, government, and research sectors aimed at promoting greater cooperation in the engineering and maintenance of a robust, scalable global Internet infrastructure.
Game-changing dimensions -- All: change rules, morph game, raise stakes.

Concept: www.bis.int. International Bureau of Internet Statistics. (start with bis.org/bis.gov in US)
Borrowing ideas from the Bureau of Labor Statistics (bls.gov) and its analogous agencies around the world, as well as the OECD (oecd.org), the BIS will assist public-private partnerships into distilling data that can be made available to various stakeholders in different privacy-respecting forms for use in developing economic, education, social and science policy.

Vision: We recognize that several other proposals for cybersecurity advances, including recent reports from the Center for Strategic and International Studies
(http://www.csis.org/media/csis/pubs/081208_securingcyberspace_44.pdf)
and the Internet Security Alliance
(http://www.isalliance.org/images/stories/The_Cyber_Security_Social_Contract_11182008.pdf)
will have to navigate data acquisition and sharing issues that have plagued cybersecurity as well as other technical efforts since the National Science Foundation left the cyberinfrastructure stewardship scene in 1995. All reports in this area have the common need for -- and lack of -- an international organization devoted to objective, neutral data on the Internet. In fact, the root of the cybersecurity challenge is the limits on trusted empirical knowledge generation imposed by economic polices that render knowledge accumulation slower than it is for our enemies.

Based on the best available data on infrastructure security, stability and economic sustainability, as well as coordinated feedback from stakeholders (workshops, etc), BIS should engage in many game-changing organizational roles:

1) identify the most important cybersecurity research questions the cybersecurity research community should pursue, and the data needed to pursue them. (change rules, raise stakes)

2) with research agencies and projects such as PREDICT (predict.org) and COMMONS (see companion proposal), help get necessary data to approved researchers (morph game)

3) promote cooperative data collectives among trusted enclaves (all three)

4) independent reports on accuracy of resource (e.g., IP addresses) ownership data from Internet registries, and other security-relevant databases. (all three)

5) track which cybersecurity strategies are working over time, (e.g., feedback on impact of DNSSEC, additional TLDs, SIDR)

Methods:
If statistics are intended to illuminate a sector, they must be designed by people who understand what aspects are important to the industry itself, and how industry processes relate to and result in measured outcomes. We painfully recognize a critical disjunction in the unfortunately intimately related financial sector, between lots of potential metrics and data and the (near absence of)

illumination, leading to a crisis we certainly cannot afford to risk -- but presently find ourselves disturbingly similarly situated -- in cybersecurity (and for the same reasons). Methods will have to include ongoing assessment and refinement of the metrics to be monitored, as well as an awareness of the limitations of statistics for improbable but catastrophic events, cf. Normal Accidents, Black Swan.

Fortunately, more building blocks for this type of effort exist now than have ever existed before. DHS's PREDICT program has learned many lessons regarding data sharing to support cybersecurity research, which could be applied to this effort. DOD's TIC program has already taken initial steps to make empirical analyses of critical cyberinfrastructure scalable and sustainable, and through its EINSTEIN effort is gaining an appreciation for the volumes of data involved, and the need for information theory as well as practical advances in data curation and management. The OECD has also developed respected methods of sensitive data acquisition, analysis, and publication. We propose leveraging experience from all of these sectors with what we have learned does and does not work, and closely tracking the effectiveness of new methods as they are tried. Other methods we propose:

1) incent participation through well-tested methods outlined in CSIS and ISA reports above (government purchasing power, research agency funding incentives), as well as regulatory tools proven effective for other critical infrastructure, and new methods geared to specific incentives and risks in cyberspace. (success of these methods assessed annually)

2) 'adaptive foresight' and 'scenario planning' workshops for public and private sector xperts to discuss what are the most important data to be collecting and collating, and how it can be collected, anonymized, and shared to satisfy security as well as privacy objectives.

3) sponsor projects such as "A Day in the Life of the Internet" (www.caida.org/projects/ditl), using federation of public and private measurement infrastructure available to support cybersecurity research, and guided by specific situational awareness questions, e.g., "how many vulnerable DNS resolvers are observable?" Retain historical data over time.

4) host workshops with legal and technology policy experts to discuss legislative updates to obsolete frameworks, with aim toward consistency across nations where sensible.

4) work with OECD and foreign government agencies to gather and improve data on cybersecurity related activity, and compare to what is available on U.S. networks.

Example macroscopic statistics the BIS could retain data for:

IP and AS topology, including coverage changes over time; BGP routing dynamics, including hijacking, e.g., PHAS; active measurement (RTT, bandwidth) gathered from research infrastructures around the world; flow statistics; trends in spam, malware, phishing, encryption, ciphers in e-commerce and other uses; IPv4 and IPv6 address space utilizatoin statistics; provisioning cost data

Dream Team: CAIDA, NSF CyberTrust, DHS S&T, NIST, security experts (whitehat teams, e.g., shadowserver), legal scholars with expertise in telecom data, e.g., Aaron Burstein.

How clear is the way forward?: To the extent that we're borrowing from already existing or proven techniques, it's clear. Whether they will work in this domain is less clear, and it is likely that legislative changes will be needed to support it. So on a scale of 1-10, it's a 5.

How high are the hurdles? Without the CSIS proposed NOC, or something like it, probably too high. With something like the NOC, hurdles are not only navigable but must be lept anyway.

per http://edocket.access.gpo.gov/2008/E8-24257.htm

Who we are: Cooperative Association for Internet Data Analysis, http://www.caida.org/ .
A collaborative undertaking among organizations in the commercial, government, and research sectors aimed at promoting greater cooperation in the engineering and maintenance of a robust, scalable global Internet infrastructure.
Game-changing dimensions -- All: change rules, morph game, raise stakes.

Concept: ***Cooperative Measurement and Modeling of Open Networked Systems (COMMONS)***

We propose to use the spare capacity recently announced on Internet2's backbone (NLR's backbone also has spare capacity) to connect select community and municipal networks to each other, and to the global Internet. Peering would be conditionally available to government entities, academic institutions, and community wireless initiatives committed to advancing the cybersecurity research agenda. The conditions for attaching networks are: (1) make some operational data available to cyberinfrastructure researchers under appropriate legal data sharing frameworks; (2) work with public safety community to develop dual-use infrastructures that give public safety authorities joint access to private or hybrid infrastructure during emergencies. (3) cooperatively develop and abide by policies, including experimental ones, based on confirmed results of data analyses.

Vision: We propose a collaboration to simultaneously solve four acute and growing problems facing the Internet: a self-reported financial crisis in the Internet infrastructure provider industry that limits investment into cybersecurity needs; a data acquisition crisis which has severely stunted the fields of cybersecurity research and network science; a fragmented and ineffective approach to public safety communications nation-wide; and a struggle for survival within emerging community and municipal networks, who are in an ideal position to assist with the first three problems but often lack resources and experience to make informed operational decisions, and are also continually threatened by incumbent-driven legislation.

The proposed project -- Cooperative Measurement and Modeling of Open Networked Systems (COMMONS) -- addresses the four highlighted problems, and without federal regulatory involvement (at least initially), which is still feared to be a cure worse than the disease(s) even by the regulators themselves. By offloading from commercial providers the responsibility for supporting Internet service delivery in unprofitable areas, we will measurably improve the financial situation of these providers. Second, COMMONS offers an unprecedented opportunity to establish standards of scientific integrity in the field of cyberinfrastructure research -- by providing rigorous empirical data against which to validate theories, models and simulations. Furthermore, because the COMMONS testbed will support public analysis of actual Internet traffic, it will inform debates on increasingly important technical, economic, policy, and social issues related to cybersecurity. Third, COMMONS infrastructure will provide an additional source of public safety communications, as well as a real-world platform for experimenting with how public safety needs can be accomodated by everyday communications infrastructure in times of emergency. Fourth, the COMMONS project not only allows struggling community networks to cost-share a financially daunting component of their operation, but it also provides a forum for the cooperating networks and the research community to share lessons learned with each other.

Ten methods COMMONS will use to improve cyberinfrastructure research capability:

1) in conjunction with representatives from IRB'S around the country and Internet2's new Network Research Review Committee (http://blog.caida.org/best_available_data/2008/10/10/internet2-launching-its-own-irb/) develop guidelines for privacy-respecting cybersecurity research, similar to the Belmont report written for human subjects research: http://ohsr.od.nih.gov/guidelines/belmont.html

2) use report developed in step (1) above to educate (a) legal scholars on how laws in different jurisdictions should be changed to support cybersecurity research; and (2) institutional research boards (IRBs) on how to update their processes to advance cybersecurity research

3) create efficient buy-in processes for regional networks to cooperate. facilitate transparent negotiation among public and corporate interests for e.g., right-of-way, spectrum sharing

4) guide participating networks in developing empirical analysis of cost, efficiency, and security of alternative ownership models, enabling a subfield of operational Internet research that does not currently exist

5) maintain repository of freely available software tools for measurement and analysis of operationally relevant network data, refine tool functionality based on feedback from users

6) through privacy-protecting projects such as PREDICT and DatCat, provide network data and meta-data to experts for independent research and analysis of security-related phenomena

7) promote cooperative research and data collectives among trusted enclaves via funding and legal support, and provide secure technologies to share lessons learned with eachother (see http://blog.caida.org/best_available_data/2007/09/18/renewing-us-telecommunications-research/ for related proposal)

8) collectively develop approaches to federated community network experimentation with new network, routing, and application technologies, using Interent2 or NLR as a backbone platform

9) support projects such as "A Day in the Life of the Observable Internet" (www.caida.org/projects/ditl) with both data and analysis targeted toward improving improving accountability and research methodologies of carriers and regulators

10) accessible outreach to educate users (i.e., public) on how they can improve their security odds in cyberspace: (including appealing material like DOD's recent 'Science of Victory' video)

Dream Team: NSF, Internet2, NLR, Internet data experts, privacy and legal scholars.

.

How clear is the way forward?: Not so clear, but we're facing a unique opportunity in current I2 and NLR conditions, and we've had two workshops discussing the idea. On scale of 1-10, it's a 5.

How high are the hurdles?

Legislative changes will be needed to protect data-sharing. Similar to the bis.int proposal, if a National Office of Cyberspace emergencies in the next administration (hopefully they won't call it that), hurdles are not only navigable but must be lept anyway.

RFI Name: RFI3 National Cyber Leap Year

Title of Concept: CERT C and CERT C++

RFI Focus Area: Morph the Game Board

Submitter's Contact Information: Robert C. Seacord, CERT / Software Engineering Institute
4500 Fifth Avenue, Pittsburgh PA 15213

Who We Are: CERT, the home of world-renowned CERT® Coordination Center, is located at Carnegie Mellon University's Software Engineering Institute. We study internet security vulnerabilities and research long-term changes in networked systems.

Concept: The majority of software vulnerabilities are caused by coding errors. For example, 64% of the vulnerabilities in the National Vulnerability Database in 2004 resulted from programming errors. The C and C++ languages are particularly prone to coding errors that result in vulnerabilities because of the lack of type-safety in these languages. CERT proposes a holistic solution to the problem that includes secure coding guidelines, program verification, static analysis, constraining the behavior of standard-conforming C and C++ implementations, and new runtime libraries. CERT C and C++ eliminate the following kinds of vulnerabilities:

- Writing outside the bounds of an object (e.g., buffer overflow)
- Reading outside the bounds of an object
- Arbitrary reads/writes (e.g., wild-pointer stores)
- Integer overflow
- Integer truncation

These problems represent the majority of vulnerabilities in C and C++. In 2007, MITRE reported that buffer overflows are still the number one issue as reported in operating system (OS) vendor advisories and that integer overflows, barely in the top 10 overall in the past few years, are number two for OS vendor advisories.

To address software security among these vendors, CERT proposes a broad approach that can significantly impact existing practice in producing secure C and C++ language programs. The buffer-overflow problem, for example, is solved using static analysis for issues that can be resolved at compile- and link-time, and dynamic analysis using highly-optimized code sequences for issues that can only be resolved at run-time.

Modern compilers for C and C++ already perform significant static analysis to understand program semantics for optimizations, especially on vector and super-scalar hardware. Furthermore, in well-written programs the array-bounds information is already maintained in variables defined by the programmer. CERT C and C++ provide a method for the compiler to track the bounds information and verify (at compile-time, link-time, or run-time) that reads and writes are valid. CERT C and C++ methods generate fatal diagnostic messages in any case where buffer overflow cannot be definitively prevented.

Integer-overflow errors are also a major source of vulnerabilities. Previous attempts to trap overflows created runtime overhead and hampered optimization. There is considerable latitude for optimization within the C and C++ standards, which permit shortcuts if the result at specified

observation points is the same as if specified semantics were followed.

In our new as-if infinitely ranged (AIR) model of integer overflow, when an observation point is reached, if overflow traps have not been disabled, and if no traps have been raised, then any integer value in the output is correctly represented ("as if infinitely ranged"). These traps are implemented using either the existing hardware traps (such as divide-by-zero) or by invoking a runtime-constraint handler. The same solution is applied to unsigned integer wrapping.

Vision: For any solution to have a significant difference in the reliability of the software infrastructure, the methods must be incorporated into tools that working programmers are using to build their applications. CERT and other members of the ISO/IEC WG14 standards committee are proposing a conditionally normative "analyzability" annex to the C1X major revision to the C standard now under development. This annex defines a security profile that requires that programs eliminate critically undefined behaviors that can result in buffer overflows, wild-pointer stores, and integer overflow as well as other security enhancements.

To encourage adoption of these methods into working compilers, CERT proposes extending ROSE to perform the program analysis and produce an advice file for the platform-dependent compiler. ROSE is an open source compiler infrastructure to build source-to-source program transformation and analysis tools for large-scale C and C++ applications. ROSE is particularly well suited for building custom tools for static analysis and software security. The generation of advice files was first proposed as a method for providing optimization advice from a front-end source-analysis tool to a platform-dependent back-end compiler.

Along with the ROSE Advisor a pre-linker is also required, to read and process the full collection of bounds-data files from all components of the application being compiled and linked.

CERT has already developed secure coding guidelines for C and C++ that will form the basis for the development of code checkers using advanced analysis techniques. CERT Secure Coding Standards are developed through a community (wiki) process, through which over 300 experts provided comment, as well as by a years-long ongoing dialog with the experts on the ISO/IEC WG14 C Standards Committee. We propose adding analysis methods to ROSE to provide assistance to projects attempting to revise their source code to enforce these guidelines.

Eventually, CERT or other organizations implementing CERT C and CERT C++ can certify that the applications are free from buffer overflows and the other classes of vulnerabilities addressed by this proposal.

Method: CERT will complete the prototypes of CERT C and C++ including an implementation of the AIR integer model using GCC, LLVM, or similar compiler infrastructures. Preliminary tests of a GCC prototype developed at CERT show that the runtime overhead of AIR integers is negligible (5.5% slowdown at –O0 running the SPEC CINT2006 benchmarks).

Portions of the checking needed for buffer overflow prevention were prototyped twice, first providing checks at the statement level, and then checking each relevant operator. These prototypes provided experience necessary to estimate the runtime overhead at well under 10%.

Dream team

David M. Keaton, Thomas Plum (founder, Plum Hall, Inc.), Dan Quinlan (LLNL), Kirk Sayre, Robert Seacord, David Svoboda.

RFI Name: RFI3 National Cyber Leap Year

Title of Concept: Secure Design Patterns for Fine-Grained Modular Security

RFI Focus Area: Morph the Game Board

Submitter's Contact Information: Robert C. Seacord, CERT / Software Engineering Institute 4500 Fifth Avenue, Pittsburgh PA 15213

Who We Are: CERT, the home of world-renowned CERT® Coordination Center, is located at Carnegie Mellon University's Software Engineering Institute. We study internet security vulnerabilities and research long-term changes in networked systems.

Concept: The use of secure design patterns for fine-grained modular security will morph the cybersecurity game board by:

- Allowing software designers to rapidly modify or completely swap out the security components of a system at a fine-grained level.
- Allow for the creation of reusable libraries of certified and verified low level security modules.
- Providing well understood design patterns providing a clear separation between security concerns and user functionality concerns in the design of secure software, leading to software that is easier to write, test, and verify.
- Provide reusable, separable designs for providing flexible fine-grained access and permission controls.

A *design pattern* is a general reusable solution to a commonly occurring problem in design of computer software. Software design patterns, such as the design patterns specified by Gamma, et. al., have received widespread adoption in the design and implementation of software. Software design patterns provide software developers with structured solutions for common software design problems that result in system designs that are modular, relatively easy to maintain and easy to understand. It is possible to extend existing, commonly used, and well understood design patterns to explicitly address security concerns in the same structured, modular manner that existing design patterns address general user functionality concerns. These extended design patterns are referred to as *secure design patterns*.

A *secure design pattern* may be created by carefully extending an existing design pattern to add additional structure and functionality to explicitly handle security considerations in a repeatable, well-defined manner. For example, the well known Visitor design pattern can be extended to create the Secure Visitor secure design pattern through the addition of locked and unlocked data nodes. Initial visits by a Visitor object are handled by locked data nodes, which only unlock their contents when explicit security requirements are met. A second example of a secure design pattern is the Secure State Machine secure design pattern, which extends the common State design pattern by explicitly breaking the state machine into two state machines, a secure state machine that handles security considerations and acts as a gatekeeper for a user functionality state machine that handles the actual user-level functionality. A secure design pattern presents the same interface and basic functionality as an existing, commonly used design pattern. Because software designers are already familiar with the use of the existing design pattern, they will be able to easily and quickly apply the corresponding secure design pattern. This characteristic of secure design patterns will greatly aid in the adoption of this new technology.

A secure design pattern also defines explicit boundaries between the portions of the design that handle security concerns and portions of the design that handle user-level functionality. Much in the same way that high level, coarse grained tools such as firewalls and secure network transport protocols

allow for the testing, verification, and swapping of different security mechanisms to be handled in a reusable, modular manner, the clear separation between security functionality and user-level functionality at the fine-grained design level will allow software developers to easily:

- Test the security functionality of the system. While it may be cost prohibitive to perform full, extremely rigorous testing of both the user-level functionality and security functionality of a portion of a software system, it may be much more feasible to perform rigorous testing of the security functionality of the system. Secure design patterns are specifically designed to make the portions of the design that implement security functionality highly separable and modular, making it easy to test the system security in isolation from the user-level functionality.
- Formally verify the security functionality of the system. While formal verification techniques can be used to provide strong assurance of the correctness of software, many formal versification techniques do not scale well to larger software systems. The modular nature of secure design patterns will make it possible to formally verify the security functionality of the design in situations where it would not be possible to verify both the user-level functionality and security functionality of the system.
- Patch or swap out the security functionality of the system. Unlike tightly couple designs that mix user-level and security functionality, the modular nature of secure design patterns makes it easy to modify or swap out system security, at a fine-grained level, in response to security flaws or changing user requirements.

Vision: Software developers would rely on their current knowledge of existing non-secure design patterns to identify system functionality that could be implemented with a non-secure design pattern. If the functionality includes security specific functionality they would then use the secure analog of the existing non-secure design pattern to implement the functionality, possibly taking the implementation of the secure portion of the secure design pattern from an existing library of tested and verified implementations of the secure portions of secure design patterns, thereby freeing the developer from having to reimplement and verify the security of the system. The steps needed to allow for the wide spread use of secure design patterns are:

1. Create a library of secure design patterns by extending existing non-secure patterns. Some initial secure design patterns have already been created by CERT.
2. Use the library of secure design patterns in the implementation of several example systems with security requirements.
3. Create a library of tested and formally verified secure portions of secure design patterns in a variety of commonly used languages that can be easily dropped in place and used by developers. This is possible due to the highly modular nature of secure design patterns.

Method: The concept of secure design patterns will be refined through the analysis and extension of exiting non-secure design patterns. The main assumption underlying the concept of secure design patterns is that the explicit separation of security functionality from user-level functionality at an understood fine-grained level will provide significant benefits in software reusability, understandability, and reliability and that the preservation of the interfaces of common existing non-secure design patterns will greatly aid in the use and adoption of secure design patterns.

Dream team:
Kirk Sayre, Robert Seacord, and Chad Dougherty all participated in the preliminary work.

Cyber Leap Year—Ipv6 Standardization for Network Fuzzing Techniques

Who We Are: Rhette (Margaret) Marsh, Routing and Switching CCIE #17476, CCDE candidate, CCIE Security Candidate, and GGSG (Global Government Solutions Group) of Cisco Systems, Inc.

Game-changing dimension: change the board and change the stakes.

Concept:

Code security has multiple models of accepted standard at the system level, but to date No comprehensive analysis has been done for network security fuzzing. Historically, Fuzzers have used single field permutations within a well-defined protocol to automate vulnerability discovery. Custom-coded fuzzers in IPv6, using Scappy6 et al. are focused on simple permutation, but also evolutionary and generation fuzzers ("intelligent fuzzing"). This proposal changes the gameboard and raises the stakes for the miscreants as it reduces their possible attack space and increases the potential complexity of a successful attack.

Vision:

This research proposal addresses how to automate testing to permutate protocols throughout architectures of interest to find threat vectors through unintended or unpredictable consequence to the transit infrastructure. The intent of the research is to close the gap of pure permutation fuzzing to include compound attack space, and then to assist in the development of an industry standard for infrastructure. Compound attack space here is defined as either more than one attack performed or attacks performed, in sequence, whose outputs could feed into the specificity of the attack for a given architecture.

Method:

For the formulation of this proposal, we referred to network fuzzing documentation for Scappy and Scappy6 (opensource), Codenomicon, Mu Dynamics, Bearing Point, and others. We examined trends in unintended cross-protocol interactions causing protocol failures or DoS conditions. This proposal is assuming and dependent on the ability to replay packet captures and fuzz on particular fields of that capture.

Dreamteam: Cisco Systems reachback, potentially manufacturer involvement from Codenomicon and Mu Dynamics, and others.

CAIDA-UCSD IPv6 Network Telescope Cyber Leap Year Proposal

Who We Are: Rhette (Margaret) Marsh, Routing and Switching CCIE #17476, CCDE candidate, CCIE Security Candidate, and GGSG (Global Government Solutions Group) of Cisco Systems.

Game-changing dimension: change the board and change the stakes.

Concept:

As IPv6 deployment becomes a global reality, so does the likelihood of zero-day attack vectors. There is a marked lack of research done in this area of IPv6 data collection and analysis leading to new identifying and mitigating techniques. Because technology is lagging behind implementation for complex inspection rules for variable length fields of IPv6, there is significant need for analysis methods of malformed packets designed to exploit the lack of detection. Likewise, worm, botnet, new protocol attack taxonomies, and attacks on infrastructure are also little understood in IPv6 and have even fewer widespread best practices. Attacks on the infrastructure itself include attacking the route cache of a router, IPv6 related denial of service against routers (core or edge), and sending malformed packets that may impact the integrity of the router. The UCSD Network Telescope for IPv6 will allow for trending, identification, and analysis of these attacks. As more applications are written natively in IPv6, propagation of malware and worms needs to be studied to help formulate best practices for minimizing risk for emergent architecture. This is a proposal for changing the gameboard by giving the whitehat side much more visibility, and will raise the stakes for the miscreants by making their movements exposed and predictable.

Vision:

This research proposal is focused at identifying propagation vectors of worms and botnets, reconnaissance techniques, routing attacks, DNS and DHCP attacks, spoofing attacks, header and fragmentation attacks, amplification attacks, routing protocol attacks, and attacks to the infrastructure itself. Previously critical participation from ISPs has reinforced global trending for new threats. CAIDA and UCSD's Network Telescope have successfully looked at these attacks in IPv4 traffic, and appear to be the logical place to propose collaboration with Cisco.

Method:

Execution of this proposal includes, but is not exclusive to: installation of a Honeynet for IPv6 at critical points in the backbone, installation of data collection for the Network Telescope for UCSD and potentially ISP points, development of techniques for IPv6 which isolate malicious backscatter from normal traffic, development of heuristics for analyzing novel types of reconnaissance (scanning methods are likely to change in IPv6), include methods against improper configuration on border routers, development of

heuristics to determine patterns in worm propagation in IPv6, development of heuristics to determine DNS attacks, and development of heuristics for attacks on routers themselves. For the formulation of this proposal, we used “Advanced Techniques to Detect and Control Global Security Threats” and updated it for emergent IPv6 threat research.

Dreamteam:

Members of the Dream team could include CAIDA, UCSD staff , and reachback within Cisco, ISP involvement, and members of the GGSG (Global Government Solutions Group).

1 Introduction

It may be possible for attackers to modify integrated circuits (ICs) to insert covert, malicious circuitry into manufactured components; a recent Department of Defense report [2] identifies several trends that contribute to this threat. First, it is infeasible economically for government-based IC suppliers to produce technology that matches the performance of commercial suppliers. These high-performance ICs provide a tactical advantage making them an indispensable resource. Second, commercial suppliers are moving more design, manufacturing, and testing of ICs to a geographically diverse set of countries in an effort to cut costs, making it infeasible to secure these steps in the IC life cycle. Together, these trends lead to an "enormous and increasing" opportunity for attack [2].

Motivated attackers will subvert the IC supply chain if doing so provides sufficient value. Since modifying an IC is an expensive attack, it is doubtful that "script kiddies" will turn their adolescent energies to malicious processors, but the same cannot be said for attackers with resources. If malicious processors are capable of running valuable attacks, governments, terrorist organizations, and so on will deploy them despite their cost. Historically, these types of organizations are experienced at covert operations, and have demonstrated considerable ingenuity in pursuing their goals. In contrast, there is little work on malicious processors.

If an attacker were able to include a malicious IC within a computer system, it would give them a fundamentally higher level of control compared to software-based attacks. While the recent SubVirt project shows that attackers can gain control over operating systems by using virtual-machine monitors (VMMs) to control the layer beneath [1], ICs occupy yet a lower layer. A malicious IC would be below all software, including VMMs, so compromising ICs gives attackers complete control over the entire software stack. This high level of control provides attackers with a fundamental advantage over defenders running above.

We propose developing intelligent malicious processors (IMPs) that run malicious services within the processor itself. Clearly simple attacks are possible (e.g., shut off the processor after a period of time), but we show that attackers can carry out sophisticated attacks using IMPs. We will design and implement example attacks and we will show that general-purpose attacks implemented using IMPs are possible, practical, and qualitatively harder to detect and defend against than current software-based attacks. We will identify fundamental perturbations to the system resulting from IMPs, and we will identify some challenges one must overcome to design, implement, and deploy general-purpose attacks using malicious processors.

We will lay the groundwork for implementing malicious services by developing several implementation strategies that highlight distinct perturbations. The perturbations inherent in our implementation strategies are analog perturbations (e.g., power consumption or temperature) due to increased transistor counts, timing perturbations due to resource contention, and visible attack states and events. Also, we consider three implementation strategies with varying levels of perturbation. Our first, *memory access*, allows unprivileged code to bypass memory protections. This uses few additional transistors (minimizing analog perturbations), but is visible within the system. The second, *hardware-only*, encodes attack logic completely in hardware using additional circuits that run in parallel with the main processor. This avoids adding attack states and events in the system, but at the cost of additional logic and analog perturbations. Our third strategy, *shadow mode*, aims to reduce analog perturbations resulting from additional logic by reusing existing circuits within the pipeline whenever possible. Like the hardware-only technique, the shadow mode

technique avoids adding attack states and events into the system, but does so without adding significant additional logic. However, reusing existing circuits introduces timing perturbations into the system by occupying pipeline resources that would otherwise be available to non-attack code.

Using these mechanisms, we will implement general-purpose malicious services to help better understand the challenges attackers may encounter when designing IMPs. Among the attacks we will implement are stealing software encryption keys, stealing passwords, a privilege escalation attack, and a service for enabling backdoor logins. When developing these attacks, two key challenges we will address are communicating with remote machines via the network and understanding high-level software abstractions (e.g., variables and functions) using only the low-level details available from within the processor (e.g., registers and instructions).

To evaluate our ideas, we will run our IMPs in hardware using a commercial processor. The Leon3 processor is an open source SPARC processor that is currently being used by the European Space Agency and in the Taiwanese ARGO satellite. We will modify the design of the processor at the hardware-descriptor-language (VHDL) level to implement our attacks, and we will synthesize and run our designs on an embedded system development board. Our processor will run on a field programmable gate array (FPGA) and the system includes many of the components found on a typical computer system such as Ethernet, USB, VGA, and PS/2, making it a realistic platform for evaluating malicious processors.

The contributions of this proposal are:

- We will be the first to design and implement general-purpose attacks (four in all) using malicious ICs, and we will be the first to run attacks on real hardware.
- We will be the first to show design tradeoffs attackers may make when designing malicious circuits.
- We will be the first to address some of the challenges of implementing practical attacks using malicious circuits.
- We will highlight what we believe are fundamental perturbations to a system infected with an IMP
- We will design and implement a defensive strategy that focuses on preventing attackers from abusing CPU caches.

2 Defending against malicious processors

Defending against malicious hardware is more difficult than protecting against malicious software. Software attacks can be averted through the control of the lower level hardware layer; with hardware attacks, there is no lower layer to leverage to regain control against malicious behavior. Nor can we simply rely on traditional fault-tolerance techniques, since the "failure" of malicious hardware will be correlated among all chips from one producer. Furthermore, the simplest denial of service based attacks (e.g., stop functioning after a period of time) are likely to be impossible to prevent. Yet if we cannot stop all attacks, then we may be able to prevent some, and make others more expensive/difficult to get away with. Therefore, for defense, our goals are twofold: to force attackers to increase the presence of anomalies between malicious and non-malicious chips, and

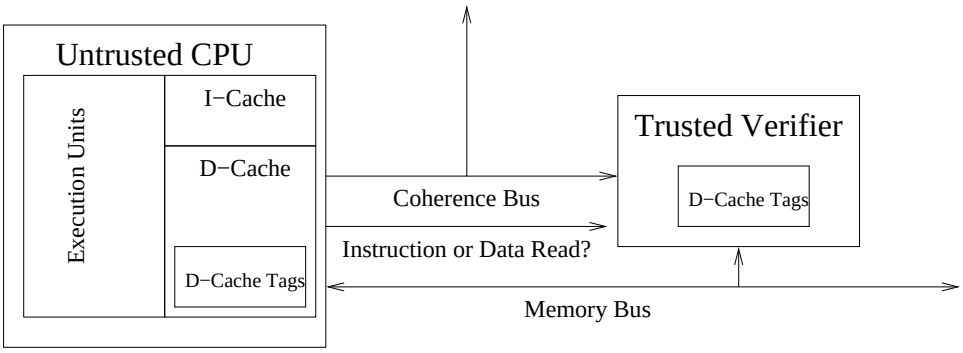

Figure 1: Trusted chip verifying an untrusted CPU.

to detect such anomalies so that malicious chips can be found. To demonstrate the feasibility of defending in such a matter, we propose a defensive technique which, by monitoring bus traffic, forces an attacker to eschew borrowing from the processor's cache or risk detection.

One of the ways an attacker can reduce the amount of malicious silicon is to reuse the already existing cache storage as a place to store malicious state. The key to our cache defense, then, is to force the attacker to avoid using the cache, or, if they do use it, to force them to become visible to higher levels. The memory necessary for "firmware" based attacks is likely to require a high transistor count (6 transistors per bit for SRAM-based cache); such larger changes are more easily detectable by inspection, and are more likely to perturb analog characteristics of the processor. Therefore, preventing malicious use of the cache is important.

Figure 1 shows an overview of our proposed defense. A trusted chip, the Trusted Verifier, duplicates the tag logic of the untrusted CPU. By duplicating tag logic, the Trusted Verifier can detect changes of the designed cache replacement algorithm from outside of the untrusted CPU. To reconstruct cache replacement state, the Trusted Verifier snoops on the memory bus and cache coherence traffic, and also requires an additional bit detailing if read traffic is for instruction fetch or data fetch.

References

[1] S. T. King, P. M. Chen, Y.-M. Wang, C. Verbowski, H. J. Wang, and J. R. Lorch. SubVirt: Implementing malware with virtual machines. In *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 314–327, May 2006.

[2] U. S. D. of Defense. Defense science board task force on high performance microchip supply. February 2005. http://www.acq.osd.mil/dsb/reports/2005-02-HPMS_Report_Final.pdf.

RFI name:	RFI-3—National Cyber Leap Year
Title of Concept:	"Shifting Sands: Operating functionally equivalent stable applications in a diverse, dynamically compiled environment"
RFI Focus area:	Randomization/Moving Target

Submitters Contact Information

Name:	Robert Grapes
Organization:	Cloakware, Inc.
Address:	84 Hines Road, Ottawa, Canada K0A 1T0
Telephone Number:	613.271.9446

Summary of who you are:

Credentials	Robert Grapes is Chief Technologist for Cloakware's Datacenter Solutions business (www.cloakware.com), maintaining the insight to the Enterprise market opportunities for Cloakware security technology. Rob has more than 17 years of professional experience in the technology sector. Prior to joining Cloakware in 2004, Rob spent many years with Entrust Technologies as a software toolkit product manager, with Cognos in vertical analyst relations, and with Allen-Bradley as a control systems automation developer. Rob's expertise on enterprise security and Governance, Risk Management and Compliance (GRC) has enabled many large government and financial service organizations to meet their audit requirements for PCI-DSS, FISMA, FERC and other regulations while reducing risk and improving operational efficiency.
Concept:	Dynamically compiled and deployed programs, or sub-programs, that incorporate multiple forms of code diversity provides a "shifting sands" approach to implementing a functionally equivalent and stable application environment. Static and dynamic analysis of software attempts to discover weaknesses in the code implementation for the purposes of creating an exploit. Automated software attacks leverage the discovered exploit for the purposes of perpetrating larger scale attacks. While encryption, obfuscation and integrity techniques remain as viable security techniques to frustrate the would be attacker, the fact remains that once a program is deployed "in the wild" it is subject to dedicated or even off-line attacks.

If a program, in its entirety or as a collection of sub-programs, can be continually compiled to include diversity and deployed into the production environment it will be possible to, in effect, push the “restart” button on any attack attempts.

Vision:

The vision is to introduce a software application deployment model that includes the necessary infrastructure components and processes for the incorporation of the automated compilation task(s) as an integral part of the deployment effort rather than as a pre-task. Introducing code diversity into the dynamically compiled and deployed applications will render automated and dedicated attacks fruitless especially as time scales between deployments is reduced, right down to hours and minutes if needed.

The fact that the compilation process will be performed on developed, tested and stable code eliminates the concern of untested software in a production environment. The “deployment” compilation process will be used only to introduce new diversity “seeds” into the compiled result…the code remains functionally equivalent.

A “head-end” will be required that has access to, or stores, the production source code and defines via policy the sources of entropy as inputs to the coefficients of the diversity function(s). This same head-end system will feed the existing software deployment system with the resultant compiled objects on a pre-defined schedule or an ad-hoc/on-demand basis. Depending on the criticality of a specific program or application it will be possible to define individual schedules and policies to meet the potential attack profile expected against that program.

As an Enterprise solution the vision for this product will include the complementary integration points with existing authentication, authorization, notification, reporting, event management and audit tools.

All of the technology to deliver on this vision exists today.

Method:

The consumer device and content distribution market have long been supporters of code diversity as an appropriate control for the mitigation of the risk of sophisticated software analysis attacks. Code diversity brings many

	security benefits not achievable by other means. Used in concert with other security techniques software diversity can help to deliver a comprehensive deployment model to defend against dedicated and automated attacks. With over one billion protected applications deployed we have many customers/partners in the device and content distribution market taking advantage of the benefits to be derived from code diversity in destroying the attackers' business model, however; the enterprise market is generally unaware of this security technique and the potential benefits to be gained through its use. Additionally, there are no existing Enterprise solutions created to deliver the "head-end" infrastructure to automate the scheduled/ad-hoc compilation and distribution of diverse software instances.
Dream Team:	None.

RFI-3 – National Cyber Leap Year

Title: Project Groundswell

RFI Focus Area – *Changing the Rules*

Submitter's Contact Information: Michael Grove, CollabWorks, www.collabworks.com, 650 El Camino Real, Suite O, Redwood City, CA; 650-346-8059; and Brian Barrett, Touchstone Consulting Group, 1920 N Street, N.W. Suite 600, Washington D.C. 20036, 202-449-7427, www.touchstone.com.

Who are we: CollabWorks provides a suite of collaborative services and infrastructure to foster among separate entities the sharing of solutions and resources to solve common problems among separate entities. Touchstone provides innovative facilitation, tools, and techniques to support collaborative activities among government agencies such as the CIO Council and the private sector.

Concept: *The core concept is to create a multi-tier network of entities that are as agile and effective as the bad guys in creating and distributing a solution.* Beginning with small pilots around real problem solving, develop collaborative "Cells" each consisting of 3-6 collaborating entities that share best practices, risk management assessment, and governance processes to effectively manage their individual cyber security threat. This process is called "Entity to Entity Collaboration™" or E2EC™. These network E2EC Cells™ would share solutions, vendor experiences, roadmaps, and expertise thereby increasing Cell effectiveness while significantly lowering costs and cycle time. A collaborative E2EC Network™ of 20-40 E2EC Cells would in turn share among themselves on a periodic basis the results, initiatives, and challenges of each E2EC Cell. Cumulative knowledge, solutions, and needs would be shared via events and a collaborative framework of processes and tools. Industry and government E2EC Forums™ composed of E2EC Networks would share a similarly structured approach but with more emphasis on policy and standards. Similarly, entity networks would be fostered in other countries, and these international E2EC Forums would foster global solution distribution. Thus, through an organic process, problem solving collaborative cells, networks, and forums would cooperate collectively to defeat the threat, much like the human body defeats a virus.

E2EC is a game changer because it addresses the root cause of the problem – an open Internet that promotes freedom of exchange, driven by consumer forces, and a plethora of devices and systems. Small entity groups exchanging knowledge, sharing resources, building trust, setting best practices and metrics is how to produce real and effective problem solving. Larger groups bog down in talk with very little walk – they just don't

scale. E2EC is a business architecture that lines up the value chain for each individual entity so decisions can be made on when, why, and how to spend resources. E2EC creates a multi-tier fabric of trusted relationships tied together in a problem solving network that succeeds because it is relevant and effective. Entities can self select which E2EC Cells to join. Competitors can avoid each other or work together.

Project Groundswell is an evolutionary process, just as the Internet and the cyber threat. By sharing solutions to problems, there are fundamentally fewer processes required to solve a problem. The greater the participation in the E2EC network, the greater value to all of the participants – this is the *network effect* working for the good guys. Further more suppliers are driven by the demand E2EC networks. Successful suppliers will line up their processes to capture E2EC network adoption. Suppliers can serve a network of customers instead of "selling each customer". Thus fewer processes are needed to generate value. The growth of the E2EC network will be driven by ruthless economics.

Vision: By 2014, more than 80% of the Global 2000 corporations, all G20 governments, and major universities will be participating in the E2EC network. They will be tied together by a collaborative infrastructure of policy, governance, compliance processes; frameworks for balancing threats, budgets, and risk; legal frameworks; rules of engagement; and a plethora of collaborative services that manage the use and flow of information. This network will evolve as the Internet did - organically. It will be driven by problem solving and economics. More than 80% of the solutions/practices that can prevent a problem are in place and constantly adapting as the threat adapts. Innovation will flow rapidly driven by user demand (as opposed to supplier inertia). Just as the body defeats a virus, the E2EC network will develop agile and effective mechanisms at the cell level. Rogue states will find themselves increasingly exposed and penalized by the huge network effect of trusted economic relationships. Bad guys will see a dramatic shift towards risk and away from reward.

Method: After two years of E2EC network research with high tech CIOs, it was clear that sustainable collaboration meant that resources spent had to be continuously relevant to each participant. Trust had to be sufficient to risk deliverables on shared execution. The 80-15-5 rule must be met – 80% people (trust), 15% process (effectiveness), 5% technology (scale). Collaborative management is essential to meet deliverables and maintain quality control. This will require a network of domain experts, collaborative leaders, and effective training and support.

Dream Team: Pilots of E2EC Cells can be created from existing cyber security initiatives, private-public partnerships, trade associates, and other groups of entities that have already established trusted relationships. The Defense Industrial Board, the Federal CIO Council, and BITS are good examples to start.

Telcordia Technologies, Inc.
One Telcordia Drive
Piscataway, NJ 08854

Configuration-Free Cyber Security

Programming the Configuration Virtual Machine

Who you are—Dr. Sanjai Narain[1], Dr. Gary Levin, Telcordia Technologies; Professor Daniel Jackson, Massachusetts Institute of Technology; Professor Sharad Malik, Princeton University; Professor Trent Jaeger, Pennsylvania State University

Game-changing dimension—Change the rules.

Concept—Eliminate the need to configure security. Security configuration languages are analogous to assembly language. No one programs in assembly language anymore. Why should we continue to do so for security? A typical machine can contain thousands of security configurations at the network, operating system and application layers and an infrastructure can contain thousands of machines. These configurations are manually crafted, and developed piecemeal, so it is impossible to provide assurance that these implement intended security goals and don't block intended functionality goals. The recent report "Securing Cyberspace for the 44th Presidency" states that security configuration remains a major operational challenge and that 80% of Air Force vulnerabilities were due to incorrect configurations.

Vision—Change the currently untenable state of affairs. We envision a world where people never have to write a single line of security configuration. They only specify holistic, high-level security goals and plans. All security configurations are automatically and correctly generated by a compiler.

What would the world look like if this were in place? We envision a **Configuration Virtual Machine**. As a first approximation, this is the union of all Reference Monitors in the infrastructure at the network, operating system and application layers. A Reference Monitor controls access of subjects to resources. Each such monitor has a configuration language for specifying low-level access-control rules. These Monitors would run over a distributed trust architecture to permit configuration across administrative domains in the absence of a centralized configuration authority. The assembly language of the Configuration Virtual Machine would be the union of configuration languages of all these Monitors. We also envision an *intuitive* security goal specification language and a compiler to compile this language into the Configuration Virtual Machine's assembly language. When these configurations are "interpreted" by their associated Monitors, the security goal will be accomplished.

The security goal specification language can be designed as follows: it contains primitive goals that capture what security assurance is provided by what configuration for what security technology. Examples of fundamental assurances are properties such as authentication, authorization, auditing, integrity and confidentiality. The language will also contain operators to compose primitive goals into system-wide ones such as Biba, Bell-LaPadula, Clark-Wilson, Chinese-Wall, least-privilege, and statements like "X-windows copy and paste between virtual machines at different security levels is disallowed". In turn, these will be composed into infrastructure-wide goals such as defense-in-depth, damage-containment and survivability. The idea is that if a human can conceptualize a good security plan (like an algorithm), he should be able to express it compactly in our language.

What makes you think this is possible? We are not proposing to invent new security technologies, only build a language, compiler and distributed infrastructure to make existing security technologies *usable*. Our goal is similar to that of programming language designers who abstracted from their computing experience with assembly language to design intuitive languages with abstractions like data structures, procedures, recursion, objects and methods. The number of fundamental security technologies is not large. Roughly, it is the number of security protocols in use. There is a wealth of

.

experience with using these so we are ready to formalize their configurations and create useful security goals and composition operators.

How would people get it, use it? The goal language and compiler would be installed on each machine that needs to be configured for security. System administrators would define their goals in this language, compile these into configurations then apply these to components under their control. When Reference Monitors on components interpret configurations, security goals would be realized.

What needs to happen for this to become real? We would work closely with organizations that address security challenges on a large and heterogeneous scale. We would understand these challenges in depth then abstract from that experience, design and implement the security goal language, compiler and distributed trust infrastructure, and evaluate it against the above deployments. We would repeat this step till mature technology results.

Which parts already exist; which parts need to be invented? The ConfigAssure project has shown significant promise in being able to formalize security goals as constraints and then, via a constraint solver, compile these into security configurations for infrastructure of realistic scale. Proofs of unsolvability are used as basis for configuration-error diagnosis and debugging. The Shamon project has created prototypes of distributed Reference Monitors for Xen/SELinux-based security. The most important item that needs to be invented is an *intuitive* language for specifying security goals for the infrastructure as a whole. A critical requirement of this language is an ability to capture relationship between security goals and functionality goals and associated configuration databases. Otherwise, concepts like "least privilege" and "need-to-know" cannot be specified and neither can one ensure that security does not interfere with functionality. Another important item that needs to be invented is the security goal compiler, particularly when the Configuration Virtual Machine is distributed across multiple components. For example, when components are added or deleted to the infrastructure, how do security configurations change to preserve infrastructure-wide goals?

Method—Our analysis is based on our experience with synthesizing and debugging security configurations for DISA's Multi National Information Sharing network, for the Securities and Exchange network, and in preliminary way, for NSA's High Assurance Platform. We have also created prototypes of distributed Reference Monitors for Xen/SELinux-based security.

Dream Team—We have interacted with the experts below in the course of developing our ideas and will try to collaborate with them. Independently, we believe the NITRD community would benefit from interacting with these experts.

1. Butler Lampson, MIT. Turing Award Winner. In his paper "Computer Security in the Real World" he states that security configuration is an outstanding problem.
2. Steve Bellovin, Columbia University. Network security and security architectures. In a paper to appear in IEEE JSAC, he also argues that configuration management is vital for security and that security is a systems property.
3. Donald Simard, NSA. Technical Director for High Assurance Platform program.
4. Peter Loscocco, NSA. Leads secure operating systems research. An author of SELinux.
5. Kevin Walker. DISA. Chief Engineer for Multi National Information Sharing program.
6. Paul Anderson, University of Edinburgh, UK. Principal Computing Officer, Develops logic-based technology to manage an infrastructure of several thousand machines at his university

Author Backgrounds

1. Daniel Jackson. Programming languages, mathematical logic, compilers, software engineering.
2. Sharad Malik. SAT-solver technology, digital systems synthesis and verification.
3. Trent Jaeger. Operating systems security, access control, source code and policy analysis
4. Gary Levin. Programming languages, mathematical logic and software development
5. Sanjai Narain. Programming languages, mathematical logic and software development

Volunteer Cyber Department (VCD)

Who you are – Concurrent Technologies Corporation *(CTC)*

CTC has been putting ideas into action since 1987; supporting a wide range of high-priority defense requirements and helping U.S. industry compete in the global market. CTC is an independent, nonprofit, applied scientific research and development professional services organization. CTC is classified as a section 501(c)(3) organization. CTC serves our client base with over 1,400 scientific, technical, and business professionals in over 50 locations across the nation.

Game-changing dimension —Change the Board, by bringing on-demand resources to your side when attacked.

Concept —The Volunteer Cyber Department (VCD), like the Volunteer Fire Department, will bring qualified defensive resources to bear against any attacker. Subscribers to this protection must volunteer their resources or pay a membership fee. VCD volunteers will have to commit to maintain computing infrastructure and processors under their control with certified cyber-defense tools. They will also have to submit to independent auditing and inspection to ensure their powers are only used for good.

The VCD members may "sound the alarm" when they detect any attack to their system or other trusted systems. VCD first responders will assess the situation and validate the cyber emergency. For validated attacks, a network of certified resources will counter the attackers with coordination of defensive, and potentially counter-offensive, means. Repeat false alarms will be charged a fee. This provides VCD members and subscribers with incentives to join for collective protection. The sacrifice of privacy and complete control over some computing/network resources must balance the perceived increase in security. Membership prices must be set at the point of value of the added security afforded to subscribers.

Vision —This world would look a lot like our US towns and cities in the face of fires. VCD member organizations would be incentivized by reduced risk and potentially improved insurance rates. Users would volunteer their computing time and resources to be available in response to attacks. They would also have the alternative of paying a fee to sustain the other resources and still gain the protection without volunteering. This would be possible where the benefits outweigh the costs for risk reduction. Many individuals would perceive the volunteering as "free" and provide a rich base of resource support. The coordination mechanisms for the VCD and the legal agreements would be a management layer that currently doesn't exist. The VCD would also have to provide an infrastructure for the certification and inspection of resources to ensure they were qualified to be members. Finally, a core of trained professionals would offer the benefit of maintaining the standards and administering the VCD. Just as actuarial science underpins the financial incentives for fire protection, research into the costs and value of cyber protection is called for to develop the agreements and quantify the risks and rewards for the VCD.

Method —This method uses the analogy of an evolved complex system. With the evolution of our firefighting infrastructure throughout the country to protect valued real estate, the VCD offers a similar construct to protect a "virtual estate." The demonstrated willingness for people to individually build the skills and sacrifice their resources to gain collective protection against

risk may be extended as the cyber domain permeates our households, and we share risk to our personal information resources.

Dream Team

As an independent nonprofit trusted advisor, *CTC* will take the lead in coordinating the following individuals and organizations for across the country to provide this innovative and vitally important capability:

- Jonathan Zittrain, Professor, Harvard Law School, Co-Founder and Faculty Co-Director, Berkman Center for Internet & Society
- Mary Ellen Hynes, Ph.D., Program Manager, HSARPA Critical Infrastructure
- Yaneer Bar-Yam, New England Complex Systems Institute
- Carnegie-Mellon CERT Coordinate Response researchers
- United States Computer Emergency Readiness Team (US-CERT) Public-Private partnership

Semantic Service Oriented Architecture (SSOA)

Who you are – Concurrent Technologies Corporation *(CTC)*

CTC has been putting ideas into action since 1987; supporting a wide range of high-priority defense requirements and helping U.S. industry compete in the global market. *CTC* is an independent, nonprofit, applied scientific research and development professional services organization. *CTC* is classified as a section 501(c)(3) organization. *CTC* serves our client base with over 1,400 scientific, technical, and business professionals in over 50 locations across the nation.

Game-changing dimension – Using a Semantic Service Oriented Architecture (SSOA) to Integrate Commercial Service Oriented Architecture (SOA) Platforms to allow for improved cyber defense across the United States critical information technology infrastructure.

Concept – SOA is currently an extremely popular systems development and software integration methodology. At its core, SOA provides a cost-effective means to integrate business processes in an enterprise. The SOA philosophy is straight forward: if all software applications or services and their associated data interactions were defined with standard interfaces, maintenance and updates could be easily accomplished and security improved. As these interfaces became prevalent, new applications could be created by simply "plugging" existing services together. Databases from various platforms with varying schemas could utilize standard interfaces for data sharing and integration.

An organization-wide, fully pervasive SOA will transform and empower users in an organization. Instead of requiring time and focused efforts of technical developers and engineers every time system or application changes are desired, a non-technical user can create a new application or data flow by connecting the existing services together in new ways. This "SOA Nirvana" will have enormously positive impacts to the overall productivity and security of the organization and have positive rippling effects across the organization.

Unfortunately, more times than not, the marketing hype revolving around SOA projects is substantially greater than the actual benefits delivered. A number of issues stand between the dream and reality that must be recognized and addressed. One of the most pressing of these issues is the focus of this proposed research:

- SOA product providers develop their individual product offerings in ways that make many of the key services and integration points to their system proprietary. As such, it can be virtually impossible to integrate these products together. This presents an enormous technical challenge that has dramatic and real business aspects. Since these products have expensive licensing fees, Government customers must either commit to a single product line or simply accept the fact that different enclaves of services on different SOA platforms will not easily communicate with each other.

The lack of SOA platform interoperability defeats one of the most compelling aspects and promises of SOA. It should be noted that the commercial vendors are motivated to continue their proprietary SOA approaches since once they have an initial deployment, there is great potential for further license sales. This results in an expensive and inefficient technical solution

for Government customers. SOA can deliver great results but the proprietary nature of leading SOA products makes true enterprise wide SOA interoperability extremely difficult.

Vision – *CTC* believes that the extensive Research and Development (R&D) experiences it has developed from key programs at NSA and AFRL over the last four years provides a very unique qualification for *CTC* to address this problem. *CTC* is confident that it can successfully increase the scope of SSOA to attempt to enable Service Discovery across proprietary SOA platforms.

The impact of this R&D could be extremely substantial for all Government agencies embracing SOA. It has been estimated that hundreds of millions of dollars are spent across the Government on SOA licensing fees. As these various SOA programs mature and deploy, they will not be able to interoperate with each other since no single SOA platform has been adopted as a standard. Each program will most likely claim success but the end result will be stove pipes of SOA platforms that can't share data and services. As a non-profit, trusted agent, CTC will extend their already existing work to begin to solve this challenging but compelling technical need.

Method – *CTC* is an R&D leader in this technical area; working for the last four years on key research projects at both NSA and AFRL that have proved the value of adding semantic technology into SOA frameworks to provide significant value in aiding in Service Discovery.

During these R&D projects, *CTC* developed new architectural components that allowed consumers (both users and applications) with means of finding services that meet their needs beyond "word of mouth" advertisement. Many complain that with the current Universal Description, Discovery and Integration (UDDI) registry model, where a service can only be "discovered" if it is already known to exist. *CTC* has developed prototype systems for NSA and AFRL that prove how semantic technologies can solve these challenges. Every service in the SOA framework has service descriptions created based on a developed service ontology that described the activities each service can complete at a task level. *CTC* then created and deployed two new key architectural objects to facilitate creation, storage and usage of these semantic descriptions. Specifically, the SSR (or Semantic Service Registry) was created to load, store and access the services and their associated semantic task level descriptions. Additionally, the SPR (or Semantic Process Registry) was created to load, store and access the workflows and their associated semantic tasks.

By creating easy to access structures (the SSR and SPR), services became very easy to discover and access. When one specific service went offline, it was straight forward for a software agent to search the SSR for other services that had similar semantic descriptions. Additionally, by monitoring the services that users typically used (and creating models of the semantic attributes), when new services were introduced into the SOA, their semantic attributes were compared to the attributes in the users models. Top matches were then pushed to the user in order to make them aware of the new services that may have been of high interest and relevancy to the user. These approaches will provide the foundation to move forward on solving the critical interoperability issues posed from proprietary SOA products and allow for improved security and SOA based cross domain solutions.

Dream team – *CTC*

Who you are – Concurrent Technologies Corporation *(CTC)*

CTC has been putting ideas into action since 1987; supporting a wide range of high-priority defense requirements and helping U.S. industry compete in the global market. *CTC* is an independent, nonprofit, applied scientific research and development professional services organization. *CTC* is classified as a section 501(c)(3) organization. *CTC* serves our client base with over 1,400 scientific, technical, and business professionals in over 50 locations across the nation.

Game-changing dimension – While we are currently putting numerous measures in place to protect our physical cyber world we have only just begun to understand the threats to virtual worlds and how that will negatively impact our real environment.. There are an ever increasing number of people using virtual worlds at a growing rate of 15% per month. It is time to level the cyberspace playing field by transforming the way we react to the potential for our adversaries to use virtual worlds to control these environments; use them as training grounds and a means to conduct illicit transactions; and to develop exploits that impact the real world. We have a current environment where our cyber professionals are not properly trained to interact with these virtual worlds and understand potential threats. This has created a safe haven for terrorists and criminals.

Concept – *CTC* is uniquely positioned for this type of effort. Our roots in academia, research, gaming and training technologies make us a perfect fit to tackle this arena. Specifically, *CTC* will make significant strides for the government by exploring virtual world environment to achieve the following objectives: identifying entities of interest; training security professionals on protocols and strategies; identifying cultural norms and tendencies; testing future technologies and techniques; providing non-attributable and cost-effective execution; develop methods of identity management and implementing classified exploitation methods.

Vision – Cyber professionals are currently exploring the virtual world in their spare time. We need to transform our way of thinking to include virtual worlds as potential areas of exploitation the same way we view physical worlds. With this change of focus, we can formalize methodologies to have cyber professionals devote resources to understanding and protecting virtual worlds in order to maintain parity with adversarial use. Some potential areas to explore include:

- Enhanced targeting of individuals and organizations. In these Cyber/Virtual Worlds, identities (avatars) are the focal point. We will dissect what is referred to as the “magic circle”, which describes the imaginary barrier between a person’s real life and their one or more avatars in the cyber world. The cyber world protects the user’s identity allowing them to cause real harm in the physical world. In 2001, Edward Castronova concluded that the value of the currency in the MMORPG Everquest was equal to 0.0107 USD, which is higher than the Yen or Lira. South Korea has begun to assess income taxes made by playing virtual games. Italy has declared it a crime to commit sex offenses against virtual children in the game *SecondLife*. One can easily translate this into an avenue for terrorist organizations to gain capital. In 2006, a woman from China became the first millionaire through profits earned entirely inside the virtual world of *SecondLife*.
- Intelligence gathering. Cultural trends and tendencies have spilled over into these virtual worlds. We now have access to unlimited information regarding how group’s around the

world act and train. For example: many countries have embassies in the game *SecondLife* such as the Maldives – located on "Diplomacy Island". Avatars can talk face-to-face with ambassadors about visas, trade, etc.

Method – *CTC*'s methodology will enable the United States to make significant strides by exploring simulated virtual world environments to achieve the following objectives:

- Create a database of entities of interest – Gather metrics and data regarding suspect individuals and organizations, resulting in more accurate and interesting targeting.
  - Develop social networks to connect previously unknown threats. Consider the electronic fingerprint left by these users: cell phones, VOIP, MAC/IP addresses, IM, emailing, etc.
- Train cyber professionals on the strict protocols and playing strategies – Enabling them to become one of the adversaries in daily life. The virtual worlds are made of cliques; for example, a small, exclusive group of members of a religious affiliation. It is to our advantage knowing the protocols and having the right level of players (not entry level users) in order to infiltrate desired rosters.
- Identify cultural norms and tendencies – Each culture plays these games differently. The government would gain critical intelligence in support of psychological operations.
- Test future technologies – Such as stealth, cloaking, and invisibility within these virtual worlds (be a fly on the wall).
- Identity management development – Provide a means to assure that a virtual user is who they identify themselves as.
- Classified Exploitation Methods – One can easily begin to imagine the classified possibilities for analysis and counterterrorism (not included in this paper).

Dream team - *CTC*'s roots in academia, research, gaming, and training technologies make us a perfect fit to tackle this arena. Utilizing our ability to partner with the following major universities, *CTC* is uniquely positioned with an unparalleled network of resident experts with vast knowledge of these Virtual Worlds and experienced avatars ready for analyst use, including: Akron, Auburn, California, Carnegie Mellon, Cincinnati, Dayton, Delaware, Denver, Drexel, East Carolina, Florida, Florida A&M, George Mason, Hawaii, Illinois, Iowa, Johns Hopkins, Lehigh, Lincoln, Maryland, Montana, Ohio, Ohio State, Old Dominion, Pennsylvania State, Rockhurst, Saint Joseph's, Seattle, Tarleton State, Temple, Utah, Villanova, Washington, Wisconsin, and Yale, as well as many others.

Secure Enterprise Networks Consortium (SEN-C)
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Cyber Attack Early Warning System (CAEWS)

Who We Are

The Secure Enterprise Networks Consortium (SEN-C) is comprised of Accenture, Los Alamos National Laboratory, Sun Microsystems, and CA, Inc. SEN-C focuses on bringing leading skills together—from thought leadership and solution development to systems integration excellence. By collaborating with government, we seek to achieve outcomes that enable CNCI initiatives and improve our nation's security.

Game-Changing Dimension: Raise the Stakes

Concept

We propose to make cyber attacks more difficult by developing a national Cyber Attack Early Warning System (CAEWS) and response capability, equivalent to NORAD, complete with:

- **Phased-Array Radar** for intrusion detection looking inward and outward at the cyber enterprise. We recommend equipping a network of watch sites with the equivalent of more modern "phased-array radar" intrusion sensors, as opposed to today's conventional Host-based and Network-based Intrusion Detection Systems (HIDS, NIDS).
- **Extended community** of Grey Hats to plot threat vectors, incursion routes, and understanding of attackers' offensive capabilities. The community connects thinking of public, commercial and academic communities. Premise: better defense comes from understanding of adversaries' offensive capabilities, shared among defenders.
- **Morphing targets** that fill an enterprise: some real, many fake, all indistinguishable from each other. Morphing targets keep attackers guessing. They decrease change-of-system disruption or compromise. They are both networks and applications that generate ghost ecosystems to diffuse threats and to monitor intrusion activity.
- **Situational awareness** to develop attack scenarios and responses that integrate network information assurance with information operations to cut off threat vectors. This will be part of a Cyber Attack Station, cyber equivalent of the "Cheyenne Mountain Air Station" of the Ballistic Missile Early Warning System (BMEWS), coordinated with data from other sensors at the CAEWS sites.

Vision

We envision a sensor-to-analyst-to-ground cyber early warning system, like the BMEWS, the first operational ballistic missile detection radar. The BMEWS provided long-range warning against Arctic ballistic missile attack. It also provided satellite-tracking data.

The CAEWS will serve similarly. It would act as a clearinghouse, sharing defense strategies and methods with public and private organizations to raise the baseline of our national security. It will help stay ahead of cyber attackers through analysis and rapid adoption of new techniques, tactics, and procedures. It will do so across a heterogeneous network landscape and through integration of offense and defense to shut down threat vectors proactively.

The CAEWS has five significant requirements for success:

1. Vastly improved collaboration across IT security, application/software, and infrastructure space that spans public and private sectors
2. True decoupling of application and infrastructure layers of systems so that application stacks and business processes can reconstitute dynamically and fluidly on different infrastructure, based on business/mission need to stay one step ahead of threats
3. Enterprise, network and open source data needs collected in greater scale and greater timeliness help operators understand threats and to respond effectively.
4. Attribution of activity is still more of art than science, with shadow networks, bot net controllers, and unwitting host computers confusing doors to stop an attack.
5. Countermeasures limited for all but most advanced IO elements of DoD. Countermeasures require R&D in advanced countermeasures to cut off threat vectors. Combined with next generation attribution, countermeasures give more immediate relief to organizations under attack.

Method

SEN-C formulated this concept with contributions to the Accenture Collaborative Innovation Solution (ACIS) by members of Accenture, Los Alamos National Labs, Sun Federal, and the University of Virginia (UVA). With ACIS, we grouped and refined similar ideas based on topic and quality rating from the broader set of participants. This concept draws on areas of experience in IT security, information management, and software engineering. Some groundwork for morphing targets has occurred in prior UVA research programs and in industry exploring mobile software agents that move from machine to machine to avoid intrusions.

Dream Team

We propose for our team: the National Cyber Security Center, the Global Information Grid IA portfolio office, chief security officers from critical infrastructure providers (customers or stakeholders of the CAEWS), the University of Virginia (which has invested R&D already), the University of Illinois (a leader in cyber modeling and simulation), Accenture, Sun Microsystems, and Los Alamos National Labs.

Leap Ahead Concept in Response to
NITRD CNCI (Comprehensive National Cybersecurity Initiative) RFI
National Cyber Leap Year

Policy-Enabled Intelligent Agent Forces

Submitted December 15,2008

Submitted by:
The Boeing Company
Boeing Phantom Works
P.O. Box 516
St Louis, MO 63166-0516

Principal Points of Contact:

Technical:
Jai Choi, Boeing Phantom Works
Phone: (425) 373-2902
Facsimile(425) 373-2969

Stephen A Ridlon, Boeing Phantom Works
Phone: (425) 373-2853
Facsimile:(425) 373-2960

Contractual: Tonya Karp
Phone: (206) 544-1327
Fax: (206) 766-5655

Boeing CAGE Number: 4JSW3
Boeing DUNS Number: 622377070

RFI Submission to NITRD National Cyber Leap Year
December 12, 2008

Who We are:

The Boeing Company:
Boeing Phantom Works
P.O. Box 516
St Louis, MO 63166

http://boeing.com/

Boeing is the world's leading aerospace company and the largest manufacturer of commercial jetliners and military aircraft combined. Additionally, Boeing designs and manufactures rotorcraft, electronic and defense systems, missiles, satellites, launch vehicles and advanced information and communication systems.
Boeing Phantom Works and Integrated Defense Systems are at the forefront of research and development and implementation of advance cyber solutions. Additionally, Boeing maintains one of the world's largest network infrastructures supporting all of its business operations with international partners.

Game changing dimension – morph the game board

Develop a policy driven intelligent-agent based cyber defense system that pro-actively prevents and minimizes attacks keeping the cyber infrastructure operating securely and effectively.

Concept:

Cyber attacks threaten the safety and well being of the United States. These attacks come from foreign militaries, intelligence services, and global civilian sources.
What if we changed the game board to anticipate and prevent these attacks as well as providing defensive mechanisms that mitigate and counter attack, thus keeping our systems running securely and efficiently - intelligent agents would penetrate every facet of the network infrastructure as a security demon, self-controlling/managing OS level service that is transparent to users, with no added cost of computing.

Vision:

Our vision is for a policy driven intelligent agent based system that anticipates attacks and that both pro-actively prevents them and minimizes their effect, keeping the cyber infrastructure operating securely and effectively. We will build a reflective and event-oriented Self-regulating Agent Operating System (SAOS) middleware for cyber security context awareness and maximized service autonomy. In such a system, intelligent security agents will autonomously collaborate among each other and create (or remove)

sub-agents on the fly, and use them to support situation acquisition and dispatch services tailored for local missions in a distributed environment. These agents will interact to monitor and detect as well as gather intelligence and synthesize information. They will then act in concert to initiate measures that prevent or mitigate attacks. A cluster of large scale agent forces can be deployed as “agent-in-a-box” or “agent-plug-in service.” These agents would be controlled via a policy based mechanism that would assure consistency of approach and compliance with regulations and policy directives, and be able to react to changes in global policy. The landscape is constantly changing in terms of friends and foes, necessitating flexibility and responsiveness in the approach.

Combined with the advancement in C2 management systems, we can achieve operational survivability and sustainability under attack with minimal impact to security transparency. We have confidence that this extension is feasible as we already have considerable experience in prototyping both intelligent-agents based and policy based cyber security components. The existing agent based systems need to be researched for scalability and evolved into a semi-autonomous state that can self-regulate based upon defined electronically captured policies. The supporting policy mechanism must be developed and integrated with multi-agent technology.

Method:

Two major R&D efforts are needed for this leap idea, involving the development of a new kind of agent operating system that will embed cyber intelligence into the SAOS middleware level. A cyber policy framework should also be developed to accommodate universal adaptation of a cyber policy language. We called together a broad spectrum of researchers and practitioners in communications and information assurance from across the Boeing Enterprise. Together we analyzed our existing capabilities and research endeavors to identify various components that would support an attack prevention and mitigation capability. We see the need to further research and develop and then integrate the underlying technologies (such as agents, policy, security management capabilities, etc.).

Dream Team –

Boeing Phantom Works and Advanced Systems professionals, who have years of experience in a broad spectrum of communications, security, and information assurance research and development would be a major part of the team.

Additionally, UIUC (the University of Illinois at Urban-Champaign) and ISU (the Iowa State University), who are both currently research partners with Boeing in cyber operations and computing security efforts, would contribute. Boeing’s strategic research partners such as IBM, with its autonomic computing systems capability, could contribute.

NITRD National Cyber Leap Year RFI Response: Filtering Spam 'At Your Leisure'

1. Submitter: Joe St Sauver, Ph.D.
PO Box 3504, Eugene, Oregon 97403

I am involved with a number of cyber security activities, including serving as a senior technical advisor for the carrier Messaging Anti-Abuse Working Group (MAAWG), however this submission is being made solely in my individual capacity and not on behalf of MAAWG or any other entity. Any opinions expressed are strictly my own.

2. Game-changing dimension: Rules

3. Concept: Currently messages are synchronously and immutably categorized as spam or ham at delivery time (or immediately thereafter). Often only very limited information is available at delivery time about the source of each message or the character of its content. Constraining message categorization to just that one miniscule moment unduly and unnecessarily hinders anti-spam processing since there will often be a period of minutes, hours or days between the receipt of a message by a mail server and knowledge of the existence of that message by its the ultimate recipient (e.g., some users only check to see if they've got email once or twice a day, or a couple of times a week).

So let's take advantage of that window. Let's imagine an alternative paradigm whereby no message is considered irrevocably "delivered" until its existence has been disclosed to the user. Until the user learns of the existence of a message, assume we're free to update the "spam status" of that message as additional information becomes available.

4. Vision: What might we learn during that interval that might make us change our mind? Well, we might learn that the sending IP address has been block listed by Spamhaus, or we might learn that a URI present in the body of the message drops malware or is a phishing site or is otherwise unsavory. Had we known that at delivery time, we might have given the message the "thumbs down" then, but life is not perfect – the true nature of a sending IP or the trustworthiness of a message payload may not be known (or knowable!) until later, far after we've made our best effort attempt at categorization.

So, then, why make a binding decision about the status of a message *only* at delivery time? At least for the many users who check their mail only on a sporadic basis, we can non-disruptively revise our assessment of a particular message's nature right up to the point where the user checks their account and that message's existence is revealed to them. Until then, I believe we can (and should!) revise the status of that message to reflect any new information we may learn.

Eliminating the "race" between the spammer and the spam filter operator means that more messages will ultimately be accurately categorized as ham or spam (relative to the current paradigm where a filter operator is only allowed to make an irrevocable snap decision based on limited knowledge at the moment a message is received at the server).

5. Method: Messages would be re-evaluated when event transitions occur. Event transitions consist of things such as (a) a connecting IP address gets listed on a DNS block list, (b) a message body URI gets listed on the SURBL or URIBL, or (c) a message checksum get listed as bad on a collaborative filtering service, among many other possible examples. As soon as an event transition of that sort occurs and that change of status becomes known, the message status of all relevant messages can be updated. Clearly, however, it isn't practical to do a sequential rescan of all messages every time an event transition occurs since event transitions occur virtually continuously (or at least every hour or few hours in the form of periodic zone file updates for DNSBLs)

To make event-transition-triggered filtering realistically possible, we need to borrow an architectural lesson from Usenet News (NNTP). For those who may not be familiar with Usenet, INN (a popular open source news server implementation) creates an "overview" database containing distilled information about each article. I'm proposing that every time we receive a mail message, providers should create a similar "mail overviews" database entry for that message. By distilling key information for each message into such a consolidated/indexed database, we can avoid the need to rescan individual messages.

Each overviews record might contain some or all of the following items: (a) message ID, (b) maildir path for the message (or an equivalent reference), (c) mail "folder" associated with the message currently (inbox, spam folder, other automatically selected folder), (d) IP address of the system that connected and handed us that message, (e) date/time the message was received, (f) message envelope sender and message body sender, (g) any message body URIs as well as the IP addresses, ASNs and name servers of those URIs, (h) message body or attachment checksums , (i) current spam status, etc.

Once your server has such an overview database in place, system access to mail messages (via web email, POP, IMAP, etc.) can occur via that mail overviews abstraction layer, but the user email presentation would be unchanged. [Note that at least one IMAP server (Dovecot) already creates indices similar to what I'm proposing, although those indicies contain less info and are not specifically intended to facilitate *post hoc* spam filtering.]

So what occurs if/when the system learns of an event transition, such as a URI that's been listed as being spammy? That event could trigger database updates to the spam status of all as-yet-unaccessed messages associated with that URI, including potentially triggering other actions such as refoldering or /dev/null'ing now-tagged-as-unwanted-messages, etc

I describe this concept and how it might be easily implemented in my publicly available 2006 MAAWG talk that's at: http://www.uoregon.edu/~joe/maawg7/maawg7.pdf

6. Dream team: Because email is stored and delivered via a variety of open source, commercial and proprietary mail software products, and this method fundamentally changes how messages are store and manipulated, implementing this approach will requires working with one or more mail server software developers. Because Dovecot is a popular open source option that's widely deployed and which already contains a version of the required indicies, I'd suggest beginning with the developers of that product.

NITRD National Cyber Leap Year RFI Response:
Who Will Clean Up The Millions of Newly Infected Personal Computers Each Month? Or, "Why We Need A Cyber CDC or Cyber World Health Organization"

1. Submitter: Joe St Sauver, Ph.D.
PO Box 3504, Eugene, Oregon 97403

I am involved with a number of cyber security activities, including having presented an invited talk on this idea at the 2007 Anti Phishing Working Group (APWG) E-Crime Summit, however this submission is being made solely on my own behalf, and not on behalf of APWG or any other entity. Any opinions expressed below are strictly my own.

2. Game-changing dimension: Morph the gameboard

3. Concept: Worldwide, there are millions of compromised (malware-infected) consumer PCs. Currently no one in the United States government is focused on helping to clean up and harden those infected systems, nor has anyone else, outside of government, stepped up and taken responsibility for this problem. I propose that the U.S. government create a "Cyber CDC" or "Cyber World Health Organization" to tackle this issue worldwide.

4. Vision: To get a sense of the magnitude of the malware crisis we face, consider the millions of IP addresses listed on the CBL anti-spam DNS blocklist (http://cbl.abuseat.org/). While the hosts listed on the CBL are there for having sent spam, most of those same botted hosts could just as readily be quickly repurposed to serve child porn images or pirated software, to conduct distributed denial of service attacks against government systems or critical civilian infrastructure, or for other nefarious purposes. This pandemic of compromised machines constitutes an international cyber crisis, albeit one that has seen surprisingly little headline coverage to date.

Part of the problem we face is that currently no one accepts responsibility for the clean up and hardening of compromised non-governmental computers -- not the owners, not the ISPs that connect those hosts to the Internet, not the vendors who made those systems or wrote the software they use, and surely not the malware writers who compromised them!

In talking to people about this problem, it is common to hear the opinion that the owner of the system should be responsible for keeping their system malware-free, and there certainly is some simple logic to that. Unfortunately a variety of practical constraints (such as a lack of expertise, a lack of time, a lack of critical specialized software tools, non-existent backups and many other factors) make it very hard for a typical non-technical computer owner to clean up their computer once it has become infected.

Moreover, as long as their system still "sort of works" for their purposes and as long as their ISP hasn't cut them off, users may not see the point in paying a professional to clean up or rebuild their system for them, particularly when it may be cheaper to just buy a new machine (selling their old and still-infected computer to someone else), and particularly when they may get "serially re-infected" immediately after they've been cleaned up!

Having exhausted all other possible avenues for relief, we thus find ourselves looking to the government as a "provider of last resort" when it comes to tackling this problem, much as the government acts as a provider of last resort to deliver emergency services during conventional disasters, such as hurricanes.

But who in government might help with this *cyber public health crisis?* Cyber public health service can't come from existing law enforcement agencies -- they're already overcommitted and in many cases infested systems may have become infested through what I sometimes refer to as "low grade online illegal behavior." For example, we know that many infections come from downloading trojan'd software or tainted music files. I don't want a user with an infested system to be reluctant to ask for help just because they've violated someone's copyright or because they're embarrassed at having become infected while visiting an online illegal gambling site. No one is saying that it is okay to break the law, but the overriding goal in this case is to get infected systems cleaned up, and that can only happen if users know that they can safely ask for help without risking action by law enforcement. Requests for cyber assistance must be privileged, at least for misdemeanor-class offenses.

So if not law enforcement, then who? It would be convenient if we could just point to an existing federal agency and say, "Ah, this would be a perfect fit for DOJ, or the FTC, or the FCC, or Interior or <fill in the blank>" but unfortunately I don't see any agency that's both appropriately focused and eager to take on the massive challenges which would be associated with delivering cyber public health services to our nation.

5. Method: I am therefore left with no option but to suggest that we need a new federal agency to deal with cyber public health. This agency should NOT be a "Department of the Internet" -- anything that all-encompassing will immediately run into a storm of knee-jerk opposition as everyone worries about what a "Department of the Internet" might do about network neutrality or whatever might be the Internet policy crisis of the day. We just want (and need!) an agency dedicated to helping us clean up our cyber mess, and the cyber mess that's also proliferating overseas.

I describe this concept and how it might be implemented in my publicly available 2007 APWG talk that's at: http://www.uoregon.edu/~joe/ecrime-summit/ecrime-summit.pdf

6. Dream team: What's needed is something that combines the "public health" skills of the CDC or World Health organization with the grass roots neighborliness and helpfulness of a Peace Corp or AmeriCorps program with the cyber acumen of US CERT and the international savoir-faire and language expertise of the State Department.

Many colleges and universities are familiar with how to scalably attack large numbers of unmanaged infested computers (since we see those systems arrive en masse each fall), however what we really need is the service delivery expertise that is homed with national scale boots-on-the-ground distributed government service agencies such as agricultural extension offices.

NITRD National Cyber Leap Year RFI Response:
Attacking the Online Underground Economy by Regulating, Licensing and Taxing Online Affiliate Programs And Supporting Services

1. Submitter: Joe St Sauver, Ph.D.
PO Box 3504, Eugene, Oregon 97403

While I am involved with a number of cyber security activities, including having participated as an invited panelist for an April 2008 RSA Conference panel session on the underground economy, this submission is being made solely on my own behalf, and not on behalf of any other entity. Any opinions expressed below are strictly my own.

2. Game-changing dimension: Raise the stakes

3. Concept: Cyber crime today is, for the most part, a business.[1]

Many cyber crimes rely on techniques originally introduced and pioneered by legitimate online businesses. For example, consider "affiliate marketing models." Affiliate marketing models can be used to drive traffic to legitimate online businesses or, equally successfully, to drive traffic to criminal web sites. If we can regulate, license and tax other businesses and marketing channels, surely we can also regulate, license and tax the online affiliate marketing programs used to promote the sale of knock off watches, the sale of illegal drugs, and other illegal goods and services. At the same time we regulate, license and tax affiliate marketing models, we also need to attack the delivery of business support services to online illegal enterprises.

Just like legitimate online businesses, online criminal enterprises routinely rely on things like credit card processors in order to accept online payments for illegal online products and services. Those credit card processors might be "niche" processors, processors who specialize in handling so-called "high risk" online sales, and criminals may pay a financial premium to get those specialized services, but criminals **can** currently get the services their online illegal enterprises require. That needs to change.

I propose attacking those circumstances by regulating, licensing and taxing online affiliate programs, and by focusing attention on miscreant-critical business support services (such as those high risk payment processors I've previously mentioned).

If we're successful in curtailing illegal affiliate programs, cyber criminals will see fewer customers to their web sites. If we're successful in attacking their ability to readily process customer payments and conduct other routine but critical business processes, we will (a) literally "increase the cost of doing business" for the cyber criminals, while (b) making it less convenient for customers to make illegal purchases, and (c) providing law enforcement with improved opportunities to "follow the money."

---

[1] See, e.g., "An Inquiry Into The Nature and Causes of the Wealth of Miscreants," http://www.icir.org/vern/papers/miscreant-wealth.ccs07.pdf

4. Vision: Consider an online pharmacy illegally engaged in selling scheduled controlled substances without a valid prescription. The benzodiazepines or steroids or narcotics being sold by that concern might be produced in Southern Asia, Eastern Europe, or the Caribbean, but those producers need to connect their supply with interested retail customers in North America, Western Europe, and elsewhere. Illegal affiliate programs currently serve to create that connection, "making a market" for those products and ultimately channeling the orders they generate upstream to a drop shipper for fulfillment. In exchange for acting as a salesman, the affiliate may be paid a comparatively large fraction of the order amount as a "commission." Now imagine a world where:

- Anonymous affiliates are a thing of the past, and each affiliate's ID is registered to his or her real identity, so that abusive promotional practices can be tied back to the affiliate (or the affiliate marketing program can be held responsible)

- Affiliate program managers can be held responsible for promoting illegal products

- Income tax and other customary government payments would be routinely withheld from affiliate compensation

- Convicted criminals could be barred from obtaining an affiliate marketing license

Similarly imagine the impact if, instead of being able to charge a credit card or use an anonymous gift debit card, online illegal drug marketers had to find some other way to receive payments for their products. We know that at least one bank card association routinely forbids its associates from processing charges for online pharmaceuticals, but why are the other bank card associations and their associates still allowing criminals to continue to take advantage of their services? Surely this, too, can be regulated if necessary?

Undercutting the basic online criminal business model will help to reduce cyber miscreants ability to make a living from illegal online activity, thereby driving down the market for things like malware and bots, while also reducing the attractiveness of "anonymous" illegal online business activities to organized crime elements.

5. Method: Because affiliate programs are a business activity affecting interstate and/or foreign commerce, they can be federally regulated/licensed/taxed to control abuses via appropriate legislation. Online business services (such as high risk payment processors) can be targeted for increased enforcement attention by financial crimes specialists. If necessary, compulsory criminal processes can be used to obtain access to high risk payment processor customer information and settlement details, but I would hope that industry self-regulation could resolve this problem instead.

6. Dream team: Work on this concept would require participation by the Department of Justice, the Federal Trade Commission, Internal Revenue, and the Department of Commerce, as well as participation by the payment card industry and other business support service providers.

The Boeing Company
Integrated Defense Systems

Masking Network Assets

In Response to:
Request for Input (RFI) – National Cyber Leap Year
15 December 2008

Submitted To:
Networking and Information Technology
Research and Development

Submitted By:
Company: The Boeing Company
Two Gateway Centre
Aurora, CO 80011

Technical Contacts: Robert Bandstra, phone: 303-307-5897, fax: 303-307-3310
Company : Security First Corp.

Bill Goodwin, phone: 949-275-4947, fax: 949-858-7092

Boeing Contracting Contact: Terri Ferrari, phone: 719-572-8112, fax: 719-637-8535

Game-Changing Dimension

Compare the cyber situation to a football game. It is our goal to change the game to a very one-sided football game where only the good guys actually ever score.

Concept

Change the dimensions of the game. If we could make it impossible for the opposing team to either tackle or even see the football players on our team it would be a big advantage. If we could create a situation where the good guys are invisible and they have the ability to instantaneously slice, dice and reassemble the ball wherever they desire we could change the dimensions of the game. When the good guys are on offense, as soon as the ball is snapped, it is instantaneously split into pieces and distributed to several different players. When any good guy crosses the end zone – the remaining pieces are reassembled. The offense can score at will because the defense simply has no chance of guessing where the ball is on the playing field until it shows up in the end zone. For the defense to prevent the score, they would be required to see the invisible good guys and collect the pieces and have the ability to reassemble the ball to keep the offense from scoring. When the good guys are on defense, the invisible good guys can make the ball appear wherever they want, say in the end zone in the hands of the good guys for another score.

By analogy, if we could make all data in our cyber infrastructure completely invisible to our adversaries and they have no chance of interfering with the data or stopping our data in transit we would change the cyber game. Many of the denial of service attacks are targeted at our ability to transfer data from one point to another and if we can deny this attack channel it would be a game-changer. Our concept would give us the ability to encrypt the data at the bit level, split the encrypted data, and send it to its destination via multiple paths where it could be reconstructed and decrypted.

Vision

With this proposed strategy, the good guys could send their information at will; the bad guys could send it when the good guys allowed it. Its use would be ubiquitous and transparent to all users. We could circumvent many of today's attack channels through a method of securing the data that simply wasn't available to us until very recently.

Method

We know this vision is possible because the technology already exists to make data in motion unusable to a potential adversary. The SecureParser® technology has been validated via the commercial standard, Federal Information Processing Standard (FIPS) 140-2. Integrated with the Unisys Stealth Solutions for Networks, this technology establishes a tunneling protocol at the link layer, requiring work group keys in order to be recognized as a participant in the network. If there is no authorized work group key presented to the Stealth-protected network there is no response provided, making the data unusable and unrecognizable to an adversary. If network data is captured and analyzed, a disjointed set of encrypted and sliced IP packets will be seen instead of a coherent stream of packets. The data bits in the packet are cryptographically split in such a way that even if a network analyst captured all of the slices, he would still need the split session key to reassemble the bits and the encryption session key to decrypt the original message.
In order to make this vision real, we will need to break down a series of Information Assurance process barriers to gain acceptance of this technology. Today's encryption techniques rely on encryption of containers or packets of data, allowing capture of these packets and possible disruption of the data flow through a number of techniques. We will need to change this paradigm so that the security community will accept a revolutionary approach to encryption and data transmission. The SecureParser technology has been tested and certified to FIPS 140-2. Technical papers regarding its validity have been written and presented and it has been issued several patents. We have been conducting further research in our labs and successfully tested the concepts in several DoD exercises since 2004 including a recent successful performance assessment in Coalition Warrior Interoperability Demonstration (CWID) '08.

Dream Team

Our Dream Team consists of Boeing as a System Integrator, Security First Corp and Unisys for the technology solutions, and a motivated governmental security organization with the vision and passion to work though the issues of gaining approval for use. NIST may be a key player in establishing standards which others can use and implement this technology.

Who you are-- Georgia State University Research Foundation--We are a nonprofit corporation, created to support research activities of the University through obtaining and administration of grants and contracts for the performance of sponsored research.

Contact: Jim D. Flowers, Interim VP External Affairs, **Game-changing dimension--** Morph the Gameboard

Concept-- There is a clear need to provide better IT security on our nation's campuses. In today's world, cyberinfrastructure provides the critical enablers and benefits of an advanced knowledge society, but unfortunately also permits intrusive and high-risk attacks on that cyberinfrastructure, threatening our nation despite adequate externally facing physical defenses. Cyberinfrastructure is not just technology but a tightly inter-dependent set of people, process and technology components. We propose to change the gameboard by enabling every citizen of the knowledge society to understand their responsibility for security of processes and technology – essentially enlisting everyone in the effort to secure our cyberinfrastructure and eliminate vulnerabilities that are otherwise exploited to cause financial, physical or societal harm.

cyberPROMPT (Cyberinfrastructure – Personal Responsibility Of Management, Process and Technology) provides an integrated, enterprise resource approach to addressing security and risk management. **cyberPROMPT** emphasizes the *personal* engagement of individuals to protect the people, process and technology components of the national cybersecurity. The **cyberPROMPT Dashboard** is the means to educate, raise awareness and provide specific metrics of those threats.

Using higher education as a testing ground due to its extremely complex nature, we are proposing to create a model where the institutional risk profile is both comprehensive yet comprehendible by all campus personnel. Dashboard metrics will provide specific risk profile information that can be aggregated from the individual perspective to successive aggregations representing the departmental and organizational layers. Individuals, departmental and organizational leadership, governance boards, state or federal legislators or oversight groups, and other funding partners will have appropriate security policy metrics to track and monitor compliance objectives.

The proposed **cyberPROMPT Dashboard**, along with a risk/compliance/cost matrix will provide a continuous approach for clearer accountability and responsibility regarding the security of higher education infrastructures. It will also correlate regulatory compliance information into a reporting mechanism to track progress and take action on risk areas, while allowing the information to roll up for every level of governance from local to federal.

Vision-- cyberPROMPT will provide a single, unified solution ensuring compliance to statutory and regulatory policy obligations at the federal and state levels. Driving the discussion of security and policy compliance to stark economic and performance terms of all people, process and technology, as proposed by the **cyberPROMPT** solution, will focus administrators, staff and students, as well as lawmakers and other policy makers, on the costs and consequences of non-compliance to physical and cyber security obligations.

The **cyberPROMPT** approach:

- Captures specific compliance profile details
- Identifies areas of risk
- Sets appropriate governance controls and risk compliance thresholds
- Monitors, aggregates and reports on ongoing compliance

The following graphic depicts the key deliverables our solution will offer. This **cyberPROMPT Dashboard** will be ubiquitously deployed and displayed at all levels of an organization with settings that reflect personal, departmental and organizational views and aggregations.

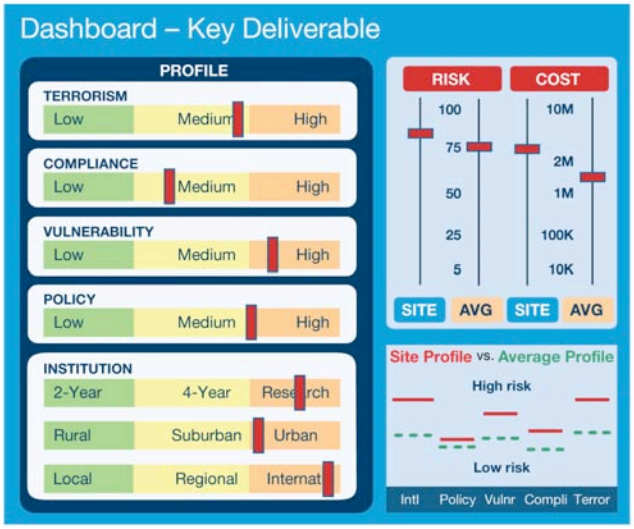

Method-- At Georgia State University we have an ongoing Information Technology (IT) Risk Management Working Group that includes information systems researchers, IT professionals, and risk management policy experts. This IT Risk Working Group has developed concepts and approaches to addressing IT security topics, including how to improve policy compliance through *proactive* monitoring and personal awareness as a way to stem intrusions and reduce costs of *reactive* responses. This working group has developed collaboration with an industry partner with expertise in the commercial practice of policy compliance and security technology. This collaboration has been productive in enabling iterative research and proactive discussion highlighting the importance of human factors in any successful IT Risk prevention.

The concept that has emerged is to provide a dashboard of key compliance and risk factors measured against actual practice. Active establishment of policy profiles, IT risk potentials and the pervasive feedback of actual performance metrics can create a proactive cultural change – with each individual becoming an engaged vector of vigilance. Such personal awareness and vigilance is expected to change the game board – no longer will our cyberinfrastructure be susceptible to intrusion because of vulnerabilities from individual failures of people, process and technology – the vulnerabilities so easily exploited with negative consequences today.

Given the complex nature of universities, we feel higher education is an excellent research and testing ground that can eventually extend results to other sectors of government and industry. The higher education venue represents an accurate statistical sampling of organizational complexity for a comprehensive preventive approach including human and technological mandated policies and best practices.

Dream team-- Georgia State University, Board of Regents of University System of Georgia (36 institutions and their Chief Information Security Officers), industry IT risk expertise (such as PricewaterhouseCoopers), Oracle information security practice or information systems, Six Sigma Management experts, State of Georgia Technology Authority, commercial vendors with policy and/or technical compliance solutions, State and Federal legislatures, Department of Homeland Security, FBI, U.S. Secret Service, Department of Education, individuals across Georgia State and the University System of Georgia.

Countering Game Changing in Next-Generation Reconnaissance and Coordination Attacks Using FireSealer and CovertDetector Defense Techniques

Who you are: (1) Ehab Al-Shaer, DePaul University, Chicago, (Director of Assurable Networking Research Center (www.arc.cs.depaul.edu), (2) Haining Wang (Colleague of William and Merry).

Game-changing dimensions: Change the board and stakes

Concept: As network attack detection and prevention systems become more accurate and effective; attackers will seek more stealthy and evasive attack vectors. Investigating novel attack strategies and tactics is crucial to develop future defense systems. In fact, it might be a pre-condition for advancing the network security research. Network attack reconnaissance and coordination are two important steps in launching a large-scale network attack. First, network reconnaissance allows the adversary to learn the victim network and service information so that malicious attacks can propagate safely and cause the maximum damage. This is usually performed by sending scanning probes from single or multiple sources. Second, as a next step after compromising machines, the adversary needs to establish a connection with those compromised machines for command-and-control or sharing information, which is essential to successfully accomplish the attackers' goals. We investigate new attacks enable an adversary to change the board of the game uncover network security configuration information and exploit covert channels as command-and-control mechanism for orchestrating attacks, in replacements of regular network scanning and traditional IRC (Internet Relay Chat) channels, respectively. On one hand, this proposal offers a significant complement to the cyber-security research community by investigating countermeasure techniques to such next-generation network reconnaissance and botnet coordination.

Vision: Our goal is counter this game change by another game change that will deceive and detect these attacks and stay in ahead of the game. We will explore effective countermeasures to thwart the next-generation network attack strategies. Leveraging our successful experience in firewalls and covert channels research we investigate novel network attack reconnaissance and coordination techniques, and seek the corresponding countermeasures. In the first part of this work, we will test the limit of our novel scan analysis and policy space navigation techniques, called FireCracker, to discover firewall rules remotely via intelligent active probing. This understanding will enable us to develop a countermeasure, called FireSealer, to be integrated in firewalls in order to perform attack detection and deception. In the second part of this project, we will investigate the implementing of an evasive timing-based covert channels, called NetPecker, and explore a series of entropy-based detection methods to capture these evasive covert channels in an accurately and timely manner.

Method: These attack techniques are game changing as they enable the adversary to navigate the network and propagate the attack quickly and safely. We expect the results of this project to enable transformative rethinking of the current network security defense mechanisms beyond traditional detection/prevention solutions. Our proposed techniques will be developed in FireCracker and NetPecker attack tools and FireSealer and CovertDetector countermeasure tools.

As investigating the attack is the first step toward an effective countermeasure, we will first investigate these already developed tools to understand the capabilities and limitations.

First, FireCracker goes beyond any scanning technique or tool by providing an intelligent analysis of probe results, and leads to the unveiling the security policy rules with high accuracy and limited probes, close to the theoretical bounds. We model the learnability of ACL firewall polices using Decision Lists and on-line learning models to establish a rigorous theoretical foundation for evaluating our technique. We use three novel algorithms for adaptive scanning: region growing, split-and-merge, and a hybrid approach to demonstrate new learning capabilities of this reconnaissance attack.

This attack on security configuration privacy reveals the entire ACL information, including (1) accessible and non-accessible destination IP addresses and ports, (2) policy structure like rule exceptions and default-deny space, and (3) topological information, which can be further used to discover defense-in-depth configuration. This can be devastating for IDS/IPS as such knowledge enables adversaries to be much more effective to initiate and propagate attacks. We will also investigate the limits of this technique in generating high stealthy scanning for fine-grain goal-oriented discovery, or distributed collaborative FireCracker agents.

In the countermeasure part, will adopt a hybrid approach of attack deception and detection using the concept of *shadow policy* as new technique to deceive attackers and distort the scanning results. Shadow policy will be integrated in the firewall to allow deceptive responses for scanning probes hitting deny address space. This will be implemented (called FiresSealer) as a software/hardware module that will be integrated as layer before the actual firewall filtering. FireSealer intercept packet going to the firewall and performs detection or deception techniques. We seek to investigated a hybrid approach that combines the advantages of detection and deception in order to minimize the attack and possibly uncover the attacker.

In the second part of the proposal (NetPecker), we propose a highly evasive timing-based covert channel for attack coordination. In the timing channels, we exploit the statistical properties of legitimate network traffic to create covert timing channels. We will design and develop a framework for building evasive covert timing channels. The framework includes four components, filter, analyzer, encoder, and transmitter. The filter profiles the legitimate traffic, and the analyzer fits the legitimate traffic behavior to a model. Then, based on the model, the encoder chooses the appropriate distribution functions from statistical tools and traffic generation libraries to create covert timing channels.

There is a contention between covert channel design and detection. To maintain the lead, researchers need to continue to improve detection methods and investigate new attacks. One goal of this proposal is to increase the understanding of more advanced covert channel design in order to develop accurate and robust countermeasure (called CovertDetector). Thus, we will propose a variety of entropy-based detection methods and reveal the relationship between detection speed and channel capacity.

Who you are – David Dittrich,[1] Affiliate Principal Scientist, Applied Physics Laboratory, University of Washington, and founding member/former-Officer of the Honeynet Project.

Game-changing dimension – Board, Rules, and Stakes.

Concept – Integration of advanced Honeypot and Honeynet capabilities for Active Response to network intrusions in production environments.

Vision – Since its inception in 2000, the Honeynet Project has proven a group of motivated individuals, with a little financial backing from a government entity, can produce a number of tools, papers, and ideas. It significantly raised awareness of security threats. Many companies today market *honeypot* or *honeynet* technologies, or use them in commercial products and services. But to-date, this effort has not resulted in a *game changing* effect in terms of defensive security operations. There are many reasons for this: there is no market, per se, for honeypot technologies alone; without better integration of features, existing tools are primarily *stand-alone* and serve very limited, mostly research-oriented goals; lastly, a general focus on basic research has a more limited benefit than might a more targeted *applied research* focus aimed at improving current security operations tasks. Missing from the honeynet/honeypot capabilities today are such elements as unified data formats, visualization capabilities,[2] decision support capabilities, integration of honeynets/honeypots with IDS/IPS/SIM/SEM systems and firewalls, secure data exchange mechanisms supporting use in confederated environments, etc. Even with a group of motivated individuals, though, efforts that are primarily voluntary in nature can only accomplish so much, so fast.

In this author's opinion, taking the current *honeypot* and *honeynet* tools and techniques to the next level of sophistication requires a much more significant, coherent, and well-funded effort than is available today. Greater support would enable development efforts to achieve a higher level of complexity, maintain a sustained focus over time, and foster a more professional research and development (R&D) methodology that can better integrate multiple components working in concert. By tying basic and applied research more closely with security operations, development of more advanced and complex tools can be accelerated and focused. Agile development methodologies will ensure requirements for efficient intrusion response are better met by allowing immediate testing and validation in the field. Leveraging the deep knowledge of these tools and how they are used will result in education and training capacity that will speed adoption in the field. Done creatively, a higher-level focus can change the board, the rules, *and* the stakes.

- Honeynet technologies can **change the board** by creating false hosts and networks. Networks could be designed to use Network Address Translation (NAT) in ways that support re-addressing of hosts, allowing them to appear to move from place to place in the network. Employing deep knowledge of how attackers and defenders interact in real intrusions, false network traffic and files could be planted such that only the attacker would find them, allowing a defender to control the attackers' perception of the state of the network.

[1] Web page: `http://staff.washington.edu/dittrich`. This submission is made as a private individual. All opinions expressed here are solely those of this author.

[2] Jed Haile developed the concept of a *Unified Data Analysis Framework* to serve these first two purposes.

- The **rules can be changed** with a deep understanding of attacker tools and tactics, and rules of engagement that support an increase in surveillance capability without violating the law or impacting the privacy rights of innocent third parties. For example: With the right knowledge and tools, it is possible to interact with the systems under the attackers' control; Carefully choosing how to interact, it would be possible to better attribute malicious acts and resources under an attacker's control; By replacing production systems, or turning them into honeypots, data can be injected to steer attackers towards non-production hosts.

- Possessing a deep understanding of attacker tools and activities, and the right tools, can enable defenders to *get inside the attackers' OODA Loop*.[3] The key is to be alert to attacks, agile in responding and reacting to them, and deliberate in countering an attack in ways that are not easily detectable by an adversary. Moving beyond static defenses and *wipe and reinstall* reactions, a more sophisticated defense can **change the stakes** for an attacker who will never know her actions are not invisible on the internet.

These goals can only be achieved in a reasonable amount of time with more funding to develop the higher-level coordinating tools and operational training that is necessary to pull existing techniques into a coherent system, and to extend and test them in production environments defending against today's advanced attack tools. Physical weapons are tested in the field to prove their efficacy and safety before being deployed on the battlefield. They are also designed to achieve specific performance goals and fit specific physical parameters in order to work with other weapons and logistic systems. They are not built by volunteers in their spare time, or by companies who only focus on short-term return on investment.

Method – The concepts put forward here were developed over more than a decade of security operations experience responding to distributed denial of service (DDoS) attacks, *botnets*, network and keyboard sniffers, various forms of cyber-crime, and the legal/ethic framework for response. Discussions and involvement in tool development as a member of the Honeynet Project also contributed. It is assumed that a more operations-focused effort to create a higher-level and more complex blending of Honeynet related technologies will accelerate current development projects, and spawn new ones, by solving discrete and urgent problems.

Dream team – As presented in a paper June 2008[4], the ideal team structure would combine: (a) experts in security and network operations from state and local government, research universities, and mid to large corporations; (b) academic research faculty and students (at undergraduate, graduate, and post-graduate levels); (c) full-time research staff to provide long-term focus, institutional memory, and serve as a bridge between other team members. They key is to focus on building tools that serve immediate *operational needs*, with a focus on *applied research* complementary to basic research and education. The result is a pipeline of learning, R&D, and public service, all working *in parallel*. This is not a model that fits the typical academic environment, nor does it fit the typical business model, but truly has the capacity to be a game changer.

[3] *OODA* stands for Observe, Orient, Decide and Act. This concept was developed by the late Air Force Col. John Boyd. For more on the OODA Loop, see http://en.wikipedia.org/wiki/OODA_Loop

[4] "On Developing Tomorrow's 'Cyber Warriors'," David Dittrich, in *Proceedings of the 12th Colloquium for Information Systems Security Education*, Dallas, Texas, June 2008

NITRD – Request for Input (RFI) – National Cyber Leap Year – December 2008

Who you are - Elastic Response Systems (ERS)
ERS is a cyber security system that focuses on multisource real-time situational awareness. The enduring long-term prospect of our game changing solution is the ability to drastically speed up multi-dimensional analysis of massive network datasets. Malicious software and other cyber threats invariably leave some trace evidence in their wake. However, the trace evidence is often lost in the volume of information the network and security devices generate. And, for new exploits the critical network data indicator is unknown, therefore the telltale sign of that intrusion is likely to be unmonitored. ERS's heritage, both the technology and the people, is from the clandestine side of intelligence. Systems with similar architecture to ERS have been deployed to find needle-in-the-haystack information in enormous datasets. At ERS, we use these successful and proven techniques to help solve the cyber security problem domain.

ERS is a compliment to traditional network defenses like firewalls and intrusion detection systems (IDS). In fact, firewall and IDS information are important datasets that ERS incorporates. ERS expertise is on early anomaly detection with an emphasis on an improvement in the quality of results. We believe early detection is important because defensive operations are constantly playing "catch up" to an ever-increasing onslaught of attacks that seem to always stay one step ahead. ERS helps anticipate and avoid threats by understanding the cyber situation, predicting adversarial actions and assessing potential effects. To protect information systems, ERS helps defeat threats with mechanisms such as adversary denial and deception.

Game changing dimension
The ERS application actively increases network security by **improving**:

Speed and Accuracy - There is a tremendous amount of information that is ignored by security systems: IDS \ Firewall log events, low granularity inspection (packet level, or lower in the network stack), netflow information, etc. This is due to the fact that existing systems do not use high-speed analytical databases or cloud computing. Often, the evidence of the intrusion was there, but was lost in the noise. The ERS platform has demonstrated a 50x-100x performance improvement in analyzing network data. Testing in the cloud has shown that adding an additional compute node offers a near linear improvement in processing.

Situational Awareness - Deployed systems are narrowly focused with no information sharing. For example, IDS and Firewalls are both dedicated to protecting the edge of internal networks, but very little automated information sharing occurs between the two (aside from manual configuration based on human discovery). ERS uses data-fusion techniques, which allows related contextual information to be combined regardless of the data's point of origin.

Efficiency - The signature\patch update cycle is inefficient. When an exploit is discovered, it must be reverse engineered, and then a patch or signature update must be applied (and often aren't applied). The ERS platform fosters self-discovery and do-it-yourself solution generation.

Collaboration - Cyber attackers act in concert and share tools and strategies that are effective. ERS allows defenders to do the same. The first ERS user to detect a threat can share the technique with others immediately. Threats should not have to be independently discovered as they are today, nor should organizational boundaries hold back the dissemination of critical information for effective decision-making. ERS allows the users of our platform to share discoveries, annotate graphs, and map threats. We are now living in the so-called web 2.0 age; security systems can utilize web 2.0 techniques. Large and small organizations can work together to offer a better defense than they can individually (ERS plans to utilize a network effect, every additional customer adds value to the rest, increasing visibility, adding novel / diverse solutions).

Visualization - The current generation of security products have user interfaces that use technology and design considerations that are over a decade old. In contrast, ERS uses a modern visualization approach that guides operators with only relevant information. Additionally, ERS user interfaces can be customized to focus on specific areas of interest.

Responsiveness - Current responses to malware attacks are too passive, generally just forensic analysis, shutting off network connectivity to the attacking subnet, or patching the system that was compromised. ERS fosters the ability to actively disrupt malicious attempts by: modifying incoming / outgoing packets, sending noise to the malicious command and control structure, making the attack economically painful. ERS's offensive capabilities are agile and can be scaled up or down within seconds as appropriate.

Subterfuge - ERS uses honeynets to virtualize entire networks of servers, applications, and clients. A few machines can be used to simulate thousands of attack points as decoys or for intelligence purposes. Virtual honeynets can act as a cyber firebreak mechanism to slow down attacks and reduce potential damage. The goal is to offer the capability to gather intelligence and inject uncertainty through strategic deception

Concept – What is the idea and why does it change the game?

ERS is a predictive analysis system that uses early anomaly-threat-failure correlation. ERS's unique intellectual property and methods addresses the speed and accuracy issues, the multi-dimensional analytics required and simplifies the complex relationship and schema management to provide the needed breakthrough for the next level of cyber defense.

Vision – Make us believe in your idea (What would the world look like if this were in place? How would people get it, use it? What makes you think this is possible? What needs to happen for this to become real? Which parts already exist; which parts need to be invented?)

If ERS was in place then early anomaly detection could be provided as SaaS (Software as a Service) on an extensible, agile, high performance platform. We believe this is possible because many of the application components are already operational in a cross section of industries in applications that have similar demands to the cyber defense problem domain. In addition, the unique intellectual property and methods that ERS brings to the table to integrate this together have been successfully prototyped and tested.

Please contact us if you are interested in learning more under circumstances of confidentiality that protects ERS' unique intellectual property and methds.

Enabling Game Changing in Security Planning and Configuration
Using Multi-Dimensional Quantitative Decision Support Sliders

Who you are: (1) Ehab Al-Shaer, DePaul University, Chicago, (Director of Assurable Networking Research Center (www.arc.cs.depaul.edu), (2) Carl Gunter, UIUC, Urbana-Champaign, IL (Director of Illinois Security Lab), (3) Ninghui Li, Purdue University, IN, (Member of CERIAS Center).

Game-changing dimensions: Change the stakes

Concept: Enterprises are often faced with the need to decide how they can attain adequate cyber-security for their operations. For most organizations security is not an end in itself but is instead risk mitigation, so ideally it should be assured with as little expense and inconvenience to the core mission as possible. Such decisions arise on a regular basis, but there is little that the security community has provided for *decision support* beyond a collection of technologies and a few general security principles. This means that system administrators and IT architects must proceed in a largely *ad hoc* manner to decide whether a security precaution makes a good tradeoff between risk mitigation, cost, and inconvenience. While there have been assorted efforts to provide security metrics as a decision support tool, these efforts have been hampered by the complexities of enterprise operations and the lack of meaningful raw data about risks from which to create plausible measures. Thus, what if we change the rules of security planning and administration to allow designing and/or configuring the security automatically based mathematically defined objectives (metrics) and constraints. This rule enables security configuration decision making to become not only accurate and robust but also provable, by eliminating design flaws, misconfigurations, guesswork and the dependency on few experts to secure our networks, as a result significantly raising the bar on intruders or inside attackers. Consider for example the sliders used to set security configurations for browsers on personal computers. The slider ranges from less to more privacy by exploiting a tradeoff with less versus more usability. It does this by automatically setting configuration rules for cookies and other browser parameters based on a pre-defined strategy. How this generalized to set up the security configuration for a whole enterprise information system? In such a case there would be sliders along three dimensions: usability, cost, and risk. Multi-Dimensional Security Configuration Support is a technique to use these tradeoffs and configuration targets to simplify the management of tradeoffs for enterprise configurations.

Vision: Effective decision support for security in enterprise systems has been a holy grail of the security research community. Providing a comparative approach for risk analysis that covers both hosts and networks and has meaningful applications would be an important step toward more ambitious aims. In addition, creating numerical metrics for *usability*, *cost*, and *risk* based on security configuration weakness that can be used to create optimal architecture and configuration and analyze tradeoffs the way enterprises address other missions though meaningful statistical control will be significant leap-ahead achievement in the field.

Our aim is to develop a framework and tools for Security Decision Support or SecDST, to automate decision support for enterprise security configuration. Using SecDST, security configuration decision making becomes not only accurate and *robust* but also *provable*. Users will be able to ask: what are different security zones (levels) exist in the network? how do they differ in term of risk and configuration? what could be reconfigured to qualify a system to communicate or migrate to a higher security zone safely? what is the impact of configuration hardening/portioning on usability and cost? what residual risk in the current architecture and configuration and how to cope with it? how to improve usability while enforcing minimum individual (host) or global (network) risk? and, what is the cost of applying security design principles: least-privilege, separation of prevailers and

privilege escalation on usability and budget?. Using the leap-ahead idea/technology, users will not struggle any more about what security architecture, strategy to adopt, what security boxes to deploy, what access control rules to use, etc. This will all be decided, justified and deployed automatically by SecDST. Alternative solutions/decision can be generated for further simulation and investigation. Security boxes will be viewed as connected organs forming one network body ready to accomplish the sated mission.

Method: Decision support is the concept of having a methodology, a model, a framework and ideally tools for deciding among these options based on a practical assessment of risk, cost, and convenience. Ideally we would like to have metrics that provide statistical-sound estimates and from which precise tradeoffs between risk and cost can be calculated. However, security is an especially slippery subject with respect to such ideals. Real-world data about attacks and their damages are notoriously hard to get. Attackers are generally humans who have a knack for doing the sneaky statistically improbable things. This makes it hard to obtain meaningful raw statistical data on which to base risk calculations. Even given such raw data, the task of assembling it into a practical overall assessment is challenged by the shear complexity of IT systems, which include diverse applications, hosts, middleware, network elements, and workflow procedures.

Given such complexities, it is almost impossible to address the full decision support problem for enterprises. In this project we would like to get a handle on it using a pair of assumptions. The first of these is that the overall problem can be addressed by modeling network and host configurations in a unified, mathematically precise framework. The second is that the decision support can be effectively provided largely through the use of partial orderings on such configurations. Both of these assumptions take on less than the full problem since there are considerations outside of access configurations and ultimately there must be at least some metrics to lay against risk and cost considerations, which are intrinsically numerical. However, the gap between current practice and a true scientific foundation is huge and progress on a meaningful treatment of configurations based on partial orders would be a significant step forward.

Our plan in the project is to create a *Distributed Vulnerability Surface (DVS)* which is precise enough and scalable enough to support a comparative optimization algorithm as a decision support tool (SecDST) in both top-down analysis that builds a new system with optimal design and bottom-up analysis that evolves an existing system in an optimal manner. We will build on prior work of the principal investigators that has demonstrated formal comparative models. On the one hand, we will create a precise concept of *Host Vulnerability Surface (HVS)* that can be used to make meaningful comparisons between the security offered by different configurations for a given operating system and even between different operating systems. On the other hand, we will design formal ways to model and analyze a *Network Access Surface (NAS)* that characterizes how and in what ways hosts can be accessed through middlebox elements such as firewalls, proxies, NATs, and VPN gateways. Next, after developing a integrated model for DVS using HVS and NAS we will create models for risk and security trade-offs, and then evaluate these models in real-life applications. In the third stage, we will explore optimization strategies that exploit the formal model to guide decision-making by calculating an anti-chain of locally optimal choices derived from a partial order induced by the DVS in order to reason about the proper network partitioning and service isolation architectures based on hosts configurations, potential risk impact, connectivity/security requirements, cost, and usability. In the fourth stage we will refine and validate this approach.

Dream Team: The PIs are planning to involve researchers from actuarial and business sciences.

Secure Enterprise Networks Consortium (SEN-C)
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Entity Passports

Who We Are

The Secure Enterprise Networks Consortium (SEN-C) is comprised of Accenture, Los Alamos National Laboratory, Sun Microsystems, and CA, Inc. SEN-C focuses on bringing leading skills together—from thought leadership and solution development to systems integration excellence. By collaborating with government, we seek to achieve outcomes that enable CNCI initiatives and improve our nation's security.

Game-Changing Dimension: Raise the Stakes

Concept

We will measure and assign trust for our users, hardware, software loads, and application processes. Using security/trust mechanisms we can build trusted distributed networks of collaborative enclaves. Using this trust infrastructure with a policy enforcement infrastructure, we would run un-trusted processes in controlled, physically separate sandboxes. Such environments would enable policy to cut off selected sandboxes and run only trusted components during times of attack. Our approach makes it harder for an adversary to insert or run malicious code in our systems. This concept receives support from the ongoing work of the Trusted Computing Group (TCG), the Internet Engineering Task Force (IETF), and the National Institute of Standards and Technology (NIST) as well as others in education and industry.

Our concept builds on existing industry technology in both hardware and software to support security trust infrastructure with component pieces to measure, register, and manage levels of security and trust across a distributed application. Some examples of available solutions that could integrate into this concept are commercial-off-the-shelf (COTS) efforts: protected random access memory (RAM) access, authenticated flash memory, virtualization, trusted platform modules, hardware full-disk encryption, network appliances, trusted execution technology, dedicated on-CPU encryption, biometrics, ARM trust zone, wireless, and mobile networks. Our concept would protect the complete security eco-system and coordinate trust between distributed enclaves without human intervention.

Vision

We would field networked enterprise solutions with distributed secure "passports" built into the infrastructure, managed by business rules and policies. "Passport" solutions would serve as standards and be able to build, expand, or contract trusted distributed networks for collaborative enclaves. Our concept starts with the ability to document and authenticate ("passport") end-point users, hardware, software, and firmware as part of the standards-based infrastructure. Such would support varied digital passport mechanisms to register and bind users, component pieces, and processes at startup. These steps would force un-trusted processes to run in a controlled, physically separate sandbox.

The passport architecture would include a policy oversight and enforcement mechanism. It also would include audit tools to monitor the life cycle support of the infrastructure (identity management, transaction logs, service level agreements). Component pieces of this passport architecture and infrastructure exist today, but integrated solutions can exploit them more fully.

Method

Our concept derives in part on analysis of industry trends and technology tipping points. It also derives from internal collaborative efforts such as discussions on security services and the Accenture Collaborative Innovation Solution (ACIS). The current industry trend is to build hardware solutions into infrastructure components. This includes use of the trusted platform module (TPM) for secure storage, key generation, secure system measurements, and security services management. We considered Intel integration of TPMs on virtualized chips, Intel's move to add a hardware crypto chip accessible to the CPU bus, and hardware full-disk encryption by several disk-drive manufacturers.

Technology tipping points include:

1. Rapid commoditizing of IT, with security/trust services built into hardware
2. De-perimeterizating, with impacts on central security
3. End-point security (network access control) standards and mechanisms are available
4. "Whack-a-Mole" Information Assurance (IA) (layered protection): necessary but expensive and insufficient
5. Commoditizing of encryption (including central key management)
6. Information constrained by policy and privacy legalities (expensive if misused)
7. Distributed Service-Oriented Architectures (SOA), able to scale for security/trust rapidly
8. Increase in business process modeling and policy measurement/enforcement points

We would establish a technical team to review COTS solutions for robustness and cutting edge. Sources would range from the SEN-C to university research and proposals for applicability. We would model and or integrate top solutions into a proof of concept test bed. We also would test scalability of the concept(s) as well as Certification & Accreditation.

Dream Team

- Trusted Computing Group (Intel Juniper Networks, Wave Systems) – fielded solutions/insight
- Select universities (technical cutting edge ideas)
- NIST/NSA (security standardization, oversight input)

RFI Name: RFI–3 – National Cyber Leap Year

The Enterprise Security API (ESAPI) Project (www.owasp.org/index.php/ESAPI)

RFI Focus Area: Morph the Gameboard

Submitter's Contact Information – Jeff Williams and Dave Wichers, Aspect Security and the OWASP Foundation, 9175 Guilford Road, Suite 300, Columbia, MD 21046, 301 804 4882

Summary of who you are – Jeff Williams and Dave Wichers founded Aspect Security, an industry leading consulting company focused exclusively on application security, and are leading contributors to the Open Web Application Security Project (OWASP) (www.owasp.org) which is an open source organization dedicated to helping the world address application security. Jeff and Dave established the OWASP Foundation, which is a 501c3 non-profit corporation to help manage OWASP's efforts. They serve on the OWASP Board of Directors and serve as OWASP's Chair and Conferences Chair, respectively. OWASP is an international community with 16,000+ participants, over 130 local chapters around the world, and 100's of application security research and tools projects, including the ESAPI project.

Concept – Building secure software is tremendously difficult today. Application software has simply become far too interconnected and far too complex for developers to write securely without a great deal of experience and skill. Over the years, we have identified a small set of root causes of these problems:

- **Missing Controls** – Many applications simply do not contain the appropriate controls to stop attacks. For example, there is no encryption control available to use on credit cards, making them easier for an attacker to extract.
- **Broken Controls** – If controls are present, they are often poorly designed or implemented. For example, an output escaping control may not properly escape the right characters, allowing injection attacks.
- **Misused Controls** – Even if controls are present and correct, developers still must use them properly and in all the right places. For example, a developer may forget to call the access control check in a critical business function, or validate form input using a standard validator, but with a weak validation pattern, allowing dangerous input.

These mistakes lead to the vast majority of software vulnerabilities affecting Cyberspace today. Eliminating these common flaws could save government and industry billions in software development and remediation costs and security losses. In addition, a more secure Cyberspace will allow organizations to innovate with confidence. To make progress in application security, we must make it easier for software developers to write secure code. Reactive approaches will simply never allow us to make progress against the ever increasing tide of interconnectivity and complexity.

We can radically simplify security for developers by ensuring that every developer in every environment has a complete set of trustworthy and easy to use security controls. We have initiated the OWASP Enterprise Security API (ESAPI) project with exactly these goals in mind. Organizations have adopted this approach with great success, saving time and money while

dramatically improving security. Standardizing security is a proven approach that has worked well for other security areas, such as encryption, where almost all developers now use NIST approved standard encryption components.

From a technical perspective, we propose to build a complete set of fundamental security building blocks, not elaborate security frameworks and systems. To achieve our goal, we propose establishing a product team to deliver a package of technical security controls for the most popular development environments. Each package will be made available under a free and open software license, and will include a library of foundational security controls, documentation, standards, and training materials.

Vision – In a world where developers are supported with strong and easy to use security controls, vulnerabilities are aberrations, not the norm. Static analysis tools can verify the proper use of these controls, instead of struggling to identify endless variations of vulnerabilities. We have already proven the approach works within large financial organizations and now it is time to take these successes to the world.

This proposal builds on thousands of hours of work that has already been done on ESAPI, which has been under serious development for over two years. The ESAPI project was developed at Aspect Security and donated to OWASP under the permissive BSD license. All of the materials, including the source code, documentation, and demonstration application are available from the OWASP website at http://www.owasp.org/index.php/ESAPI. ESAPI is featured at conferences such as the OWASP AppSec series, JavaOne, QCon, DHS SWA Forum, Jazoon, and more. Many companies are using or evaluating ESAPI for their organization, including Sun Microsystems, Oracle, Lockheed Martin, Infinite Campus, UBS, and numerous financial organizations.

Method – For the last 10 years, we have focused on making the developer's security job as simple as possible. We arrived at the conclusion that the best way to convey security knowledge was with an API, and have spent the last 2 years refining the API to be as simple and easy to use as possible. The conclusions we draw are fully supported within large development teams, we are assuming that we can bring these benefits to the rest of the software development community.

While the ESAPI project is thriving and making progress, the urgency of the application security challenge facing the country makes it critical for us to grow the project quickly. We cannot do this with volunteer effort. We propose to create an organization to manage the ESAPI project across all software environments. The project will establish an organization and become self-sustaining quickly by providing paid support to users of the ESAPI libraries.

Dream team – Existing ESAPI development team, the OWASP community, NIST, Rome Labs, Sun Microsystems, Microsoft, and the Federal Government.

Labeling of Proprietary Information – A key aspect of changing the game with ESAPI is that everything associated with ESAPI is free and open source. As such, there is nothing in this proposal that is proprietary, and no proprietary results will be produced by this effort.

From: RDavis
Sent: Friday, December 05, 2008 6:15 PM
To: Leapyear
Subject: Leap Year Idea

Who you are:

Russell Davis, D.Sc., CEO Femtosecond

Game-changing dimension:

People enter networks using the addressing of the addressing of the location they enter at. With IPv6 entering into the federal space concurrent with the OMB approach to limit internet access points to approximately 50, this will necessitate greater use of tunneling. With the preponderance of malware, software exclusive controls are no longer adequate.

Concept:

Every federal employee is required to have a Personal Identity Verification (PIV) card as part of HSPD-12. There is much discussion as to including an IPv6 address in the Global UID (GUID) field located in the Card Holder UID (CHUID). If this is done, the PIV cards can be combined to establish a solid VPN that is tied to strong identification and authentication (I&A). Without strong I&A, there is no basis for security. With bi-directional PIV VPN established using the PIV cards, man-in-the-middle attacks should shrink to insignificance.

Vision

The IPv6 address assigned to the card would allow cleaner auditing of a person's activities and provide a higher assurance for end-to-end security. In effect the PIV card becomes part of the VPN. Applied to all connections, internal and external, this will mitigate against vulnerabilities with the current used with PIV/CAC/VPN authentication that will likely become known shortly. We would in effect be authentication the network and the card holder.

Method

Software needs to integrate into the current PIV (or the PIV designed needs adjusting) to allow trusted network establishment. Ideally, a certified interface on the PIV card would allow one program to execute on a host regardless of the malware resident. By forcing the VPN to include part of the setup through the PIV card there is a class of currently unexploited vulnerabilities that will be closed before they can become problems. The malware attacks are not likely to abate any time soon.

Currently, software is used for the VPN with the PIV card used for I&A purposes. By integrating a low bandwidth component of the session to say continuously update a session key the server main network is sure the PIV card is still connected. In future renditions, this trusted channel could allow the connected location a safe location to peek at the local host without any malware risk. Dynamically adapting to new threat environments could be addressed by downloading a tight applet to the PIV card extended area through the trusted channel. In effect, providing a sensor as needed.

Dream team

PIV Card manufactures (those that write the applets), network vendor, and a standards team.

From: RDavis
Sent: Saturday, December 06, 2008 3:28 PM
To: Leapyear
Subject: Leap Year Submission 2

I am sending another submission (this is in addition to the previous submission).

Who you are:

Russell Davis, D.Sc., CEO Femtosecond

Game-changing dimension:

Malicious software (malware) is more prevalent than ever. When combined with zero-date attacks, existing systems are at a significant disadvantage. The concept presented can provide a tool in determining the state of a current host computer based on the binaries resident. It can answer questions such as: is the software a test version; what version and build are in use; what patches are installed; and is the binary corrupt. This is must know information

Concept:

In 1990, I wrote the paper “Software Checking with the Auditor’s Aid,” (Proceedings of the Sixth Annual Computer Security Applications Conference, IEEE, 1990). I believe the risk from rogue programs and human error is justification to revisit the concept. At the agency end, all approved binaries are hashed and the resulting configuration information entered into the database. At the client end the approach is to examine the binary objects, hash the objects, and use the resulting hash value as a database lookup to ascertain what version of the object exists.

Vision

Recently, 3,000 TWIC enrollments were lost when test software was inadvertently used to enroll new people. There are cases of significant losses when test software (or outdated software) is used in production environments. By using a controlled environment to hash the known good binary objects (executables, pictures, evidence, and the like) these values along with the configuration information can be maintained in a database. In my last Leap year suggestion, I posited adding a capability to the next generation PIV card to allow a secure channel between the card and the agency connected to. I suggested code could be moved to the card to run programs to examine the host environment. The suggestion presented here is one such application. Consider the DOD recently banned the use of USB memory devices to combat a work. In an environment with significant sneaker net, what is the cost of such an approach? Having the ability to determine if binaries are correct and not corrupted by malware, without using host software, is what this can provide.

Method

The approach would be to use a SHA-256 hash (instead of the 1990 paper’s polynomial checksum) and a next generation PIV card with expanded memory and processing power (Hashing is currently done in software, it needs to be done on card for the hardware assurance).at the agency level, a current database could be constructed along with a standard object schema (XML). Before critical applications are executed or binary objects distributed, the object can be hashed and the hash value can be checked against a database where the hash value is used as the index to the information on the binary object. If there is no match, this implies the object has been corrupted and should not be trusted. If there is a match, it will determine the version, patch version, build and other information regarding the object. When combined with a trusted hardware connection, a small PIV card loaded program could pull up objects,

hash them on the card, and pass the value back to the connected location. The Agency then could determine if the object is the correct version or if it was corrupted. Moreover, the card could be run as a background task thereby providing continuous monitoring of the local host.

Dream team

PIV Card manufactures (those that write the applets) and database developer.

A Whitepaper

Fixing the Internet:
A Security Solution

by Roger A. Grimes
InfoWorld Magazine security columnist

February 19, 2009

Version 2.0

Table of Contents

Abstract

The Problem

The Internet has a significant amount of malicious activities and security risk. Cybercrime is high reward coupled with low risk, and continuing to increase in scope and sophistication. Society is ever increasing its reliance on the Internet for everyday activities, including mission-critical applications that should not be utilized over untrusted networks like the Internet. The increasing (or even sustained) level of maliciousness is colliding with society's increasing dependence on the Internet for legitimate needs, creating an unacceptable risk. If left unaddressed, it could lead to a tipping point event; or at the very least is already causing a significant opportunity loss (e.g. money spent on inadequate defenses, slower computing performance, and people unwilling to conduct meaningful transactions over the Internet, etc.). Few currently available or advertised solutions (known to me) appear ready to change that fact over the next 5 to 10 years.

The Solution

Although many in society think we must learn to live with the current level of maliciousness, there are open standard solutions that can significantly minimize Internet maliciousness in the mid- and long-term. It will take a global, coordinated, community-based effort, along with accepting an increased managed (both centralized and private) control presence. This sort of infrastructure maturation has occurred throughout history, turning nomadic peoples and uncivilized societies into collaborative, productive centers where all citizens (participating or not) benefit.

Fixing the Internet's security problems will require a two-fold approach: a world-wide, global "dream" team of security experts working in concert to solve the systemic problems; and a global Internet security infrastructure solution that addresses and provides security protections holistically.

The solution(s) to fixing the Internet must:

- Use Open Standards
- Vendor and platform neutral
- Use an Open and Transparent Process
- Be Voluntary, Opt-In
- Be Performance Neutral
- Integrate with Legacy Systems
- Not Be Disruptive to Users and Services

As difficult and complex as this seems at first, it can be accomplished. Contrary to established, knowledgeable critics, this goal is readily achievable, today, using already existing open standards.

This paper will present the underlying security problems with the Internet, provide a global framework upon which to build stronger, longer lasting Internet security solutions, and ends by presenting two possible solutions that could fit within that framework. Readers are invited to critique, support, or reject.

[Note: The ideas and recommendations contained in this paper are solely the responsibility of Roger A. Grimes. No vendor or sponsor has been involved in the creation, editing, or approval of this whitepaper.]

Revision History

Date	Author	Version	Change reference
5/1/08	Roger A. Grimes	0.9	Initial draft, reviewed by key people
5/8/08	Roger A. Grimes	1.0	First draft released publically to the Internet and introduced in InfoWorld security column
5/12/08	Roger A. Grimes	1.1	Added FAQ section to answer most commonly asked questions
2/19/09	Roger A. Grimes	2.0	1. Updated various text portions to be more inclusive 2. Integrated WS-* open standards into text 3. Added additional authentication scenario

Author Bio

Roger A. Grimes

- 22-year anti-malware researcher (started fighting malware on Apple IIs, Commodores, Amigas, and DOS)
- Author of 4 books, co-author of 1 book, and contributing author of 3 books, and over 200 national magazine articles on computer security
- Expertise in securing Microsoft Windows, but also runs OpenBSD and various Linux distros at home
- InfoWorld security columnist for over 3 years
- Runs 8 honeypots tracking malware and hacker behavior
- Former Ultimate Hacking Expert instructor for Foundstone
- 9 years of experience as a penetration tester
    - Have compromised every company, site, and device I've been hired to pen test, in every instance in 1-hour or less (with one exception that took 3 hours)
- Frequently invited guest speaker or keynote at national security conferences
- Currently a security architect at a large OS and application vender
- Has participated in designing large scale and national identity metasystems

Introduction

Fact #1 – The Internet is full of malicious behavior which is not expected to decrease significantly over the next 5 to 10 years

The Internet is over two decades old[1], and unfortunately, rife with malware and malicious activities. Spam currently compromises over 70 - 90% of all email traffic[2]. Some experts estimate that malicious activities compromise 2-6% of all Internet traffic today[3]. Phishing attacks are becoming more targeted, and successfully compromising both casual home computer users and Fortune 100 executives[4] alike. Hundreds of thousands to millions of malicious bots control vulnerable computers[5] - conducting identity theft, adware redirection, distributed denial of service (DDoS) attacks, privacy invasions, corporate and government espionage, extortion, child pornography, and other malicious objectives. Malware is getting increasingly sophisticated (e.g. fast-fluxing[6], server-side polymorphism[7], generic one-offs that will never exist again) and propagated through legitimate (compromised) web sites[8]. People's bank accounts and stock portfolios have been emptied. Over a quarter of all U.S. adults had their financial identity information compromised in one year alone[9].

It is highly likely that millions of dollars are being stolen on the Internet every day[10], not to mention credit histories ruined, and legitimate operations and people's lives disrupted. We almost never catch the criminals. Of the major crimeware gangs (e.g. Russian Business Network[11], Rockphish[12], etc.) we have never identified, much less caught and prosecuted a single member. Internet crime is high yield and low risk. Current anti-malware defenses are being challenged like never before to accurately respond, and it is highly unlikely that most of the traditional solutions will significantly reduce malware over the mid- and long-term.

Fact #2 – Society Is becoming increasingly reliant on the Internet for basic and mission-critical services

At the same time, more and more of society's activities are moving online. What starts out as a public service convenience, turns into the primary way business is conducted, and leads to the only way business can be conducted. These include traditional commercial transactions (e.g. airplane tickets, concert tickets, paying bills, requesting services, etc.), as well newly evolving mission-critical applications that were never intended for an unsafe transport mediums. These include online healthcare records, software-as-a-service applications, university emergency alert systems, remote workers, online banking, television, and Voice-over-IP telephony.

Many mission-critical applications that the general public would never imagine were hosted on the Internet are. For example, the SQL Slammer[13] worm in 2003 compromised tens of thousands of unpatched SQL servers in under 10 minutes[14]. Hundreds of banks, including many of the world's largest banks were compromised and shutdown during that outbreak.

Since then, the incredibly appealing low price point of using the Internet as a VPN transportation pipe versus other alternatives has attracted more mission-critical applications to the Internet, not less. Many large-scale, city and regional supporting infrastructures are dependent on the Internet, and are being compromised over the Internet. Even the highest-risk, mission-critical applications (e.g. 911 response systems, public utilities, police systems, nuclear management facilties, etc.) that we are told aren't connected to the Internet, can easily be affected by Internet performance issues because they share strategic "choke points" along transmission lines and within telecom facilities.

Malware outbreaks affecting non-online public and private services is nothing new. Several past malware outbreaks (e.g. Iloveyou worm[15], Blaster worm[16], etc.) in the early 2000's affected integrated resources, causing telephone, pager, and cellphone disruptions, network news delays, and even the late delivery of basic goods and services. We are so inter-connected now with the Internet, that a single, widespread online attack will always impact the physical world. This fact isn't new. We have been lulled into a sense of complacency because no big Internet attacks have happened over the last few years. What is disturbing is the increased reliance on the Internet and what a new widespread disturbance would mean today, or in 5 years, or 10 years?

Fact #3 – Current Computer Defenses Are Inadequate

Current anti-malware defenses (e.g. antivirus, anti-spam, anti-spyware, firewalls, etc.) have proven mostly inadequate over the last twenty years and are ever decreasing in effectiveness. Whatever computer defenses vendors have come up with have been easily circumvented by faster reacting malicious hackers. Unfortunately, even though the current, traditional defenses are imperfect, end-users and business entities are forced to accept them (and their expense) because there is nothing else out there currently is better. Many vendors are trying to develop stronger, longer-lasting, harder-to-defeat defenses, but they are many years away from production release or require global adoption to work.

This brings up many important questions, including:

Why do we keep creating the same types of traditional point defenses against malicious computer activities when they so obviously don't work (with over two decades of largely imperfect anti-malware history as proof)?

How many people will not conduct legitimate business over the Internet (i.e. opportunity cost) today because of realistic, appropriate fear?

How much legitimate business does not happen over the Internet (i.e. opportunity cost) today because of realistic, appropriate fear?

How can we possibly be looking to put our personal medical records online[17] with the Internet's stability and security so much in question?

We are looking to improve the overall conditions of the world using the power of the Internet (e.g. One Laptop Per Child[18], etc.), and yet we are inviting these technologically new users into an inherently unsafe place.

It is these colliding realities, rampant maliciousness and increasing reliance on the Internet, which makes the improved security of the Internet of vital importance.

A Solution Framework

Solving the Internet's security problems will require a global, community effort.

No Single Vendor Solution Is the Answer

No single vendor solution can make the Internet more secure, for the following reasons:

- The substantial security problems of the Internet are not a "product" or "protocol" problem. The underlying problems are systemic and affect every vendor, every product, every protocol; and which if fixed, would make the other point solutions more successful (or unneeded).
- Most vendor security responses are acute, point-specific, in nature, not focusing on the underlying strategic problem; resulting in inefficient "whack-a-mole" defense solutions. When one hole is closed, the hacker attacks another weak link.
- Security defenses evolve slower than malicious attacker techniques.
- Every network packet is exposed to the same levels of scrutiny (or lack of scrutiny) and given the same speed of delivery regardless of the demonstrated historical trust of the originating gateway (e.g. a packet from a trusted partner is treated identically to a packet from an untrusted source). For examples, traffic from the Russian Business Network IP space travels around the Internet and to your Internet egress point at the same speed as a long-time, trusted business partner or loved one.
- Most global security events (e.g. large bot DDoS attack, phishing and spam floods) are only noticed by a small set of selected vendors. If the data was shared faster, globally with everyone, the benefit would be greater.
- No globally accepted security initiative addresses the systemic problems.
- No global Internet servers or services address security broadly.
- No Internet global body has a charter focused solely on malicious prevention, although we have dozens covering response.
- Few currently proposed solutions (that I am aware of, with one exception covered below) will make a significant decrease in malicious attacks over the next 5 to 10 years.
- End-user education is highly overvalued (many end-users are not technologically sophisticated enough to recognize malicious events). We need to develop solutions that minimize asking the end-user to make trust decisions.
- There is no accountability for poor security or poor coding (e.g. some of the poorest security performers are gaining market share, and origination sites with consistently poor trust records can access destination resources at the same quality of service as proven trustworthy providers).

Real Solutions

Fixing the Internet's security problems will require a two-fold approach:

- A world-wide, global "dream" team of tactical security experts working in concert to design systems and protocols to solve the systemic problems
- A global Internet security infrastructure service (likened to DNS) that addresses and provides holistic security protections

Global Security Dream Team

The Internet is full of very bright, sometimes popular and accomplished, sometimes relatively known, security experts with good solutions to the Internet security problems. Unfortunately, their good ideas languish inside of their respective employers (due to competing self-interests), on Internet discussion lists only known to the lists participants, or in research papers left largely unread.

We have many times in the past, when faced with a seemingly insurmountable problem, gathered together the world's best minds in their respective disciplines and solved the unsolvable. Examples abound: clean water, vaccines, computers, nuclear energy, outer space programs, and ending wars. This point in time requires that we build another team dedicated to significantly improving Internet security.

Selected top vendors (open source and commercial) and independent security experts should be brought together for a period of 6 months to 2 years to debate the problems of the Internet and recommend strategic and tactical solutions. An open and transparent consortium should be created to facilitate these expert meetings, and participants should agree to work toward common, agreed upon objectives.

Note: Many existing national commissions already exist and have had the ear of high-level politicians, even the United States President. Unfortunately, all of the prior committees have been heavy on executive and strategic thinkers, but missing senior security tactical thinkers and technicians. While strategic and political committees are absolutely necessary, there has yet to be one designated to design and deliver actual solutions.

Team Makeup and Responsibilities

There should be a different, independent team created for each critical core component, which naturally seems to lie somewhat along the OSI model's definitions (e.g. Physical, Logical, Network, Session, Application, etc.). To that idea we should add other shared necessary components, such as Cryptography, Identity, ISP, IANA, Legal, Global Considerations, Privacy, Open source, End-User, etc.

There should be a larger, more strategic Executive committee team that helps coordinate and integrate the various lower component teams and provides strategic direction to each component committee. There should be a team leader (with only 1 vote and chosen by a majority of participants (each also with 1 vote)) of each component committee. Each component committee will be responsible for developing the tactical ideas to be passed along to the technical participants under each component. The technical participants will be responsible for coming up with technical solutions and standards to meet the tactical assignments. The technical (and end-user, public, and shared committees) will also be tasked with providing technical guidance to the higher committees (i.e. can the tactic be realistically implemented).

The figure below shows the basic consortium design along with the component committees.

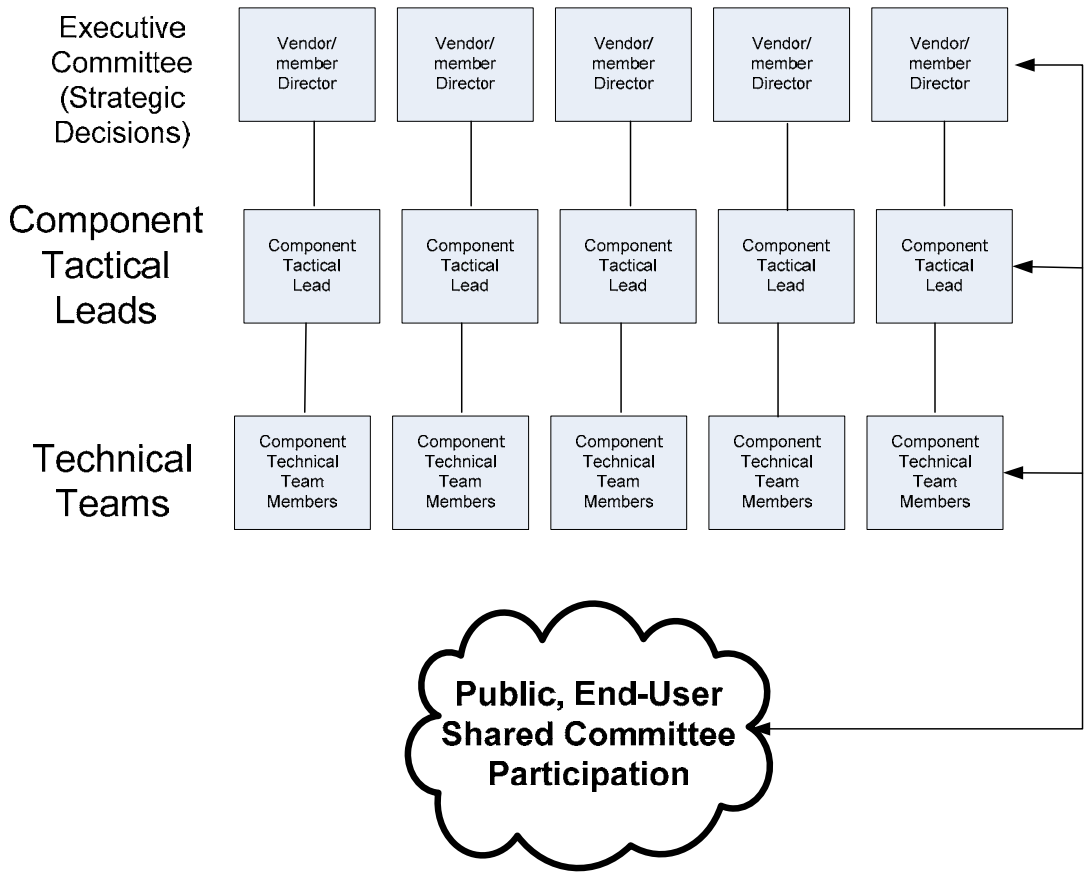

How Big of a Team?

Although any number I pick now is arbitrary, 10-20 participating members on the Executive committee and 10-20 members on each component committee seems a realistic starting number. Invite 5-10 vendor leaders into each component committee, and another 5-10 independent field experts. Initial component members (no more than half of the total members) could be chosen by the Executive committee, and additional members voted on by the original members by majority rule.

Example Vendor Participant Members

Participating vendors would have to dedicate and fund multiple original committee members, including:

- Senior Management (responsible for selecting Execute Director representative, non-voting)
- Execute Director (voting member, responsible for coordinating member's response and vote)
- Assistant to Director (logistics, minutes recording, etc.)
- Technical Lead for each tactical component the vendor is involved with
- Senior Technical Staff under each technical lead (although who participates here can vary according to need)

Thereafter, the community-based consortium would require ongoing, permanent (but revolving) members to address standard updates, either to address improvements, additional coverage, or to respond to vulnerabilities.

The Hardest Part

The hardest part of solving the Internet's security problems is not generating the technical security solutions. If you can solve the hardest part, the technical solutions will come easily. Getting vendors and independent experts to dedicate 6 months to 2 years of their life to a single, societal goal is among, if not the hardest part, of solving the Internet's security problems. Natural sustainability (usually revenue or earnings) dictates that members work on their own self interests to maximize revenue. How do we get vendors and individuals to give up potential, immediately recognizable revenue gains to concentrate on the greater good, which ultimately benefits themselves and the commons? I'm not sure. I'm hoping that if enough end-user interest is generated by this idea, governments will call upon citizens to do their civic duty and vendors will volunteer to participate as much is possible for a reasonable period of time. Or perhaps, a grant of some type could be awarded to offset the revenue reduction to benefit the greater good. This is a tough issue to solve.

Transparent and Open Submissions

It is important that every single consortium word, decision, and result would be posted on the Internet to be as transparent as possible. In order to get a world-wide community solution we need the community's trust. It must be supported by open source and commercial concerns.

One Member, One Vote, Public Participation

Each participating member would be given one, equal vote on all proposals, and additional members could only be added by majority vote. In order for this idea to be successful we must guarantee to participants (many of whom will be hesitant otherwise) that this is not vendor-specific dominated initiative. Multiple public and private participants can be present and engage in debate, but only one vote is allowed in a particular component committee.

The public will be invited to participate at multiple points and their comments and submissions reviewed (by a sub- or full-committee as each component committee deems appropriate); although in order for any idea or issue to be voted upon it must be brought into full committee by a voting member. All proposals must receive an up or down vote, and majority rules.

Why the 6 months to 2 Years Timeline?

I believe that time is of the essence, not because we don't have time necessarily, but we need to use time as a tool to minimize members debating details to death and getting lost in the weeds, and forgetting the overall goal. I propose the following time schedule:

- 2-3 months to organize the effort
- 6 months for the teams to meet and discuss possible solutions
- 6 months for public review and discussion
- 6 months for technical review and decisions, and the final vote on document 1.0
- 6 months to document decisions and release new security proposals to all vendors

Additionally, one of the primary problems why we have not solved the Internet's security problems is the relative speed at which malicious hackers move as compared to the security problem solvers. By proving that we can move quickly within a naturally bureaucratic system, it will provide some measurable disincentive to future malware writers. Plus we can use the lessons learned to move even more rapidly in the future, when responding to new challenges.

If we created a global consortium to concentrate on resolving the Internet's security problems, two years from now we would have new global, community supported Internet security standards, which

could be implemented by participating vendors and individuals. At the end of two years, vendors and individuals could then take the time they need to implement the standards in their own way (or reject them and not participate directly). Legacy devices and software must be able co-exist and function with the newer devices. If done appropriately, no one is deprived of legitimate service, except the malicious hackers.

Other Solution Ideas:

- What the committees can't agree on will be tabled or split (for just that issue) so we can get the overall, strategic and tactical goals met. Let's vote on what we agree on.
- Solutions must be opt-in, with more "carrot" and little "stick". People choosing not to opt-in are only disadvantaged by not directly participating in better security.
- Solution must address all computer platforms (PCs, PDAs, cell phones, media players, TVs, etc.)
- Any response to hacker vulnerabilities against the new standards must be rapid. We want to demoralize the current and potential hackers, and show that the defenders can respond as quickly as the bad guys.
- There are human, process, and education elements to consider.
- We need strong global participation for global acceptance.
- Optional idea: Funding for the long-term community consortium members can be collected through some minor (voluntary) monthly or manual minimal fee collected at the Internet's egress points.

Challenges/Questions

- The normal issues associated with global, strategic direction without explicit authority.
- How to create enough self-interest to motivate major vendors and other needed participants to meet?
- How to be quickly responsive to changing malware tactics…must be built into process.
- Balkanization of committees, objectives, or protocols (that's why we will table and split when needed), but majority rules; let good ideas emerge, even if differing.
- Should this be an entirely new committee or rolled into some existing body (e.g. IANA, IETF, CERT, Trusted Computing Group, etc.)?

Global Internet Security Infrastructure Service

The Internet's major security problems cannot be solved by a single vendor or a vendor-specific solution. Whatever the solutions are coming from the above mentioned Internet security consortium, the outcomes will be global and require global, coordinate participation (in most cases). The Internet lacks any service or infrastructure dedicated to coordinating/advertising/publishing security services (again, think DNS). Accordingly, I propose building a global Internet infrastructure service to provide coordination, advertising, and publication of the various global security initiatives.

This idea is similar to an imagined cross between the global DNS infrastructure , a web services' Universal Description Discovery and Integration, UDDI[19] service, and the Trusted Computing Group's new IF-MAP standard, applied globally. The diagram below re-summarizes the concept.

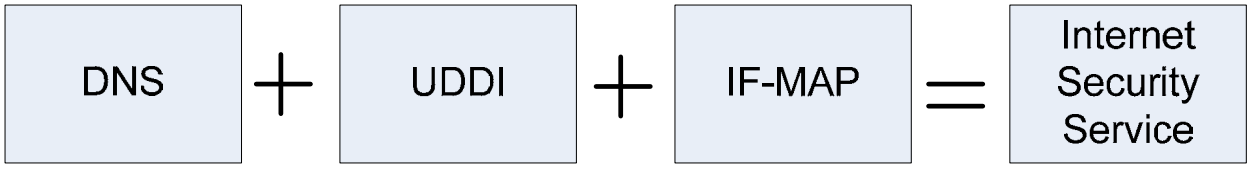

The new global Internet security infrastructure service should DNS-like in that there would be fault-tolerant, distributed "root" servers dedicated to directing querying clients to the appropriate security service server(s). It would be UDDI-like in that each participating global, sub-root server would to serve up IP addresses to the corresponding needed security services (and to advertise and publish such services). It would be IF-MAP-like in that the existing sub-root servers would allow participating members to report and respond in a global, holistic, multi-service manner.

If you are not familiar with IF-MAP, in a nutshell, the new Trusted Computing Group's (www.trustedcomputingroup.org) IF-MAP standard (https://www.trustedcomputinggroup.org/specs/TNC/IFMAP_FAQ_april_28.pdf) allows participating devices to report security events and receive notifications from other security devices to be able to respond in a coordinated fashion.

For example, if a firewall notes an unauthorized outbound stream that it recognizes as a bot spam stream, the firewall can contact the IF-MAP service, which can then contact a policy server that contacts another service that shunts the offending device off the network. The Internet security service would be similar to IF-MAP in that it would allow the coordination (i.e. reporting, advertising, direction, and response) of multiple disparate services, but be global in scale. Currently, the IF-MAP standard focuses on coordination within a single control domain. The Internet security service would be available for global coordination and direction, and should be integrated with private IF-MAP devices. The global Internet security service would have to be resilient, fault-tolerant, and cryptographically sound.

The following diagram gives an example of what the infrastructure might look like:

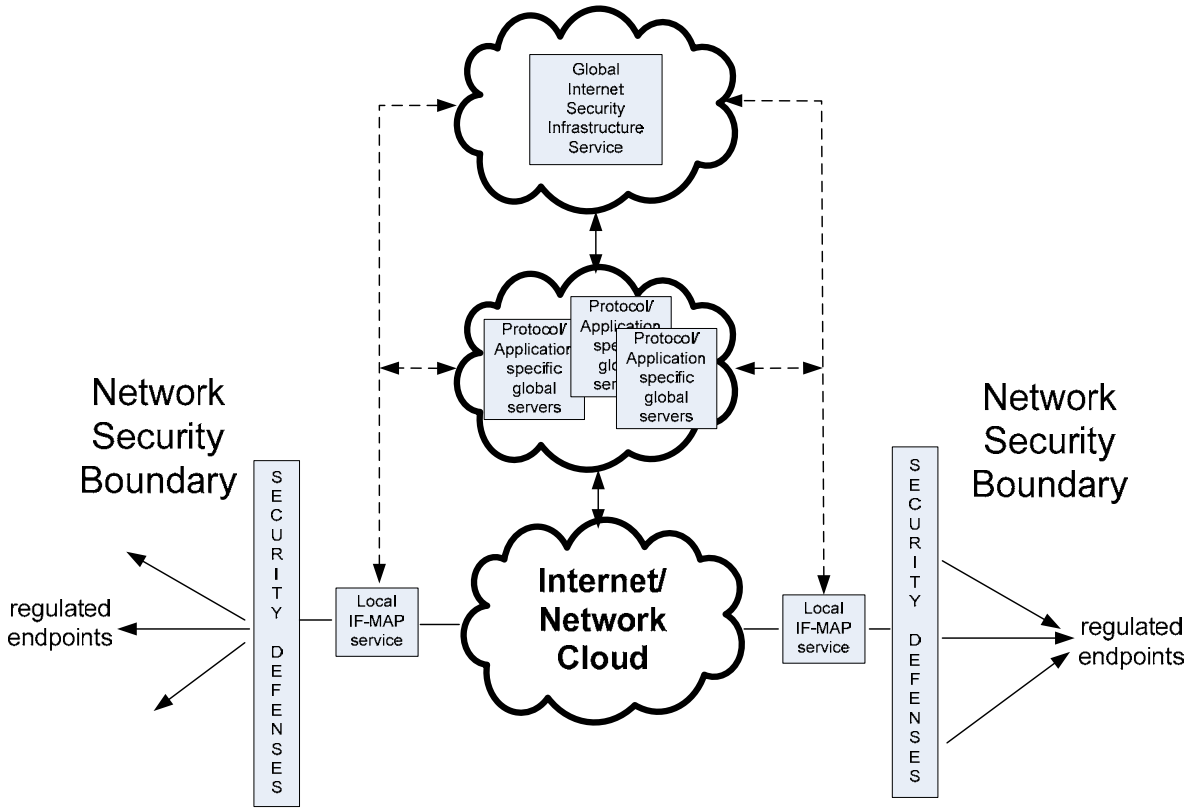

The local IF-MAP services could take advantage of the global Internet security service, and be better able to respond (and report) threats. This would allow local security domains to respond quicker to threats noted by other partners, and be able to report local threats to other partners for their benefit. This sort of cooperative coordination has so far only realized in commercial, private, and more narrowly-focused public projects.

For example, several large anti-malware vendors (e.g. Symantec, Microsoft, McAfee, etc.) are able to capture and respond to large global threats because they have millions of participating nodes collecting and reporting statistics. Several open source and commercial anti-malware black lists have been around and used publicly for over a decade, albeit limited to a few uses (e.g. anti-spam, anti-phishing, etc.). There are several private groups, often led by anti-malware researchers, which collect and disseminate information to its members. Other groups, like SANS (www.dshield.org) collect limited information from participating members, and share the collected information publicly. These are all laudable goals, but suffer from limited membership or focus. A global Internet security service could collect information on a broader scope and its wider information used by more people. If global threat information was publicly communicated instantly, each participating entity, and the Internet, in general, would greatly benefit. Malicious hackers depend on the lack of global coordination to be successful. Let's take that strategic advantage way.

Example Scenarios Benefitting from a Global Internet Security Infrastructure Service

- Your network or web server comes under attack by a DDoS attack. Your local IF-MAP security device could connect to a root Internet security server and get directed to one or more services to allow an efficient response and defense to the attack. Your network could get subscribed on-

the-fly to an anti-DDoS service, fire up additional availability resources on new IP spaces, or lead all the other participating networks into shunting off the offending bot-infected computers.

- Your company participates in a global whitelist/blacklist of IP addresses. Your company's whitelist/blacklist servers/service could contact the global root servers to get instantaneous updates of the Russian Business Networks' changing IP address space.
- Your anti-spam device or anti-phishing filter can learn instantly when a massive new spam or phishing attack occurs instead of waiting for a vendor update or allowing only the already existing global email servicers to learn about the attack.
- Supposed a MySQL-based Slammer type, zero-day, worm gets launched that can be successful against all existing, contactable MySQL servers on the Internet. Your firewall could be notified of the zero day attack and shut down the port until a better remedy is provided.

Regarding the last example. The original MS-SQL Slammer worm went off in the early morning weekend hours (in the United States). The majority of compromised servers occurred in under 10 minutes. Not only did the attack start and essentially end in under 10 minutes, but it was six to eight hours before the vast majority of the waking up Internet users (in the U.S.) learned of the attack, and began to respond. It seems unusually risky that we do not have devices ready to automatically respond to instantaneous global threats and are still relying on humans (which on average are asleep one-third of the time) to implement reactive solutions. It would be better if we had widespread, global early warning systems with rule-triggered IF-MAP devices to handle the initial response.

The following diagram shows some example coordinated services and proposed connection points.

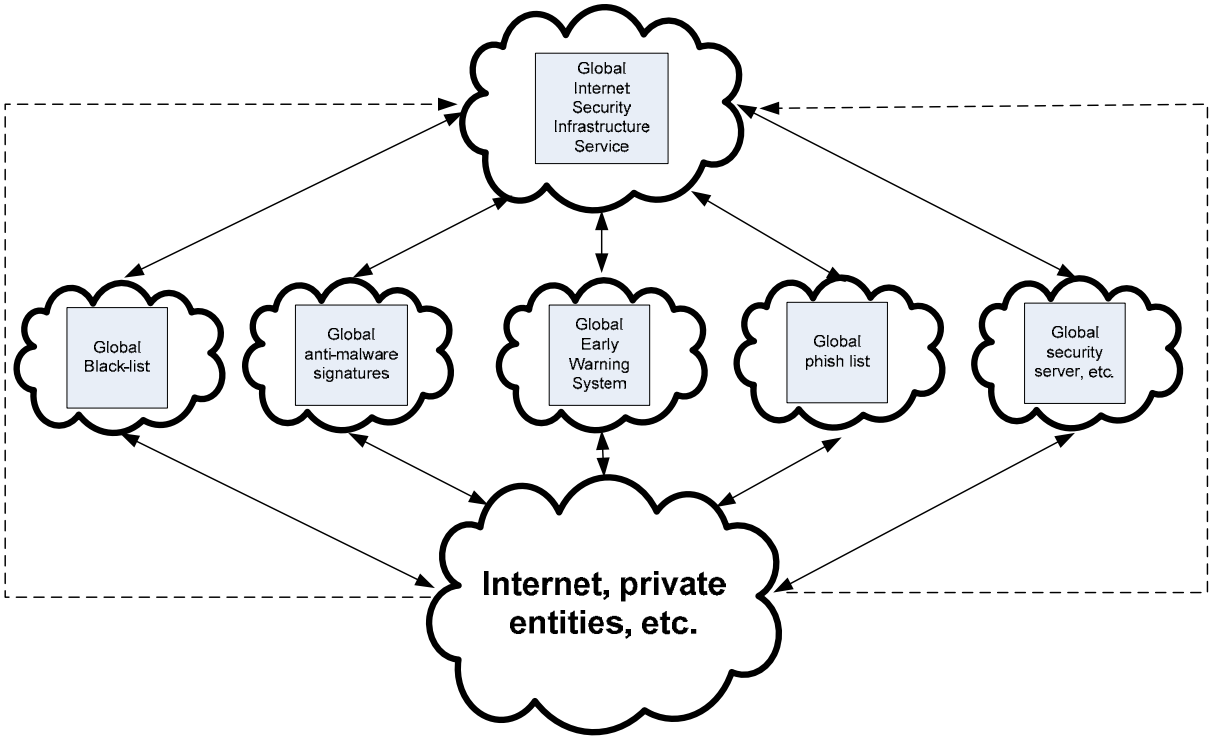

We know we need global, coordinated security early warning and responses, but we do not have a global security infrastructure to support this need. This type of solution would be in the end-users self-interests because it provides better, holistic solutions, and provide lower cost and better performance

as Internet maliciousness decreases. It would be in the vendors self-interests because they get to develop a new stream of products and defense responses they haven't even considered, yielding new customers and better solutions in an otherwise staid space.

Possible Solution #1– Replace Default Anonymity with Pervasive Identity and Integrity

by Roger A. Grimes

Abstract

The major underlying Internet security issue that is preventing a significant reduction in malicious behavior is the pervasiveness of default anonymity on the Internet. Because we can't identify malicious hackers with a high degree of confidence we cannot identify or hold them accountable. Internet crime is high-yield and low risk. If the Internet's model of default anonymity was replaced with default identity and integrity, the amount of maliciousness would significantly decrease.

I propose that every participating Internet component, hardware and software, be modified to provide increased identity and integrity assurance. Participating devices and users would provide improved levels of trust and be treated appropriately. All participating network traffic would be cryptographically tagged with a "trust level", which could be evaluated and acted upon accordingly. Each participating security domain would be responsible for assuring the trust and labeling of its egress traffic and responsible for acting upon tagged ingress traffic (and be held accountable for its attestations).

A security domain gateway device (called a " trust gateway") would perform the necessary trust labeling and evaluation. Every component (e.g. hardware, OS, network devices and pathway, identity, etc.) would end up being evaluated and assigned a numerical trust rating. Levels of trust, and how to obtain them, would be determined by a consortium of computer security experts, and published in an open, transparent manner.

Increased assurance levels would result in higher trust level ratings. For example, a user logging on with non-complex, short password would result in a lower trust rating than a user using two-factor authentication. Identity of participating nodes and users must be assured, but does not necessarily mean that each unique identity translates to a specific entity or user (i.e. user's real name).

All participating traffic would be encrypted and authenticated from origination to destination trust gateway end-points. Participating nodes and network traffic, demonstrating increased reliance and assurance, would undergo less inspection and given an increased quality of service. Nodes wishing not to participate would still be accepted and evaluated exactly as they are today, albeit with a lesser quality of service as compared to participating nodes. The solution would be vendor independent, transparent, open, voluntary opt-in, performance neutral, with least service and end-user interruption as possible, and driven by user and vendor self-interests.

Note1: The ideas and recommendations contained in this paper are solely the responsibility of Roger A. Grimes. No vendor or sponsor has been involved in the creation, editing, or approval of this whitepaper.

Note2: I accept that this particular solution will not make everyone happy. I'm bound to have critics that strongly disagree with it. However, it is my hope that this part of the whitepaper (*Possible Solution#1 – Replace Default Anonymity with Pervasive Identity and Integrity*) stands alone and is evaluated separately from the solution framework provided in the first part of the paper above.

What's Wrong With The Internet?

To understand how to improve Internet security you have to ask why things are as bad as they are. Most people when asked this question respond with problems (and solutions) that are pain point-specific (e.g. anti-virus technologies aren't accurate enough, we have to patch too often, software is always insecure, the end-user is the problem, etc.), but don't always focus on the strategic, underlying issues.

Security issues and solutions can be broken down into the CIA triad components: Confidentiality, Integrity, and Availability. All are important. But if you ask which one, if solved, would significantly decrease Internet maliciousness? It is without a doubt, Integrity. If we could confirm that the email is from who it says it is, we would end all spam and phishing. If we could confirm that the offered security patch is really from the vendor who says it is, we would not install malware. If we could identify the origination of released malware, we could track the hackers. If we could identify malicious hackers, we could arrest them. In fact, I can't think of a single significant, remaining Internet problem that isn't an identity or integrity issue.

Most of the Internet's infrastructure and its components run with default anonymity making it difficult to hold the majority of malicious participants accountable. Why do malicious hackers hack? Because they can do it with near impunity. Without greatly improved identity, integrity, and accountability, there can be no significant reduction in malicious Internet activity.

Solution

Build into the Internet pervasive, reliable, trustworthy identification and integrity into participating components and transactions, from source to destination. This will require a world-wide, community-based approach and the strengthening of every core component (called "trust components") along the OSI model, including:

- Hardware
- OS Boot Process and Loading
- Device and User Identity
- Network Stack and Protocols
- Applications
- Network Transmission Devices and Packets
- Communication Sessions

And it must be accomplished vendor independent, voluntary, opt-in, performance neutral, and with least service and end-user interruption as possible. An accepted solution must integrate legacy components while providing (voluntary) compelling reasons for consumers, vendors, and service providers to adopt solution-compatible components. I propose doing this by making each Internet egress network responsible and accountable for the security and trust of the endpoints in their network.

This applies to corporate environments, as well as, ISPs being responsible for the security of their end-user clients (to a variable degree). Each egress network access point would be known as a "trust network", and the management and technical teams responsible and accountable for implementing improved security trust mechanisms (e.g. egress filtering, two-factor authentication, anti-malware, secure coding practices, etc.).

A world-wide community consortium of computer security experts would transparently decide what levels of trust are assigned to the various trust components and how various trust networks earn increasing levels of trust. Egress points with poorly demonstrated levels of security will be given a low trust rating, and that rating known to all participants (e.g. world-wide trust rating list). This should encourage trust networks to improve their security to be rated higher, and at the same time hold accountable questionable networks (e.g. Russian Business Network's malicious IP space).

These global trust ratings would be sharable and available to each communicating trust network. Each receiving trust network can decide how to treat incoming traffic based on the originator's trust rating; and even provide custom trust ratings to trusted private trading partners (regardless of the packet's tagged trust). Traffic with higher ratings of trust should be inspected less and be delivered faster to end-points.

Trust Gateways

Each trust gateway should implement a trust gateway device (which can be a separate component or integrated into other egress/ingress point devices and software (e.g. ISA server). The trust gateway device is responsible for tagging egress traffic with a community decided upon trust rating, and appropriately handling (and handing off) incoming traffic based upon the trust rating with which it is marked.

Community-Based Trust Rating Server

A participating Trust Network's trust will be registered on a community-based Trust Rating Server. Trust gateways can periodically query the Trust Rating Server and download the trust ratings for various trust networks. This way we can update trust ratings and track when the bad guy networks move, and communicate that move to all participants. All network ratings, good and bad, will be readily available for inspection. We will have to build a process for rating and updating, efficiently. If a trust rating cannot be updated quickly and with integrity, the whole system breaks down. At first this may seem like an alien idea, but we have many such community-based servers, but none focusing on holistic trust.

How Trust Is Determined?

Every defined trust component (e.g. hardware, boot, OS, identity, software, network, etc.) contributes to the overall trust rating of the packet leaving or traversing a trust gateway device. Each trust component receives its own trust rating, and culmination of all trust component ratings leads to an overall packet trust rating. Each participating network transmission device is also assigned a trust rating, and the transmission path of each network packet from source to destination adds an additional network pathway trust rating. Thus packets sent along trusted network pathways are given higher levels of trust than those traversing lesser secured routers and devices.

For example, one-factor identity gets a lower rating as compared to two-factor, and so on. There will be a network device rating. Network routers without source routing enabled, fully patched, with strong passwords, without known vulnerable scripts, etc. will be given a strong rating.

The diagram below shows a logical representation of two packets with trust ratings, showing their individual component trust rankings and the overall packet trust ranking.

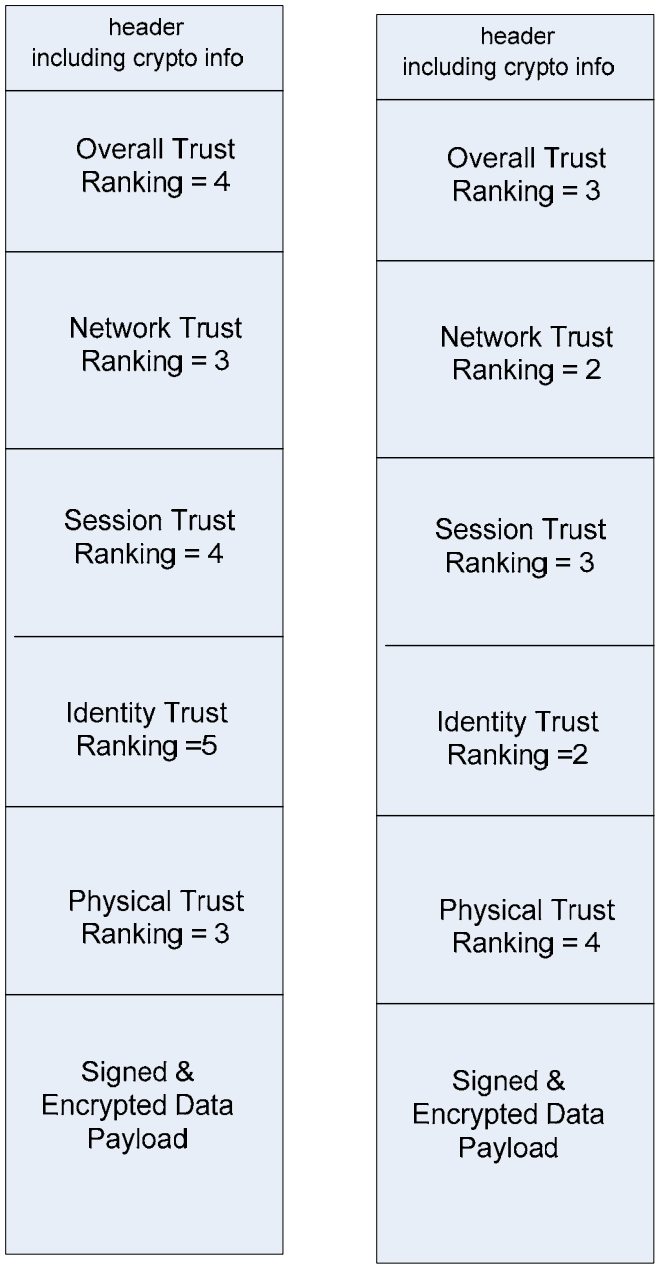

How a component is ranked will be determined by community-based decisions, and documented in a transparent, public-accessible document. It will be a common-criteria sort of document, but based on real, implemented security best practices. Most other common-criteria sort of guides are flawed because they end up being paper exercises and don't translate to real improvements in security. This document will be immediately usable and help all users and networks to improve security.

All component ratings end up generating the packet's overall trust rating, and both component and overall trust ratings are built into the network protocol for inspection by intermediate and final destination trust devices.

Ingress trust devices can treat network traffic differently depending on individual component trust rankings or rely solely on the packet's cumulative rating. This gives flexibility to ingress points that require different security policies (e.g. an online bank requires higher identity ranking while a network peering partner requires higher levels of network trust). Legacy devices will ignore the trust component, but pass along the trust components unmodified.

Trust ratings will be tagged into the traffic, and securely protected against unauthorized modification. Ingress trust gateways can rely upon the packet's attestation level and/or query a global community trust rating server to confirm the incoming security domain's historical trust ranking. If a particular security domain ends up being recognized as a poor trust decision, then the global trust rating servers can deliver that message to the ingress gateway device.

My idea is summarized in the diagram below.

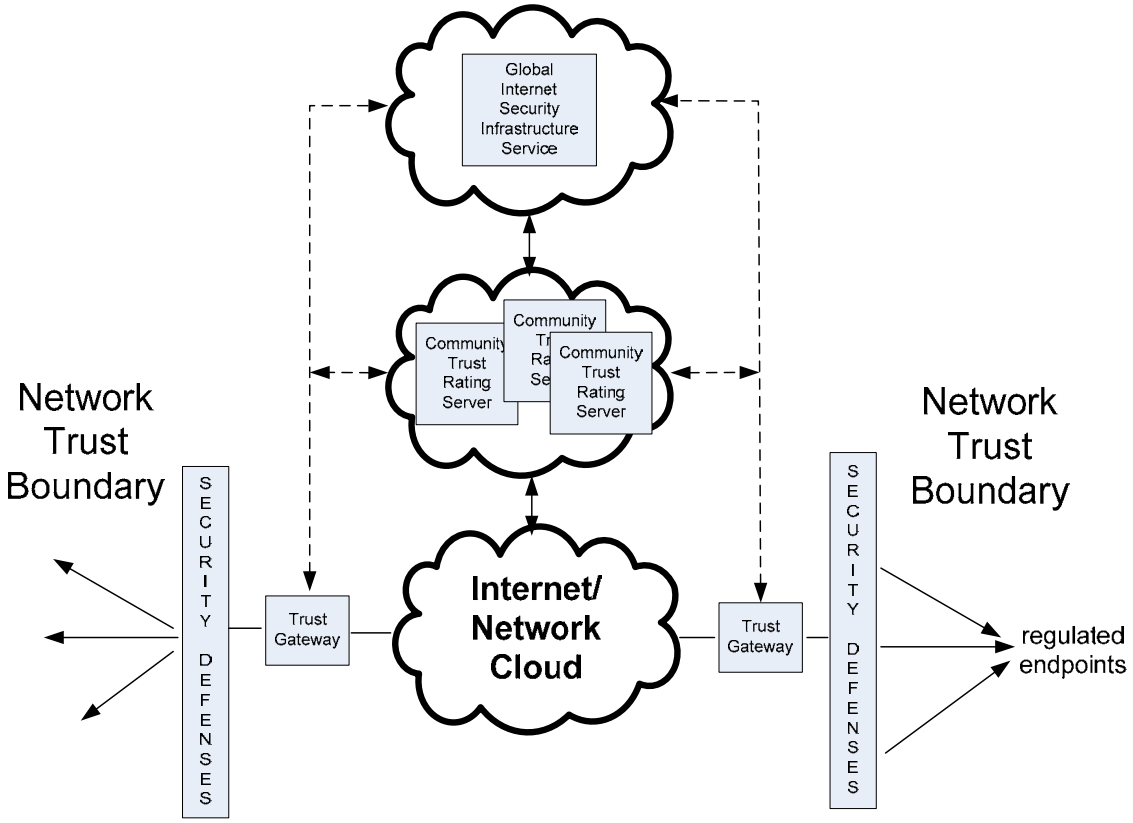

Thus, a roving malware network, constantly changing IP addresses could be tracked and identified by the global trust servers. No longer could malware writers hide behind fast-fluxing IP and DNS domain name changes. Another example, could be a previously highly trusted network or web site becomes infiltrated by malware. During the active attack, the compromised network or host could be assigned a lower trust rating, and that lower trust rating communicated to all participating parties. Once the malware was cleaned up and the network or host running clean again, its trust rating could be

improved, maybe slowly at first. But certainly after a set period of time, it could regain its original trust rating, or actually improve it beyond the original if newer, more secure practices were used. Currently, there is no way for the Internet community to be aware that a particular, popular host or network is compromised. With more and more legitimate sites being used to host malware, we need some sort of warning system.

Integrity and Identity Without Personal Identification

Privacy proponents, of which I am one, might decline this solution on its face because of the forced identification to participate. It is true that in order for this solution to work, that the destination network must be able to rely on the identity of the originator. But this doesn't necessarily mean that the destination network knows the originator's true identity. There are mechanisms and companies dedicated to the idea of identity without personal identification. The idea is that I can prove my real identity to a trusted third party, who then gives me a global token that I can use on behalf of myself…or perhaps multiple tokens, unique to each use, so I can't be tracked or identified by anyone. This is known as *pseudo-anonymity*. Thus, Internet participants can choose to be truly anonymous, pseudo-anonymous, or authenticated along various levels of increasing trust assurance. True privacy advocates can choose not to be identified (i.e. remain anonymous), but it doesn't have to be a binary decision.

The destination network/host can choose whether to require the originator's real identity, or just a reliable proxy identity, or to accept truly anonymous connections, and treat received traffic accordingly. Originators may choose whether or not to participate with a destination network depending on the destination network's identity requirement. For example, my destination network may choose to drop traffic without a real person's identity attached to it, or just treat it differently than personally identified traffic. The idea is that right now all networks must accepted poorly authenticated traffic as the same level as more trusted traffic. This new solution would give both origination and destination networks a choice to handle trusted and untrusted traffic differently.

Cryptographically Sound

This solution requires that open cryptographic standards be employed to ensure that all participating transactions are secure, confidential, and have integrity. The participating, chained components in the trust pathway must cryptographically verify the next participating component (much like is done in the Trusted Platform Module chip today). Device and user identity must be cryptographically verified and attested. Each trust component and its trust ranking must be cryptographically verified and attested. Network traffic must be tagged in a cryptographically sound manner that detects unauthorized modification. Lastly, information sent is cryptographically protected (encrypted and signed) by default, and can only be read or verified by the destination network. Default encryption and signing of data is not required for this solution to work, but is encouraged to prevent unauthorized viewing and manipulation.

How To Satisfy the Remaining Critics and Non-Participants

This solution takes into account that initially some large portion of critics and end-point nodes will choose not to participate. This solution is an opt-in solution. If it provides a compelling reason to join, we can expect some of the critics to join as success is demonstrated. End-nodes not participating are not harmed beyond their current service levels and expectations, other than being given a lower quality of

service rating as compared to more trusted traffic. If this solution significantly decreases malicious traffic on the Internet (rated at 2-6% of overall Internet traffic), even non-participants should benefit from increased performance, or at worst be performance neutral.

Possible Solution #2 – Global Identity Metasystem

by Roger A. Grimes

This solution proposes creating global infrastructure layers to provide one or more identity/authentication pairs to end-users from one or more Authentication Providers (APs) for use by content and service providers (let's call them Content Providers to simplify). Essentially, an End-User would request one or more identity/authentication pair from one or more Authentication Providers. Authentication Providers could provide password services, biometric identities, two-factor authentication tokens, smart cards, or whatever identity/authentication pair they want to offer- each with a defined trust assurance level.

Trust Assurance Levels (TALs) would be defined globally, published, and available for anyone to see. All participating Authentication Providers would have to build their identity/authentication pairs to meet a certain level of assurance as predefined in the TAL table. Example TAL table might look something like this:

TAL Value	Assurance Level	Authentication Type
0	None	Unknown connection
1	None	True Anonymous Connection
100	Low Assurance	Simple password, made up identity
500	Medium Assurance	Pseudo-anonymous identity using InfoCard, complex password and registered, verified identity
1000	Medium Assurance	Smart card, two-factor, identity verified by local proxy
65000	High Assurance	Three factor biometric identity, verified in person by certified representative, background investigation, etc.

End-Users would be free to obtain identity/authentication pairs from any participating Authentication Provider, and could have multiple identity/authentication pairs, and submit different ones to different Content Providers.

Authentication Providers would be audited by a central authority and given their own trust assurance level. Authentication Providers could not assign identity/authentication pairs above their own trust

assurance level. Abuses by an Authentication Provider might result in censure or decrease in their trust assurance level.

Each participating Content Provider would re-code their site or applications to work with participating Authentication Providers. A Content Provider would designate what minimum level of assurance is needed for an end-user to connect to their content or service. It could be done at a domain or site level, down to as granular as a specific object. For example, a payroll processing company would allow anyone, anonymous or not, to connect and download public documents. However, to see individual paycheck results might require medium assurance. To withdraw payroll money might require high assurance.

When an End-User connects to the Content Provider's site, the Content Provider prompts the user for their identity/authentication pair, along with the minimum level of assurance needed. The End-User's computer would then securely supply the appropriate identity/authentication pair to the Content Provider's web site/application to begin authentication. In most cases, the Content Provider would not ever see the End-User's authentication token, just enough to identify the End-User's identity, the type of authentication token used, and the originating Authentication Provider's identity.

The Content Provider could pass along the submitted identity/authentication pair to an Authentication Provider for authentication. The Authentication Provider would approve or deny the identity/authentication pair, which the Content Provider could handle accordingly.

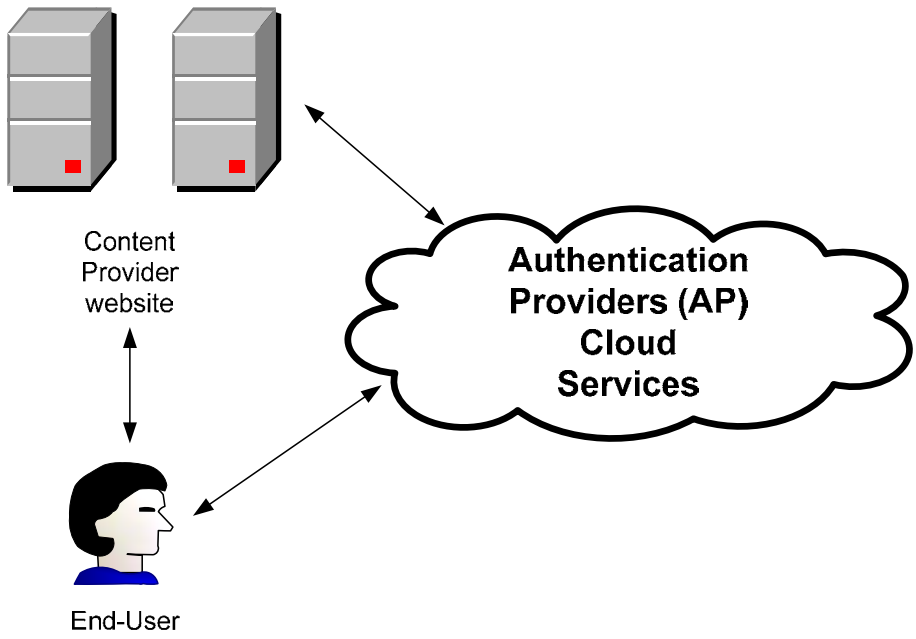

Essentially, you would have three, independent, but inter-connected layers, as shown below.

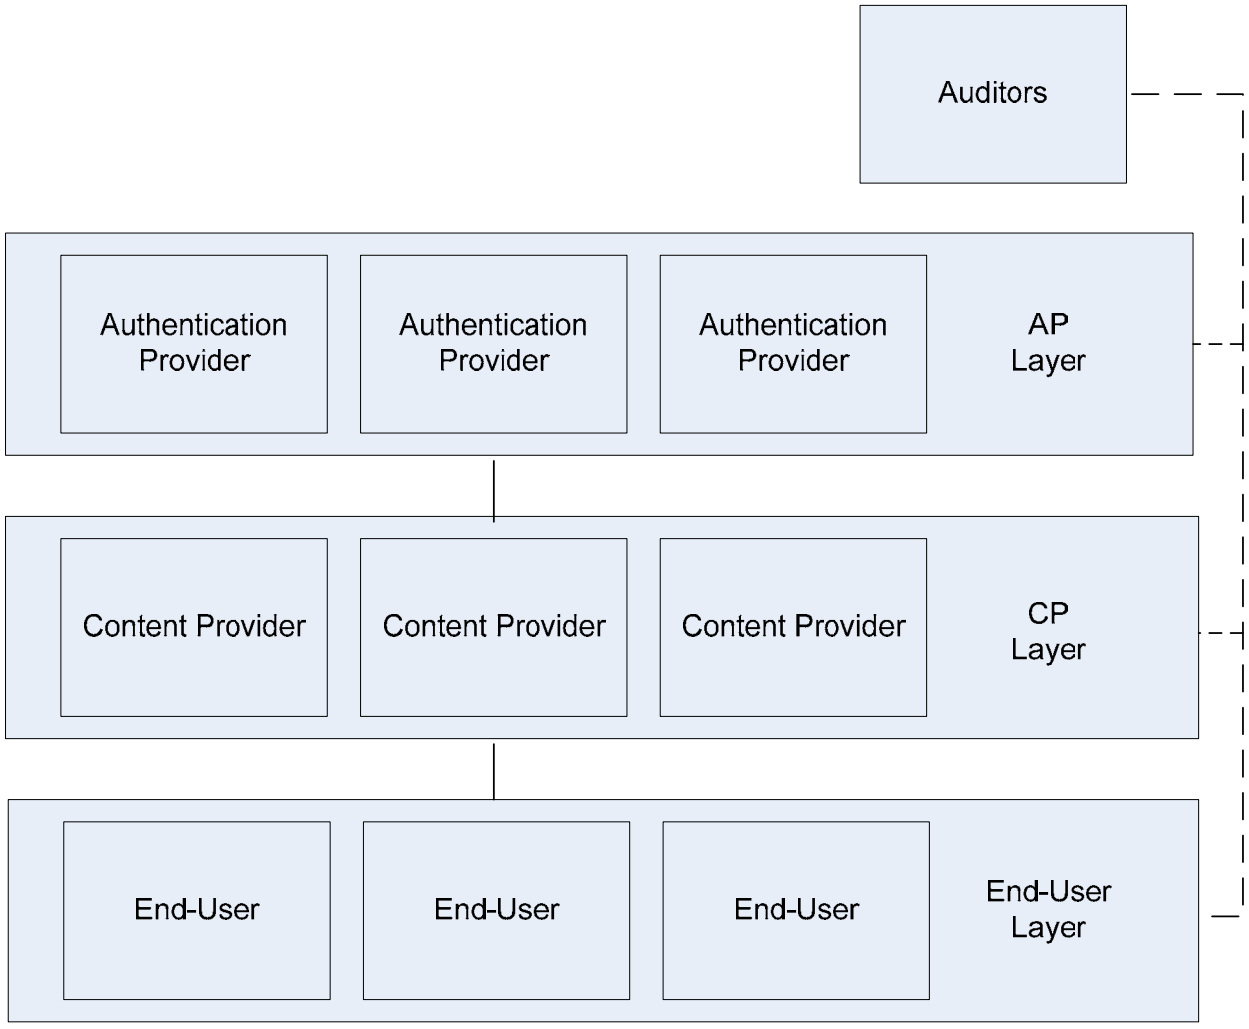

Connecting Existing Identity Systems

Each home computer user, business, enterprise, Internet Service Provider, and in some cases, entire countries have their own identity systems. This solution allows each individual identity system to be connected to the large identity metasystem using the appropriate protocols and coding, gateway, or service. With a gateway server or services, the Content Provider doesn't have to modify all their applications to take advantage of the global identity metasystem. See the diagram below.

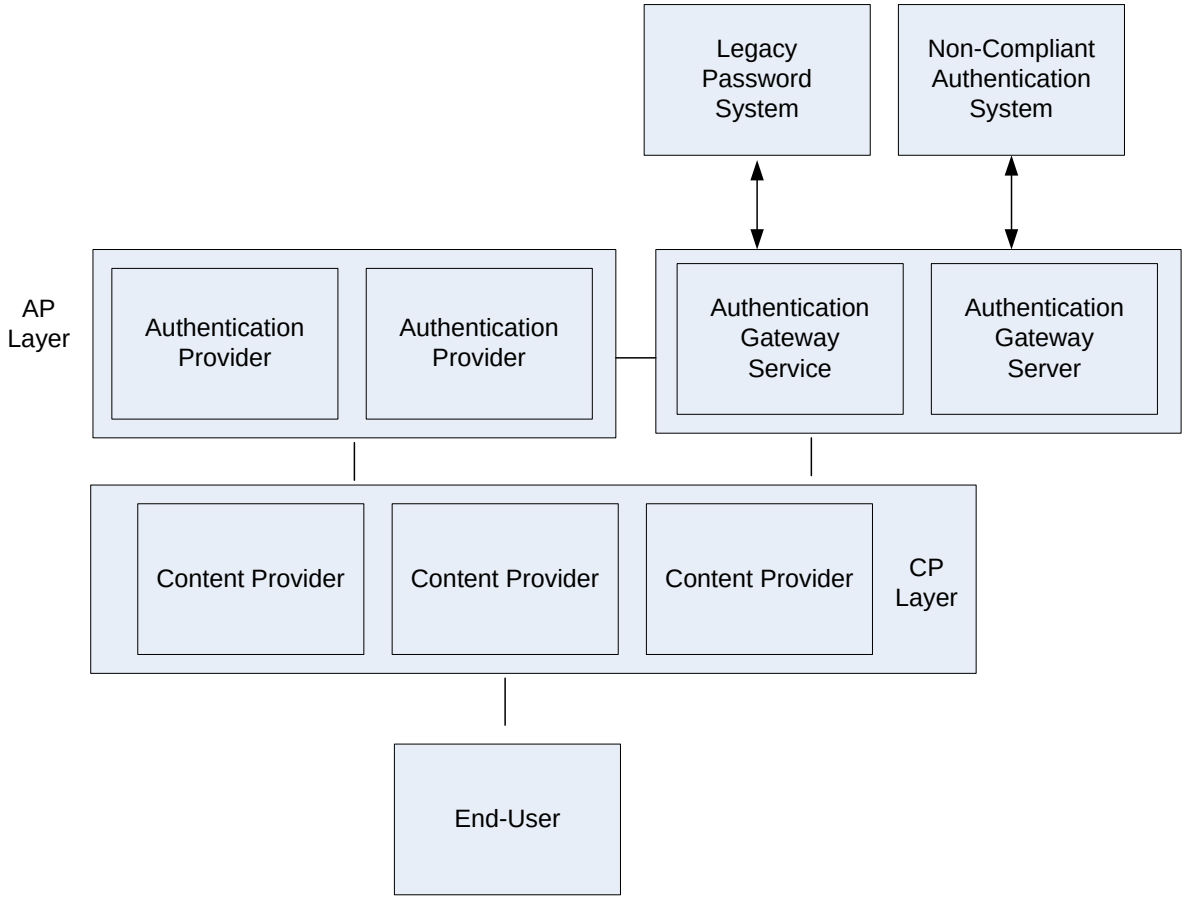

This open standards solution model has already been developed by vendors, and products are shipping (with some components in beta form, as of 2/10/09).

Open Standards Exist Today To Support Better Solutions

Today, there are enough existing open standards to support better solutions, including the possible solutions proposed above. These standards already have major industry support and products which implement them already exist. Those standards include:

- TCP/IP, especially IPv6
- Web Services (WS)
- Web Service Extensions (WS-*)
- WS – Trust
- WS – Federation
- Security Assertion Markup Language 2.0 (SAML 2.0)
- InfoCard
- DNSSec
- x.509 Digital Certificate Formats
- x.500 LDAP Directories
- Trusted Network Connect
- Network Access Control
- Trusted Platform Module chip

These standards and protocols can be used to make a more secure Internet.

FAQs

1. **Your solution decreases individual privacy. I'm completely against what you propose.**
   A: I, too, am a big privacy advocate, but I don't know of a solution that can significantly secure the Internet that doesn't involve improved, default, authentication and integrity. Let me know if you can think of one. Plus, my solution doesn't require that someone give up their anonymity, if they want to maintain it. Individuals can choose what level of identity or anonymity to give to a particular destination network. And the destination network can choose how to treat inbound traffic based upon the level of identity contained. Some hosts might choose to drop traffic, while others (I suspect the vast majority) will simply inspect the traffic more; while strongly authenticated traffic is given less inspection and faster transmission.
2. **Do you expect for your solution(s) to be adopted anytime soon?**
   A: It's highly unlikely in the near future, but dare to dream. I'm fairly confident that something along the lines of an Internet security service infrastructure will develop, because it is the only reasonable solution I can see for fighting larger, polymorphic threats. But overall, no, I don't think the world's vendors and security experts will come together to solve the Internet's big security problems until a tipping point event happens or the world's biggest governments get involved. Society, in general, is great at being reactive, and not so good at being proactive.
3. **Do any companies or entities currently support any part of your solution(s)?**
   A: Yes and no. No single company supports my exact solutions, but several already support similar ideas (or sometimes the exact concepts). Part of the reason I wrote this paper is that many of the ideas that I've been promoting for years, publicly and privately, are starting to become mainstream recommendations (e.g. Microsoft's End-to-End Trust initiative, Trusted Computing Group's IF-MAP standard, etc.), and I've been more right than wrong about the evolving threats. So, I thought by sharing more of my ideas in larger forums that people and companies with similar visions can come together and try to make a difference before the tipping point event happens.
4. **Would you be open to a public-private partnership, like what created the Internet?**
   A: Absolutely. If this solution is able to be accomplished, it will like involve participate from both sides.
5. **Do any of your solutions offer enough significant advantages that vendors will be forced to adopt them?**
   A: No. That is a very large problem. How do you induce individuals and individual companies to act against their natural self-interests to do something for the great good? I'm hoping that significantly improved security, improved performance, and custom demand is enough to entice the initial players into the solution. After the big players and names are on board, the rest of the world should follow.
6. **Other Internet protocols, like DNSSec, SenderID, and IPSec, offer significant security improvements to the Internet, but haven't taken hold. How do you expect your idea to be any different?**

A: The problem with these other laudable protocols are that they are too limited in scope. Everyone knows that if you fix DNS, fix email, etc. that you are just fixing a point issue, and malicious Internet behavior will continue nearly unabated. My solution "fixes" all protocols. Fix the plumbing pipes and you don't have to fix nearly as much of the traffic in the pipes.

7. **You mention that there are few (i.e. meaning you know of some) defenses being developed that you think can significantly improve computer security. What are they?**
8. A: First, DNSSec, SenderID, and IPSec are current protocols, that if adopted more completely would significantly improve security. Plus there are many new emerging defenses and protocols (End-to-End Trust, IF-MAP, any Trusted Computing Group standard, application signing, application and content whitelisting, Extended Validation SSL, Dshield-like data collection points, etc.) that appear to be very advantageous.
9. **Doesn't any solution of this type naturally discriminate against smaller companies and individuals who can't afford the newer stuff required to support the decision?**
   A: Yes, at least to some limited extent, whether intentional or not. It's like requiring a photo ID to vote. There's a valid reason to require a photo ID (i.e. voter fraud), but people who do not have easy access to a photo ID are discriminated against. There are many such decisions in the world (e.g. driver's license, social security card, passport, etc.). But with that said, it is my hope that the opt-in nature will allow, and very little discrimination (after all people not joining in will only be subject to the same scrutiny that they are today), will prove to be more like people moving from analog phone lines to broadband for Internet access (i.e. something people want to do). Many of the solution components (e.g. InfoCard, etc.) are zero cost.
10. **How can you realistically expect to increase security and not impact performance?**
    A: This is a difficult challenge, but with 2-6% of the Internet and 70-90% of all email being malicious in nature, if we can reduce those levels to near zero, it gives us a lot of room to play with before it actually slows down overall computing.
11. **Doesn't your solutions erode people's privacy?**
    A: Yes and no. Yes, at least a little, if you want improved security. No, if you choose not to participate. It's not a binary decision. We give up privacy all the time for more security (e.g. driver's license, city and community laws, etc.). And if you want to participate in better security without giving away your real identity, go pseudo-anonymous. Security and privacy are not completely exclusive of each other. Privacy isn't a binary choice anymore than security is.
12. **You propose that network traffic be encrypted and signed end-to-end. Won't many governments oppose this on the grounds that they need to inspect the traffic?**
    A: Probably, but so far there is no law that says I have to let the government read my information. Most governments have all sorts of rights and laws to try and read our traffic, but I don't know of any government law (I'm sure there are some) that requires people to let the government read it. For example, the U.S. government may sometimes have the legal right to capture your network traffic or listen on your phone calls, but there is no law saying that I cannot encrypt my phone call or network traffic further so they can't read it. Personally, I would strongly fight any law that says I have to show the government my information for basic services, and without a court order.

13. **You propose creating a new "dream team" consortium to solve the Internet's major security issues. Doesn't IETF, IANA, CERT, TCG, (or whoever), already exist to protect the Internet?**
    A: Nearly so. The Trusted Computing Group (TCG) is the closest model to what is needed. I'm open to imagining the security dream team as part of one of the aforementioned groups, as long as the team can act quickly (not something these former groups are always known for).
14. **How would the community trust rating server get populated with security domain trust rankings (i.e. would it be possible for a malicious person to maliciously malign my host or network in order to lower its trust rating)?**
    A: I'm not sure exactly how it would work, but yes, there would have to be protections in place to prevent malicious manipulation. This sort of thing is done for all sorts of services already with varying degrees of success.
15. **Why do you mention DNS in your solution as an example technology when DNS is so insecure?**
    A: For two reasons. First, it's mainly mentioned as an example of a global, redundant, distributed infrastructure service. Attackers will try to take down any global security service, so we need to mention that it is possible, as demonstrated by DNS, to do it globally and secure. Second, DNSSec is one of three technologies (the other two being IF-MAP and Sender ID) that are true security solutions. I consider most other things security theater.
16. **How would you do X and XX in your solution?**
    A: I don't have all the answers. That's why I propose bringing together a dream team of experts under each component discipline to solve the tough technical challenges.

If you've reached this part of the paper, I thank you for your time and participation.

Bibliography

[1] History of the Internet, Wikipedia, http://en.wikipedia.org/wiki/History_of_the_Internet

[2] MessageLabs Intelligence Report April 2008, MessageLabs, http://www.messagelabs.com/mlireport/MLI_Report_April_2008.pdf

[3] Internet has a trash problem, InfoWorld Magazine, April 1, 2008, http://www.infoworld.com/article/08/04/01/Internet-has-a-trash-problem_1.html; and Up to 3 percent of traffic is malicious says researcher, CSO Online, April 1, 2008, http://www.csoonline.com/article/326013/Up_to_Three_Percent_of_Internet_Traffic_is_Malicious_Researcher_Says; and A Peek at ISP DDOS, Spam Traffic Trends, DarkReading, April 1, 2008, http://www.darkreading.com/document.asp?doc_id=149866&WT.svl=news1_1; and 2% of Internet Traffic Raw Sewage, Arbor Networks, March 31, 2008, http://asert.arbornetworks.com/2008/03/2-of-internet-traffic-raw-sewage.

[4] Hackers harpoon executives, The Sydney Morning Herald, May 6, 2008, http://www.smh.com.au/news/security/going-after-the-big-phish/2008/05/06/1209839606696.html; and History and current status of phishing, Wikipedia, http://en.wikipedia.org/wiki/Phishing; and US Federal Court Subpoena Phish, Junkfax.org, http://www.junkfax.org/fax/phish/uscourtsPhish.htm; and U.S. District Court Subpoena, Snopes.com, http://www.snopes.com/fraud/phishing/subpoena.asp.

[5] Top Spam Bots Exposed, Secure Works, April 8, 2008, http://www.secureworks.com/research/threats/topbotnets; and Top botnets control 1M hijacked computers, ComputerWorld magazine, April 9, 2008, http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9076278; and Storm botnet, Wikipedia, http://en.wikipedia.org/wiki/Storm_botnet.

[6] From Fast Fluxing to Rockphish – Part 1, McAfee AVERT Labs, November 30, 2007, http://www.avertlabs.com/research/blog/index.php/2007/11/30/from-fast-flux-to-rockphish-part-1.

[7] What is server-side polymorphism?, Anti-Virus Rants, August 10, 2007, http://anti-virus-rants.blogspot.com/2007/08/what-is-server-side-polymorphism.html

[8] Compromised web sites serve more malware than malicious ones, Ars Technica, January 22, 2008, http://arstechnica.com/news.ars/post/20080122-compromised-websites-serve-more-malware-than-malicious-ones.html; and Malware spikes in 1Q as hackers increasingly infect web sites, InformationWeek, April 24, 2007, http://www.informationweek.com/news/internet/showArticle.jhtml?articleID=199201032; and Massive Site Compromise: The Siege Continues, TrendMicro Malware Blog, April 3, 2008.

[9] Privacy Protection: The Government is No Help, InfoWorld magazine, June 16, 2006, http://www.infoworld.com/article/06/06/16/79260_25OPsecadvise_1.html

[10] 2007 Internet Crime Report, FBI, http://www.ic3.gov/media/annualreport/2007_IC3Report.pdf and CSI 2007 Survey, FBI, http://i.cmpnet.com/v2.gocsi.com/pdf/CSISurvey2007.pdf

[11] Russian Business Network, Wikipedia, http://en.wikipedia.org/wiki/Russian_Business_Network.

[12] Rockphish, EU Spam Trackers, http://spamtrackers.eu/wiki/index.php?title=Rockphish.

[13] CERT Advisory CA-2003-04 MS-SQL Server worm, CERT, January 25, 2003, http://www.cert.org/advisories/CA-2003-04.html.

[14] SQL Slammer (computer worm), Wikipedia, http://en.wikipedia.org/wiki/SQL_slammer_worm.

[15] Iloveyou worm, Wikipedia, http://en.wikipedia.org/wiki/ILOVEYOU.

[16] Blaster (computer worm), Wikipedia, http://en.wikipedia.org/wiki/Blaster_worm.

[17] Google Health: A First Look, Official Google Blog, 2/28/2008, http://googleblog.blogspot.com/2008/02/google-health-first-look.html.

[18] One Laptop Per Child, One Laptop Per Child website, http://laptop.org.

[19] Universal Description Discovery and Integration service, Wikipedia, http://en.wikipedia.org/wiki/Universal_Description_Discovery_and_Integration.

20 Trusted Network Connect IF-MAP Announcement FAQ, Trusted Computing Group, https://www.trustedcomputinggroup.org/specs/TNC/IFMAP_FAQ_april_28.pdf.

Separate the Concerns

Michael Franz, University of California, Irvine

Who you are – I am a tenured professor at *UC Irvine,* designated by the National Security Agency as a National Center of Academic Excellence in Information Assurance Research (CAE-R). I have led numerous federally-funded research projects with an emphasis on systems-level software and cyber security. My close relationship with the open-source community has led to a major technology transfer success: my research at UC Irvine created the (patent pending) just-in-time compilation technology that forms the basis of the "TraceMonkey" JavaScript engine in the *Mozilla Firefox* web browser (300 Million deployments). Additionally, I am collaborating with *Adobe* to integrate my invention into *Flash* (deployed on more than 2 Billion devices) and I am working with *Sun Microsystems* on a new *Java* virtual machine based on my academic research. I am also in talks with *Microsoft* to integrate the technology into their *Dot Net, Silverlight,* and *Azure* platforms.

Game-changing dimension – Morph the gameboard / change the rules.

Concept – Computers are starting to be very inexpensive. Capable "netbook" computers can be purchased for under $300 and initiatives such as "one laptop per child" are targeting a $100 computer. At this hardware price level, it no longer makes sense to treat the "personal" computer as a precious resource that must simultaneously be able to run user-installed programs, surf the open web, as well as perform critical functions such as electronic banking and e-commerce.

I propose a separation of concerns, in which we differentiate between user-managed PCs on one hand, and create an entirely separate category of Trustworthy Information Appliances (TIAs) on the other hand. These new devices would be so cheap that banks could give them away for free and governments could subsidize distribution to low-income households. Trustworthy Information Appliances would become the gateways to a new National Trusted Cyber-Infrastructure (NTCI) that is parallel to and independent of the open Internet, and tightly managed end-to-end.

Many people would probably continue to own "personal" computers and use them for discretionary activities. However, as time progressed, "critical" services would increasingly migrate to the "managed" NTCI and be accessible only via devices following the TIA standard.

By design, a transaction on the NTCI is end-to-end between a single TIA and a single remote party such as a financial institution or a service aggregator. For example, a citizen would insert her banking card into her TIA, turning the TIA into a dedicated banking terminal until the card is removed. The TIA would connect via a VPN tunnel directly to the issuer of the smart card, using VPN credentials stored on the card (and not even requiring a DNS lookup). The smart card also provides the only form of persistent storage in this architecture; when it is removed, the TIA resets itself, erasing all of its memory and returning to a generic device state.

Such an architecture would be game-changing, as it would render most existing client-side vulnerabilities irrelevant. Citizens would eventually conduct all their critical online transactions via secure point-to-point channels and no longer commingle them with discretionary activities. The "Internet of trust" could then evolve separately from the "Internet of everything else."

Vision – A TIA is a low-cost "thin client computer" that is specifically targeted towards managed "Web 2.0" services. A laptop-like device, it uses secure boot techniques to enter a remotely measurable trusted state. Rather than providing the flexibility of a full laptop, the TIA provides a

computing platform based on existing open HTML, XML, and JavaScript standards, with added information-flow controls. The trusted code base for this thin Web 2.0 client is expected to be orders of magnitude smaller than that for a standard PC, enabling a full human audit.

Note that the point-to-point approach does not necessarily imply a loss of convenience for the end user. I envision a situation in which the service-side endpoints of the managed infrastructure will be mostly "aggregators" that in turn provide access to merchants. For example, a credit card company or an ISP could set up an electronic "shopping mall" through which a variety of services can be accessed, with service costs paid by the merchants. Alternatively, an aggregator could offer subscription-based "trusted e-commerce services," in which merchants are carefully vetted and certain assurances (and perhaps even financial indemnity) are provided as part of the package.

E-government services, on the other hand, would probably be very welcoming of the point-to-point approach, using government-issued digital credentials (e.g., smart cards issued by DMVs). A big advantage of the proposed architecture is that *all* NTCI connections become tunnels with decentralized key management and out of band key distribution (via physical smart cards).

Most of the technologies required to build such an infrastructure exist already. What is required is a clear leadership for establishing a standard that stays clear of any vested commercial interests to become acceptable for the majority of stakeholders.

Method – We are at an inflection point at which three things are happening simultaneously: 1) The "cyber crime tax" that we are paying is becoming so significant that people will adopt new practices if these are demonstrably safer. 2) The Web 2.0 environment is so rich that the line between traditional desktop applications and "cloud computing" services is increasingly blurring. For example, *Gmail* has virtually all of the functions of a traditional email program. Few, if any, client-side functionality would need to be excluded from our approach because they cannot be made to run inside a Web 2.0 browser. 3) Hardware costs have dropped to the point where it is no longer preposterous to suggest that users own a separate device for the purpose of accessing trustworthy hosted services.

Taken together, I believe the time is now to stop investing huge amounts of money and effort into fixing an existing home computing infrastructure that was never designed for trustworthy computing. The needs of private owners of home PCs that are used for gaming, the kids' social networking, and discretionary web browsing are fundamentally irreconcilable with the requirements of a national trustworthy e-commerce and e-government infrastructure. It is time that we separate these concerns and create a separate, managed ecosystem for trusted e-activities to complement the unregulated general Internet that is so vulnerable.

Dream team – In a first phase, form a small research team composed of top-notch academic and government researchers in conjunction with an established open-source organization such as Mozilla, to create an architecture and build a working system. In a second phase, form a non-profit consortium encompassing government, ISPs, hardware manufacturers, and the financial industry, to standardize and deploy the solution.

The goal is to judiciously combine existing and emergent technologies into a coherent, robust whole. The challenge is as much political as technical, because this plan will encroach on entrenched interests. The best way of guaranteeing success is to define this system independent of vested commercial interests, with the help of the open-source community and under academic leadership.

Financial Services Sector Coordinating Council
for Critical Infrastructure Protection and Homeland Security

Submitted via *http://www.nitrd.gov/leapyear/* and via e-mail: *leapyear@nitrd.gov*.

December 15, 2008

Dear Sir/Madam,

Thank your for the opportunity to provide input to the Networking and Information Technology Research and Development (NITRD)'s request on "Cyber Leap Year" which appeared in the Federal Register (FR Doc. E8-24257) on October 14, 2008. As members of the Research and Development Committee of the Financial Services Sector Coordinating Council for Critical Infrastructure Protection and Homeland Security (FSSCC), we would like to submit the following three "game changing" technologies:

1. Five-Star Rating for Software Security
2. Self Healing Application Services Framework
3. Rapid Reconstitution Resiliency

These suggestions are included in the FSSCC R&D Agenda which was released publicly in September (see: https://www.fsscc.org/fsscc/reports/2008/RD_Agenda-FINAL.pdf) and will appear in Sector Specific Plan for the Banking and Finance Sector.

Sincerely,

FSSCC R&D Committee:

Alexander Abramov, JPMorgan Chase
Warren Axelrod, Financial Services Technology Consortium
Andy Bach, Securities Industry Automation Corporation
John Carlson, BITS/Financial Services Roundtable (chairman)
Frank Castelluccio, The Options Clearing Corporation
Dan DeWaal, The Options Clearing Corporation
Eric Guerrino, Bank of New York Mellon Corporation
Mark Merkow, American Express Company
William Nelson, Financial Services Information Sharing and Analysis Center
Dan Schutzer, Financial Services Technology Consortium
Robert Vitali, MorganStanley
Brian Peretti, U.S. Department of the Treasury (Public Sector Representative)

1. *FIVE-STAR RATING SYSTEM FOR SOFTWARE SECURITY*

Who you are - Financial Services Sector Coordinating Council (FSSCC) R&D Committee (www.fsscc.org). Experts in technology, information security, and risk management
Group Members:

- Alexander Abramov, JPMorgan Chase
- Warren Axelrod, Financial Services Technology Consortium
- Andy Bach, Securities Industry Automation Corporation
- John Carlson, BITS/Financial Services Roundtable (chairman)
- Frank Castelluccio, The Options Clearing Corporation
- Dan DeWaal, The Options Clearing Corporation
- Eric Guerrino, Bank of New York Mellon Corporation
- Mark Merkow, American Express Company
- William Nelson, Financial Services Information Sharing and Analysis Center
- Dan Schutzer, Financial Services Technology Consortium
- Robert Vitali, MorganStanley
- Brian Peretti, U.S. Department of the Treasury (Public Sector Representative)

Game-changing Dimension - Change the rules

Concept - A five-star rating system has been very effective in improving the safety and quality of vehicles over the last several years. We are calling for a similar five-star rating system for software security. This five star rating will provide an indicator of the assurance related to software's "*degree of protection*" or "*level of resistance*" against *known threats* to application and system (infrastructure) software.

This rating system may be mandated by government and commercial procurements, made transparent and readily available to consumers and enterprises to improve decision-making and comparability across similar software genre.

Vision - The vision is to implement a reliable, transparent, reasonably fast and inexpensive standardized methodology to test and rate software security on a simple five point scale thereby removing the drawbacks of existing complex, proprietary, and expensive testing schemes, such as the Common Criteria (ISO/IEC 15408).

The rating system quantifies the protection measures built into the design, development, and deployment of the software against all known threats and to an extent, against unknown threats (zero-day attacks). The rating system could provide a uniform/normalized score based on results from automated and/or manual analysis on the source code and deployed software in different usage scenarios:

- Default out-of-the-box installation
- Maximum Security Configuration
- Typical Deployment Configuration - fully operational with all interfacing systems
- When the software, hardware, network and/or interfacing systems fail from security vulnerabilities

Each star rating in the system could be established to describe the strength of the software's controls/protection against known and unknown attack conditions – measuring the software resistance to attacks in terms of time or alerting-capabilities

A more stringent rating system will need to be developed and enforced for "embedded software" – e.g. medical equipment, automobiles, and other critical devices or application.

This rating information about software products should then be made available on all "shrink-wrap" boxes, manufacturer and vendor's Web sites description, and in advertisements for these products.

Method – The idea for a 5-Star Rating System emanated from an ongoing initiative by the FSSCC R&D Committee to publish and maintain a prioritized list of Research Challenges to improve cybersecurity across the Finance and Banking Sector. The top rated challenge is Advancing the State of the Art in Designing and Testing Secure Applications and this project helps to meet several of the objectives described within the challenge. The FSSCC Research Agenda and Challenges Document may be located at https://www.fsscc.org/fsscc/reports/2008/RD_Agenda-FINAL.pdf

Dream Team

- Selected members of the FSSCC R&D Committee
- SANS (Mason Brown, Allan Paller, David Rice)
- Microsoft (Michael Howard, Steve Lipner)
- Cigital (Gary McGraw)
- Stonewall Software (John Viega)
- KRvW Associates (Ken van Wyk)
- OWASP (Tom Brennan)

2. SELF HEALING APPLICATION SERVICES FRAMEWORK

Who you are - Financial Services Sector Coordinating Council (FSSCC) R&D Committee (www.fsscc.org). Experts in technology, information security, and risk management
Group Members:

- Alexander Abramov, JPMorgan Chase
- Warren Axelrod, Financial Services Technology Consortium
- Andy Bach, Securities Industry Automation Corporation
- John Carlson, BITS/Financial Services Roundtable (chairman)
- Frank Castelluccio, The Options Clearing Corporation
- Dan DeWaal, The Options Clearing Corporation
- Eric Guerrino, Bank of New York Mellon Corporation
- Mark Merkow, American Express Company
- William Nelson, Financial Services Information Sharing and Analysis Center
- Dan Schutzer, Financial Services Technology Consortium
- Robert Vitali, MorganStanley
- Brian Peretti, U.S. Department of the Treasury (Public Sector Representative)

Game-changing Dimension – Morph the gameboard

Concept – The *de rigueur* methods for creating secure systems based upon insecure or unreliable application and infrastructure software is fundamentally flawed. These systems often result in layers upon layers of protection mechanisms to overcome the weaknesses or unknown risk of component parts used in construction. Now, more than ever, there is an increasing need to squeeze out development costs while improving the security and quality of applications.

While applications will continue to be stitched together for cost-effectiveness and meeting the pressures of time-to-market, there is a widespread need to create a self healing application framework.. We are looking for a self reliant application (or a set of applications) that could detect the pattern of attack and quickly respond triggering an automated cycle of remediate – test – redeploy.

Vision - The vision is to create a reliable, self healing application services framework that would allow any developer to build an inexpensive, yet secure application ready for deployment in any environment, without being dependent on a developer to build all the safety and security mechanisms in each and every application. The self healing application service framework would provide:

- Simple coding techniques to hook into any components of code;
- Runtime self-healing capabilities to:
    - Monitor and Detect the pattern / abnormal behavior
    - Validate and confirm the potential security threat
    - Evaluate candidate services to address
    - Activate / Request services to avoid the incident
    - Increase level to monitoring to collect additional information
- Trigger automated cycle to analyze, rebuild, test and redeploy application / services or component(s)

A more desirable feature would be the self learning capability to be environment "aware" to tune the self healing services to the desired level, based on external factors. All commercial/open world pieces of code would need to be "self healing ready" certified for use with other self healing application software.

Method: The idea for a Self-Healing Application Services Framework originated from an ongoing initiative by the FSSCC R&D Committee to publish and maintain a prioritized list of Research Challenges to improve cybersecurity across the Finance and Banking Sector. The top rated challenge is Advancing the State of the Art in Designing and Testing Secure Applications and this project helps to meet several of the objectives described within the challenge. The FSSCC Research Agenda and Challenges Document may be located at https://www.fsscc.org/fsscc/reports/2008/RD_Agenda-FINAL.pdf

Dream team

- Selected members of the FSSCC R&D Committee
- SANS (Alan Paller)
- Microsoft (Michael Howard, Steve Lipner)
- Cigital (Gary McGraw)
- Stonewall Software (John Viega)
- KRvW Associates (Ken van Wyk)
- OWASP (Tom Brennan)

3. RAPID RECONSTITUTION RESILIENCY

Who you are - Financial Services Sector Coordinating Council (FSSCC) R&D Committee (www.fsscc.org). Experts in technology, information security, and risk management
Group Members:

- Alexander Abramov, JPMorgan Chase
- Warren Axelrod, Financial Services Technology Consortium
- Andy Bach, Securities Industry Automation Corporation
- John Carlson, BITS/Financial Services Roundtable (chairman)
- Frank Castelluccio, The Options Clearing Corporation
- Dan DeWaal, The Options Clearing Corporation
- Eric Guerrino, Bank of New York Mellon Corporation
- Mark Merkow, American Express Company
- William Nelson, Financial Services Information Sharing and Analysis Center
- Dan Schutzer, Financial Services Technology Consortium
- Robert Vitali, MorganStanley
- Brian Peretti, U.S. Department of the Treasury (Public Sector Representative)

Game-changing Dimension - Change the rules

Concept – Much attention has been focused on improving an enterprise resiliency by duplicating and backing up data centers and power systems (including locating them sufficient distance apart), and procuring redundant and diverse telecommunications. This concept includes the ability to reconstruct rapidly and dynamically at appropriate locations depending on the situation. The concept includes the discovery of remaining remnants of the enterprise system and support infrastructure to automatically and quickly rebuild any missing functionality and capability.

Organizations would achieve high levels of resilience if employees could carry or buy off-the-shelf portable computers, communications and power components with them, that are capable of rapid installation and set-up without loss of data or transactions, even when there has been some downtime. This would enable organizations no matter how hard it is hit to re-assemble remaining employees who are appropriately trained, to reconstitute the enterprise processes and systems out of portable available components they can carry or easily transport with them.

Vision - The vision is to establish a set of processes, procedures and supporting architecture that permits employees to securely transport with them the necessary components and data to allow reassembly and reconstitution of the financial institution's key enterprise systems. The goal is to continue operations at minimum essential levels until the main support infrastructure is restored.

This concept could be achieved by leveraging the great strides in the miniaturization of computing and communications equipment, in the area of rapid installation of wireless communications nodes, and advances in military communications and computing. With the right planning, training and architecture, a secure set of portable components could be built that is parallel and distributed.

Method – The idea for a rapid reconstitution resiliency approach originated from an ongoing initiative by the FSSCC R&D Committee to publish and maintain a prioritized list of Research Challenges to improve cybersecurity across the Finance and Banking Sector. Since cybersecurity attacks, combined with physical attacks, can take out of service vital computer, communications and potentially power systems, providing this sort of rapid reconstitution capability is highly desirable and discussed in the second highest rated challenge in the report: More Secure and Resilient Financial Transaction Systems. This project helps to meet several of the objectives described within the challenge. The FSSCC Research Agenda and Challenges Document may be located at https://www.fsscc.org/fsscc/reports/2008/RD_Agenda-FINAL.pdf

Dream Team

- Selected members of the FSSCC R&D Committee
- Associations (e.g., SANS Institute)
- Software companies (e.g., Microsoft, Oracle)
- University and think tanks (e.g., Carnegie Mellon University, RAND, MITRE, SAIC),

Game Changing Cyber Security Concept: Personal biometric identity verification that works with existing access control/security infrastructure enabling high assurance verification without requiring the addition of specialized equipment or modification of existing systems.

Who We Are: Privaris, Inc. is a small, privately held technology company pioneering new approaches to the use of biometrics for identity verification in IT and physical security. Privaris implements the concept of "personal biometrics" whereby biometric identity verification is accomplished via a small device carried by the user. www.privaris.com.

Contact: John Petze, President & CEO

Game Changing Dimension: Change the rules related to credentials used for access to cyber systems by creating a solution that insures high reliability identity verification via biometrics, supports widely adopted standards for secure credentials and is easily deployed.

Concept: The proposed concept is to implement biometric identity verification on a small key-fob sized device carried by the user. The user verifies their identity via a one-to-one match of their live fingerprint to the fingerprint template placed on their device during enrollment. Upon a successful match by the rightful owner, the device then outputs standard credentials, such as digital certificates, that work with existing systems. In this way, biometric identity verification can be added to existing systems without requiring the installation of specialized equipment on all individual assets.

Vision: Establishing identity of the individual attempting to access IT resources, whether a computer, network or application, is a critical element of any security protocol. Reliable identity verification is the foundation on which subsequent actions are taken, i.e., allowing access to systems and determining privileges within those systems, etc. In other words, it is essential to know who is at the door or the keyboard.

Typical methods of identity authentication such as passwords and cards have numerous disadvantages including: inherent insecurity due to the ability to be shared or used if lost or stolen; complexity for the user, e.g., the need to remember passwords that must be changed on a regular basis.

Biometrics Can Help Addresses These Issues. Biometrics, the measure of a unique physical characteristic of the user, such as a fingerprint, provides a much higher assurance of identity than cards and passwords. Biometrics eliminates the need for users to remember passwords and change them over time. Biometrics cannot be lost, stolen or shared.

How to Overcome the Limitations of Past Approaches to Biometrics. Historically the use of biometrics has required the installation of expensive, "fixed mount" readers on every asset to be protected. This is costly and cumbersome and has significant capital expense implications. This, and other drawbacks of "fixed mount" biometrics, have limited use. Biometric identity solutions can only become widely used when the need to install infrastructure is eliminated. The proposed approach accomplishes this.

Method for Implementing Personal Biometrics: The concept of personal biometrics is achieved in a small, key fob sized device which contains a fingerprint sensor, secure microprocessor, re-chargeable battery, and multi-mode wireless communications capabilities. Credential delivery formats supported include: legacy RFID (125KHz) and the latest RFID standards based on ISO-14443 and 15693 (which are the basis for the HSPD12 cards such as PIV, TWIC and CAC), Bluetooth, 802.15.4, and USB. Because the device performs all biometric matching locally on the device, and communicates widely accepted standard cryptographic credentials it does not require expensive modifications to existing systems. This allows high reliability biometric identity verification to deployed rapidly across large populations of users.

Dream Team: DoD, Dept of State, Verisign and/or RSA, a major bank. A consortium of players across Government and industry can provide the framework for secure issuance of personal biometric credentials and demonstrate acceptance which would then be emulated across Government agencies and commercial entities.

Capability, Status, and Opportunity for Advancement of the Concept
Privaris has developed a first generation of fully functional, secure personal biometric identity verification devices, which are being successfully used in Government and commercial applications. The opportunity for advancement of the concept is twofold:

> First, take the lessons learned to date and develop 2nd generation devices that implement personal identity verification and multi-factor credential delivery at a significantly lower cost, thereby increasing the range of applications that can be addressed cost effectively. Second, by assembling a dream team to support the adoption and acceptance of the personal biometrics model it will be possible to create an ecosystem supporting secure, "cooperative issuance" of these identity credentials. This will enable users to receive a secure identity verification device from an authorized issuer. Once vetted and enrolled in their personal identity device they would then be able to use it across multiple agencies and applications, with each accepting agency able to manage the specific credentials used for their own applications.

Closing Thoughts: Biometric technology provides high assurance identity verification for transactions of all types. The challenge to more widespread use has been the cost and complexity of deploying "fixed mount" biometric readers on the individual assets to be protected, along with the need to enroll in each individual system. This infrastructure-centric approach has limited the adoption of biometrics to only high security applications. Continued focus on the infrastructure-centric approach is the primary impediment to biometrics becoming universally applicable across business and society. The concept we propose is to move to a distributed model where biometric verification is a capability carried by the individual user. We submit that just as telephone communications moved from fixed infrastructure (phone booths) to personal, mobile communications (cell phones), identity verification too can only become a ubiquitous, fully integrated national approach through the implementation of such a a distributed model. The technology is here today and proven with first generation product. What is needed achieve the end goal is a consortium of key players that have the reach and authority to issue and accept personal identity verification devices.

Who you are-- Georgia State University Research Foundation--We are a nonprofit corporation, created to support research activities of the University through securing gifts, contributions and grants from individuals, private organizations, and public agencies and in obtaining contracts for the performance of sponsored research, development, or other programs.

Contact: Jim D. Flowers, Interim VP External Affairs

Game-changing dimension—Morphing the game board – integrating security compliance into organizational structure – incentivizing 100 percent of organizational members to make the organization secure – analogous to arming the inhabitants inside a secured perimeter.

Concept-- This compliance program uses a framework to organize and direct research and development projects toward an integrated, enterprise resource planning (ERP) approach to addressing security and risk management in a wide variety of organizations. This framework cuts across silos of information and practices that arise naturally within moder, complex, computer-based organizations. This proposed framework is multidisciplinary, leveraging knowledge and experience of private sector practice and public sector research.

Vision-- The compliance framework uses simple concepts to organize and coordinate normally disparate security and policy compliance activities. It acts fundamentally as both a "wrapper" and an extensible interface. As a wrapper, it envelopes an organization's information technology infrastructure including communications, database, service bus, web services, etc. Based on the compliance governance model, it provides a form of universal protocol to interface with a wide variety of organizational security and policy compliance elements. Examples of such elements are denoted as "golf tees" components that optionally plug into the backplane. These optional elements connect the backplane to the dashboard. The dashboard is an element that provides a human interface to the organizational systems, stakeholders, managers, and policy makers. (Figure 1 provides a graphical depiction of the Compliance Model.)

Method-- The Board of Regents of the University System of Georgia (USG) is operating a program that seeks to address these security threats on its own campuses within a framework that is largely exportable to many other organizational forms. The USG effort, guided by principal investigators representing multiple disciplines at Georgia State University and teamed with industry experts in risk management and enterprise information systems, will research, develop and implement an integrated solution. The approach will not only be suitable for US colleges and universities, it will be suitable for many other kinds of government, commercial and service organizations. In addition to common physical and technical security aspects, organizational and individual behaviors must be addressed through compliance training and focused organizational change.

This compliance program aims to achieve a single, unified solution ensuring compliance to statutory and regulatory policy obligations at the federal and state levels. By focusing the program of security on policy compliance, each program element helps focus organizational members and decision makers, as well as lawmakers and other policy makers, on the costs and consequences of non-compliance to physical and technical security obligations.

The framework in this compliance program provides an integrated enterprise approach to security. Potential elements include:

- A compliance policy framework;

- Economic and financial models to determine best solution fit given an organization's mission, budget resources and weighted risks and liabilities;
- A means for continuous update of methods, practices and compliance obligations as other institutions across the nation add their experiences to the knowledge base – such a dynamic system represents a "living" deliverable that will increase in value over time.
- Definitions of the people, process and technical capabilities necessary to gain an accurate view of institutional characteristics and vulnerabilities;
- Audit, legal, regulatory and/or law enforcement bodies access to critical and relevant information;
- Management planning tools (i.e., integrated security dashboard) to correct and adjust risk profiles as required by laws and regulations.
- A methodology, model and tools for delivering risk assessments and compliance analyses that will significantly improve the safety, security and policy compliance of a wide variety of organizational forms.

Dream team— Georgia State University, Board of Regents of University System of Georgia (36 institutions and their Chief Information Security Officers), industry IT risk expertise (such as PwC), Oracle information security practice or information systems, Six Sigma Management experts, State of Georgia technology authority, commercial vendors with policy and/or technical compliance solutions, State and Federal legislatures, Department of Homeland Security, Department of Education, individuals across Georgia State and the University System of Georgia.

Figure 1. Security and Policy Complance Framework

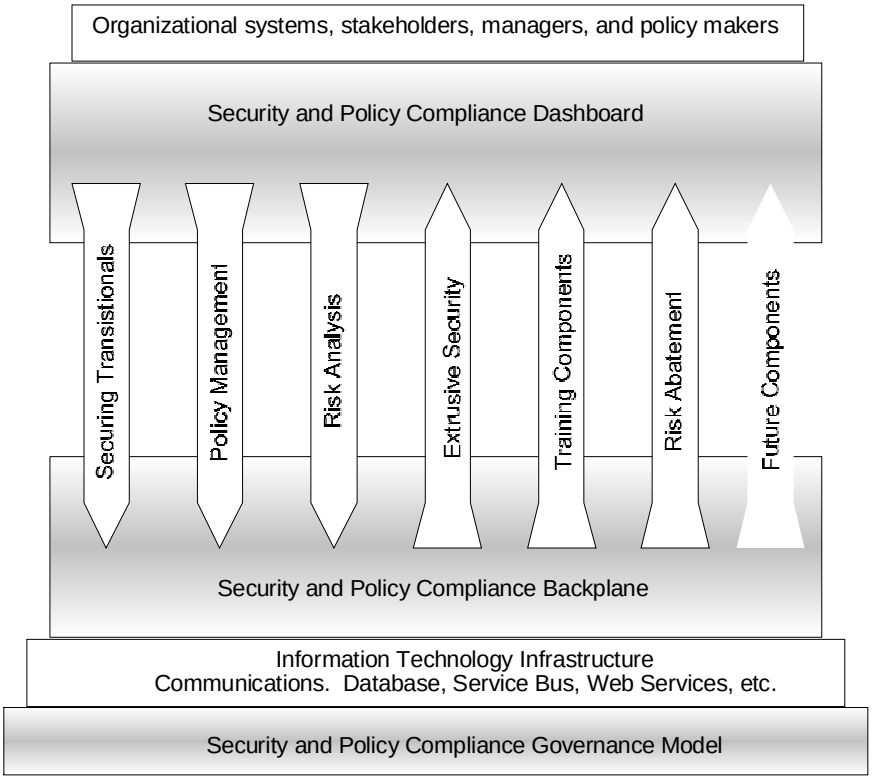

Who you are-- Georgia State University Research Foundation--We are a nonprofit corporation, created to support research activities of the University through securing gifts, contributions and grants from individuals, private organizations, and public agencies and in obtaining contracts for the performance of sponsored research, development, or other programs.

Contact: Jim D. Flowers, Interim VP External Affairs **Game-changing**

dimension—Raising the stakes

Concept--. Extrusive Information Security Systems ("Extrusives") are information technology (IT) security systems that extend beyond the boundaries of the organizations that own and control the systems. Extrusives impose a common security architecture on information trading partners that exchange information across a common organizational boundary.

Extrusives operate on the principle that a host organization's information exchanges involve "domestic" and "foreign" IT. Domestic IT are systems under the control of the host organization, and subject to its information security management principles and technology. Foreign IT are systems that are under the control of other organizations, and not subject to the host organization's security management principles and technology Extrusives "change the game" in security by forcibly imposing an organization's security management principles beyond its boundaries and into the external IT systems to which it connects. Extrusives sense foreign IT security vulnerabilities and "invade" the foreign IT in order to close the vulnerabilities.

While potentially forceful and potentially invasive, Extrusives are designed such that they will not operate in an extrusive mode unless permitted by authority of foreign IT ownership. Such authority might be established by installation of Extrusives client-server IT on foreign IT. As a consequence, Extrusives could operate cooperatively with compliant security components in foreign systems, such as a situation where companion Extrusives are mounted by both trading partners in an information exchange setting.

Vision-- Extrusives would offer a solution to a number of intractable problems confronting many different kinds of IT systems. Below are some examples

1. Business trading partners
   A setting where businesses exchange information.
2. IT offshore outsourcing partners
   A setting where information crosses borders, leaving legal jurisdictions
3. Telecommuting, home computer use by employees
   A setting where home computers attain networked device status within a controlled network.
4. Federated organizations
   A setting where centralized control of security is limited by the basic design of distributed organizational control
5. Portable computing devices introduced into corporate networks (laptops, phones, pdas)
   A setting where wireless or wireline devices enter, exit, and reenter a controlled network environment at will
6. Educational student computer use

> A setting where inexpensive, vulnerable IT devices enter, exit, and reenter a controlled network environment at will

The large presence of home computing as a potential venue for Extrusives makes the concept a potential great-leap advance for better management of Internet-connected home computers. Many home computers act as unwilling hosts for Botnets and traffic relays. Extrusives technology might provide an avenue whereby networks, such as the Internet itself, can protect itself against corrupted or vulnerable network nodes. In this way, Extrusives might not only enable Universities and Businesses to protect themselves from corrupted student or employee computers, but also enable a national network infrastructure to protect itself from widespread corruption of vastly distributed control nodes.

Method-- This approach to research involves a constructive approach to knowledge, basically learning about problems and solutions through the scientific act of designing solutions for problems. Essentially acts of "designing" and "theorizing" are merged into a single science process that results in the construction of a solution artifact. Acts of "artifact evaluation" in the context of problem-solving and "field experiment" are merged into another scientific process. These fundamental processes define a design-centric approach to science that directly engages the intractable problems of societies with the scientists capable of creating innovative solutions for these problems.

Dream team-- Georgia State University, Board of Regents of University System of Georgia (36 institutions and their Chief Information Security Officers), industry IT risk expertise (such as PwC), Oracle information security practice or information systems, Six Sigma Management experts, State of Georgia technology authority, commercial vendors with policy and/or technical compliance solutions, State and Federal legislatures, Department of Homeland Security, Department of Education, individuals across Georgia State and the University System of Georgia.

Who you are-- Georgia State University Research Foundation--We are a nonprofit corporation, created to support research activities of the University through securing gifts, contributions and grants from individuals, private organizations, and public agencies and in obtaining contracts for the performance of sponsored research, development, or other programs.

Contact: Jim D. Flowers, Interim VP External Affairs

Game-changing dimension—Changing the rules – using the dynamics of transitionals to create and adaptive cyberculture which changes as the threats change

Concept--. Transitional Information Systems ("Transitionals") are information technology (IT) systems undergoing a passage in their fundamental situation. These systems are comprised of component units that are autonomously changing from one state, stage, or form and developing or evolving into another such state, stage, or form. These changes are continuous and non-synchronous.

A Transitional presents intractable problems for information security for two fundamental reasons. First, system security is a weak-link phenomenon. If one unit in a system is not-secured, it provides an exploit path into all other units in the system. Where units are continuously changing, a unit may change itself into a configuration that unravels the security across the system. Further multiple units may change in ways that retain individual security, yet end up in a configuration where the security of the whole system has been defeated. Second, centralized control is an unstated premise of security governance approaches. Security governance and management schemes assume that a central security organization can impose regulation on system units. However, the nature of Transitionals precludes centralized control. Transitional units are often autonomous and can overrule centralized control.

These two problems, the weak-link phenomenon and decentralized governance, combine to make security in a Transitional problematic. In order to be effective, the principles of security management in a Transitional must be fundamentally different than those assumed by much of the security literature.

Vision— Dynamic settings that initiate transitional systems are commonplace and there are many examples that include ad hoc that are often cobbled together on-the-fly in order to meet unexpected or unplanned information requirements. These systems are rarely developed from scratch. Usually these systems are composed by linking together preexisting, dissimilar, and/or autonomous systems that may be crudely integrated as component sub systems. The use of programming techniques such as glue code or wrappers can be engaged to butt together systems never originally intended to interoperate. Examples of settings in which ad hoc transitional information systems commonly arise include Military coalitions, emergency response, and initial merger and acquisition systems.

> *Military coalitions* usually arise in response to unexpected unified operations involving forces of different states or nations. While there may be (or may not be) some common preparation for interoperability of systems, the actual exercise of IT–based command, control, and intelligence across organizational boundaries may raise unexpected and unplanned information flows.

Emergency response coalitions usually arise in response to large scale disasters. These coalitions may involve different kinds of organizations (police, fire, medical, military, volunteer, etc.) assembling from different municipalities, states or nations. As with military coalitions, there may be (or may not be) some common preparation for interoperability of systems. But the actual exercise of IT–based command, control, and intelligence across these organizational boundaries may similarly raise unexpected and unplanned information flows.

Initial merger and acquisition systems may similarly throw dissimilar systems together in a setting where information across organizationally boundaries is badly wanted.

The nature of Transitionals leads to at least four different kinds of security inhibitors for which there are innovative theoretical bases available as foundations for solutions.

1. *Central versus distributed control* is addressable from alternatives like autopoiesis, self-referencing and self-organizing systems theory.
2. *Social and political pressures* are addressable from alternatives like communities of practice theory.
3. *Application level incompatibilities* are addressable from complexity theory.
4. *Emergence driven incompatibilities* are addressable from process and variance theory and theories of information business decisions.

Method-- This approach to research involves a constructive approach to knowledge, basically learning about problems and solutions through the scientific act of designing solutions for problems. Essentially acts of "designing" and "theorizing" are merged into a single science process that results in the construction of a solution artifact. Acts of "artifact evaluation" in the context of problem-solving and "field experiment" are merged into another scientific process. These fundamental processes define a design-centric approach to science that directly engages the intractable problems of societies with the scientists capable of creating innovative solutions for these problems.

Dream team-- Georgia State University, Board of Regents of University System of Georgia (36 institutions and their Chief Information Security Officers), industry IT risk expertise (such as PwC), Oracle information security practice or information systems, Six Sigma Management experts, State of Georgia technology authority, commercial vendors with policy and/or technical compliance solutions, State and Federal legislatures, Department of Homeland Security, Department of Education, individuals across Georgia State and the University System of Georgia.

Who we are:

- GMU Center for Secure Information Systems led by Dr. Anup Ghosh with Dr. Sushil Jajodia and Dr. Angelos Stavrou.
- BAE Systems led by Dr. Srikanta Kumar
- GMU's Critical Infrastructure Protection Program led by LTG (R) Mick Kicklighter
- GMU's Center for Air Transportation Systems Research led by Dr. George Donohue and Dr. Lance Sherry
- GMU's Center for Geospatial Intelligence led by Dr. Peggy Agouris

Game-changing dimension – Board

- Change computing infrastructure from largely unmanaged desktop systems to diskless thin clients that connect to cloud.
- Change target space by presenting different & diverse virtualized servers that changes with requests

Concept:

Currently users tend to be the greatest threat to the enterprise. Simply surfing the Web, downloading and playing multi-media content, and opening email is sufficient to compromise desktop systems. We propose two approaches to change the game for desktop and server computing in order to make it significantly and quantifiably more difficult for the adversary:

1. To protect the user environment, employ diskless thin clients connected to managed virtual desktops in the cloud instead of workstations that users control. The virtual desktops reset to their pristine condition after each use
2. To protect critical network services, change the attack surface area by presenting a different server image for each request. The server image presented is chosen randomly from a pool of functionally equivalent servers. The uncertainty in which server is presented will thwart prepared attacks. Each presented server is presented in pristine condition. Change attack surface area for servers by presenting diversified software stack that is functionally equivalent.

Vision:

In our vision of computing, users use diskless thin clients to access virtual desktops in the cloud. Furthermore, the virtual desktops are stateless in the sense that persistent changes are not made to the operating system and services. Instead, any persistent data is stored separately in networked file servers as non-executable data. Any changes made to the virtual desktop are removed after the user terminates his session.

The solution would be rolled out to enterprise users initially, simply because of the high bandwidth, network services, and managed desktops enterprise users typically get. Eventually, the service can be rolled out to home users via ISPs, especially has high bandwidth to the home becomes more prevalent outside of urban/suburban areas.

Current virtualization technologies and cloud computing services support the needed

technical capabilities already. For this to become real, enterprises will have to plan for migrating from current desktop-centric hosts, to cloud-centric computing. This is a trend that certain application service providers (e.g., Software as a Service (Saas) vendors) have already successfully pushed.

For our server-based solution, we need to develop approaches to automatically diversify software images to be structurally different, but functionally equivalent. Prior work funded by DARPA can be leveraged. In addition, we can leverage our own work in trustworthy feedback control of servers to manage different virtualized server entities to handle different requests.

Method:

We've developed related concepts around presenting pristine operating systems for each application session. We expanded this concept here to diskless computers that connect to a cloud computing infrastructure. The assumptions that underlie our method is that sufficient network bandwidth is available to make connecting to virtual desktops seamless for users. We depend on the network to continue to stay up.

Dream team:

We have assembled a multi-disciplinary university team representing different sectors, technologies, and solutions, including the Center for Secure Information Systems (CSIS), Critical Infrastructure Protection Program (CIPP), Center for Air Transportation Systems Research (CATSR), and the Center for Geo-Spatial Intelligence.

In addition, on our dream team, we believe it is important to have key performers in the system integrator space to implement these solutions on the Federal side, a key government advocate, an entrepreneurial start-up to develop new ideas and bring them to market, a venture capital firm to provide investment capital and market guidance, and major commercial players to promote wide-spread adoption. Our dream team members are listed below.

System Integrator: BAE Systems
Venture Capital: Grotech, Novak-Biddle
Entrepreneurs: Secure Command
Govenmentt Advocates: DHS Infrastructure Protection, JTF-GNO, NSA
Industry: Verizon, Google, Microsoft, Fortune 500

We believe having all of these partners work together completes the value chain for developing and bringing game-changing technologies to market and enabling wide-spread adoption. We are fortunate to have relationships with these organizations listed above.

RFI Name: RFI – 3 – National Cyber Leap Year

Title of Concept: Phone-Based Authentication Helps Hunt Cyber Criminals

RFI Focus Area: Morph the Game Board

Submitter's Contact Information:

Name: Girish Kothapati

ID: 215509 University: Governors State University.

Concept:

Informs when a hacker tries to login: A server-based software, say GUARD, rings an employee's phone as part of User authentication and requires the employee to enter a PIN number to access their e-mail, intranet, SAP or any other server-based application.

Catches the Hacker: When confirmed that it's a false login, the software traps the hacker inside a mock-up version of its site, even populating the site with false information to keep the hacker thinking he's just struck gold. With the hacker still logged in, the security team swoops in, trying to pinpoint the hacker in a bid to aid law enforcement and figure out exactly who and where the hacker is.

Vision:

The vision is to hunt for Hackers with the help of a server-based software, which when installed, checks for a two-point authentication. While at first blush this seems like any other possible piece of User authentication, a particularly interesting scenario makes this Software for a mobile worker backed by a thoughtful security team.

Method:

Say an employee is out of the office, whether on vacation or driving to work. A hacker logs into his company's software using the employee's login and password. GUARD Software, which is installed as an agent watching logins to the server and also is integrated with Active Directory, calls the employee on his mobile phone, tells him via voice prompt that someone has logged into such-and-such application, and asks for his PIN number.

Some online services, like Pay Pal, have begun doing similar things, like sending a PIN number via SMS to a user who then enters that PIN into a browser form.

In this case, here's where the interesting part comes in. If the second factor of a two-point authentication is a thumbprint or a USB key, the process would likely be over at this point, because it would likely require too much work for a hacker to reproduce a specific thumbprint or also steal a relevant USB key. However, since the second point here is pro-active, the employee whose login is being used has the chance to notify his employer that someone is falsely using his login. The employee just hit #9 on his phone keypad and that tells the employer that the login is fraudulent.

Innovative security organizations can take this one step further. For example, a financial institution sets up a honey pot whenever there's a false login, trapping the hacker inside a mock-up version of its site, even populating the site with false financial figures to keep the hacker thinking he's just struck gold. With the hacker still logged in, the security team swoops in, trying to pinpoint the hacker in a bid to aid law enforcement and figure out exactly who and where the hacker is.

Dream Team:

Strong Security Organizations, Secure mobile & Internet service providers, Oracle or any other database company, GPS, Federal Communications commission, etc

One of the most important elements associated with this two point authentication software is how it can be applied or integrated to any client server application: thus a hard challenge can be issued inside any application: Remote access, Secure FTP, SSO, Terminal Service application, Online Financial Services or even to another Saas application delivered in the cloud. If used in conjunction with strong security policy, this software can eliminate fraud in all areas of businesses and commerce.

Global Uni-Docs Corporation
P.O. Box 7123
Dallas, TX 75209-0123

Name: RFI–3 – National Cyber Leap Year, Global Uni-Docs submission #1

Title of Concept: "*Mitigating Counterfeit Parts and Document Manipulation using a Content-Centric Security (CCS) Approach for Product Life-Cycle Certifications*"

RFI Focus Area/Game-Changing Dimension: We morph the game-board by efficiently assuring that both sender and receiver are authenticated and that electronic content has not been modified. Imagine a world in which part certification documentation actually authenticates technicians, mechanics, QA personnel, auditors and others who participate in the events associated with each unique part.

Submitter's Contact Information – David M. Shaw, Global Uni-Docs Corporation, 214-718-0325, PO Box 7123, Dallas, TX 75209-0123

Who we are: Global Uni-Docs Corporation, an information integrity start-up pioneering Content-Centric Security (CCS). CCS is an approach that will change the rules on how we certify product life-cycle events. www.gud.cc

Concept: The threat to our cyber-infrastructure fundamentally results from an inability to enforce information integrity. Increasingly sophisticated attackers are successfully beating existing IT security measures. The root cause of many of today's cyber problems is that the tools we use for protecting content are separated from the information itself. But what if we interwove information and its protective shield together such that they could not be separated? What if we could build a content life-cycle that held multi-factor authentication rule-sets, created an audit trail, and had behavioral controls concerning how the content could be utilized? Existing approaches are network-centric, Our approach enables peer-to-peer files transfers with the content protecting itself after release. Attackers must act at a granular level which radically changes how the game is played.

Vision: The GUD team will place authentication procedures and behavioral controls into electronic content itself. Unlike the current "exoskeleton" model in which soft information is protected by cyber-walls, GUD will harden content with an "endoskeleton" so that it can determine for itself who can do what and when they can do it. Like a train with pneumatic brakes, information will default to safe, instead of defaulting to at-risk. Consequently, the content remains secure and is continuously self-protecting, without dependency on a central server, and regardless of where it may be intentionally or accidentally sent.

Initially we envision the construction of authenticated histories of critical or high-value manufactured parts. For the first time authenticated histories, bound by a unique ID, will use controlling rule-sets for user profile trust-levels, compliance requirements (export, environmental, QA, etc.), business process requirements, and other factors. Each life-

cycle event certification will be captured within the content. With each event/access, user profile information is aggregated within the content, creating an unimpeachable audit trail. Since the process is not dependent upon a central server to authenticate each exchange, the part certification documentation is able to self-protect after release and move freely across domains to the end-nodes of the Global Information Grid.

Method: We proposa a pilot study to demonstrate the protection of high-value information in a collaborative environment. The application domain we will consider is ensuring the accuracy of manufacturing certifications in the DoD, NASA or other supply chains, thereby greatly reducing the risk of counterfeit parts being introduced into service.

Unlike existing alternatives, the game-changing deliverables are to provide authentication within the content itself and eliminate the need for continuous communication with third-party central servers, create an audit trail that is built within the content itself, and demonstrate a series of peer-to-peer exchanges in which the content self-governs after release. We envision a pilot being based upon a client-server model using peer-to-peer exchanges. Simple rule-sets will be authored into the content to replicate participants in a high-value supply chain certifying various manufacturing stages. The approach sits on top of the existing network architecture and complements existing security regimes. Since it is based upon a one-way algorithm it promises to efficiently scale across domains that currently are disparate and inhibit collaboration.

Out team has extensive expertise in securing electronic exchange. Many years ago, our founders recognized that a train-wreck was emerging due to the dependency on central authorities to authenticate exchanges. A recent study showed that the cost of cleaning up a single computer after a major attack runs $5,000 to $7,000.[1] For about the cost of cleaning up 200 computers, we can pioneer “default to safe” information protection.

Dream Team: Robert Smith, Founder and CEO, Global Uni-Docs, Corp. (Veteran-owned, Small Business) ; **Will Bralick, Ph.D**., President, Paladin Logic, LLC (Service Disabled Veteran-owned, Small Business) and Adjunct Professor of Computer Science, Southern Methodist University; **Sos Agaian, Ph.D**., Peter Flawn Distinguished Professor, Department of Electrical and Computer Engineering, The University of Texas at San Antonio; **Alain Bensoussan, Ph.D.,** Director, International Center for Decision & Risk Management and Distinguished Research Professor, Risk Management, Operations Management, The University of Texas at Dallas; **James Smith**, Principal, Critical1 Consulting; **Mike Grove**, Founder and CEO, CollabWorks™, **Alan Greenberg**, Boeing, Tech Director Cyber Business Unit, N&IS IA Lead, **Others** – DoD, NASA, DLA, Air Force, Army, Navy, NSF, etc.

[1] *Growing threat from cyber attacks,* http://www.breitbart.com/article.php?id=CNG.69709d80015b6cd5c5a20c040aa7e0b6.11b1&show_article=1

Global Uni-Docs Corporation
P.O. Box 7123
Dallas, TX 75209-0123

Name: RFI–3 – National Cyber Leap Year, Global Uni-Docs submission #2

Title of Concept: ***"A Game-Changing Approach to Protecting Patient Records and Ensuring Privacy"***

RFI Focus Area: Change the Rules

Submitter's Contact Information – David M. Shaw, Global Uni-Docs Corporation, 214-718-0325, PO Box 7123, Dallas, TX 75209-0123

Who we are: Global Uni-Docs Corporation, an information integrity start-up pioneering Content-Centric Security (CCS), an approach that will change the rules on how we protect medical records and individual patient privacy[1]. www.gud.cc

Concept – Instead of using fragmented back-end servers that limit the ability of multiple participants to securely move medical records across domains, we are going to look at placing user profiles into medical records so that the record itself self-governs who is authorized to gain access and under what conditions. Users may include doctors, nurses, medical technicians, administrators, insurance personnel and, or course the patient. The promise is a radical method to protect content and the privacy of those authenticated to access medical records.

This approach changes the rules because CCS enables the author of a medical record to use a one-way algorithm to determine who can access the record. The benefit of CCS enabled medical records is that they self-protect after release, regardless of where the record is sent, who attempts to access it, and once authorized what participants are permitted to do with the record after they've secured access.

Unlike other methodologies that continuously depend on a back-end central authority for authentication, our rule-changing approach allows the content to determine who is authorized. For the first time content including image, document, text, audio or video files will be able to self-protect ensuring the privacy of both the record and those who have accessed it.

Vision – GUD envisions deploying CCS using a client server model to protect medical records and ensure the privacy of the patients. We will model a pilot based upon the creation of a simple electronic medical record. Unlike existing alternatives this rule-changing pilot will consider multiple participants whose access rights are authored into the initial record based upon their trust level, actual unique ID or other rule-sets. The record will consider simple "need-to-know" access rules. A key rule-changing element of this approach will be to demonstrate how the secure record moves across various

[1] "One year later: Five lessons learned from the VA data breach", http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9022678

domains. Only authorized users will only be able to gain entry to the covert portion of the record. Content may include digital documents, text, images, audio or video files.

Method – We propose a pilot that will differentiate CCS from existing alternatives and their root cause failure points which typically are due to the separation of content from authentication procedures. Network-centric approaches to protecting content, especially in environments with high-value content in highly-collaborative systems break-down due to the dependency upon a centralized authority to perfect an exchange. Problems associated with the protection of medical records and patient privacy are likely to grow as more content, devices, and applications continue to grow with more fragmented infrastructure and varying trust levels of users.

Unlike existing alternatives, our approach will provide authentication within the medical record itself and eliminate the need for continuous communication with third-party central servers, create an audit trail that is built within the content itself, and demonstrate a exchanges in which the content self-governs after release. We envision a pilot being based upon a client-server model using peer-to-peer exchanges. Simple rule-sets will be covertly authored into a record to replicate participants in a medical record environment. The approach sits on top of the existing network architecture and complements existing security regimes. Since it is based upon a one-way algorithm it promises to efficiently scale across domains that currently are disparate and inhibit collaboration

Out team has extensive expertise in securing electronic exchange. The dependency on central authorities to authenticate exchanges has cost enterprises significantly. A recent study showed that the cost of cleaning up a single computer after a major attack runs $5,000 to $7,000.[2] For about the cost of cleaning up 200 computers, we can pioneer "default to safe" information protection. Unlike existing alternatives we change the rules by creating a pro-active methodology to provide a higher-level

Dream Team: Robert Smith, Founder and CEO, Global Uni-Docs, Corp. (Veteran-owned, Small Business) ; **Will Bralick, Ph.D**., President, Paladin Logic, LLC (Service Disabled Veteran-owned, Small Business) and Adjunct Professor of Computer Science, Southern Methodist University; **Sos Agaian, Ph.D**., Peter Flawn Distinguished Professor, Department of Electrical and Computer Engineering, The University of Texas at San Antonio; **Alain Bensoussan, Ph.D.,** Director, International Center for Decision & Risk Management and Distinguished Research Professor, Risk Management, Operations Management, The University of Texas at Dallas; **James Smith**, Principal, Critical1 Consulting; **Mike Grove**, Founder and CEO, CollabWorks™, **Alan Greenberg**, Boeing, Tech Director Cyber Business Unit, N&IS IA Lead, **Others** – NIH, NSF, NIST, VA, major hospital or healthcare network provider, etc.

[2] *Growing threat from cyber attacks*
http://www.breitbart.com/article.php?id=CNG.69709d80015b6cd5c5a20c040aa7e0b6.11b1&show_article=1

Global Uni-Docs Corporation
P.O. Box 7123
Dallas, TX 75209-0123

RFI Name: RFI–3 – National Cyber Leap Year, Global Uni-Docs submission #3

Title of Concept: "***How Content-Centric Security (CCS) Content Controls Changes the Economic Costs for Rogue Actors***"

RFI Focus Area: Raise the stakes

Submitter's Contact Information – David M. Shaw, Global Uni-Docs Corporation, 214-718-0325, PO Box 7123, Dallas, TX 75209-0123

Who we are: Global Uni-Docs Corporation is an information integrity start-up pioneering Content-Centric Security, (CCS) an approach that will proactively raise the stakes and discourage cyber attackers by increasing their costs and decreasing their odds of successfully disrupting our critical infrastructure[1]. www.gud.cc

Concept – Placing rule-sets into content so that the digital file itself self-protects and self-governs. The content is self-protecting based upon the value of the information, the nature of the collaboration participants, and the content life-cycle. The idea is to use CCS as a methodology that creates economic disincentives to attackers. Current methods typically fail due to the network-centric perimeterization of data protection. Once the crab shell is broken, all of the meat is vulnerable and the entire system is impugned. As such, the economic reward for an attacker can be very large and justify the time/effort expended.

Our approach flips that model and binds content to the authentication rules. Current methodologies enable one hacker to get into a network and have access to thousands of files. The costs to defend against these attacks, to remedy known failures, and to engineer solutions are growing and are reactive. With CCS, each organ within the crab self-protects so regardless of whether the crab shell is broken or not, the critical organs are safe.

Left out of the usual analysis is the value of the content which is compromised. Reactive solutions are of little value if the compromised content is of unique value. The cost of losing one's social security number pales compared to the forced failure of the power grid, or impugning a network by manipulating its content. CCS forces attackers to expend time/money attempting to hack what is essentially robust, object-oriented content. The granularity of our approach creates a game-changing economic disincentive for rogue parties. Unlike other alternatives, we re-define multi-factor authentication at the content level. Security factors may include user profiles, network addresses, GPS coordinates, time stamps, device fingerprints, etc.

[1] Hiroshima 2.0, Stephens, Brett, The Wall Street Journal, 04/14/09, http://online.wsj.com/article/SB123966785804815355.html#mod=djemEditorialPage

Insider trust issues create other significant problems that currently are largely addressed in a reactive way. Forensic techniques are used to detect variances after–the-fact. Those are of little good if the power grid is down, the financial payments system is impugned or munitions/hazardous material have been stolen. The CCS approach creates an environment that is similar to having a camera on top of a traffic light. If rogue insiders understand their user profile is captured in an unimpeachable fashion, it will deter hostile actions and proactively mitigate insider trust risk.

Vision – One person hacked the University of Texas network, escaping with over 200,000 personal data files. Using our approach, that same person would have had to try to hack 200,000 independent files protected with our proprietary CCS technology. Recently, Citigroup had thousands of data tapes stolen as they were shipped by UPS to Experian. The rogue parties modified the manifests to redirect the shipments. Our approach, with proper rule-sets engrained into the manifests, would have kept those tapes from being stolen. In short, our approach renders content too costly to pursue, and even if successful the attacker may end up with something that is of little value.

Method – We propose establishing a pilot test using a simple document exchange over a peer-to-peer platform. We can scale the technology for larger systems such as power grids or ad-hoc mobile networks. Engraining factors such as multiple user profiles based upon trust-level, a GPS coordinate, time-stamp, device fingerprint, etc., based upon the required business process, will provide multiple factors that are unique to a particular event. The approach sits on top of the existing network architecture and complements existing security regimes. It promises to efficiently scale across domains that currently are disparate and inhibit collaboration.

A recent study showed that the cost of cleaning up a single computer after a major attack runs $5,000 to $7,000.[2] For about the cost of cleaning up 200 computers, we can pioneer “default to safe” information protection and radically raise the stakes and increase the economic disincentives for rogue actors.

Dream Team: Robert Smith, Founder and CEO, Global Uni-Docs, Corp. (Veteran-owned, Small Business) ; **Will Bralick, Ph.D**., President, Paladin Logic, LLC and Adjunct Professor of Computer Science, Southern Methodist University; **Sos Agaian, Ph.D**., Peter Flawn Distinguished Professor, Department of Electrical and Computer Engineering, The University of Texas at San Antonio; **Alain Bensoussan, Ph.D.,** Director, International Center for Decision & Risk Management and Distinguished Research Professor, Risk Management, Operations Management, The University of Texas at Dallas; **James Smith**, Principal, Critical1 Consulting; **Mike Grove**, Founder and CEO, CollabWorks™, **Alan Greenberg**, Boeing, Tech Director Cyber Business Unit, N&IS IA Lead, **Others** –DoD, NASA, NIH, DLA, Air Force, Army, Navy, etc.d

[2] *Growing threat from cyber attacks*
http://www.breitbart.com/article.php?id=CNG.69709d80015b6cd5c5a20c040aa7e0b6.11b1&show_article=1

From: Hans-Werner Braun
Sent: Tuesday, October 21, 2008 11:45 AM
To: Leapyear
Subject: Leap Year Idea

From my childhood on I learned that if you thin out a forest so only one kind of tree (i.e., the species you like) remains, the whole forest might keel over if that remaining species get infected.
If you keep many species, the risk of a systemic catastrophic failure is greatly reduced.

Irrespective of whether you like Microsoft Windows or not, if there is an almost singular reliance on a specific (especially proprietary) operating system, or even just a single (especially proprietary) document format like Word and Powerpoint, there is a vulnerability risk strong enough to be a threat to our national security. E.g., even leaving alone malicious viruses and so, and especially without open source software, even just the potential of a disgruntled software worker leaving a trigger-able subroutine behind scares me.

I think we should seriously consider to strongly encourage more diversity in our computational vehicles, and, if possible, rely more on publicly reviewable open source software for our critical underlying infrastructure.

Hans-Werner Braun

From: Hans-Werner Braun
Sent: Tuesday, October 21, 2008 12:53 PM
To: Leapyear
Subject: Leap Year Idea

Second suggestion of mine:

Common sense dictates that you need to be able to lock your door if you think an intruder may take advantage of your house. Whether they are entirely successful or not, Border Patrol and US Customs have the responsibility to protect the "US house." Within the US, US law is enforcible. What is not US-enforcible is crimes committed in other countries, even if it impacts US citizens.

I believe there should be a low-overhead mechanism (e.g., via DNS/UDP lookup) to determine whether IP traffic originates from a source where US law is enforcible. From logs at my servers I see many daily attacks from foreign countries, and it would sure be nice if I could do at least one of:

. prohibit the traffic if it originates from a non-US source
. exclude/include specific countries
. prohibit the traffic if it originates from a location where US law is not enforcible

Not that such a filter would have to typically be in place, but, specifically, do we have the capability to close our borders on zero notice for our cyberinfrastructure if the need arises. And, even more so, can we give everyone the ability to do so as they please? I am specifically wondering about the impact on junk email if we could elect to only accept it if at least the last-hop traffic originates from a location that has to abide by US law.

Hans-Werner Braun

ECONOMIC SIGNALS:
THE KEY TO INVULERABLE SOFTWARE

Who you are

HP, the world's largest technology company, provides printing and personal computing products and IT services, software and solutions that simplify the technology experience for consumers and businesses. HP has been protecting the majority of the world's financial data for 35 years. Information about HP is available at http://www.hp.com.

Game-changing dimension

The current game between vendors and purchasers of software and systems and adversaries who wish to exploit flaws in those products is one in which the players are largely blindfolded. Purchasers have no reliable way of determining the security quality of a product (e.g. number and severity of exploitable security defects present in released products). Vendors are under pressure to release software and systems with more and more features, with security often a secondary consideration. Even when vendors make a concerted effort to institute quality security engineering practices, they have no ability to prove to their customers that their software is secure or even relatively secure compared to another vendor. As a result, there is often little motivation for vendors to expend resources on security quality. Adversaries take advantage of the situation, developing exploits before the legitimate parties are even aware of a problem.

What if the security quality of products could be reliably measured? Purchasers could make wiser decisions, increase pressure on vendors to improve the quality of their products, and, as a result, increase the difficulty for attackers to exploit these products.

This is a game-changing proposal that raises the stakes by forcing significantly more transparency and objective evaluation of the security quality of software and systems. In other markets, reliable economic signals (e.g. CarFAX, Consumer Reports, UL) have transformed market dynamics for the better. We believe it can be done for security, too.

Concept

The proposal is to create mechanisms that provide transparent security quality signals to fuel a marketplace for high quality security products. An industry which permits many thousands of vulnerabilities is the result of an economic problem, not a technical one (technical solutions follow once the economic incentives work correctly). Any attempt to solve the problem technically without also tackling the economics is doomed to failure. We envision at least two mechanisms to fix the economics.

First, vendors would receive certification that they have an effective secure product development process integral to their product development. Transparency into this process, by public disclosure of the process, or certification by a trusted third party, would provide purchasers with a signal indicating vendors are following secure software development practices, increasing the aggregate security of products in the market.

Second, we propose creating something like an Underwriter's Laboratory (UL) for security products. Such an organization would be able to objectively evaluate the security quality of products. To be practical, the evaluation process would be low-cost, focusing on using automated security testing tools. Technologies exist today that could be leveraged, but can only signal poor security quality. With the appropriate investment in research, advancements could be made to make such testing more robust, gradually raising the bar for measures of good security quality.

These are just two examples of the types of security quality signals we envision. The broader agenda of this proposal is to develop effective signals that can lead to greater transparency and objective evaluation of the security of products in the market.

Vision

We believe that transparency and objective evaluation are the keys to improving the overall security of products in the market and does not exist today. With this approach, there will be key incentives for organizations to continually improve the security of their products. Some major organizations, including HP and Microsoft, have adopted secure product development lifecycles, but the industry as a whole is still far behind.

Method

We arrived at the concept from a combination of observation of the pervasively poor state of security quality across the industry, by our internal efforts to institute and measure secure product development, and review of publications on the economics of security.

That effective signals can and will be found is an assumption, but the ideas above are promising. That other industries have been transformed by the creation of such signals (new and used cars, UL listing for electrical appliances, etc) provides good reason to believe such signals can be found and can transform the Cyber Security landscape.

Dream team

The dream team would include broad participation from major IT vendors, especially those who have experience deploying secure product development lifecycles; organizations like UL, who have vast experience in testing and certifying products; economists and insurance companies, who know how to develop risk models based on testing outcomes; security researchers, especially those who perform automated security penetration testing and program analysis; and public sector organizations, such as NIST, NSA, and DISA who develop security standards and have experience with Common Criteria Certification.

BRIDGING THE IT GOVERNANCE-OPERATIONS GAP

Who you are

HP, the world's largest technology company, provides printing and personal computing products and IT services, software and solutions that simplify the technology experience for consumers and businesses. HP has been protecting the majority of the world's financial data for 35 years. More information is available at http://www.hp.com.

Game-changing dimension

Security is an asymmetric game that favors the bad guys. The bad guys need only find one vulnerable path along which to launch an attack. The good guys, on the other hand, have to get everything right: governance, configuration, monitoring and troubleshooting. The scale and complexity of our systems further amplifies the asymmetry of the game. IT organizations today are preoccupied supporting operations in a way that does not scale, while hackers are constantly inventing new ways to exploit systems in ways we never envisioned.

Enterprise security policies govern the process for delivering IT services in a consistent and cost-effective manner, backed up by accountable decision making; the intent is that sound governance will improve security. In reality, system security depends crucially on correctly configuring low-level settings on thousands of components across the entire technology stack – network devices, infrastructure servers, and software applications. But coordinating all these configurations is an ad hoc process, independent of any systematic top-down policy. It is a time consuming and error-prone manual process. In the absence of guiding principles and abstractions, those in a position of defining policy have no ability to influence the real security of their systems, because they have no way to specify policies that can be translated into configurations. Security failures arising from misconfiguration will remain rampant as long as IT activities remain scattered across poorly coordinated IT silos.

Concept

Our idea is to develop abstractions and principles for defining high-level system-wide security policies that will drive the security configurations of individual components. This is a game changing proposal that reverses the asymmetry in favor of the good guys. By defining system security policies at a high level of abstraction, overall system security could be automated; security will no longer remain an amorphous property whose control is distributed among thousands of systems administrators.

Attackers can no longer rely on exploiting known vulnerabilities and configuration errors – instead, they must discover new and previously unknown vulnerabilities. They must work harder. On the other hand, the good guys will free up considerable resources that can be used to develop stronger defenses and keep proactively ahead of the bad guys.

Vision

We need new ways to think about high-level policy abstractions that can bridge the gap between governance and operations. We need new ways to define how software is configured that will be more amenable to being guided by policy. The challenge is developing policies that can be defined with business goals in mind, yet can be used to guide operations.

We have identified several technical challenges, which if solved, would enable significant progress. First, we need more robust ways to define and represent policies. Technical policy definitions today are limited to being prescriptive in nature, and refer only to controls on individual entities and processes. This approach is awkward in constraining system-wide behavior; distributed policies tend to be brittle and must be rewritten when the control technology changes. Second, we need ways to automatically translate high-level policies into actionable control policies on low-level components, as well as policy checking and remediation mechanisms. Third, we need ways to reason about global properties of systems that are in flux. Finally, since security properties depend both on the functional behaviors of individual components and in the interconnections between them, we need standard models of software components that capture information flows; this will require cooperation from vendors.

Significant research advances in recent years, although mostly focused on specific problems in isolation, as well as current trends in industry towards centralized enterprise-wide configurations and common representation standards lead us to believe that our vision is attainable.

Method

We arrived at the concept from observations of the current state of IT security service practices, the evolution of IT management software, and progress in several related, but disjoint, research areas ranging from network security to distributed deduction. We have worked on related problems for some time and, while the overarching vision is admittedly difficult, we are confident that these challenges can be overcome by a combination of research advances and standardization efforts within industry.

Dream team

NIST, Major vendors of network, applications and management software, Security and Systems Integrations specialists, Academic and Industrial researchers in areas of software specification, modeling, security, and verification.

SECURITY BY RANDOMIZATION

Who you are

HP, the world's largest technology company, provides printing and personal computing products and IT services, software and solutions that simplify the technology experience for consumers and businesses. HP has been protecting the majority of the world's financial data for 35 years. More information about HP is available at http://www.hp.com.

Game-changing dimension

Any potential game changer has to recognize two fundamental facts. First, the adversary has the initial advantage because his job is to only find one (or a few) paths to successful compromise, whereas the defender's job is to thwart all or most attacks. A reactive security strategy that depends on guessing the adversary's exact moves (such as virus testing) is bound to eventually fail. Second, current system management philosophy depends on keeping our systems in static configurations that allow the adversary to know exactly what state the system is in. What if we could build our systems to keep changing constantly so that they look orderly to the legitimate users but chaotic to the adversary? Rather than incrementally improve our systems security, can we take the board away from the attacker by building our system security to be independent of the adversary's strategy?

Concept

The key idea of our game-changing approach is randomization. Randomization is a useful theoretical tool: it breaks symmetries that can cause deadlocks in distributed systems and it can make systems uniformly hard to break, such as in cryptography. A randomized system can be designed so that (i) the system behaves unpredictably when attacked by pushing it beyond its boundaries and (ii) provides a constantly moving target that negates fixed attack strategies but (iii) a legitimate user can remain automatically "sync'ed" so she never sees any difference. We believe that this simple and hence scalable idea can be applied across the board to increase security.

For example, in information protection, the current game between vendors, purchaser and adversaries is one in which vendors create software, adversaries obtain a copy of the software, reverse engineer it, discover a vulnerability, and create an exploit which is then leveraged against a purchaser's copy of that software. What if every copy of software was randomly different in some fundamental way, so that even if a hacker developed an exploit against his own copy of the software, that exploit would not work against the purchaser's copy of the software? Such an approach could prevent a wide range of attacks currently waged against software.

Another example would be *IP Address Hopping*, in which the IP address of a particularly sensitive target is changed periodically. This approach has already been used to thwart certain kinds of attacks in which an adversary uses a static IP address to attack a system. What if the IP address hopping of a machine is hooked to the DNS service that is used on an Internet scale to find machines using simple names such as www.hp.com? This address hopping scheme can be further carried out in an orderly fashion in large enterprise networks and university networks so that most of the Internet will look like it is in constant chaos to any attacker that is based on hard-wired or cached IP addresses, thereby eliminating a huge fraction of today's attacks in one stroke.

These are just two examples of the ways in which randomization could be used to improve security. The broader agenda of this proposal is to develop these and other ideas that leverage randomization to eliminate entire classes of security flaws.

Vision

Our idea is not entirely new – we see isolated uses of randomness already but the key of this proposal is that the idea needs to be more widely used and integrated into larger systems. For example, randomization of the address space layout inside a computer program (i.e. where the executable, libraries, heap and stacks are located in memory), an old idea now implemented in Vista, is an effective means to thwart buffer overflow and stack smashing attacks. Single-use credit cards eliminate entire classes of credit card fraud. Randomized mechanisms are used to avoid collisions in legitimate requests, e.g. random port assignment and session ids. All these techniques have the property that the good guy, following the protocol, sees no difference in behavior. However, anyone stepping outside the protocol either sees unpredictable behavior or gets no service at all.

Method

We arrived at the concept of randomization as a result of our own observations from a number of research projects. We started from the question "What would it take to make the Internet totally chaotic for an adversary and totally orderly for a legitimate user?" It was clear from our research in cryptography, systems security, and software testing that just a small but pervasive introduction of randomization to the entire computing and networking fabric has the potential to eliminate a huge fraction of attacks on the Internet by simply changing the rules profoundly for attackers.

Dream team

The dream team would include broad participation from IT vendors, especially those who sell basic infrastructure such as servers and routers as well as academic researchers, some of whom have come up with these randomization ideas in the first place. We will also need the active participation of standards bodies such as DNS and ICANN.

Leap-ahead Technology Concept

Who you are – Dr. Yi Hu (Assistant Professor, Computer Science Department, Northern Kentucky University, Highland Heights, KY, 41099 **Dr. Brajendra Panda** (Professor, Computer Science and Computer Engineering Department, University of Arkansas, Fayetteville, AR 72701

Game-changing dimension – Change stakes

Concept – Cyber attack is a real threat to survivability of information systems. With the increasing number of databases that are accessible through the Internet, database security is a growing concern not only for government installations with critical data but also to businesses, industries and civilian sectors. These databases may store sensitive information such as social security numbers, credit card numbers, and other financial or medical information. Attackers take advantage of any vulnerability that may exist in the database system to steal or corrupt critical information. Organizations spend significant amount time and money in securing the network and host computer, but applications such as database are often ignored.

Existing post information warfare damage assessment and recovery protocols use a passive approach. This is highly undesirable. These damage assessment and recovery schemes are designed to act after the detection of attacks. Organizations have no knowledge of either how vulnerable their databases are to cyber attacks or how damage would propagate through the system, before an information attack actually occurs. Moreover, when database developers design a database, most often they do not consider the potential of a cyber attack as a factor. The result is that in most cases no extra fault tolerant mechanisms are developed to protect the most vulnerable data in the database system. Data located at distributed sites can densely couple with each other. When the database is hit by the attack, it is extremely time-consuming to recover damaged data since the damage assessment process has to be correlated among different sites. Without knowing the vulnerabilities of a database and how quickly damage may propagate, before attack even occurs, it is difficult to develop a survivable information system that can withstand information warfare. The main purpose of this research is to construct a database vulnerability assessment model that is capable of generating vulnerability profile and projected damage profile for the database system prior to an attack. We would develop a method to find out what data items form the weakest points in the system that would cause a major part of the system to malfunction if damaged, what data are more likely to cause damage to spread, and what data we may not want to recover if damaged because if these data are heavily damaged, it may not worth the cost of recovery.

Vision – Our ultimate goal is to employ our research ideas to create a database vulnerability profile visualization engine and also a projected damage profile visualization engine. The vulnerability assessment and visualization tool can help decision makers in evaluating the risks faced by information warfare better, retrofitting the databases to reduce vulnerabilities, deploying a customized intrusion detection system, and choosing the most appropriate damage assessment and recovery strategy.

The goal of the database vulnerability profile visualization engine is to generate an interactive graphical user interface that facilitates understanding data coupling information, temporal data accessing information, and the vulnerabilities inflicted by the data dependencies. It would be designed to help users identify data relationships and vulnerabilities at different granularity levels. Users can employ this tool to understand how databases hosted at different sites are related to each other, how data items are densely connected with each other inside a clique, which represents a group of data that are closely related to each other in term of data dependencies, how tables at one site are coupled with others at the same or other sites etc. It will also point out the vulnerabilities at different levels so that users can zoom in to a particular part of the database system, such as a clique, to watch such vulnerability information more closely.

The projected damage visualization engine is responsible for generating a visualized damage profile so that users can have a concrete impression of the projected damage and damage propagation trend before information attack ever happens. It can give users different views of the damage at different granularity levels. Users would be able to zoom in to a selected damaged site to explore the databases or cliques/meta-cliques that are damaged. Then they may zoom in to one particular database followed by to a damaged table to find out which tables and which rows might have been affected. The probability of each projected damaged item will also be made available in the user interface.

Method – We expect to achieve our goal by employing a data mining approach that generates concise data dependencies and temporal accessing rules that are employed to find out what data are vulnerable to cyber attacks. Our model is also expected to generate a projected damage profile if the set of assumed-damaged data and assumed intrusion detection latencies are given.

In the proposed model, a data mining engine will be responsible for discovering intra-transaction and inter-transaction data dependencies and temporal data accessing correlations during normal database operation phase. These dependencies and temporal correlations will be described as concise rules, which would be employed to generate the vulnerability profile of the database system. The vulnerability assessment unit will determine the set of data items in the system that are vulnerable to potential attacks and then generate an on-demand vulnerability profile. The vulnerability assessment unit would be exclusively dependent on the compact rules mined instead of the database log. Thus, it can generate a vulnerability profile very quickly. The vulnerability profile visualization engine would be responsible for rendering a visualized vulnerability profile that can clearly illustrate how data are coupled with each other and what data are more vulnerable to attacks than others. The projected damage generation engine will be responsible for generating a projected damage profile by assuming some data are damaged and the attack is detected at certain detection latency. It provides the organization a concrete idea on damage propagation in case of an attack. This can be thought of as a fire drill in the scenario of cyber defense. The damage visualization engine will offer a potential damage profile of data on screen that can illustrate damage at different granularities and allow users to explore damage at different parts interactively.

Dream team – Dr. Yi Hu (Assistant Professor, Computer Science Department, Northern Kentucky University), Dr. Brajendra Panda (Professor, Computer Science and Computer Engineering Department, University of Arkansas)

Who we are – Intelligent Automation, Inc. (www.i-a-i.com): We are a woman-owned firm with over 100 researchers and technical staffs dedicated to the innovative research and product development.

Game-changing dimension – Morph the game board

Concept – Software is man-made, and therefore imperfect. Before attacking a network service, attackers need to know at least two things: (1) where is the sever that provides the network service? (2) Which versions are the operating system and the application software and the associated vulnerabilities? The attacks may become much harder if we randomize the network address of the sever and obfuscate the software code.

Vision – The vision is to minimize the knowledge about the protected network services by introducing randomization into the network topology and software execution. Normal users should be able to gain various network services without the need to know where the server is and how the service is implemented. For attackers, it is difficult to collect the information such as IP, port, version of the applications and launch subsequent attacks, because that information will be changed due to randomization. Even if attackers can successfully compromise one server by utilizing some specific software vulnerabilities, there is little opportunity for them to reuse the same code to automatically compromise a large number of the same types of servers, because the code on each server is somewhat different (and different enough) due to code obfuscation, plus the fact that the code will be running in different memory space address.

Method – Morph everything we can. We could introduce multiple layers of randomness to protect the network services. First, we can randomize the IP address and Port number to make the location of the service unpredicted to attackers. The idea is that the attacker's information could be rendered stale if network services are forced to frequently change their IP addresses and Port numbers. The connection on servers may be disconnected if we change their addresses and ports when connections are still in progress. We can solve this problem by using Network Address Translation (NAT).

Second, because the bad guys go after commercial software that is pervasive and vulnerable, we can modify that software using the same sorts of tools that the hackers use to obfuscate their own code. If every copy of Operating System or application code is just a little bit different than the next, even though the base code is the same, such obfuscation makes it significantly more difficult to write malware that could infect enough machines so that it is financially viable for the bad guys.

Third, it has been observed that most attacks use absolute memory addresses during memory corruption attacks. Address space randomization techniques randomize the layout of process memory, thereby making the critical memory addresses unpredictable and breaking the hard-coded address assumption. By randomly shifting critical memory regions at process initialization time, address space randomization converts an otherwise successful malicious attack into a benign process crash. We should provide sufficient randomness to prevent successful brute-force attacks without noticeable performance degradation.

Dream Team – Government agencies as purchasers, at least in near term; law makers on liability issues in a long-term; software vendors to change mindset for software products; R&D companies and academia for innovations and solutions.

Randomized Defense Strategies to Proactively Protect Servers

Arun Sood

Who we are:

PI: Arun Sood, International Cyber Center (ICC), George Mason University (GMU), Fairfax, VA.

ICC mission is to facilitate strategic collaboration and information sharing to better identify and address global cyber issues from policy, technical, and research perspectives. We do this by (a) building a community of interest around priority issues by organizing workshops and conferences, and (b) creating and leveraging interdisciplinary teams of Computer Scientists, Engineers, Economists (especially Experimental Economists), Risk Managers, Legal experts, and Public Policy analysts. To facilitate the process of building such inter-disciplinary teams, ICC is organized as one of the few centers at GMU that reports directly to the Provost.

Game-changing dimension:

Our goal is to make it orders of magnitude harder for an intruder to succeed in achieving the purpose of the intrusion. We plan to develop an array of evasive maneuvers so that it will become progressively harder and more costly for an intruder – even one using a zero-day vulnerability – to do damage. Our first goal is to increase the difficulty of doing such damage by an order of magnitude (at least ten times). Our next goal is to raise the difficulty by an additional order of magnitude. We achieve this by trading off a small reduction in computing power in exchange for a significant increase in security. Our success will be measured not by the intrusions we prevent, but by how difficult we make it for miscreants who successfully intrude, to do measurable damage.

Concept:

To protect physical human assets, randomness and constant change are often used to protect the asset. The focus of this proposal is to develop Randomized Defense Strategies that will protect high-value computing resources. The techniques of hide, obscure, move, alter, and speed are some of the strategies that are used to protect physical assets. We plan to develop similar approaches to protect computing resources, like servers and client stations.

Vision:

Cybersecurity is a hard problem, one that hackers are winning! In spite of large investments in computer security, attackers continue to evade the most advanced intrusion prevention and detection systems. Current systems are able to detect less than 50 percent of the malware. In 2008 more than 30 million consumer records have been comprised, and data breaches are costing in the millions per breach. The problem stems in large part from (1) the constant innovation and evolution of attack techniques, (2) rapid development of exploits based on recently discovered software vulnerabilities, and (3) most defense capability needs specifics (e.g., signatures) of the attack to create a defense that will work. Incredibly, the major focus of cyber defense is based on a model in which there is no defense against most attacks until AFTER successful attacks have taken place and been detected.

The current reactive methods are inadequate because the bad guys are always one step ahead. We conclude that a robust defense-in-depth strategy requires a new, additional layer of defense that makes one critical assumption -- that *intrusions are inevitable* – and institutes measures to minimize the damages from all intrusions. In this research we propose Randomized

Defense Strategies (RDS) to protect servers located in the DMZ. . By exploiting virtualization technology, our approach has application to the cloud computing strategies that are based on delivering clients a broad range of services using virtual machines.

We have already developed techniques that reduce the exposure time of the servers located in the DMZ, and believe that by extending these techniques we can develop systems that cover all the five aspects – hide, obscure, move, alter, and speed – underlying the RDS approach. Our approach is not a substitute for removing vulnerabilities, but an additional layer of defense to manage the negative impact of inevitable intrusion. We encourage research in developing good programming techniques, but until these are available, RDS focuses on containing the losses from an intrusion even in the absence of knowledge that an intrusion has occurred.

Method;

In the last six years the GMU researchers have developed a new approach to Intrusion Tolerance – characterized as an effort to contain losses from intrusions. As compared to other intrusion tolerance techniques that emerged from DARPA's OASIS project, our approach does not rely on intrusion detection. Self Cleansing Intrusion Tolerance (SCIT) has resulted in several publications and four patent applications and links to these are available from http://cs.gmu.edu/~asood/scit. This research has been funded by federal funds and more recently there has been funding and testing by SUN, Lockheed Martin, and Northrop Grumman. Significantly, this experience shows that SCIT is able to delete surreptitiously installed malware every minute or so, without the requirement that the malware even have been detected!

Current servers are persistent, connected to the internet for long duration, and are almost like sitting ducks. SCIT converts the servers into agile, dynamic environments. SCIT achieves this **without** changing the application software. However, in this approach we are not able to defend against some attacks, for example, the Denial of Service attack. SCIT is a server-based strategy, and DOS or DDOS are often network based attacks. The RDS approach will be tailored to defend against a DOS or DDOS attack, and reduce the damage that will take place. Once again, we will not prevent attacks, or even successful intrusions, but we will reduce the damage that can occur. A critical aspect of RDS is that there are costs incurred, leading to performance tradeoffs. In this project, it is our goal to identify the cost elements, and prepare a robust cost-benefit analysis. We believe that additional security costs and this cost needs to be justified before the security approach is widely accepted. Our general philosophy is to shift from a risk-prevention focus to broader approach of risk management, and this requires collecting substantial experimental data. For this reason, a multidisciplinary team is essential.

Dream Team:

We will assemble a multidisciplinary team to work on this effort. This will include theoretical and experimental researchers and will be guided by a practical viewpoint. To meet this objective, we will constitute a project advisory board that will bring the practical experience to bear on this problem. We note that ICC has an Advisory Board with membership of leading corporations, current and former government executives and GMU faculty. So assembling a team of advisors for this project should be readily achievable. The disciplines relevant for taking this from concept to widespread application are Computer Scientists, Engineers, and Experimental Economists with Risk Management focus, Social Scientists and Legal Experts. For this expertise we will work with the researchers at GMU and where necessary get support from researchers at other institutions.

4121 WILSON BOULEVARD, SUITE 101 • ARLINGTON, VIRGINIA 22203 • www.i-lawgroup.com • tel.703.243.8100 • fax.703.243.8162

Who We Are – Internet Law Group is a law firm that specializes in tracking cyber fraudsters and their enablers, and implementing strategic offensive actions that can be brought against them. Our law firm has successfully implemented the process outlined below.

Game-Changing Dimension – Our ideas change the game in all three dimensions – by morphing the game board, changing the rules, and raising the stakes.

Concept – Build a cost-efficient, robust civil investigative and legal enforcement "engine" that systematically collects cyber crime forensic data, and processes that data through formal and informal civil legal processes to identify and pursue major cyber criminals targeting American businesses, governments and consumers. This civil enforcement "engine" will bridge the gap between technical defensive measures (that attempt to prevent cyber crime in the moment) and traditional criminal law enforcement (which lacks the resources and procedural efficiencies required to move quickly in response to cyber crime). The only reason this "engine" does not currently exist is because the costs of cyber crime fall largely on the public commons. Overcoming this tragedy of the commons is the key to success.

Vision -- Unlike street crime, cyber crime depends almost entirely on an infrastructure of legitimate service providers to support it. For example, illegal online pharmacies advertise their goods through email, ad words and SEO. They host their drugstore webpages on computers manufactured by well known makers, and those machines connect to the Internet through traditional ISPs and access providers. Cyber drug dealers have merchant accounts within the credit card system that allow them to charge consumer credit cards for the deadly drugs they sell. They ship their fake drugs using traditional international and domestic shippers, including the US Postal Service. They communicate with each other using cell phones just like you and me, and they operate toll free call centers using VOIP numbers they lease from well known telecommunication companies. And most importantly, the proceeds of their illegal enterprise move through the financial world using long-developed traditional methods. Cyber crime's dependence on these enablers is its Achille's Heel. We can largely stop cyber crime by systematically cutting off its access to any one aspect of these critical enabling services.

The civil enforcement "engine" we envision will cut cyber criminals off from the enabling services they need. It is also technically simple to operate. It begins with simple data collection and analysis tools that already exist. Complaints about spam, for example, can be collected from any number of data sources. Collectively, these sources identify millions of web pages a day. Existing technology tools can capture the source code of these web pages along with the technical data underlying their architecture (WhoIs, DNS, A hosts, etc.).

While the data is voluminous, even a cursory review of it reveals a remarkable convergence in the underlying "fingerprints." This convergence suggests a very small number of criminals are responsible for the vast majority of cyber crime.[1] Much of this convergence can be observed on the very first day of analysis. For example, a simple comparison of WhoIs look ups reveals the vast majority of domains hosting spam-advertised web pages are sold to the criminal community by a mere handful of registrars. This convergence provides strategic opportunities to identify and cut off enabling services.

The key to acting strategically is to conduct "triage" on the data collected, and create a civil law/investigative process that can bring a broad wave of low-level investigative and legal pressure on enablers across the board, and increase that pressure against those who either fail to respond appropriately or who are found to be in the best position to cut off key enabling services.

Civil law enforcers have many tools at their disposal. Automated notices to enablers can be issued at very low cost. Such notices leverage investigative resources by challenging enablers to investigate and take action on their own, and report back on the results of their investigations. In addition, undercover buys can be used to acquire financial information on crime rings. Most importantly, civil subpoenas can be issued to witnesses and enablers to acquire information that cannot be obtained informally.

Armed with all that can be discovered about cyber criminals and their enablers, strategic actions can be brought against them, their assets and the systems on which they depend.

Unlike information developed through government and criminal investigations, all the information acquired through this civil enforcement engine can be shared under appropriate circumstances across the victim base (to ensure criminals are not provided enabling services in the future), and with government law enforcers.

Method – the process outlined above is one our law firm has already successfully implemented to identify and stop cyber crime. It is a proven strategy that merits broader consideration and support.

Dream Team – In addition to the resources and skills our law firm can contribute, the process outlined above will require funding from public and private sector sources. Additional support will be needed from a small number of players in the IT community and/or academia to refine the data collection and analysis tools needed to identify points of convergence in the data. The team needed to implement this plan is small, buy in across a broad spectrum is not necessary, and no changes in law or legal processes are required for success.

---

[1] This observation is demonstrably true, and was anecdotally demonstrated most recently by the tremendous drop in spam volumes that resulted for a time when one US-based ISP was taken off line in response to an investigative story written by Brian Krebs with the WashingtonPost.com. See Major Source of Online Scams and Spams Knocked Offline, November 11, 2008 <http://voices.washingtonpost.com/securityfix/2008/11/major_source_of_online_scams_a.html> (reporting 75% drop in spam following McColo disconnection from the Internet).

Submission to RFI

Who we are

We are the security informatics research group in the School of Informatics at Indiana University. In our group, Professor L. Jean Camp's core interest is in the trust in context: security and privacy in commerce and finance; Dr. Wang has extensive research experiences in information security and incentive engineering; PhD candidate Debin Liu has been focusing on the analysis and modeling of information risks.

Game-changing dimension

Raise the stakes

Concept

From the perspective of information security economics, many security failures are not purely technical but rather caused by the misalignment of incentives. Hence, there is a clear and profound need for effective risk communication. However, it's always a challenge to effectively communicate risk information to users. What if we change this situation by "visualizing the risk" and "prompting appropriate response"? The basic idea is to quantify the risk associated with each action users could take and ask them to pay for their actions.

The quantified risk could be presented as *action price* via risk tokens or risk points, which visualize the risk and convey the perfect truth of the risk information to the users. By assigning an initial endowment of risk points and asking users to pay for their actions in the terms of risk points, the total risk a user is bearing is then visualized by her *risk budget*. Moreover, we add *punishment* to the users who exhaust their risk budget so that users can be expected with appropriate responses against unnecessary risk-seeking actions.

Vision

With the speed meter, a driver can easily know how fast she is driving. With the gas meter, a driver can be easily aware of how farther she can go. Similarly, a user is able to effectively obtain information of the risk of her action and the total risk she bear by means of the *action price* and *risk budget* respectively.

Understanding that exhaustion of gasoline would cause a huge cost such as vehicle pulled over and gasoline refill, a driver is then motivated against gasoline abuse and unnecessary vehicle usage. In a similar way, a user is prompted to take appropriate

action responding to risk communication due to the possible punishment resulted from the exhaustion of her risk budget.

Previous related researches have proposed several methods to quantify estimate of risk associated with an action request. Then given quantified risk, we now can apply Human-Computer Interaction (HCI) design theory to visualize the risk and effectively convey the information to users. Furthermore, we need to carefully design and build an incentive-based mechanism to punish users' risk-seeking behaviors and prompt appropriate response actions.

Method

We conducted human-subject experiments to discover the misalignment between the mental models of the risk communication designers and the receivers. We modeled and analyzed the incentives and motivations of malicious insiders, inadvertent insiders and common legitimate users using game theory and incentive engineering. The PhD candidate in our research group has participated in several HCI interface design projects and accumulated experience in visualization design. We also completed human-subject experiments to address users' risk behaviors. Moreover we proved that we are able to effectively and positively affect users' response decisions with incentives.

Our previous works have been published in several peer-reviewed conferences and journeys.

Dream team

The Human-Computer Interaction group in the School of Informatics at Indiana University, Watson Research Center at Hawthorne.

Secure Enterprise Networks Consortium (SEN-C)
National Science Foundation National Cyber Leap Year Request for Input (RFI)
December 15, 2008

Information Assurance Ecosystem

Who We Are

The Secure Enterprise Networks Consortium (SEN-C) is comprised of Accenture, Los Alamos National Laboratory, Sun Microsystems, and CA, Inc. SEN-C focuses on bringing leading skills together—from thought leadership and solution development to systems integration excellence. By collaborating with government, we seek to achieve outcomes that enable CNCI initiatives and improve our nation's security.

Game-Changing Dimension: Morph the Board

Concept

Success requires both goal(s) and a roadmap. Our concept defines a solution model inclusive of all aspects of a defined security/trust ecosystem that provides operations security (OPSEC). It incorporates:

- Basic, agreed-upon Technical Reference Model (TRM) for security/trust/information assurance (IA) distributed information system assets
- Ontology (security/trust/IA relationships)
- Security/trust/IA ecosystem taxonomy
- Lexicon(s) and semantics
- Stated standards for security, trust and information assurance
- Distributed policy enforcement model that is secure, trusted, and assured in current distributed information system environment
- Component Security for broad inclusion in systems and software
- Link-level encryption

Our concept recognizes multiple stakeholders for secure/trusted information technology solutions architected within an enterprise or between enterprises. This drives the need to understand the security/trust ecosystem, the nature and relationships of component pieces, taxonomy/terminology, and semantics used by stakeholders with different views. These relationships will be captured in an interactive model/tool using social/collaborative interaction between stakeholders.

Vision

The component pieces of the trust/security ecosystem are a finite set of definable entities. These include:

- Identity (management, verification, authentication, audit)
- Trust (attestation, non-repudiation, electronic signing, access, audit)
- Data (capture, management, electronic signing, storage, encryption, transmission, audit)

- Business usage (applications, rule sets, policy enforcement, service level agreements, audit)
- Lifecycle support (backup, disaster recovery, continuity-of-operations planning (COOP), remediation, end-of-life)

We believe mapping can make sense from different stakeholder perceptions. A common language can help discuss and model the entities. Both would help lead to an acceptable TRM to define architecture and use case examples. They also could be a factor in needed technical modeling. With an agreed-upon, lower-level model, stakeholders would be a better position to discuss and address complex issues of governance (Infrastructure management, user needs, business needs, security posture, policy enforcement, training, certification, and their evolution). We would have a security/trust infrastructure built in as part of the solution rather than added ad hoc.

Using currently available social engineering and collaboration tools, a cross section of stakeholders could engage periodically. A focused technical team could document findings and put use case information into modeling tools for analysis and reporting. The project's security/trust roadmap would guide this spiral approach.

Method

We recommend collaborative tools (Cisco MeetingPlace audio conferencing, wikis, Microsoft SharePoint, Accenture Collaborative Innovation Service (ACIS)). A core technical team will lead discussion, capturing progress, summarizing results, addressing stakeholder feedback, and maintaining schedule. The first task after stakeholders are identified will be to define and establish an eco-system roadmap supported by a basic taxonomy, definitions, and semantics for discussing the eco-system and roadmap. Next, they will identify low-hanging fruit and bleeding–edge technology. They also will consider ongoing standards work and affiliation of our ideas and concepts with other groups interested in moving data into existing or new standards. Further, they will move ecosystem data into a modeling tool and select demonstrations to test the validity of this approach. Success will be measured by yet-to-be-determined factors that equate results to stakeholder satisfaction and broader community acceptance of the work.

This approach is based on work we participated in for DOD related efforts to capture data and model interaction of Maritime Domain Awareness fusion approaches.

Dream Team

- Stakeholders from multiple domains (minimal time/week, periodic audio meetings)
- Participants from the academic community, standards groups, and government committees

Who we are – A partnership between Netronome and Concurrent Technologies Corporation

Netronome is a leading supplier of highly programmable semiconductor products that are used for intelligent flow processing in network and communications devices. Netronome's solutions include tightly integrated network flow processors and acceleration cards that scale to more than 20 Gbps. They are used in carrier-grade and enterprise-class communications products, as well as virtualized servers and appliances that require deep packet inspection, flow analysis and content processing, all at very high speeds for millions of simultaneous flows. Netronome is headquartered in Pittsburgh, Pennsylvania, with core operations in San Jose, California and Boxborough, Massachusetts. – www.netronome.com

Concurrent Technologies Corporation (*CTC*) has been putting ideas into action since 1987, supporting a wide range of high-priority defense requirements and helping U.S. industry compete in the global market. *CTC* is an independent, nonprofit, applied scientific research and development professional services organization. *CTC* is classified as a section 501(c)(3) organization. *CTC* serves our client base with over 1,400 scientific, technical, and business professionals in over 50 locations across the nation. – www.ctc.com

Game-changing dimension – Network users (malicious and non-malicious) understand where to find the blind spots in the network. Those blind spots will be exploited unless something is done to understand what is happening in previously un-monitored areas of the network. A complete tool set is needed that will provide cyber defenders with visibility into data in motion.

Concept –Our concept will provide that visibility and also the tools needed to analyze and determine whether improper or malicious content is in the stream. The issue at hand is the rogue contractor or employee who looks like they are trusted individual, but may be enticed to leak sensitive information for lucrative payouts or other reasons (job loss, etc). The demographic of this type of person could be anyone; there is no stereo-typical characteristic defining them. End users are aware that encrypted transmissions leaving the enterprise are often disregarded by network security devices because of the inability to decrypt traffic. This traffic, defined by the Secure Sockets Layer (SSL), is now becoming a ubiquitous protocol because of its very nature – encryption. Insiders can now tunnel sensitive information over SSL encrypted applications freely without detection.

Vision – The vision is to provide a method by which network security personnel would have visibility into encrypted (SSL) outbound network traffic streams. SSL encrypted communications have become a favorite attack vehicle for hackers to infiltrate computing resources but security architects have no reliable way to "see into" or inspect the plaintext of these encrypted flows. SSL has become the ubiquitous choice to secure web-based transactions as well as other applications such as secure email and SSL-based VPNs/extranets. As such, SSL-encrypted traffic has grown to constitute a significant percentage of data transmitted to and from the enterprise or home user. This encrypted traffic poses a security risk, though, as SSL also provides a mechanism for more nefarious applications. Network-based threats, such as spam, spyware and viruses—not to mention phishing, identity theft, and other forms of cyber-crime have become commonplace. SSL can make it difficult or impossible for network administrators

to locate these threats and enforce corporate acceptable use policies and to ensure that threats, like viruses, spam and malware, are stopped before they reach individual users.

Method – Many security companies have attempted to solve this problem, but only SSL proxy solutions have been available for deployment. In such a case, users would know there is an active monitoring device and would be thwarted in their attempt to leak data. The ideal solution would be with a transparent SSL decryption device like the one from Netronome Systems, Inc. Because of its transparency, users inside the enterprise would never know it is there, thus becoming susceptible to having their traffic decrypted, monitored, and then re-encrypted. Netronome has teamed with *CTC* to provide the total suite of tools and services needed to provide visibility and analytical capability to combat this threat. All network transmissions, including SSL, will be clearly understood.

Dream team – *CTC*, Netronome, NetWitness

RFI-3 – National Cyber Leap Year

Title: Alternative communication networks for disaster recovery

Submitter: Kevin Fall, Intel Research Berkeley, 2150 Shattuck Ave., #1300, Berkeley CA 94704

Organization: Intel Corporation's research lab in Berkeley, CA.
http://berkeley.intel-research.net

Game-changing dimension: change the rules, morph the game board

Concept: Using store-carry-forward network routers deployed in mobile vehicles (e.g., cars, boats, UASs) with wireless micro base stations, a viable alternative communications infrastructure can be employed in times where the conventional infrastructure has either failed or become unavailable due to congestion, and this infrastructure can be used by existing portable devices already familiar to first responders. This has the potential to avoid many of the communications issues raised during natural disasters such as Katrina in which alternative communication options were either unavailable or unfamiliar to first responders (effectively precluding the use of many of these alternative communications options such as satellite phones). A vehicle-based network could potentially act as a back-up route for the conventional infrastructure if it experiences failures in its backhaul connectivity.

Vision: Today, cellular phones and radios require an operating infrastructure to function effectively over a moderate geographic area. Failure of infrastructure can lead to C2 failures endangering life and property. Allowing such devices to continue operation (perhaps somewhat degraded, if necessary) even when the traditional infrastructure has failed can be achieved using vehicle-hosted store-carry-forward gateways that allow both traditional (low-delay) routing of data, but also the ability to physically carry and replicate messages from place to place as a consequence of physical movement. Vehicles (e.g. cars) often have greater endurance and redundancy than single handsets due to onboard power storage and generation.

Method: This approach is not the same as previous work on Mobile Ad-Hoc Networks (MANETs) that focused on routing when continuous connectivity is possible due to an assumption of high node density. This concept has evolved after experience in designing and developing a new network architecture which focuses on moving data objects across intermittent and sparse networks (called Delay or Disruption Tolerant Networking, DTN, an effort that has received DARPA funding over the last few years). DTN moves data "bundles" (objects) among store and forward routers that may cache data for extended periods of time. Such nodes may elect to delay transmitting data for a variety of reasons (saving energy, avoiding detection, waiting for pre-scheduled connectivity opportunities for other mobile nodes such as LEO satellites). Previous research studies in routing over such fabrics (e.g., Grossglauser's "Island Hopping") suggest the viability of the approach.

This approach relies on an assumption that cooperation can be established between the owners of various vehicles and that effective control algorithms can be created. There have already been small-scale demonstrations of the store-carry-forward capability (e.g., in developing countries as an alternative to conventional Internet connectivity). In addition, an effective system must consider longevity of handheld devices as a function of what functions are provided. For example, a "disaster mode" for cell phones could be created that permits peer-to-peer operations (augmented by vehicular gateways) for extended periods of time without recharging if certain features are disabled or modified (e.g., tower scan intervals).

Dream Team: Experts are required in embedded computing devices, mobile handset architecture, routing algorithms, and fleet control. The following organizations would be desirable: Intel, Nokia, DHS (for air and marine assets), Fedex or UPS (fleet operators).

Proprietary Information: none claimed for this document

RFI-3 – National Cyber Leap Year

Title: Securing data instead of channels in computer networks

Submitter: Kevin Fall, Intel Research Berkeley, 2150 Shattuck Ave., #1300, Berkeley CA 94704

Organization: Intel Corporation's research lab in Berkeley, CA.
http://berkeley.intel-research.net

Game-changing dimension: change the rules

Concept: By associating cryptographically-bound tags to data objects (chunks of data useful to applications) and changing to a network protocol stack that understands such tags, controlled sharing, provenance, and find-grained use controls of data can be achieved across multiple security domains in a common network. As data can be maintained in a secure state even while at rest, it can be stored and transported even by untrusted agents. Consequently, host/ server system compromise would be more of an annoyance than catastrophic failure, myriad network protocols that each implement encryption and security could be radically simplified, and network infrastructure devices (firewalls, cross-domain solutions) can make use of this approach to implement much more intuitive, easily-deployed, and secure controls for supporting information routing, sharing and dissemination in current and future computer networks.

Vision: If this approach were made common practice and brought to fruition, sensitive data could be authored, processed and distributed among applications with multi-level security and fine-grain data use controls that are enforced by host systems, network infrastructure, and storage devices. Data use controls are bound cryptographically to the data they describe, so re-authorship or modification becomes an intentional, attributable and loggable operation. Because the security of data is divorced from the channel it is carried on, secure transfer of data across a wide variety of transfer devices (e.g., USB thumb drives) is acceptable and convenient. This approach also has the potential advantage of avoiding the creation of multiple independent computer networks for carrying data belonging to different security classifications and use controls, as is common practice today. It also substantially reduces the risks associated with device loss and/or theft (esp when containing sensitive material).

Method: This concept has evolved after experience in designing and developing a new network architecture which focuses on moving data objects across intermittent networks (called Delay or Disruption Tolerant Networking, an effort that has received DARPA funding over the last few years). This architecture focuses on moving data "bundles" (objects) among store and forward routers that may cache data for extended periods of time. Bundles are a convenient unit of granularity for applying cryptographic controls.

This approach relies on an assumption that the important and challenging problem of scalable key management in a large multi-administrative-domain network can be solved and that appropriate cryptographic mechanisms can be employed efficiently. There is reason for some

level of optimism. Key management is performed today by a number of agreed-authority public key infrastructures (PKIs), on both private and public networks. The cryptographic basis for securing objects is already established in the area of digital rights management (DRM). However, the DRM model typically posits the user as an adversary whereas in many cases the user in this context will be cooperative. In addition, the work on store-and-forward networking, content delivery networks and distributed cache systems suggests networking may evolve to become much more content oriented. That said, significant challenges remain and an industry-government collaboration would be well-positioned to tackle the task.

Dream Team: Experts are required in computer/host architecture, networking, cryptography, and data sharing policy. The following organizations would be desirable: Intel, Microsoft (Helen Wang, security researcher), Sandia National Lab (Ed Talbot, manager for Computer and Network Security), Cisco, MITRE (Robert Durst, research manager), NSA, ODNI, OSD, DHS.

Proprietary Information: none claimed for material contained in this document

Who you are – Invertix Corporation. We are a growing government contractor that provides wireless communications and related technology development, solutions, and services to the government, defense, and intelligence communities.

Game-changing dimension – Change the rules

Concept – The current ubiquitous email communications system can leave people vulnerable to information overload and spam, viruses and worms, eavesdropping and identity theft. What if an alternative and open email standard were developed that enforces and thereby guarantees identity and security in a distributed, traceable, and global fashion?

Vision – The vision is a new alternative and open peer to peer standard for email, which we'll refer to as 'trustmail,' that enforces the use of Public Key Infrastructure (PKI) principals and encryption leveraging trusted authorities. The new trustmail protocol will also address other current email shortcomings by including standardized message tracking and message recall capabilities, instant messaging, as well as a common flexibility to support multiple languages and document formats.

How might trustmail operate?

- Prior to use, trustmail users will be required to register with one of a set of trusted authorities, and valid IDs will be required to register. A trusted authority will be analogous to VeriSign or Thawte, but can include public agencies.
- Users can elect to publish their trustmail addresses publicly through the trusted authority databases, or they can elect not to publish at all.
- Once authenticated and registered, trustmail user messages, while peer-to-peer in nature, will require key validation by the trusted authorities on a periodic or per-use bases, through digital signatures, certificates, or online verification, thereby ensuring the identity of all parties in the message chain as well as ensuring message privacy through encryption – private even from the trusted authorities.
- Software that will be used to support trustmail messaging must also be validated and signed by the trusted authorities. Without validation, software will not function.
- If abuse is reported to a trusted authority, offender's keys may be revoked. When revoked, offenders cannot send trustmail messages.

What would the world look like if this were in place?

- As more and more users elect to register on and become comfortable with the alternative and open trustmail network, these same users can use trustmail for an ever increasing proportion of their electronic communications.
- Being open, trustmail has the capability to spread by word of mouth (or email) in a viral fashion (as in rapid social-based growth, rather than in the negative computer virus sense).
- If trustmail is adopted on a large scale for all electronic communications, then trustmail would safeguard its users entirely from information overload, spam, and email-based viruses, worms, eavesdropping and identity theft.

How would people use it?

- Users will configure a passphrase and possibly set up biometric authentication for their identities upon registration.
- Trustmail software will be freely available for download, and a user can easily set up the software to send and receive using their identity.
- Users can share their trustmail addresses by mouth, card, phone, or email, or by accepting registration with an optional global public trustmail address book, which will be synchronized between all trusted authorities.

What makes you think this is possible?

- Enhancements to existing email capabilities have been difficult to standardize and adopt as new add-ons must be marketed and installed by all communicators, and often a large number of competing options are available.
- Current PKI solutions, while highly useful, are difficult to use, and trusted authorities are not required.
- We believe it will be easier to build a new and improved system from scratch, and offer it as a free and open alternative to email, than to provide yet another possible enhancement to email.
- Obviously, a lot more thought is required, but we believe this is possible and valuable!

What are the challenges?

- Developing an standard that addresses all desired capabilities as well as ease of use, privacy concerns, and business viability.
- Potential users will have to be convinced that their privacy is protected, and that the cost of registration (in time / money / or both) is worth the trouble.
- Users will have to remember their passwords and keep them secure.
- Developing a business model to support the operation of trusted authorities, which must be vigilant in guaranteeing identities, validating software, storing accessible databases, and protecting their own private keys from possible subversion. Possible business models include registration fees, advertizing in software, public funding, among other possibilities.
- Some elements of this PKI model may be protected by patents, which must be addressed in the standardization process. However, the trustmail standard must be open for full adoption as openness can lead to competition between trust and software providers.

What's the way forward?

- Begin defining a draft standard for comment and input by relevant parties / dream team.

Method – Invertix held internal brainstorming sessions to discuss the broad challenges and shortcomings of current IT practices, technologies, and infrastructure. Ideas were solicited, discussed, and refined in an iterative fashion. Research was conducted using primary and secondary sources.

Dream team – Internet Engineering Task Force (IETF), World Wide Web Consortium (W3C), Federal Communications Commission (FCC), Department of Homeland Security (DHS), Defense Advanced Research Projects Agency (DARPA), Trusted Authorities (e.g. Verisign), Software Providers (e.g. Mozilla Foundation).

Who you are – Invertix Corporation. We are a growing government contractor that provides wireless communications and related technology development, solutions, and services to the government, defense, and intelligence communities.

Game-changing dimension – Raise the stakes

Concept – Cyber attackers can attack networks and data from all levels in the OSI model or Internet Protocol stack, from the application layer all the way down to the physical/link layer. What if an attacker's characteristic "fingerprint" in any mode of interaction can be recorded for later identification and reprisal?

Vision – We envision a system that characterizes the fingerprints of attackers.

- Fingerprints can be measured through characteristic user/device interactions and qualities at any level.
- At the application through session layers layer this can include characteristics of transmitted data (e.g. possibly through spectral analysis of transmitted data such as word patterns)
- At the data link / transport layers this can include analysis of data transmission characteristics (e.g. jitter / latency characteristics).
- At the physical layer this can include a wider array of data transmission qualities for wired networks; and for wireless networks, this can include RF fingerprinting (i.e. analyzing the precise waveform characteristics of a transmitter).
- Fingerprints can be recorded and even shared between trusting parties, such as between government agencies.
- In filtering mode this is analogous to MAC address filtering.

What would the world look like if this were in place?

- If a network could identify and track the entities that interact with its network, it can prevent access and possibly counter-attack known offenders.
- If attackers know their identities can be detected, the risk to them in attacking a network is increased and the probability of attacks should therefore decline.

How would people get it / use it?

- For network-related fingerprinting, we envision an appliance and algorithms that monitor network traffic characteristics – possibly within a specialized firewall.
- For RF fingerprinting, we envision a passive wideband receiver that monitors all spectral bands of interest in real time for threat identification and direction finding. Rather than necessarily decoding RF communications, the characteristic spectral pattern of target waveforms may be analyzed for characteristic discriminatory features (e.g. RF envelope structure, offsets, jitter).

What makes you think this is possible?

- Pattern recognition technologies are improving with regard to fingerprint detection by enabling the identification and exploitation of invariant features in the face of variability and noise.

- These techniques can be applied to any information and data channel, be it network based or RF.

What are the challenges?

- Although we are confident that fingerprints exist at all levels, a thorough analysis must be performed on the target network layers and RF waveforms to verify the accuracy and true efficacy of such an approach.

What's the way forward?

- Analyze both network traffic and similarly RF waveform structures for accurate and environmentally invariant fingerprint characteristic using all leading-edge pattern recognition, such as using linear and non-linear classifiers, Gaussian mixture models, graph diffusion, etc.

Method – Invertix held internal brainstorming sessions to discuss the broad challenges and shortcomings of current IT practices, technologies, and infrastructure. Ideas were solicited, discussed, and refined in an iterative fashion. Invertix experience with pattern recognition and classification technologies, communications protocols, and RF technology development lead to the identification and analysis of this idea.

Dream team – DoD / Intelligence communities, Department of Homeland Security, Law Enforcement, Universities / IEEE.

Who you are – Invertix Corporation. We are a growing government contractor that provides wireless communications and related technology development, solutions, and services to the government, defense, and intelligence communities.

Game-changing dimension – Morph the gameboard

Concept – Small Radio-Frequency (RF) communications transceivers, such as cellular phones and wireless hand-held netbooks, are continually enhancing our abilities to communicate large amounts of information wirelessly. However, as these technologies develop they also pose an increasing threat to Cybersecurity, for when such devices are smuggled into secure areas they can be used to record and transmit sensitive data and information out. To date the preferred approach for detecting transceivers at gateways and checkpoints, which include manual searches or metal detectors, have proven ineffective (especially when transceivers are powered off). What if a system for detecting communications transceivers or components were available, even when such components are switched off?

Vision – The vision is an accurate, low-cost, safe, unobtrusive, and fast system for detecting radio transceivers, whether the transceivers are powered on or off.

- Such a system can prevent both intentional and unintentional smuggling of transceivers.
- Such a system can take the form of a simple walk-through portal or possibly a less obtrusive hidden wall mounted system.
- An operator or automatic process can monitor the output from such a system and can take any security actions appropriate for the facility if a transceiver is detected.

What would the world look like if this were in place?

- Threat of critical data theft through the use of intentionally or unintentionally smuggled RF transceivers will be reduced considerably.
- Potential beneficiaries can include Defense, Intelligence, and other government agencies, as well as secure corporate facilities.

How would people get it / use it?

- Such a device should be available as a specialized security product.

What makes you think this is possible?

- The threat is real, and there is already investment from the Defense community.
- For the benefit to be fully realized it is important to collect and factor the requirements from all potential end users, including both government and industry.

Method – Invertix developed a method for transceiver detection in its internal R&D efforts - an active method that detects transceivers whether they are powered on or off. The Army is currently funding Invertix to develop this technology into a short-range transceiver detector to provide this Cybersecurity capability.

Dream team – Army / Department of Defense (DoD) / Intelligence agencies, Department of Homeland Security (DHS), Federal Communications Commission (FCC).

Comprehensive Sanitization of Untrusted Inputs

Who are you Professor Trent Jaeger [1] and Professor Swarat Chaudhuri
The Pennsylvania State University
University Park, PA, 16802

Game-changing dimension Morphing the gameboard (change the defensive terrain)

Concept Despite the invention of formal integrity models (e.g., Biba and Clark-Wilson), attackers still have a significant advantage as current systems do not protect themselves from untrusted data. We plan to change the defensive terrain by ensuring that processes that access untrusted inputs only access sanitized inputs. The resulting system will prevent many types of attacks based on malicious input, ranging from buffer overflows to SQL injections to cross-site scripting, through labeling of untrusted and potentially malicious inputs, system mechanisms to enforce sanitization on access, and application-derived techniques to sanitize such input.

Vision The vision is that the systems and program collaborate to ensure sanitization. First, a mandatory access control operating system (e.g., SELinux from the NSA) mediates all data accesses. If data is from an untrusted source (i.e., either a local or remote process whose integrity is not trusted), then the data will be assigned a label indicating its low integrity. When a process requests access to such untrusted data, the reference monitor will see this request and require sanitization. The sanitization requirements will be determined from the program code that makes the request for the untrusted data. Based on satisfaction of these sanitization requirements, the system allows the process to access the sanitized data.

- **What would the world look like if this were in place?** With this approach, systems would apply program-based sanitizations to inputs from untrusted sources prior to delivery to that running program. Reference monitors (e.g., SELinux) would be enabled with the ability to identify when an untrusted input is delivered (based on mandatory access control policy) and to ensure that the appropriate sanitizations of such data have taken place (based on program analysis and simple specifications). The result is that all inputs must be high integrity (i.e., from high integrity processes), meet sanitization requirements, or are discarded, analogously to the Clark-Wilson integrity model.
- **What makes you think that this is possible?** We believe that practical, comprehensive sanitization of input is possible because the systems can identify that inputs are from untrusted sources, there are only a modest number of interfaces to protect, and program analysis techniques have become adept at identifying input formats, test cases, and filters for known malicious inputs. With comprehensive, mandatory access control (MAC) enforcement in systems (e.g., SELinux by the NSA), it is possible to identify that a particular program interface is accessing untrusted data. MAC policy analysis can identify the objects (and their labels) that can be modified by untrusted processes. When a program tries

to access such untrusted data, the system can determine whether the sanitization requirements for the specific program interface performing the access have been fulfilled. Beyond taint analysis, this approach applies program analysis for identifying formats from program code, filtering interfaces, and test cases, to extract an understanding of how program code is related to inputs to ensure that all untrusted inputs are sanitized or discarded.

- **What needs to happen for this to become real?** We envision that programmers would have to identify the places in their programs where they expect to receive untrusted inputs and provide some high-level declarative specification regarding their sanitization. Preferably, the program analysis tools described above could be extended to generate such specifications from code, but some manual guidance for ad hoc features (e.g., data value requirements) and higher-level features (e.g., operation order) is probably required. Reference monitors can check that such sanitization specifications are met whenever running programs request an untrusted input, as identified by the system's mandatory access control (MAC) policy. While some sanitizations may be applied on the fly, we envision that sanitization may be done asynchronously, to head off possible errors and to enable safe use by multiple parties.
- **Which parts already exist? Which parts need to be invented?** A variety of components already exist, including comprehensive MAC enforcement in conventional systems (SELinux by NSA) and MAC policy analysis tools (our work, Tresys). Further, we envision that a variety of program analysis work will be useful, including building filtering interfaces from malicious inputs (Bouncer from MSR), reverse engineering of expected formats from program code (Tupni from MSR), and automated, high-coverage test generation (KLEE from Stanford).

  We will need to invent an approach by which the application programmers can work with their programs to define sanitizations. While researchers have had success defining how to prevent specific attacks, we have done little to describe what is legal. It would seem that work in high coverage test generation might be most useful as it aims to tease all program paths. Providing an approach and tools to assist programmers will be the main challenge. Tools will also be necessary to assist MAC policy designers to integrate sanitizations.

Method We will explore this problem from the operating systems and programming language levels. From the system, we will use SELinux as the basis for identifying integrity problems in programs and describing how these problems should be addressed satisfactorily by sanitization. From the program, we will perform analyses to guide the programmer, define languages for the programmer to state sanitization requirements, and generate test cases to determine consistency between the sanitization requirements and the actual code.

Dream Team Weidong Cui, Microsoft Research, or others in format identification
Marcus Peinado, Microsoft Research, or others in format identification and filter generation
Peter Loscocco, NSA, SELinux
Frank Mayer, Tresys, SELinux policy analysis
Dawson Engler, Stanford, or others in program analysis for test generation

Author Backgrounds Trent Jaeger, operating systems security, source code and policy analysis
Swarat Chaudhuri, program analysis, software model checking

Request for Input (RFI) - National Cyber Leap Year

Who you are - The Johns Hopkins University Information Security Institute (JHUISI). We are the University's focal point for research and education in information security, assurance and privacy. Members of our Institute involved with this proposal include Dr. Gerald Masson (Director and Electrical Engineering), Dr. Michael Lavine (Digital Forensics), Eoghan Casey (Digital Forensics), Bryan Hoffman (Systems Engineer), and Dr. Jorge Vasconcelos (Embedded Security Architectures).

Game-changing dimension - Change the Board by creating offensive capabilities in the field of mobile device forensics

Concept - Remote Forensics for Small-Scale Digital Devices

Our concept is to provide a unified way to retrieve all digital information from Small-Scale Digital Devices (SSDD) covertly over the network for forensic, incident response, and intelligence purposes. The SSDD landscape has rapidly evolved over recent years, making devices like the Blackberry, iPhone, and G1 the digital hub of most owners' lives. These portable devices have functionality comparable to current personal computers, often including an e-mail client, Web browser, digital camera, global positioning system (GPS), short message service (SMS), video/voice recorder, media player, text editors, and document viewers. Individuals store personal data on their iPhones, parents use GPS enabled devices to track their children, hospitals use handhelds to access medical data and support patient care, and some companies give each employee a Blackberry to support their business.

SSDDs are a double-edged sword, creating new security risks while providing valuable sources of evidence. Insiders (e.g. Robert Hanssen) or outside attackers can use these devices to steal data or cause other damage to an organization. The bombs in the 2004 train bombings in Madrid apparently used mobile phones as timers. The terrorists in the recent Mumbai attacks communicated using satellite phones. Drug dealers are heavily dependent on mobile phones. Sex offenders have video taped their crimes using mobile phones.

To address the risks associated with SSDDs and fully exploit their evidentiary value, it is vital to be able to perform forensics on these devices. However, the current approach requires special methodologies, tools and devices, designed for specific model devices and under certain circumstances not all of the data may be retrieved due to proprietary hardware and software.

Vision - "Reach Out and Touch Someone"

An employee is stealing data from a secure facility, a gym teacher is taking videos of naked children in the locker room, a terrorist is planning an attack and we can prove it by remotely acquiring all data from their Small-Scale Digital Device before they have a chance to delete the evidence and cause further harm. When surveillance is called for, we can manage the device remotely to activate the microphone on the device or take photographs.

It is only a matter of time before criminals develop methods for gaining unauthorized access to SSDDs. Our concept would enable digital investigators to respond to such device intrusions and track down the culprit. Current, service providers have some limited capabilities to perform remote administrative functions on SSDDs (e.g., reset password, erase device). However, these methods do not provide the forensic or intelligence gathering capabilities envisioned in this proposal.

The component parts of this concept exist: 1) forensic acquisition of data on SSDDs, and 2) data transfer from SSDDs over the cellular network. In addition to bringing these two processes together, we need to address limitations in current forensic acquisition methods and we need to ensure that the remote access mechanism is both secure and undetectable by the user.

Although some data on SSDDs will be readily accessible via the cellular network, acquiring the full contents of physical memory (including deleted data) is generally more difficult. Therefore, the full realization of our vision may require enhancements to SSDDs such as integrating a SoC (System-on-a-chip). The SoC we envision will have the ability to access all data in physical memory and on removable media (e.g., MicroSD), and transfer acquired data over the network via a secure encrypted channel to a centralized collection system. The SoC will also contain tamper resistant protection that will disable mobile device if a user attempts to bypass/disable the SoC. Additionally, centralized server software and management console is to be developed in order to perform the following:

- Obtain heartbeat information from SoC. (Active user identifiable information, physical location, and network addresses)
- Manage SoC remotely.
- Automatically perform data mining functionality when the Homeland Security Advisory System is at a particular Treat Level. (Example: A primary target is under investigation and being monitored from the SoC. When the primary target is called or text messaged, the individual that placed the call or sent the message automatically has their SoC activated and content automatically uploaded and tagged with the primary targets information.)

Method - From our applied research and practical investigative experience, we identified the major shortcomings and challenges within the Small-Scale Digital Device forensics field. We then researched various SSDDs to better understand potential solutions to performing remote forensics and live on-device surveillance. In addition to a SoC, we considered using existing capabilities of certain devices such as running SSH on an iPhone to acquire data. We then had additional meetings with Computer Scientists, Embedded systems developers and Systems Engineers to discuss additional functionality that could be incorporated if full control of the device was possible from a remote location.

Dream team - Federal Communications Commission, constitutional lawyer, Cellular Network Providers, Cellular Hardware and Software manufacturers. Dr. Richard Mislan at Purdue is a preferred member of the team.

JOHNS HOPKINS
UNIVERSITY

Applied Physics Laboratory

11100 Johns Hopkins Road
Laurel MD 20723-6099
240-228-5000 / Washington
443-778-5000 / Baltimore

Please refer to:
AISD-09-079

20 February 2009

VIA E-MAIL
NITRD
Suite II-405
4201 Wilson Boulevard
Arlington, VA 22230

Attention: www.nitrd.gov/leapyear/

Subject: Johns Hopkins University Applied Physics Laboratory's Submission for NITRD RFI-2 National Cyber Leap Year

Enclosures: (1) Defend Missions, Not Systems
(2) A New Vantage Point for Defense

Gentlemen/Ladies:

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) is pleased to submit the attached white papers in response to NITRD RFI-2 National Cyber Leap Year.

JHU/APL greatly appreciates NITRD's review of the attached white papers. If you have any questions concerning this proposal, please contact Mr. John A. Piorkowski at 443-778-6372.

Respectfully,

Original signed by:

John A. Piorkowski
Deputy BAE for IO

JAP/dah

Distribution (*with enclosure):
www.nitrd.gov/leapyear/ (electronic)

JOHNS HOPKINS UNIVERSITY Applied Physics Laboratory
Laurel MD 20723-6099

Enclosure (1) to:
AISD-09-079

Defend Missions, Not Systems

Who we are: We are the Applied Physics Laboratory (APL), a not-for-profit center for engineering, research, and development, and a division of one of the world's premier research universities, The Johns Hopkins University. With our outstanding staff, augmented by world-class facilities, we work on more than 400 programs that protect our homeland and advance the nation's vision in research and space science.

APL solves complex problems that present critical challenges to the nation. The expertise we bring includes advanced technology; highly qualified technically diverse teams; hands-on operational knowledge of the military and security environments; and rigorous systems engineering.

Game-changing dimension: Our idea **morphs the gameboard** by changing the definition of winning from suppressing network attacks to assuring national missions. We move cyberdefense from "World of Warcraft" hack-and-slash to "Civilization"—a strategy-and-influence model of cyberdefense.

Concept: Current network defense is based on a paradigm of preventing and responding to successful attacks. Defenders are neither aware, nor motivated by the impact their actions have on the mission. They have no choice: the intersection of doctrine, mission, attack, and defense has no common expression or model to define defensive actions and consequences in mission terms. Our concept centers around a unified model of mission and attacks that changes the definition of winning from *stop the attacker*, to *ensure mission success*.

Vision: We envision a model-driven operational defensive system that blends the goals of doctrine and missions to automatically select and execute a full range of approved defensive actions. The actions work in concert to execute a course of action designed to guarantee overall assurance of the highest priority missions. We achieve this strategic capability through predictive analysis, using the model to forecast which course of action has the greatest likelihood of ensuring mission success, then automatically executing that course. Freed from the workload of detecting and reacting to individual attacks, analysts turn their creativity toward analyzing and integrating new situations, missions, doctrine, and threats into the operational model.

Method: Today's defenders play the game by attempting to make it harder for adversaries to achieve their immediate goals, under the assumption that increasing the adversary's "work factor" will reduce the level of risk to our critical systems. A typical course of action in the context of this game is generated using attack graphs that enumerate all sequences of steps that lead to goals attackers might pursue. Mitigation strategies are then developed that block as many paths as possible from initial state to goal. A fundamental limitation of this approach is that it fails to address the creative and adaptive nature of the real-world adversary. This static approach does not directly support our real goal, which is to defend our critical missions. Indeed, it is possible for defenders playing today's game to inadvertently take actions that cause greater harm to an overall mission than an adversary would have caused if unchecked. This is one side effect of an analysis that focuses on adversary defeat, rather than on mission impac.

We propose to create and apply a unified modeling formalism that represents attacks, active defense measures, doctrine on when these measures are allowed, and the prioritized missions to be protected. Our methodology is based on the use of Petri nets. JHU/APL has already developed a model for network attacks based on Petri nets that provides an intuitive and straightforward representation of the attack space. Unlike attack graphs, Petri nets can be used to model an adversary's acquisition of resources as an attack proceeds, independent of any specific goal. Moreover, there is substantial literature concerning the use of Petri nets for representing and analyzing industrial processes and systems. Petri nets are typically used in this context to determine fault probabilities and mitigation effectiveness. JHU/APL is currently working on adapting these industrial process models to the case of computer networks. Combining the attack model with the process model will enable us to assess the impact of an adversary's acquired resources on mission success, and thereby support the development of more robust mitigation strategies.

The conditions of the model we envision will be updated automatically at line speed by deployed sensors. In response to changes in the system state, the attack layer of the Petri net model will use coverability analysis to determine the set of resources that the adversary can control. On the basis of this analysis, the mission layer will forecast the adversary's potential impact on mission success. Authorized defensive measures will then be deployed automatically. As described above, these defensive measures will be focused on reducing mission risk, which may or may not include mitigating the current attack. It will be necessary to pursue further research on elements such as faster-than-real-time evaluation and distributed operational command and control in order to realize this vision.

The use of Petri nets as a common framework linking doctrine, mission, attack, and defensive actions will enable us to shift our defensive focus from defeating attackers to protecting critical national missions.

Dream team: To make this vision a reality, a number of years of collaborative research with participation from many nationally recognized centers is required. These include JHU/APL, Duke and other academic centers, and specialists in national cybersecurity doctrine.

JOHNS HOPKINS
UNIVERSITY

Applied Physics Laboratory

11100 Johns Hopkins Road
Laurel MD 20723-6099
240-228-5000 / Washington
443-778-5000 / Baltimore

Please refer to:
AISD-09-079

20 February 2009

VIA E-MAIL
NITRD
Suite II-405
4201 Wilson Boulevard
Arlington, VA 22230

Attention: www.nitrd.gov/leapyear/

Subject: Johns Hopkins University Applied Physics Laboratory's Submission for NITRD RFI-2 National Cyber Leap Year

Enclosures: (1) Defend Missions, Not Systems
(2) A New Vantage Point for Defense

Gentlemen/Ladies:

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) is pleased to submit the attached white papers in response to NITRD RFI-2 National Cyber Leap Year.

JHU/APL greatly appreciates NITRD's review of the attached white papers. If you have any questions concerning this proposal, please contact Mr. John A. Piorkowski at 443-778-6372.

Respectfully,

Original signed by:

John A. Piorkowski
Deputy BAE for IO

JAP/dah

Distribution (*with enclosure):
www.nitrd.gov/leapyear/ (electronic)

JOHNS HOPKINS
UNIVERSITY
Applied Physics Laboratory
Laurel MD 20723-6099

A New Vantage Point for Defense

Who we are: We are the Applied Physics Laboratory (APL), a not-for-profit center for engineering, research, and development, and a division of one of the world's premier research universities, The Johns Hopkins University. We work on more than 400 programs that protect our homeland and advance the nation's vision in research and space science.

APL solves complex problems that present critical challenges to the nation. The expertise we bring includes advanced technology; highly qualified technically diverse teams; hands-on operational knowledge of the military and security environments; and a rigorous systems engineering approach.

We offer an outstanding creative staff, augmented by world-class facilities, and the ability to develop effective solutions to difficult problems.

Game-changing dimension: Morph the game board

Concept: Archimedes said "Give me a lever and a place to stand, and I will move the world." Defenders today are attempting the equivalent of standing on the world while they try to move it. Host-based Intrusion Detection Systems, virus checkers, and the like occupy the same memory space, and often depend on the services of, the vulnerable operating systems that attackers manipulate. Attackers are able to out-maneuver the defense before it makes a move. Our concept **morphs the game board** by giving the defenders a new, and invulnerable, vantage point from which to play.

Vision: Because any defended system may be compromised, observations and defensive actions must be performed from a Protected Vantage Point (PVP) within the host. We envision the PVP as a logically protected execution environment for defensive software that is able to access and control all host system resources, but cannot be modified by the commodity software running on the host. The PVP can observe the same space as the host system, but the host system cannot observe the behavior of the PVP. The PVPs will be of minimal complexity and functionality to permit verification of crucial security properties through formal methods, simplify integration with commodity components, and minimally impact normal system performance.

A PVP could be implemented in a number of ways, even with today's existing systems. At APL, for example, we have used the widespread x86 ring architecture as the basis for creating a Protected Vantage Point (PVP) for defensive operations. The x86 architecture implements four levels of memory protection, referred to as rings. Ring 0 is the most privileged ring; it can access all system resources including the memory in Rings 1-3. Today, applications run in Ring 3 (the least privileged ring), and the vulnerable commodity operating system (OS) occupies Ring 0, giving attackers complete control of the system once they have compromised the OS. We have moved the unmodified commodity OS to Ring 2, and placed a secure microkernel in Ring 0. While the OS in Ring 2 "believes" that it is controlling the lower-level platform resources, in reality it is passing control commands to the secure microkernel. Security services such as host-

base IDS and virus checking could run in Ring 1, using the services of the secure microkernel instead of the vulnerable – and possibly compromised - commodity OS. The microkernel is small and simple enough for formal verification, and steps in that direction have already been taken.

Not only is the PVP concept feasible technically, the APL work indicates that it can be made feasible operationally. The ring-based PVP can either be booted with the system, or deployed on a running system without disturbing the user. The performance impact on the user is within the acceptable range, even for the rough, prototype system. Other locations for a PVP exist in commodity hardware today, and the recent move towards more hardware guarantees for security, exemplified by the shipment of commodity computing platforms with the Trusted Computing Group's Trusted Platform Module and Intel's Trusted Execution Technology, provides the assurance that the capability needed to create a PVP will exist well into the future.

There are a number of models by which the PVP could be employed, from incorporation into everyday security tools sold to security-conscious individuals to implementation on critical infrastructure nodes where additional protection is required. Within an Enterprise, a large scale deployment of PVPs with out-of-band command and control would provide unparalleled security for host configuration and security management. In a world where defenders operate from a PVP, we can envision the following advanced capabilities unachievable today:

- Deploy state-of-the-art host-based defenses that the adversary cannot disable,
- Make observations on host activity unbiased by adversary manipulation,
- Safely observe adversary activity on a host, while protecting the rest of the network from infection,
- Manage both ends of an adversary-to-tool communication to reverse engineer malware and gather information for traceback,
- Quickly ascertain the extent of adversary penetration of a network without tipping the defensive hand,
- Create and deploy a coordinated defense across an Enterprise without adversary knowledge.

In summary, today defenders and attackers are on a "level playing field". The PVP concept lets the defense occupy the cyber high ground, overlooking and secure from the attackers.

Method: The PVP concept grew out of 10 years of research and development in creating secure monitoring systems for high assurance systems. This research was inspired by, and largely funded by, customers in the Intelligence Community seeking the same fail-safe security for software that they had provided for years in the provision of security-related hardware systems.

Dream team: The ideal team to make this vision a reality would include:

- Firms and universities pushing the boundaries in the application of formal methods,
- Industry members of the Trusted Computing Group, especially processor vendors such as Intel and AMD,
- Commercial security software vendors,
- The National Security Agency,
- APL.

JOHNS HOPKINS
UNIVERSITY

Applied Physics Laboratory

11100 Johns Hopkins Road
Laurel MD 20723-6099
240-228-5000 / Washington
443-778-5000 / Baltimore

Please refer to:
AISD-08-845

15 December 2008

VIA E-MAIL
NITRD
Suite II-405
4201 Wilson Boulevard
Arlington, VA 22230

Attention: www.nitrd.gov/leapyear/

Subject: Johns Hopkins University Applied Physics Laboratory's Submission for RFI – National Cyber Leap Year

Enclosures:
(1) Development Framework for Systems Impervious To Computer Network Attack
(2) Survivable System Research and Development
(3) Provable Information Flow Security In Java

Gentlemen/Ladies:

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) is pleased to submit the attached white papers in response to RFI – National Cyber Leap Year.

JHU/APL greatly appreciates NITRD's review of the attached white papers. If you have any questions concerning this proposal, please contact Mr. John A. Piorkowski at 443-778-6372.

Respectfully,

John A. Piorkowski
Deputy BAE for IO

JAP/dah

Distribution (*with enclosure):
www.nitrd.gov/leapyear/ (electronic)

JOHNS HOPKINS UNIVERSITY Applied Physics Laboratory
Laurel MD 20723-6099

Enclosure (1) to:
AISD-08-845

Development Framework For Systems Impervious To Computer Network Attack

December 15, 2008

Who we are: We are the Applied Physics Laboratory (APL), a not-for-profit center for engineering, research and development, and a division of one of the world's premier research universities, Johns Hopkins. We work on more than 400 programs that protect our homeland and advance the nation's vision in research and space science.

APL solves complex problems that present critical challenges to the nation. The expertise we bring includes advanced technology; highly qualified, technically diverse teams; hands-on operational knowledge of the military and security environments; and a basic systems engineering approach.

We offer an outstanding creative staff, augmented by world-class facilities, and the ability to develop effective solutions to difficult problems.

Game-changing dimension: Our idea morphs the gameboard.

Concept: Current day general computing systems are extremely vulnerable to being attacked, subverted and exploited. This gives adversaries the opportunity to access U.S. government data and information resources.

We propose to morph the gameboard by developing an approach to creating systems that are, in a very broad sense, invulnerable to computer network attacks. We propose to do this by careful redesign of critical pieces of the system and judicious application of formal methods. Systems created with this methodology will lack the flaws and vulnerabilites on which attacks depend. Attackers will simply not want to play the "computer network attack" game any more because they will always lose.

This will create a completely new game for the attacker that is much more expensive. In this new game, attackers will have to either physically access and modify a computer's hardware or subvert the theorem proving tools used during the development of the system in order to take control of it.

Vision: We propose to address the shortcomings and vulnerabilities of modern day systems by developing a framework that enables the creation of useful, cost-efficient, ultra high-assurance computing systems (evaluatable at the EAL7 level) focused on system security.

First, to address the lack of protective features in modern-day architectures, we propose to re-engineer the hardware architecture, creating an environment that is designed to prevent and deter compromise.

Second, we propose to carefully prioritize design principles that remove unnecessary complexity to produce a minimal system that still retains all the power and flexibility of a modern

computing environment. These simplifications will allow us to accurately model the system and aid in developing rigorous mathematical proofs of correct operation.

Third, we propose to develop a novel approach to building a zero-fault general computing system that will simply not be vulnerable the way our systems are today. Our strategy will actively prevent the introduction of flaws into the system during its design, guaranteeing its correct operation through layers of rigorous mathematical proofs that rest on axioms rooted in a trustworthy hardware base.

The central technical challenge that needs to be overcome is that we need to scale existing mathematical proof techniques and compose them to produce guarantees about a full-featured computer system.

Method: APL's vision for this revolutionary computing platform is based on the notion of a "high-level language" computer. Our vision for the execution engine of this computer is based on the SECD abstract machine, a well-studied architecture for executing lambda calculus-like expressions.

We have chosen this architecture because the simplicity of this machine and the large body of work tying its semantics to the theoretical foundations of computing make it an attractive starting point for defining a trustworthy hardware base. By modifying it to allow concurrent processes and combining features like tagged memory and hardware garbage collection, it will provide a simple, capable platform and make whole classes of computer network attacks (e.g. those based on buffer overflows, code injection and dangling pointers) impossible. Being an order of magnitude less complicated than conventional x86-based systems facilitates the construction of proofs and guarantees about the hardware and any software it executes. The SECD machine's Turing completeness guarantees that there are no restrictions on our ability to do arbitrary computations.

We have developed a model of this hardware architecture using the Coq proof assistant. The model serves as both a software emulator that allows us to study the system, and a model of the hardware about which we can begin developing rigorous guarantees and proofs about its behavior and the software it executes.

We have also developed a strategy that will enable us to apply formal methods to produce guarantees of correct operation through a series of rigorous mathematical proofs. The strategy we have adopted will ensure that the proofs will be both compositional and that they will span the development process (requirements, design, implementation). This will allow us to provide strong, mathematical guarantees that the platform will function correctly and not allow an adversary to control its resources.

We will build upon the decades of work and take advautage of newly developed tools and technologies (Coq, ACL2, SPIN, etc...), combining them with lessons learned from past efforts.

Dream Team: To execute this idea, it would require a number of years and a wide variety of skill sets. We would seek experts in the area of formal methods to help develop models and prove theorems of correctness about the system, and people with expertise in integrated circuit layout and fabrication, hardware architecture development, and operating system design. These teams of people would need to be carefully coordinated.

The first phase would be design and feasibility exploration, to develop the capability to both simplify the system and scale the formal methods to deal with the complexity of a general computing environment.

JOHNS HOPKINS UNIVERSITY **Applied Physics Laboratory**
Laurel MD 20723-6099

Enclosure (2) to:
AISD-08-845

Survivable System Research and Development

December 15, 2008

Who we are: We are the Applied Physics Laboratory, a not-for-profit center for engineering, research and development, and a division of one of the world's premier research universities, Johns Hopkins. We work on more than 400 programs that protect our homeland and advance the nation's vision in research and space science.

APL solves complex problems that present critical challenges to the nation. The expertise we bring includes advanced technology; highly qualified, technically diverse teams; hands-on operational knowledge of the military and security environments; and a basic systems engineering approach.

We offer an outstanding creative staff, augmented by world-class facilities, and the ability to develop effective solutions to difficult problems.

Game-changing dimension: Our idea raises the stakes.

Concept: There is a trend within the Department of Defense of IP-enabling various legacy systems that have historically been isolated and connecting them together in wide area networks. Commercial off-the-shelf systems are being deeply embedded in the communications and control infrastructure. The proposed direction of the Army's Future Combat Systems program is an example of this type of evolution.

Along with the advantages of interconnecting these systems come significant unintended consequences and vulnerabilities. An IP-enabled system is not only suceptible to destruction by an enemy, it is susceptible to compromise and control. If the adversary is clever enough, they can use our own system against us. A hypothetical example is an adversary taking over a critical node that controls the codes that launch missiles – they could remotely cause a great deal of havoc.

Our idea raises the stakes required to play the game of computer network attack by making it much more difficult for an adversary to successfully control or subvert a computer system. The concept is to develop a new paradigm for general computing: a *survivable* computing environment that will meet the computational needs of its users, but at the same time, continue to function correctly when some parts of it have been compromised and are under the control of an adversary.

Vision: We propose to adopt a distributed, decentralized design as our architectural model. However, rather than designing it so that it is a network of individual computers, the inter-

connected web of nodes, collectively, is a single, multi-user computing environment. Data processing and storage are not done by a single node but by a collection of them working in concert. This allows redundancy and allows individual nodes to be compromised without loss of overall functionality.

This research would build on the significant body of published literature already available in ad-hoc routing protocols, Byzantine-attack resistant communications strategies, distributed filesystems, parallel computing, and swarm intelligence.

Adapting the ideas behind swarm intelligence to general computing is what distinguishes this system from other distributed, Byzantine attack-resistant systems. Each node will have an extremely simple set of responses that collectively give rise to more complex behavior – in this case data manipulation, storage, and retrieval, as well as the ability to transmit information to other users of the system. A primary goal of this architecture is to ensure that the system design is as simple as it can be, while still providing a fully functional computing environent. We see this simplicity as a critical factor in enabling verification of the architectural design as well as the implementation. This would facilitate certification of the system as extremely high assurance (EAL7) and enabling its use in mission-critical applications.

Method:

In this research, APL will define a set of desirable characteristics for a distributed system and metrics that measure these characteristics. These will include standard measures of utility as well as more specialized ones relevant to this effort: computational lag, measures of redundancy, simplicity of implementation, and robustness to Byzantine attacks. We want to be able to answer questions like: "What is the probability of system malfunction if the adversary compromises N nodes out of a total of M, given a topology from class T," and we want to be able to make design choices that reduce this probability.

We propose to simulate the performance of this system both through modeling and in a testbed, and quantify the various tradeoffs. The tradeoff of various performance characteristics with respect to simplicity of design of each system node is particularly important, as a simple design will enable program verification of each node to ensure that it has no implementation vulnerabilites.

A large part of this work will be the integration of existing technologies, and making design decisions based on the various tradeoffs that we are able to quantify, in order to create a solidly engineered solution to the problem of ensuring that a general computing system survives compromise by an adversary.

Dream Team: To execute this idea will require a team of people that include experts on ad-hoc routing protocols, peer-to-peer algorithms, and computer scientists with an interest in distributed computing.

JOHNS HOPKINS UNIVERSITY Applied Physics Laboratory
Laurel MD 20723-6099

Enclosure (3) to:
AISD-08-845

Provable Information Flow Security in Java

December 15, 2008

Who we are: We are the Applied Physics Laboratory (APL), a not-for-profit center for engineering, research and development, and a division of one of the world's premier research universities, Johns Hopkins. We work on more than 400 programs that protect our homeland and advance the nation's vision in research and space science.

APL solves complex problems that present critical challenges to the nation. The expertise we bring includes advanced technology; highly qualified, technically diverse teams; hands-on operational knowledge of the military and security environments; and a basic systems engineering approach. We offer an outstanding creative staff, augmented by world-class facilities, and the ability to develop effective solutions to difficult problems.

Game-changing dimension: Morph the gameboard by helping programmers write programs that are provably secure and cannot unintentionally leak sensitive information.

Concept: Programmers lack adequate tools to write programs that enforce information security properties, leaving these programs with security flaws that may leak sensitive data. We propose to implement a static information flow type system for Java that provably enforces a security policy on programs that pass the type checker. Programs written using this security-enhanced Java can be guaranteed not to mix data of different classification levels, and certain covert channels will be entirely eliminated. Programmers will not need to learn a completely new programming language, but will integrate information security into the software development life cycle. In addition, the security type system will provide a mechanism for upgrading legacy Java programs to achieve the same provable guarantee.

The ability to assert that a program is guaranteed to respect a security policy disables an attacker's ability to exploit errors in a system that may cause information leaks. If a program passes the security type checker, it *must* be secure with respect to the policy. The security type system is designed with programmer usability in mind, in order to ensure that the programmer can correctly use the security system and be able to easily write accurate policies. With the correct policy, an attacker will be unable to learn anything about sensitive information that is manipulated by well-typed programs.

Vision: Current computing needs require the creation of trustworthy systems that manipulate data with different security requirements, and today's security policies are much more fine-grained than just classified and unclassified. Guards that are used to pass information from one system to another must have strong guarantees about how they handle information. Systems that handle information with a variety of security policies must be sure to enforce the correct security policy on each piece of sensitive information.

Examples of information leaks occur frequently, as the programs that handle sensitive data does not always do so properly. For example, an email containing top-secret nuclear weapons information of the United States was sent over insecure networks by board members of Los Alamos National Security, LLC (LANS) in 2007[1]. If the user had an email system that could enforce security policies on the messages being sent, this leak could have been prevented.

We have designed a secure information flow type inference system for Java that provably guarantees that programs will not leak high security information to low output channels. A focus on Input/Output channels and programmer usability presents a simpler system design; programs need not be annotated, only policies defined, making it simpler to make information security a part of software development. It also becomes much easier to add security to legacy programs. Type inference provides increased expressiveness and flexibility over other approaches.

The security type inference system has already been designed and proved correct for a significant portion of Java[2]. It remains to be implemented and expanded to the full Java language. The system could then be evaluated at EAL7, and any programs written using the system would provide proven security guarantees.

Method: The foundations of programming language-based information flow security have been in place for some time, though they have yet to be applied in a practical fashion to languages that programmers use. By leveraging techniques similar to those that have been used to ensure functional correctness of programs, we can also achieve security correctness.

This programming language approach assumes a trusted computing base of the hardware and operating system (OS). An attacker who can subvert the hardware or OS to e.g. read from arbitrary portions of memory will be able to bypass this security mechanism.

Dream team: Implementation of this system would require expertise in programming language-based security, including that of the system designer Mark Thober. Experience with implementing Java extensions would also be useful.

---

[1] Adam Zagorin. Anger over nuclear secrets leak. Time Magazine, June 14, 2007.

[2] Mark Thober. End-to-end information flow security for Java. Ph.D. Thesis. Johns Hopkins University, October 2007.

Compositionality in Cyber Defense

Who we are Computing and Information Sciences Department, Kansas State University.
Technical contact: Dr. Xinming Ou (http://people.cis.ksu.edu/~xou/)

Game-changing dimension: *Raise the stakes*

Concept The asymmetric nature of cyber warfare determines that the defender's task is much harder than the attacker. While an attacker only needs to find one path to get into a system, the defender needs to look at all possibilities of penetration and must make sure everything works right at all times. The security tools we have today can only find localized problems and cannot be combined to multiply the protection. For example, if one deploys a firewall and an IDS, the number of attacks the two systems together can prevent is at best the sum of each individual system. This linear increase in protection power does not work since there are combinatorially large number of possible attack scenarios given the diversity and size of possible enterprise network configurations, and it is not hard to circumvent any single protection layer. To fundamentally change the landscape of cyber warfare towards defenders' advantage, the security protections must have the *compositionality* property, such that when a new piece of defense tool is added, the protection power will be *multiplied* as opposed to linearly increased.

Vision Technologies that provide compositionality in cyber defense will enable security systems to interact, communicate and cooperate, instead of sitting at one vantage point and operating in isolation. Many systems have been successfully compromised even though they were running "appropriate security software", usually Norton's Internet Security Suite. NIS is known to be susceptible to zero day attacks because the "bad guys" purchase NIS and develop their products to get around them. NIS's 55% market share makes them the prime target. With our new technologies, the new-generation security suite will communicate and cooperate with other security tools on the system such as firewalls, network-based IDS, file-system integrity checkers, *etc.* in a framework that provides the power of compositionality. The information from the various observation vantage points will be reasoned about all together in an efficient manner, and the protection achieved by combining these tools will be combinatorially stronger than using them separately. The combined tools will detect security problems not detectable by any one of them alone: if you look at multiple places of your system and reason about what you see, chances are you will have a much better understanding on what is going on than if you only look at one place. The number of attack scenarios the combined system can prevent will be combinatorially larger than the sum of what each individual system can. This framework will also enable system administrators in the trenches to publish knowledge they used in identifying security problems, which any other organization can plug into its own protection system to identify related problems. Like adding a new security tool, adding a new piece of knowledge will also provide the compositional effect — it will multiply the power already existent in the current system by what the new knowledge enabled. It is analogous to bringing all the top-breed security experts into your system who work together to secure it. One-time experience can be generalized, replicated, and applied to a large number of future scenarios, which greatly increases the agility of security defense tools as well as dramatically reducing security administration cost. This framework can be further extended so that people can share not only knowledge, but also information regarding emerging threats after one system is compromised. This can enable other systems to predict the imminent attacks even before they happen.

Method There are two key technical challenges to fulfill the vision: 1) a generic language that can be used to express the information returned by various types of security tools and the high-level knowledge used to reason about them; 2) efficient reasoning methodologies that can digest the input information to quickly identify critical and non-obvious security problems. There has already been significant progress on both fronts for a subset of the full security defense domain. On challenge 1), MITRE and NIST have designed and put to use a number of standard languages for communicating security advisories, such as OVAL, CWE, CAPEC, CPE, and so on.[1] The open-source IDS tool Snort is a good example where a simple language with clear semantics can help the security community to develop useful tools to capture real-time events that may indicate on-going attacks. But there has not been a standard language that enable these tools to communicate their observation and enable high-level reasoning. One major difficulty is the inherent uncertainty in these observations: what does an SNMP probing to a service really mean? Ou *et al.* conducted preliminary research on how to express such uncertainty in a logical framework that enables reasoning about uncertainty.[2] But significant research is needed to identify the taxonomy of security-relevant information and how to teach the security administrators in the trenches to categorize data so the reasoning model can ingest them and provide meaningful responses. On challenge 2), there has been positive result on application of logic-based approaches to identifying security problems in an enterprise network. Examples include Ou's MulVAL attack-graph[3], Telcordia's ConfigAssure project[4], and HP's Vantage project. However, significant break-throughs are needed in how to handle the inherent uncertainty in cyber security in a logical framework. Ou's work[2] started the investigation of this problem but it remains to be seen how efficient and effective the reasoning can be conducted. If these two technical challenges can be successfully addressed, and the right level of education and cooperation can happen, this will fundamentally shift the game of cyber warfare towards the defenders' advantage.

Dream team Besides the CIS department at K-State, we would like to have collaborators who have expertise to tackle the two main technical challenges above, and can help bring the technology to the security practitioners in the trenches.

- *Telcordia Technologies* (Dr. Sanjai Narain), whose ConfigAssure project is mentioned above.
- *HP Labs* (Dr. Raj Rajagopalan), who has on-going collaboration with K-State on reasoning techniques for handling intrusion events.
- *Idaho National Laboratory* (Dr. Wayne Boyer and Miles McQueen), who leads the nation's effort on critical-infrastructure protection.
- *NIST* (Dr. Anoop Singhal) and *MITRE* (Dr. Todd Wittbold), who have significant work on standardizing security-relevant information.
- *Security system vendors*, examples are SourceFire, OSSEC, Snort, Symantec, McAfee, Norton, *etc.*
- *Small-business security consulting companies*, who can benefit from the research and help form a community-based cyber defense knowledge-base using the compositionality techniques.

[1] http://measurablesecurity.mitre.org/
[2] http://people.cis.ksu.edu/~xou/publications/tr_ou_1108.pdf
[3] http://people.cis.ksu.edu/~xou/mulval/
[4] http://www.argreenhouse.com/papers/narain/TelcordiaConfigAssureOnePager.pdf

S3: Securing Sensitive Stuff

Who We Are: Sachin Katti (ICSI/U.C.Berkeley), Andrey Ermolinskiy (U.C.Berkeley), Martin Casado (Nicira Networks), Scott Shenker (U.C.Berkeley) and Hari Balakrishnan (MIT)

Problem Statement: One of the highest cybersecurity goals is the protection of sensitive data: both users and administrators would like to restrict the flow of data, making sure that it is seen only by authorized users. The widespread vulnerabilities in commercial operating systems and applications leave most sensitive data vulnerable to outside attacks. Recent academic efforts to deal with this problem in a more fundamental way have proposed new clean-slate operating systems (such as Histar [4] and Asbestos [2]) and programming libraries (such as Flume [3]); however requiring the porting or replacing of legacy production systems is unlikely to happen in the near future, so while of deep intellectual interest we dont see these efforts as having the potential for near-term impact. The purpose of our work is to enforce high-level policies on the flow of sensitive data without requiring modifications of existing OSes or applications.

Game-changing Dimension: Our proposal will morph the gameboard since we can retrofit sensitive data protection to deployed legacy operating systems and applications, making it significantly harder for attackers to steal classified information.

Vision: We propose a low-level security substrate called **S3** to track the flow of sensitive data within an OS and across an enterprise. S3 provides the following interface for each word in memory and disk: "Was this memory word computed based on sensitive data, and if so which pieces of sensitive data?" We then build security policies to prevent exfiltration (by which we mean transmission off the host) of sensitive data on top of this interface. The most immediate benefit of S3 is that it will not require the modification or trust of existing operating systems and applications. Further, S3 will leverage the widespread adoption of hardware support for virtualization and multiple cores to implement the low-level security substrate without significant performance penalties.

Status: We have a prototype of S3 implemented inside the Xen [1] hypervisor using the Qemu emulator and are now testing its flexibility in implementing various exfiltration policies. We are also simultaneously implementing mechanisms to exploit hardware-assisted virtualization and multiple cores to improve the performance of S3. We expect to have a fully-optimized system by April, 2009.

Method: S3's proposed architecture is similar to hypervisors such as Xen. The operating system with which the user interacts, referred to as the "user OS" is run within a guest VM. There is also a minimal host operating system (similar to dom0 in Xen) which offers limited services to S3 and the guest VMs. S3 incorporates three techniques:

1. **Policy Specification:** Users or operators will specify how information can be handled by attaching policies to files (i.e., tainting files). For applications such as email where the user will have difficulty in pinpointing the exact file where sensitive information resides, S3 will provide helper applications which help bridge the semantic gap between the users view of sensitive data (such as an email) and the underlying file object. In S3, all files are maintained in the host OS, the user OS runs as a disk-less machine. Thus even if the user OS is compromised, file taint information is not. Policies can dictate who can receive/read the data in a file (by name or group) or whether the document can be forwarded or written to. For instance, a user can send email to someone and apply a "no-forward" policy to that email, and this prevents the receiving user from sending any email (or transferring a file) that contains any data

from that particular email. A user can apply a "cannot leave the enterprise" policy on a file, and no email or file that has been tainted by information from that file can leave the enterprise.

2. **Shadow Memory:** S3 maintains a shadow memory data structure which maintains taint information for each word in the user OS's memory address space. The data structure is a two-level hierarchy, starting at the page level and indirecting to another page if the corresponding page in user OS memory has different taints within the page. Since we expect that contiguous regions of memory are likely to have the same taint, the space cost of shadow memory will not be onerous. When a user opens a file with sensitive content, S3 taints the memory associated with that open file. It also sets up permission bits on the physical pages in memory such that when tainted information is accessed, the user OS traps to the hypervisor.
3. **Taint Tracking:** When the user OS touches the physical memory with sensitive contents, the hypervisor traps the access, and from that point tracks execution of the user OS such that it keeps track of the information contained in the tainted physical memory (as it is copied, computed upon, written over or deleted). When the user OS attempts to send out a packet over the network, the hypervisor checks the taint status of the data in the packet, and verifies if that data can actually be exported to the packet's destination according to the security policy. S3 proposes two techniques to make taint tracking efficient: speculative user OS execution and parallelized passive taint tracking. Since S3 only needs to passively track taint flow, it can speculatively allow the user OS to execute, while keeping a log of non-deterministic events. It then leverages extra computing available due to multiple cores to replay the log and update the taint information in the shadow memory structure. OS execution is only suspended when something bad could possibly occur (e.g., when the OS accesses the network or an unsecured peripheral like a USB key), so that taint information in the shadow memory is bought up-to-date and appropriate policy is enforced.

Conceptually, S3 is different from systems like Histar etc. in two aspects. First, S3 tracks computation and data flow, prior works [4, 2] track data flow between processes. These systems therefore have to be conservative in tainting, for example, once a piece of sensitive data is touched by a process, all subsequent data produced by that process is also tainted with the taint of the sensitive data, even if the output did not depend on the sensitive data. Hence they require users/application developers to write declassifiers, which have to figure out what data is safe to declassify and let out of the system. S3, due to its ability to track computation at the instruction level, allows the user to know if any word in memory has been derived from sensitive data, and therefore eases security policy specification. Second, S3 does not require OS modifications or trust but at the cost of low bandwidth covert channels, an acceptable tradeoff in many cases.

Dream Team: Interdisciplinary team of network architects, security researchers and systems engineers.

REFERENCES

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the art of virtualization. In *SOSP*, New York, NY, 2003.

[2] P. Efstathopoulos, M. Krohn, S. VanDeBogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris. Labels and Event Processes in the Asbestos Operating System. In *SOSP*, Brighton, UK, October 2005.

[3] M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris. Information Flow Control for Standard OS Abstractions. In *OSDI*, Stevenson, WA, October 2007.

[4] N. Zeldovich, S. Boyd-Wickizer, E. Kohler, and D. Mazieres. Making information flow explicit in HiStar. In *OSDI*, Seattle, WA, December 2006.

December 15, 2008

National Cyber Leap Year
Comprehensive National Cybersecurity Initiative
Homeland Security Presidential Directive -23

Dear sirs,

I am Ivan Krstić, an independent systems security and architecture specialist. Until recently, I was Director of Security Architecture at One Laptop per Child (OLPC), a non-profit organization started by MIT professor Nicholas Negroponte that aimed to develop the so-called "$100 laptop" as a means of improving education for children in the developing world. Described by Wired magazine as a "security guru", in 2007 I received the MIT TR-35 prize naming me one of the world's top innovators under the age of 35 for my work on the OLPC security platform, Bitfrost, which delivers exceptionally strong protection to the computer's user while obviating anti-virus, anti-malware and anti-spyware software. In 2008, eWEEK magazine editors declared me one of the top three most influential thinkers in modern computer security, and along with editors of CIO Insight and Baseline magazines, named me one of the top 100 most influential people in all of modern IT.

The **game-changing dimension** of this proposal is morphing the game board.

The **concept** is a reinvention of computer operating system security in order to shift the focus from making it *harder* for intruders to break into machines, to making it more technically and economically *useless* to do so. What if we can make it so that breaking into a machine provides neither access to the user's private data, nor facilities to spy on the user or interfere with regular use of the machine (including for sensitive operations such as electronic banking), nor the ability to turn the machine into a remotely-controlled zombie?

The **vision** is ubiquitous use of adaptive sandboxes for end-user software executing within mainstream operating systems. Instead of fighting a losing battle in trying to make sure malicious, untrusted software never executes – as we are doing today with anti-virus, anti-spyware and anti-malware software – adaptive sandboxes would forcibly apply the principle of least authority (POLA) on executing software, essentially treating all code as potentially malicious. Privilege elevation in an adaptive sandbox model occurs through implicit user action, and thus no indecipherable security prompts are presented to the end user, who, as is obvious both from the available research and casual observation, cannot be relied on to make informed systems security decisions. In this model, adaptive sandboxes, which can be implemented in various ways ranging from kernel isolation mechanisms (so-called "jails" or "zones") to full-blown virtual machines, drastically strengthen security restrictions enforced on executing software, making it extraordinarily difficult for attackers to perform data theft or assimilate the machine into a command and control network even after successfully compromising the machine by exploiting a piece of vulnerable software. Adaptive sandboxes both render most known classes of software attacks useless, *and* drastically reduce the available attack space which future attackers must compromise to devise new classes of attacks.

The **method** is an advanced adaptive sandbox design, such as the one pioneered for the One Laptop per Child's Bitfrost security system. That system's specification is public and available at `http://wiki.laptop.org/go/OLPC_Bitfrost`, and a paper documenting the approach was presented at the 2007 ACM Symposium on Usable Privacy and Security, the premier peer-reviewed conference in the emerging field of HCI-SEC, or Human Computer Interaction Security. The work has won strong acclaim from both industry and academia, and has been presented at numerous top-tier security conferences and universities, including MIT, the Harvard Law School, the Harvard Faculty of Arts and Sciences, and the University of California at Berkeley. While the Bitfrost work showed most of the technical challenges can indeed be solved successfully in practice on a Linux platform, further research work would be required in porting the approach to other mainstream operating systems, chiefly Microsoft Windows and Apple Mac OS X. Conversations with technical experts from both platforms indicate, however, that large parts of the Bitfrost work could be re-used with little modification, while adapting the remaining components poses a formidable but eminently solvable challenge.

The **dream team** for this proposal is a group of operating systems and security decision makers from the three major platforms: Microsoft Windows, Apple Mac OS X, and Linux. Specifically, Scott Charney (Corporate VP, Trustworthy Computing, Microsoft), Bertrand Serlet (Senior VP, Software Engineering, Apple) and Mark Shuttleworth (CEO, Canonical Ltd., makers of Ubuntu Linux) would be a good group to create the political will to introduce such significant security measures into mainstream operating systems despite some application compatibility problems that would inevitably result. The three of them could then propose the right technical people within their individual organizations, such as George Stathakopolous (GM Security, Microsoft) and Matt Zimmerman (CTO for Ubuntu Linux, Canonical) with whom research and implementation work can be discussed.

Kind regards,

Ivan Krstić
Drage Gervaisa 9
10000 Zagreb
Croatia, Europe

t: +385 1 3893 806

w: http://radian.org

"Advanced Inference Approach for Risk Estimation"

Who We Are: The idea proposed here is the team effort (see "Method", below), led by Russell Cameron Thomas, Principal at Meritology. Meritology is a boutique consultancy based in Burlingame, California, specializing in measuring and modeling the business value and risks associated with information technology. (http://meritology.com).

Game-changing dimension: "Raise the Stakes", or perhaps "Change the (meta) Game Board".

Concept: *treat information risk metrics as an Artificial Intelligence (AI) problem rather than as a straight calculation or statistical estimation problem.*

Specifically, the we propose an "Advanced Inference" approach that uses *advanced modeling/simulation*, *inference* and *plausible reasoning* methods to estimate *overarching risk metrics* based on a wide range of ground truth data (either historical, forecast, or projected). It also has built-in methods for learning and self-improvement. This is in contrast to traditional mathematical methods of calculating risk as a function of operational variables or estimating it using standard statistical methods from risk indicators (e.g. actuarial models).

The heart of the Advanced Inference approach is to use many estimation and inference methods at once to estimate risk via "triangulation" or "weight of the evidence". Here is a partial list of candidate inference methods we have considered:

- **Bayesian Networks**
- **Prediction Markets**
- **Agent-based and Swarm Simulations**
- **Neural Networks**
- **Stochastic Dominance**
- **Process and Capability Modeling** (e.g. Pi Calculus, computational org. theory)

The biggest research challenge will be to design a computational process to integrate and resolve inferences from many different perspectives and levels of detail, including conflicting inferences, so that a consensus estimate can emerge. Like all AI methods, this computational process would apply insights from human decision-making while leveraging the speed and power of computing technology. Like human reasoning, it doesn't guarantee an answer but instead is a "best effort".

The resulting risk metrics could be on an ordinal, interval, or ratio scale, depending on the quality of available data and the uncertainties involved. If updated frequently and continually improving, these risk metrics can support rational decision-making, saying essentially this:

> "*This is the best estimate of economic information risk possible given the available data and collective knowledge, and it's consistent with the other risk estimates in your cyber world. Rationally speaking, bets for and against this risk estimate have equal payoffs.*[*] "

[*] i.e. investors would be indifferent to being either the insurer or the insured at this risk estimate.

Most important, these overarching risk metrics can serve as the basis for effective incentive systems – e.g. risk sharing pools, risk-based pricing or service tariffs, cyber insurance or re-insurance, management and employee incentives, etc. Finally, for modeling and simulation purposes, the same methods can be used to estimate the risks and incentives for "bad guys".

If successful, this would be a game-changer because it will facilitate rational investment and design decision-making in the face of intrinsic uncertainty and rapidly changing environments. Moreover, risk metrics and incentive systems can serve as a *force multiplier* for every other aspect of information security – technology, resources, and policies.

Vision: *to help the "good guys" be as agile in their defenses as the "bad guys" are in their attacks over the time horizon of investments and architectures (6 months to 10 years).*

We believe the "Advanced Inference" approach is feasible because the methods have been developed and tested in other fields but have not yet been applied to information security risk management. It remains an open research question as to whether these methods are sufficient and what the best combination of methods is best. At the very least, the Advance Inference approach offers a completely new approach to breaking through the computational and informational barriers to this "grand challenge".

To make it real and prove viability, researchers should probably focus first on a specific organization context, and focus on a subset of the information security problem. The same methods should generalize to other contexts and problems. Idealized theoretical models and toy simulations will also be good starting points. To accelerate progress, a research challenge could be defined specifically to support the implementation of national cyber security recommendations by the CSIS Commission on Cyber Security for the 44th Presidency (http://www.csis.org/tech/cyber/ , e.g. identity management, Chapter 5, starting page 61.)

Method: In 2007, Meritology led a consortium of seven organizations on an advanced research proposal to DHS on the topic of cyber security metrics. The proposal was selected as a finalist and was deemed "of particular interest" by DHS Science & Technology reviewers. The proposal was not funded, however. Team members included Cigital, Risk Management Insights, and RTI International, with consulting support from experts in the economics of information security: Dan Geer, Jean Camp (Indiana University), Ray Kaplan, Patrick Amon (Ecole Polytechnique Federale de Lausanne) , and Bob Austin (Kore Logic). In addition, these ideas have been debated on the securitymetrics.org mailing list and at their conferences (Metricon, associated with Usenix).

Dream Team: 1) thought leaders on Economics of Information Security (e.g. team members listed above, and also WEIS conference and securitymetrics.org participants); 2) specialists in artificial intelligence, computational organization theory, and other modeling/inference methods applied to intelligence and risk management problems (e.g. RAND, SRI, RTI, plus academics); 3) information security *research* leaders from major ICT vendors (Microsoft, Google, IBM, HP, etc.); 4) risk management leaders/sponsors from critical infrastructure industries (e.g. Verizon Business Services, JP Morgan Chase, E-Bay, AFCYBER, etc.); and 5) International collaborators, particularly in Europe (e.g. ENISA).

"Incentive-based Cyber Trust"

Who We Are: Meritology is a boutique consultancy based in Burlingame, California (http://meritology.com).

Game-changing dimension: "Change the (meta) Game Board".

Concept: *put a price on cyber risks to support incentive instruments for all stakeholders (analogous to business and consumer credit rating).*

Many problems in cyber security exist at least partially because the people and institutions involved are not properly motivated to solve them, or that one party's "solution" increases risks or costs for others. In essence, the incentives for stakeholders are often perverse, misaligned, or missing. The incentive-based approach creates economically meaningful metrics or prices for risk, which can then be reflected in various incentive instruments, including cyber insurance or self-insurance, product and service prices, surety bonds, risk pooling contracts, and so on. ("Risk" is probabilistic measure of losses or total costs related to security.) The incentive-based approach works by sharing the gains (benefits) of cyber trust outcomes in order to align the interests of all stakeholders and mobilize their collective intelligence and creativity, ideally for the benefit of all. Like other market mechanisms, it has the potential to yield solutions that are substantially more efficient and effective than existing approaches in isolation – security technologies, mandates/regulations, penalties, and politics (antitrust) – while also serving as a complement to them. If successful, this would be a game-changer because it will directly address the economic problems and failures that are a major barrier to cyber security. Moreover, risk metrics and incentive systems can serve as a *force multiplier* for every other aspect of information security – technology, people, and policies.

Vision: *to help the "good guys" be as agile in their protection capabilities as the "bad guys" are in their attack capabilities.* Specific examples of incentive instruments in practice include:

1. Information and computing technology (ICT) supply chain risk-pooling instruments –some form of forward contract on predefined cash flows from both ICT vendors and customers, approximating their cyber security self-insurance costs. This would provide compelling incentives for the ICT vendors and customers to share cyber trust information and work more cooperatively to implement cost-effective cyber trust solutions.

2. Real-time cyber risk dashboard for end-users or consumers **–** a dashboard or other animated display that provides risk feedback in real-time as the consumer or individual is making use of the ICT devices and services. The most important information to give the consumer/user is relative expected value changes for alternative courses of action (e.g. visit the site vs. not).

3. Enterprise Total Cost of (In)Security[*] **–** To guide investments and decision-making, new managerial accounting methods and decision support tools are needed to measure the Total Cost of Security (or Insecurity). It would also serve as the basis for risk sharing and other incentive instruments, and also allow meaningful public disclosure in stakeholder reports.

4. Incentive funds for vulnerability research and resolution **–** such as stakeholder contribution schemes and/or completion bonds. The benefit of this approach is that it provides funding up front for speculative but socially valuable activities (i.e. vulnerability research) and it makes the economic incentives more visible.

---

[*] For more detail, see: http://meritology.com/resources/Total%20Cost%20of%20Cyber%20(In)security.ppt

Method: Here is a framework* that lays out the essential elements where research is needed:

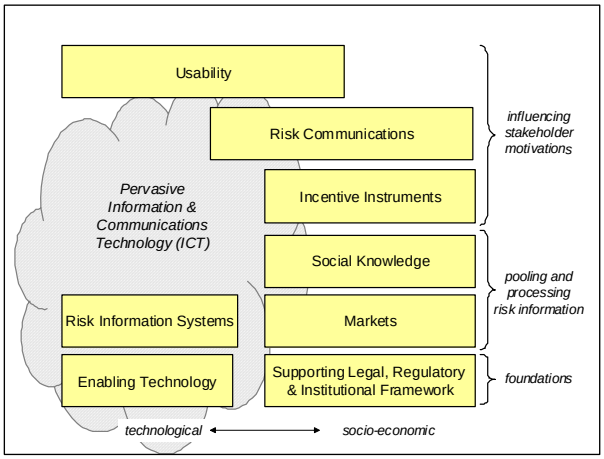

- **Usability** – Personal incentives are essentially embedded in the design of cyber trust systems, and especially the usability aspects. This includes technology, people, and processes.
- **Risk information systems** – to continuously collect and aggregate operational information related to cyber trust, and then to analyze that data to discover cause-effect relationships between operational metrics and stakeholder value.
- **Risk communication** – Cyber trust and risks should be presented in ways stakeholders can understand and act on, given their perceptions, biases, and level of understanding. This could include anything from simple disclosures to sophistication visualization.
- **Social knowledge** – including reputation systems, peer-to-peer support and sharing, and other products of social networks. Includes certification and ratings of trusted third parties.
- **Markets** – mechanisms to draw out information, to discover prices, and to support incentive instruments. Examples that have been suggested include "cap and trade" markets (similar to pollution rights markets), "Zero-day" vulnerability auctions, and prediction markets.
- **Incentive instruments** – including cyber insurance, risk sharing pools, risk-based pricing and other contingent payments, bounties, vulnerability auctions, and rights-based licensing.
- **Enabling technology** – cyber trust incentive systems should be widely distributed and embedded in the pervasive computing and communication systems.
- **Supporting legal, regulatory, and institutional framework** – supporting structures to encourage fairness and systemic trust, and to enforce self-regulation and transparency.

Dream Team: 1) Thought leaders on Economics of Information Security (e.g. WEIS conference and securitymetrics.org participants). In the past, Meritology has partnered with: Cigital, Risk Management Insights, and RTI International, Dan Geer, Jean Camp (Indiana University), Ray Kaplan, Patrick Amon (Ecole Polytechnique Federale de Lausanne) , Bob Austin (KoreLogic), and Michael R. Grimaila (Air Force Institute of Technology).

2) Specialists in artificial intelligence, computational organization theory, and other modeling/inference methods applied to risk management problems (e.g. RAND, SRI, RTI, Risk Management Solutions (RMS)).

3) Information security *research* leaders from major ICT vendors (Microsoft, Google, IBM, HP).

4) Risk management leaders/sponsors from critical infrastructure industries (e.g. Verizon Business Services, JP Morgan Chase, e-Bay, AFCYBER, etc.);

5) International collaborators, particularly in Europe (e.g. ENISA);

6) Key government agencies – DHS S&T, NIST, DoD, NSF, Treasury, NSA, and others.

* From: http://meritology.com/resources/Incentive-based%20Cyber%20Trust%20Initiative%20v3.5.pdf

Sent: Monday, December 15, 2008 4:11 PM
To: Leapyear
Subject: Leap Year Idea

RFI –Request for Information from October 14, 2008 Federal Register volume 73, number 199j

Submission to:

The National Coordination Office for Networking and Information Technologgy Research and Development (NITRD)

The National Science Foundation,

Who are we?

My company has developed transformational technologies for cyber security and multi level security information sharing for the past 10 years.

The technologies have been presented over the years in meetings to the Federal Reserve to avoid melt down of digital money streams, to NSA Information Assurance Directorate , to the Office of the Director of National Intelligence ODNI,- CIO office, to Joint Chiefs of Staff (to the former J6 General, Director of Command Control Communications and Computers ,at his invitation),to NII- Deputy CIO, to Joint Forces command- presentation 2 hour 40 participants including 2 generals, To NorthCom, To CetCom,J6,to John Hopkins Physics Laboratories, Precision Lab, and other parties.

Game Changing dimension, vision and method:

A transformational technology enabling following simultaneous results:

1) Automatic Multi Level security (MLS) information sharing (different variation of the Bell Lapdula Model)

2) Automatic defenses for the Network Centric Operations (GIG) for survivability, continuity and resiliency (defense against different types of Nation State cyber attacks including EMP attacks)

3) Automatic security

3) Automatic privacy,

4) Automatic regulatory compliance

More presentation of the method needs to take place after:

1) An appropriate legal framework for protection of intellectual property is set in place.

2) An acquisition framework is in place establishing that if government uses our technology, the company will get licensing fees from government. Furthermore if integrators or contractors are involved in any way including a bid they will need to pay royalties to the company.

Stating the “Catch 22” problem with this RFI and its continuation process:

The cyber security situation is critical, andwhat was said in testimony before Congress by Dr. Sami Saydjari about the situation is correct.

Dr. Saydjari testified on 4.25.07 that if a concentrated attack takes place as he described in his Congressional testimony “We’ve gone from a superpower to a third world nation practically overnight”.

According to the “Securing Cyberspace for the 44th Presidency” report, government should not expect too much from industry. Because as is stated in page 12: “Over reliance on the market has not produced success. As a result there has been immense damage to the national interest.”

Government has given no business motivation to industry to deliver basic good security technologies, let alone “Transformational Technologies”.

Government has to make changes as per items #1) and 2) above to create any motivation for those who have worked hard for many years to build breakthrough discontinuing technologies to enter a dialogue with government.

For companies like us it is not about getting grants of $100,000 or $1,000,000, we have survived for a long time and have the resources and will power to continue for the long term.

Its all about letting us benefit fairly in a fair acquisition process where we can get licensing fees if we deliver to government good solutions that government wants. Furthermore its about being part of a trusted process where our intellectual property is respected and protected, with no leakage to competitors or to secret black projects.

Government can take our ideas into secret projects, that we will not know about, but that is not the way for government to get the desired “Transformational” results. People who were dedicated many years to develop “transformational projects” have a lot of critical knowledge in their heads that will take a lot of time to reproduce. Bringing some of our ideas to an open bid , RFPs to integrators will only yield low level results.

The bottom line if government treats my company fairly government will be handled on a silver platter solutions that are critical to our national security.

The Dream Team:

IBM, Microsoft, Cisco, Oracle and Google + NITRD, NSA, ODNI S&T, Joint Chiefs J6,and NII

My company is in the process of establishing our “Transformational project” for commercial products, that will be comprised of independent joint ventures

Those joint ventures we are working on will be established with the leading it vendors such as IBM, Microsoft, Cisco, Oracle and Google. Each joint venture will entail further development, and integration and implementation of our technologies into existing and new commercial products of such vendors.

It’s a must for a company like ours to work with a large IT vendor to accelerate the pace of acceptance of a transformational technology.

The assets of a big vendor as well as the credibility are the only guarantee that the government will have a finished product within a short time.

Therefore Government can play a major role and help us in establishing a fair playing field with such big companies.

Government can help us innovation companies by creating a framework in which smaller company divulges its intellectual property to the big IT vendor, the IP will be respected.

Setting up an “Accelerated Protected Process for Joint Venture Development”:

Government does not have to invest money in that, but just set up a fair playing field. For example: inviting my company to a meeting with a big vendor that is of importance for a joint venture. In such a meeting the vendor agrees in writing not to undermine the intellectual property, or reengineer the technology of the smaller company. Perhaps big vendors will agree to such terms, not because of national security but just because government is the biggest client for IT.

If government states as part of the “Accelerated Protected Process for Joint Venture Development” that it will not buy products that were compromised or tainted in the Accelerated process government will benefit from the best breakthrough transactional technologies in this country.

Furthermore government by being an honest broker will get the biggest IT vendors to invest in such “Transformational technologies” and give the commercial end products the testing and credibility that will enable fast adoption in DOD, DHS, ODNI etc.

If government deals in a fair, legal, and correct way with my company as stated above:

1) An appropriate legal framework for protection of intellectual property is set in place and

2) An acquisition framework is in place establishing that if government uses technology the company will get licensing fees from government, and if integrators or contractors are involved in any way including a bid they will need to pay royalties to the company.

Then government will be able to benefit from the success of such ventures.

Those are not like “Spirals”, those are Joint Ventures that will have the best minds of the big IT vendors.

If this takes place we will cause commercial developments in such joint ventures to focus on the most critical needs of government.

If government plays fair game with my company and moves on items #1) and 2) to our satisfaction, we would like to see as part of the dream team the all or a combination of the following government players. NITRD,NSA, ODNI S&T, J6,and NII in an advisory role, or possibly contributing R&D funding.

The governmental bodies will be privy to new developments as they are developed as long as those government bodies respect that this is a commercial product development effort and we do not want to them creating competition to us.

We would see such government players as players who can contribute advice direction to the joint ventures.

The joint ventures will take place in Utah. There we will bring in young people and get them involved for the long term in what we consider a new paradigm shift that is needed for basic digital survivability of commercia1 enterprises.

Thank you

Changing the Board with On-line Privacy Manager

Who am I: Arif Ghafoor, CTO, iPrivacyManager, Inc., (women-owned small business); 25 years of R&D experience in information and Web-service security, multimedia and distributed systems; Fellow of the IEEE; Recipient of the IEEE Computer Society Technical Achievement Award. Consultant to US DoD, GE, UNDP, AT&T Bell Labs.

Change the board: Today, increasing number of users are turning to the Internet to manage their personal information regarding finances, credit, healthcare, investments, employment history, etc. This trend is being fueled by an ever-growing number of companies and government agencies such as banks, hospitals and employers that are managing users' personal information through online databases. The aim is to save time and money, by streamlining access to and manipulation of information online using the internet/intranet both in a fixed and mobile environment. However, the primary barrier to wider use of such applications is the inability of the users to define context-aware disclosure and sharing rules for their on-line data assets, in a user friendly and consistent manner. Context is defined as "any information that can be used to characterize the situation of an entity[1]." For example time of day of a certain activity is a context parameter for that activity. Similarly, location of activity is another of context parameter. The key challenge is to empower users to control their private information not only in terms of management and access but also allow the sharing of their information in a private, secure and confidential environment with others whom they authorize. The key tenet of such information sharing is that the decision to disclose personal information should entirely rest with the user.

Concept: iPrivacyManger, through the support of Purdue University, has developed a novel and intelligent Internet-based system, known as iPM[2], that allows Internet users and enterprises to manage and share their personal information (profiles, pictures, business data, etc.) within and across enterprises and through social networking sites (e.g. MySpace, Facebook). iPM can be easily integrated with the existing information store of an online user profile (via the open XML standard) and provides a unique, graphically interactive mechanism for the user to define criteria (also known as "context-aware disclosure rules") that dictate *who* can see *what* under *what* contextual conditions. Examples of contextual conditions include: time of the day, days of the week, duration (i.e. how long), location from where access to information is permissible (from a company computer or from home), events of interest, agenda, and environmental circumstances. In essence, iPM allows composition of the context-aware disclosure rules and grants users more secure access in controlling the privacy of their online information. Specifically, it provides a mechanism for the users not only to define these rules in a user friendly manner but also assists them to compose a consistent and verifiable disclosure policy using an intelligent feedback mechanism. The resulting rules are consistent and conform to predefined information disclosure standards. The primary features of the iPM technology are summarized as follows:

- A user-friendly and intuitive interface allowing intelligent feedback and simplicity for the user in creating and managing context-aware disclosure rules
- An underlying intelligent conflict resolution mechanism which allows composition of conflict free and verifiable disclosure rules

[1] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness." in HUC, 1999, pp. 304–307.

[2] A demo prototype of iPM is available at: www.iprivacymanager.com

- Both simple, predefined disclosure rules for quick and easy privacy settings and more complex, granular options for robust usage

Vision: In an era where information access is both ubiquitous and social, iPM offers a way to access and share personal online information in a secure and private manner. The application of security and privacy is based on identity, location, time and other context parameters. And these parameters are controlled by the owner of the information via disclosure rules that provide a new level of information security and privacy that is otherwise unavailable. As it stands now, online privacy management is more like a light switch with a simple on and off decision. But this is woefully inadequate. For example, in the offline world, are *all* friends/business partners of an entity given access to the personal information in the same way? No. In contrast, in the online world, the options for sharing are, more often than not, restricted to a very limiting yes-or-no choice. But this has to change as more government regulation is stepping in. A case in point is the emerging Personal Health Record technology which allows users the full-ownership of their Electronic Health Records in terms of access, management and sharing of their data across multiple healthcare providers (e.g. clinical practices, hospitals, pharmacies, etc). Another example of storage and use of an individual's personal information by a large number of users is financial information. While an individual's financial information is mostly private, some parts may still be shared with financial institutions, government agencies, advertisers etc. Data held by credit bureaus include name, social security number, bank account information, credit card accounts, financial history, etc. By utilizing the iPM technology, a user (owner of information) may define varying levels of privileges on all his/her financial information, consequently safeguarding his/her privacy. The key challenge behind such applications is to empower users to control their private information not only in terms of management and access but also allowing the sharing of the information with others whom they authorize, in a private, secure and confidential environment. iPM is expected to fill this void and will allow vendors and companies that store personal information to maintain government compliance.

iPM is easily deployable by users and vendors as a standardized API with the existing on-line information stores and allows users of a given website (i.e., social network, healthcare portal, etc.) to compose their disclosure rules for online assets they decide to protect on that site. For social networking, for example, iPM can be integrated seamlessly as an optional configuration page with an easily accessible option from the privacy section of the user's account. The generic XML-based design of iPM allows portability of the user policy to all platforms.

Method: The iPM technology has been developed using a major extension of the well-known Role-Based Access Control (RBAC) model. The extension includes provision for context-driven rules expressible that are automatically translated into an XML-based policy specification language. An intelligent policy verification engine of iPM assists the user to compose a consistent design for the underlying policy using a feedback mechanism. The technology binds contributor, consumers and owners of data and designates their roles as a part of the policy specification. Both simple, predefined disclosure rules for quick and easy privacy settings and more complex, granular options for robust usage are provided through an intuitive GUI.

Dream Team: Arif Ghafoor (iPM architect), social networking vendors (currently iPrivacyManger is engaged in discussing potential partnership with various white label social network vendors), healthcare IT vendors (in particular VA and Medicaid/Medicare healthcare) and government's financial/business sector providing social & welfare services, and IRS.

LEXISNEXIS SPECIAL SERVICES INC.

NATIONAL CYBER LEAP YEAR RFI RESPONSE

NETWORKING AND INFORMATION TECHNOLOGY RESEARCH AND DEVELOPMENT (NITRD) PROGRAM

December 15, 2008

Steve Nguyen
Vice President, Government Enterprise Solutions

Who we are
LexisNexis Special Services Inc. (LNSSI) is the government solutions arm of LexisNexis that was chartered to support classified government programs. LNSSI leverages the rich LexisNexis heritage of data expertise and innovative technologies to provide government customers with global sources of data, data fusion technology and advanced analytics that address their most challenging analytical and decisioning needs in the areas of investigation, intelligence analysis, cyber security and screening and identity verification.

Game-changing dimension
Utilizing our ability to fully index and correlate all of an organization's inbound and outbound network traffic, over the period of many months or even years, we propose a system that makes it nearly impossible to hide malicious online activity. Our ability to correlate these massive amounts of current and historical internal and external network data traffic enables the detection of previously unknown threat signatures and activity related to a multitude of ever-evolving cyber threats. Existing COTS Security Event Monitoring (SEM) and Intrusion Detection System (IDS) technologies cannot scale to handle the dozens or hundreds of terabytes (or petabytes) of data that network sensors are now routinely generating and thus cannot efficiently analyze more than a small short-term snapshot of the data. This constrains their ability to uncover new and evolving threat signatures because they are limited to detecting previously identified threat signatures and behavior and can only process a relatively small amount of netflow information. As a result, existing COTS technologies will always be a step behind new and evolving cyber threats.

LexisNexis is also proposing a game-changing shift by providing the ability to correlate malicious network activity with other person, company and threat centric data types (once legal predicate has been established to undertake such correlation). This could include lists of suspected or known nefarious IP addresses, linking web traffic to web hosting services that may be connected to known "bad actors" and correlating the origin of network traffic to known areas of criminal activity.

Concept
Use of the LexisNexis Data Analytic Supercomputer (DAS), a massively parallel, cluster-based, high performance computing platform that leverages a patented, flexible, declarative data analysis language called ECL, to enable the fusion and correlation of *all* available network security data in near real-time. With the benefit of many years of incremental improvements and evolution, the DAS platform and its Enterprise Control Language (ECL) operating system are uniquely suited to provide the capability to analyze the massive amounts of netflow data generated in today's network-centric world. The DAS platform provides the ability to perform rapid analysis of both current netflow data from a large array of sensors as well as historical network traffic spanning back months, or even years, in order to uncover previously undetectable malicious activity.

The DAS and ECL search language were developed by LexisNexis to integrate and correlate billions of records from over 10,000 disparate sources of public records data and is the core technology platform that supports the company’s multi-billion dollar risk and information analytics business. By utilizing the DAS and ECL, government information assurance

professionals will have an unprecedented ability to integrate and analyze terabytes of netflow data to discover non-obvious relationships and anomalous patterns in network activity to uncover computer "social networks", data exfiltration, botnet beacons and other potential threats.

Vision

Our vision for a proof of concept pilot is for a DAS to be deployed at each of the government's cyber security situational awareness centers. Each of the centers would feed newly identified threat signatures uncovered by the DAS from their respective Internet domain area of responsibility (.ic, .mil and .gov) to the National Cyber Security Center (NCSC) at DHS where they could be further disambiguated, matched and correlated with similar signatures from the other domains. NCSC would identify and coordinate the sharing of threat signatures with all three centers as well as the private sector.

Government cyber security analysts at the three situational awareness centers would be trained to use ECL to enable them to craft flexible and adaptive queries to enable detection of unknown threat signatures. NCSC could set up a wiki to allow analysts to publish and share their custom ECL search algorithms with analysts across the community. Analysts would translate and re-publish existing algorithms, such as those written in SiLK, to ECL.

To enable this leap-ahead solution, policies and mechanisms for information sharing between the IC, DOD, US-CERT and NCSC must be established.

The DAS is currently operational in multiple classified programs and a proof of concept prototype that demonstrates this cyber analysis capability on six months of firewall log data from an actual federal agency currently exists and has uncovered previously undetected malicious network activity that would have otherwise gone undetected. The DAS platform's current ingestion processes for multi-terabyte data sets would support ingestion of netflow data on an intra-day basis for the first spiral of a proof of concept pilot supported by NITRD. As part of a subsequent spiral effort, LNSSI would need to further partner with NITRD in order to create a data ingestion process to enable a true real-time analytical capability to further compress the time between identification and mitigation of cyber attacks.

Method

For Spiral One of a pilot, load three individual DAS appliances with data from NTOC, US-CERT and JTF-GNO and train analysts from each organization in ECL. Perform analysis on all available historical netflow data from each government domain to discover previously undetected malware and threats. Profiles and meta data for newly identified threat signatures could then be pushed to NCSC so other situational awareness centers can rapidly incorporate into their and IDS systems and ECL algorithm libraries to enable implementation of new queries on their own DAS. Spiral Two would establish real-time ingestion processes for each DAS to enable persistent analysis of network activity to detect anomalous patterns in real time.

Dream Team

LexisNexis Special Services, NCSC, US-CERT, NTOC, JTF-GNO, NIST, Carnegie Mellon CERT, National Cyber Forensics Training Alliance, Sandia National Lab, Internet Storm Center,.

NSF RFI Cyber Security Leap Year
"Play With Intelligent Pieces"

Mark G. Graff, Lawrence Livermore National Laboratory

Who we are

Lawrence Livermore National Laboratory (LLNL) is a premier research and development institution for science and technology applied to national security. Our national security mission requires special multidisciplinary capabilities that are also used to pursue programs in advanced defense technologies, energy, environment, biosciences, and basic science to meet important national needs. The Laboratory pursues research and development in areas of enduring importance to the nation, seeking challenges that reinforce its national security mission and have the potential for high-payoff results.

Principal Investigator Mark Graff is Chief Cyber Security Strategist at LLNL and a leading cyber security practitioner and thinker. He has testified as an expert witness before both Congress and the Presidential Commission on Infrastructure Protection, and served as an expert witness for the state of California. Mr. Graff has lectured on risk analysis, the future of privacy, and other security-related topics before the AAAS, the FCC, the Pentagon, and many other U.S. national security facilities and "think tanks." His most recent book, *Secure Coding: Principles and Practices* (co-authored with Ken van Wyk), is used at dozens of universities around the world to teach how to design and build secure software-based systems. A new book from Addison-Wesley is due in 2009.

Game-changing dimension

Change the way the game is played by building intelligence into the pieces.

Concept

Today, when thieves overwhelm network defenses, the point is often to steal data. Once they get it off site, they can use it or sell it to a third party. What if data were theft-proof--if it "died", say, or turned to digital dust, once removed from the place it was created? If the game were chess, we would be playing with pieces that could, say, refuse to be captured alive. Alternatively, they might look out for their own welfare, warning us if they are threatened, blocked, under-utilized, or about to be captured,

Vision

The vision is data that cannot be stolen, that can only live in a place its owner approves. And we are not speaking merely of sheltering data behind a firewall inside a network enclave. In the worlds of Web 2.0 and Web 3.0, information is smeared around the globe, and the network is the operating system. We want a way to restrict the use of data we own even after it has passed through all the doors we control.

Some parts of the solution already exist. Whole disk encryption, for example, protects against the case that an entire computer is stolen. However, the data owner presumably will decrypt the information in order to make use of it, and a compromised system could watch for that step and send the information off in plaintext. Conversely, some products today (one is PGP Desktop for Windows) automatically encrypt files and messages as they traverse a computer or network boundary, in preparation for decryption by the designated recipient. Combining this capability with whole disk encryption increases protection; but still, once the file is decrypted, it can be stolen by compromising the *recipient's* computer—or *any subsequent repository* it is copied to.

The root of this propagation problem lies in the very nature of digital information. When the Mona Lisa was stolen from the Louvre in 1911, the painting was gone from the museum until it was returned in 1913. But today, stealing a copy of a novel in Word format is theft of *the book itself*, creating an instantiation of the object indistinguishable from what we think of as the "original", while the object remains itself unchanged.

Cisco may have identified a way out from this conundrum recently with their use of IEEE "ethertypes" to influence the routing of data. Embedding the provenance and handling requirements of information in the data stream itself is a promising step.

How do we move forward to build intelligence into the game pieces—our information? Well, how about redefining the way data is stored on computers? Decades ago, the computer industry settled on ASCII character codes for text representation and 32-bit, four-byte words for data storage. Can we revisit those design decisions today--taking into account vast improvements in processor speed, the advent of public-key encryption, and ubiquitous security threats--and build provenance and access control into the very fabric of digital data, restoring the traditional meaning of "ownership"?

Method

This possibility arose out of consultations with experts in cyber security and computer architecture at Lawrence Livermore National Laboratory, an institution that has been a leader in these fields for decades and a place where many of the critical historical design decisions were made or influenced.

Dream team

In addition to computer architecture and cyber security experts, we would need intellectual property attorneys, economists, and even legislators to evaluate the impact and feasibility of such a change. Program managers and technocrats would be needed, too, to think through how such a sweeping technical change could be phased in.

Building security into the pieces on the network gameboard would still leave us vulnerable to social engineering scams that fool us into ceding access to information. Collaboration mechanisms will need to be rethought, too. But revisiting early design decisions in order to change the game is a step worth considering.

NSF RFI Cyber Security Leap Year
"Conceal the Board, the Pieces, and the Rules"

Mark G. Graff, Lawrence Livermore National Laboratory

Who we are

Lawrence Livermore National Laboratory (LLNL) is a premier research and development institution for science and technology applied to national security. Our national security mission requires special multidisciplinary capabilities that are also used to pursue programs in advanced defense technologies, energy, environment, biosciences, and basic science to meet important national needs. The Laboratory pursues research and development in areas of enduring importance to the nation, seeking challenges that reinforce its national security mission and have the potential for high-payoff results.

Principal Investigator Mark Graff is Chief Cyber Security Strategist at LLNL and a leading cyber security practitioner and thinker. He has testified as an expert witness before both Congress and the Presidential Commission on Infrastructure Protection, and served as an expert witness for the state of California. Mr. Graff has lectured on risk analysis, the future of privacy, and other security-related topics before the AAAS, the FCC, the Pentagon, and many other U.S. national security facilities and "think tanks." His most recent book, *Secure Coding: Principles and Practices* (co-authored with Ken van Wyk), is used at dozens of universities around the world to teach how to design and build secure software-based systems. A new book from Addison-Wesley is due in 2009.

Game-changing dimension

Morph the gameboard? Change the rules? Raise the stakes? How about all three, at different times, unpredictably, and to the sole advantage of the legitimate user?

Concept

Why not use modern psychology (mixed in, perhaps, with a little sleight-of-hand) to turn an everyday computer system into a hall of mirrors for intruders?

Vision

Today, one computer system works much like any other. Like automobiles, computers come in different models, but the operating principles are the same: hit the gas pedal to make the car go; click on an icon to start a program. Standard, predictable interfaces reduce production costs, training time, and the chance for accidents.

These same standards and similarities however, make life easier for intruders and interlopers who find a way to sit (in a cyber sense) in our seat. What would the computer world be like, and how much harder would it be to steal our stuff, if the way a

personal computer looked and reacted to stimuli was tailored to its individual user? Can we combine the tools of technology, psychology, and even magic to create a computer usable to one mind but maddening and misleading--even impenetrable--to others?

Consider, as a simple example, the "RSA SiteKey" authentication mechanism. Logging in, the user is sometimes (and sometimes not) confronted with a personally selected icon as verification that the website has not been spoofed in a "phishing" scam. That is, the login experience *varies*. One could easily imagine requiring the user to make a *series* of choices--which simulated room to enter into, say, or what color key to use to unlock a virtual file cabinet. The system might base a dynamic chain of authentication steps either on an explicit list of preferred motifs ("I like baseball and basketball, but not football") or one derived algorithmically from a psychological profile.

Still more promising might be the prospect of deceiving intruders, by means of misdirection, into erroneously evaluating either the defensive posture of the system or the extent to which security had already been compromised. We are aware of at least one major system designed along principles of deception and misdirection. It was deployed in the 1980's to defend some of the largest museums in Europe.

We therefore suggest exploring a new security paradigm combining the tactics of deception (against human malefactors) and unpredictability (against automatons). Research may confirm that the rightful human operator of a properly variegated user interface could navigate it in ways neither sort of attacker ever could. To put it another way, humans are well equipped to win a varying game whose pieces, gameboard, and rules are always especially designed to suit their individual traits and quirks.

Method

This approach is based in part on a twenty-year-old European system one of us helped design. The cascade of weird barriers and deceptions central to it was in turn inspired by "Rogue Moon", a novella by Algis Budrys written in the 1950's.

Dream team

To build such a system, one would need a motley crew: cyber experts familiar with attackers' targets, techniques, and mindsets, of course; algorithmicists experienced in squeezing the most out of the hardware; statisticians; and human factors engineers. But you also would want to draw in industrial psychologists, magicians, and other creative thinkers who could conceive afresh, uninfluenced by the way computers work under the hood, how to tailor an interface to one person while leading intruders astray.

One might imagine that today's faster processors would have made it easier for computers to defend themselves. That has not proved true—after all, the processors used by attackers have sped up just as much. But can we use the new processor power, combined with our burgeoning understanding of the human mind and brain, to customize computer behavior for individual users? That is a possibility worth exploring.

NSF RFI Cyber Security Leap Year
"Identify the Players and Enforce the Rules"

Mark G. Graff, Lawrence Livermore National Laboratory

Who we are

Lawrence Livermore National Laboratory (LLNL) is a premier research and development institution for science and technology applied to national security. Our national security mission requires special multidisciplinary capabilities that are also used to pursue programs in advanced defense technologies, energy, environment, biosciences, and basic science to meet important national needs. The Laboratory pursues research and development in areas of enduring importance to the nation, seeking challenges that reinforce its national security mission and have the potential for high-payoff results.

Principal Investigator Mark Graff is Chief Cyber Security Strategist at LLNL and a leading cyber security practitioner and thinker. He has testified as an expert witness before both Congress and the Presidential Commission on Infrastructure Protection, and served as an expert witness for the state of California. Mr. Graff has lectured on risk analysis, the future of privacy, and other security-related topics before the AAAS, the FCC, the Pentagon, and many other U.S. national security facilities and "think tanks." His most recent book, *Secure Coding: Principles and Practices* (co-authored with Ken van Wyk), is used at dozens of universities around the world to teach how to design and build secure software-based systems. A new book from Addison-Wesley is due in 2009.

Game-changing dimension

Don't change the rules; make it possible to *enforce* them.

Concept

Don't take what doesn't belong to you; don't go where you're not allowed. Take care what you say about others in public. In the everyday world, a person who breaks rules like these must anticipate punishment, as social mores and the law dictate. What if we could enforce in cyberspace those rules we are all expected to follow in "real life"?

Vision

What would it take to create an Internet where spam is no more common than junk mail, cyber vandalism or theft is routinely traced back and punished, and anonymous slander is a minor concern? Putting it another way, what factors are present in everyday life but absent from cyberspace that regulate behavior and enforce social norms? We suggest that it comes down to a single, venerable control: *attribution*. When something bad happens in cyberspace, we need to know *who did it*.

If we were talking about a chess tournament, we would say that players must register at the door, and must obey time limits and other restrictions or be disqualified. In sports like tennis or golf, participants also obey rules about equipment, submitting it for inspection upon demand. (And in the sports we play, by the way, concealing one's identity behind a mask while playing is considered poor form.)

Switching metaphors, we will point out that it was approximately 100 years ago that the first driver's licenses were issued in the United States. (One candidate: Missouri, 1903.) Vehicle registration soon followed, as did vehicle codes and, in most states, a legal requirement for automobile insurance. Few persons contest the need today for these controls. Do similar societal interests in personal responsibility, public safety, and a modicum of order apply in cyber space? We will not flesh out the arguments here; the need for this debate, however, is upon us.

The days of reusable, static passwords as a suitable means of authentication are long past. Two-factor and biometric methods are sounder technically, but founder ultimately because the fabric of the Internet obscures or even suppresses identity and point-of-origin information. Universal adoption of IPv6 would clear up some of the muddle in packet routing, but many identity problems would remain. To really change the game, we would need to know—with assurance, and also appropriate exceptions for anonymous protests--who the players are. And if that is the cyber world we want to live in, we will need immutable Internet identities at least as strong as the identity measures we live with in the everyday world. The technical challenge here is formidable. Not only will protocols need to be improved, but also operating systems, processors and peripherals; routers, switches, firewalls and intrusion detection systems; database software, e-mail and other communications technologies; and even methods of authorship and collaboration.

Method

The debate about what to do about the flaws in the current authentication techniques dates back decades, predating even the venerable discussions within the Internet Engineering Task Force that gave birth to IPv6. There have also been many proposals over the past decades to create secure enclaves on the Internet (some believed that "Internet2" could fulfill that need). We participated in some of those discussions, and have tried to frame the issues clearly here.

Dream team

Protocol specialists and other engineers can provide technical leadership, and regulatory and standards bodies such as IETF, ANSI, and ISO will play a pivotal role. We will need help from intellectual property experts—open source advocates as well as industry representatives. Also, since perfected Internet attribution would be a political event as much as a technical one, we'll need to bring in legislators, and civil liberties groups like the ACLU and the Electronic Freedom Foundation. We'll need diplomats, too: international cooperation is essential to knowing who Internet "players" really are.

Who we are – Los Alamos National Laboratory is a premier national security research and development institution, delivering scientific and engineering solutions for the nation's most crucial and complex problems. The Lab enhances our nation's security by developing and applying broad, multi-disciplinary scientific and technical capabilities to today's threats.

Game changing dimension – Change the rules

Concept – A major aspect of why the adversary is currently winning is that theft and misuse of sensitive data are not properly restricted. Unfortunately, the good guys must create and manipulate sensitive and innocuous data simultaneously using commodity software and off-the-shelf hardware. This provides ample opportunity for attackers to harvest and disseminate sensitive data from any node they have compromised. Furthermore, the good guys do not do not have an effective way of enforcing the policies associated with their data and may inadvertently leak sensitive data because of the complex interactions of the software they use.

Content-based systems watching outgoing traffic are not adequate as a result of complex data compositions, myriad formats and encodings, encryption and covert channels. To leap ahead of this threat, we propose that data be *indelibly* associated with its policy as it moves through *standard* applications, operating systems, and networks. This allows data to retain its policy no matter *who* manipulates it. When information reaches a policy enforcement boundary (e.g., a network perimeter or an external disk), we no longer need to consider whether this is an external attack, an honest mistake, or an insider. All the players in the game are now equal, and data policy rules drive how data are used and move through the system.

Vision – Next generation government computer systems with data policy management will enable automated policy enforcement while retaining compatibility with existing applications and platforms. Sensitive data can be automatically encrypted at rest or blocked from leaving the system or a network perimeter based on policy. Attackers will be limited to extremely low bit-rate covert channels for data exfiltration. Managers and operators can track and audit the provenance of sensitive data. Users interact with data policies using standard abstractions with which they are familiar like file permissions and classification levels.

We envision the software system that controls the data flow monitoring be implemented inside of operating system kernels, virtual machine monitors, or in hardware extensions. These systems will provide an interface for controlling and creating policy and a standard format for representing it. To enforce policy restrictions, there will be active policy wardens at protection boundaries. Several examples of such wardens include: disk controller based encryption module that automatically encrypts data blocks, a policy-aware outbound network firewall that can block data from exiting, and an operating system module that controls which applications may access certain data.

Existing research (both by us and others) has shown that these active wardens are feasible, however further study is required to understand how to efficiently design them. We must investigate how data flow tracking can still be fine grained enough for effective data policies while not requiring byte-by-byte instruction level emulation/tracking. We also have not yet fully understood how information flows through a system under standard user workloads and at what

Contact Information: Curtis V. Canada, 505-665-7453

granularity. Lastly we need to understand the limitation of this approach by answering the following: How often is innocuous data incorrectly marked as sensitive? How do we apply policies on mixed or fragmented data?

Method – We have done some preliminary work on byte-level information flow tracking and shown that we can track data through standard OSes and applications. We have experimented with using a virtual machine monitor (VMM) and on-demand emulation to enable data taint tracking (a technique borrowed from intrusion detection). In essence, taint tracking assigns a label to each byte of memory and propagates that label to CPU registers and memory as programs execute. This allows the system to follow each byte as it moves through the system at the instruction level. We have studied this preliminary system and found it to be compelling yet still undeveloped. We have also prototyped and investigated several different active policy-enforcement wardens. These systems show great promise for combating existing and future security threats, however they must be married to data policy tracking system.

To formulate and refine our ideas, we first started by soliciting ideas and concepts from the diverse body of researchers at Los Alamos. We also collected empirical observations and realistic threat models from the computer and network operations group at Los Alamos who run a large heterogeneous network that is under constant external attack. Critical to the success of this research is collaboration with the academic research community. Therefore, we consulted with researchers from the University of Illinois Information Trust Institute as well as the National Center for Supercomputing Applications to hone and advance our ideas. We also worked with integrators Accenture and Computer Associates to discuss future partnerships that might commoditize the results of our research for rapid deployment.

Dream team – To tackle a research project of such large scope, we envision engaging partners in academia, industry, and government research labs. We have had preliminary discussions aimed at enlarging existing collaborations with academic institutions specializing in system security research including the University of Illinois and the University of California at San Diego. We have also had discussions with research laboratories including the National Center for Supercomputing Applications who have experience developing and deploying large-scale security solutions. Lastly, as the research progresses we will need commercial industry partners to assist getting products to users. Such partners are available through our involvement in the Secure Enterprise Network Consortium (SEN-C) as well as through other avenues. SEN-C, composed of Accenture; Los Alamos National Laboratory; Sun Microsystems, Inc.; CA, Inc.; and Cisco Systems, Inc., is focused is on bringing leading skills together—from thought leadership and solution development to systems integration excellence—to collaborate with government and to achieve outcomes that enable CNCI initiatives and improve the Nation's security.

Contact Information: Curtis V. Canada, 505-665-7453

Advanced Simulation & Knowledge Integration Technology

Who We Are: Los Alamos National Laboratory (LANL) is a national resource of basic-science capabilities, with over 3000 Ph.D. scientists and engineers focused on national security. LANL is geared towards "big science", national-scale problems that have the combined challenges of scientific research, systems engineering, national decision making, and the nation's strategic technical capabilities and defense. LANL has had a core competency in modeling and simulation since its inception in the 1940's, has a unique expertise in decision support systems and uncertainty quantification, and is accustomed to thinking and operating in the realm of terabytes and petabytes—the same scale as the national cyber domain.

Game-changing Dimension: "Morph the gameboard" by fundamentally changing defenders

Concept: We must acknowledge that it will never be possible to have a completely secure cyber infrastructure, especially in today's environment where US citizens do not produce most of the hardware or software. We need to be able to best defend and protect ourselves, even in highly compromised situations. Imagine our edge if we could find and put together the diverse pieces of the cyber intelligence puzzle *in real-time*, enabling timely situation awareness, efficient and effective attack mitigation, attacker identification, and counter-attack. Without a similar capability, our adversaries would not be able to keep up with our knowledge-assisted defenders.

Vision: The ultimate goal is to enhance preparedness, protection, response, mitigation, and recovery activities with a scientifically defensible and reliable predictive capability. The game-changing capability should give our defenders a leap ahead with

- Validation of hypotheses explaining real-world events, with quantified confidence
- Optimal response and mitigation, with or without detection by the adversary
- Faster and more complete forensics—taking a few-month process to only hours
- With sufficient data and computing, undirected search and anomaly detection
- Preemptively identify vulnerabilities, and associate with emerging attacker capabilities
- Better training, intuition-building, response and logistics planning
- Comprehensive risk assessment and design of robust infrastructure/protocols/etc.

The CNCI defines an explicit and central role for research & development activities "to transform the cyber infrastructure so that critical national interests are protected from catastrophic damage and our society can confidently adopt new technological advances". This virtual test-bed can also be used for assessing the effectiveness, efficiency (collateral damage), and cost (both monetary and reputation) of future cyber security technologies and policies.

This capability can be achieved through a combination of modern information science for all-source information integration, and modeling and simulation to constrain the space to those supported by data and physics (that is, reality).

Modeling and Simulation (M&S) is a resource that can be used as a virtual laboratory for experimentation and constrained data fusion, especially in data-poor environments, where real-world experiments are often prohibited or impossible. M&S affords the possibility of

investigating hypothetical past and future situations constrained by both the available data and physical/logical constraints of reality. It can also facilitate the understanding of complex events and the resulting nonlinear superposition of impacts and their associated never-before-seen phenomenology, and thus is ideally suited for studying new cyber attacks.

M&S capabilities today include the ability to model systems with billions of nodes with realistic mobility and demand patters, cross-network substitution effects, and real-world infrastructure interdependencies. It can have a multi-resolution modeling paradigm to deal with variable response timelines, including tunable level-of-detail. It already provides a capability for network contingency analysis, cyber-attack analysis, situation assessment, and course-of-action analysis and optimization. Nevertheless, further development of this core technology is still desireable to fully meet the goal of leaping well ahead of our attackers.

The knowledge engine will have the task of seamlessly integrating the data being provided by network sensors, users, analysts, automated systems, and other heterogeneous sources of data including information being added by simulations. This data will be mined for the indicators and evidence supporting various hypothesized situations. The knowledge engine must be able to achieve actionable knowledge through information integration and the related uncertainty propagation. It is important to note that this is not total information awareness, but directing analysts to the most relevant information to a particular problem. Formal methods exist to do this, and even to go in reverse, predicting the piece of information, if measured next, that will most reduce the uncertainty in the top-level hypothesis. Techniques like Bayesian networks and semantic graphs are well explored, but need new basic research to cover aspects of our problem like geo-locations and temporal variability.

Some of the biggest hurdles in realizing this vision over the next decade will be

- Quantifying the interactions of political, social, economic and technical systems
- Empirically-based computational social science does not yet exist
- Comprehensive network uncertainty quantification and propagation does not yet exist
- We have large, complex data sets, but are still in data poor environments overall
- Problem spans multiple simultaneous scales and resolutions (e.g. packets to petabytes)
- Non-local, non-intuitive and interdependency effects will require study
- Predictive science requires significant calibration, verification, and validation
- The integration task is highly nontrivial.

Method: This concept has been developed over the course of the last two decades, through a series of internally- and externally-funded R&D and operational activities. Various methodologies and technologies have been investigated, with the vision described in this document emerging and being refined over the last two years.

Dream Team: Achieving this vision will require a strongly integrated multi-disciplinary team, including expertise in M&S, HPC, networking, data analysis, information science, social science, machine learning, decision making, hardware, and visualization. Team members could include LANL, David Nicol and UIUC, Don Tousley at UMass, Clifff Zou in FL, the Berkeley DETER project, Telcordia, Cisco, and IBM.

Response to National Science Foundation RFI — National Cyber Leap Year

Who We Are — LGS is a subsidiary of Alcatel-Lucent and was formerly the government solutions division of Bell Laboratories. LGS has been conducting applied research for the US government for decades. Currently approximately 200 LGS scientists and engineers work on government R&D contracts. Most have TS//SCI security clearances. A significant amount of this work is for the intelligence community; most of the rest supports the DoD.

POC: Dr. William T. Wroblicka, Technical Manager, LGS, 15 Vreeland Road, Florham Park, NJ, 07932

Game Changing Dimension — Raise the stakes.

Concept — A computer's operating system and applications contain many megabytes of programs that are virtually certain to contain unpublicized, exploitable vulnerabilities, which are open to a wide range of attacks. This problem will become worse as organizations move increasingly toward software monocultures where most hosts share some subset of common software identical in both version and patch level. In such a monoculture an attack that is successful against a particular program on one host will be successful against the instances of that program on most other hosts. Given the inevitability of attack, some measure must be taken to 1) drastically reduce the number of computers that are vulnerable to a particular attack and 2) drastically increase the cost of developing an effective attack.

Recent research has focused on diversity as a method to protect individual hosts and guard against large-scale attacks. Instruction set randomization (ISR) is used within software dynamic translation systems, providing process-specific instruction sets to protect against code injection attacks. Address randomization (AR) techniques guard against memory error exploits by randomizing the location of segments (stack, heap, and code) within process memory space, and by randomizing the order and spacing of stack and heap variables within their segments. ISR and AR techniques automatically create numerous diverse programs. But all the instances of a given program produced by these techniques have the same code structure. Since research suggests that the greater the diversity the more effective the defense, it follows that varying the code structure, which can provide an even greater number of variants, will be more effective than ISR and AR alone at thwarting attacks.

We have developed a novel approach to software diversity that randomizes code structure. We propose to develop an operationally useful tool based on this approach that creates functionally equivalent copies of a program (from its executable) that are diverse with respect to code structure and variable ordering. We propose to prove the correctness of the algorithm and its prototype implementation in the tool. This is to guarantee the correctness of the transformed programs. Furthermore, we propose to provide evidence that this new diversity technique is safe in that it does not introduce any new threat that an adversary can rely on to compromise systems.

Our approach is game-changing because it creates programs that are diverse with respect to their structure. Since new instances will have a different number of functions, the functions will vary in their parameters from the original. Furthermore, the functions will vary with respect to the stack and heap variables they create. Our approach is also unique in its method of obtaining diversity. Ours is a language-theoretic approach that induces a context-free grammar from a program, performs random transformations on the grammar, then constructs a new program from the transformed grammar.

Vision — A tool using our approach can be applied once to the software installed on each host in a network rendering each installation of the software unique with respect to code structure. The tool could also be applied to a program after each use, producing a unique instance for each invocation. Consider the following example: A worm attacks a network of computers. It gains access to the first host via a buffer overflow attack, which overflows an input buffer on the stack into another stack variable (a control variable) used in a branch condition. The worm attempts to propagate to the next host, but fails because the buffer and control variable are not in the same function, and thus it is impossible for the worm to overwrite the control variable. The worm continues to attack each host in the system, but is successful only on the very small percentage of the hosts where the two variables are present in the same function and located in the necessary order. Ultimately the stakes are raised because an attacker must expend much more effort to exploit large-scale systems; the value of a single exploit (to the attacker) is greatly reduced.

Method — The focus of this research will be the implementation and evaluation of a diversity tool based on a context-free grammar transformation algorithm. The tool will create functionally equivalent copies of a program that differ in their program structure (i.e., in the set of functions, and the distribution of function parameters, stack, and heap variables across those functions), and that have differing stack and possibly heap variable ordering.

The method is based on the observation that a change in program structure can change code and memory in ways that render some attacks ineffective. Program structure can be described by a context-free grammar, so grammar transformations are a natural way to describe restructuring. The functions and control structures are the nonterminals of the grammar. The branch conditions and atomic statements constitute the terminals of the grammar.

The basic steps of our method are to induce a context-free grammar from a program, to randomly transform the grammar, and then to construct a new program from the transformed grammar. The context-free grammar transformation creates a new program that has a different set of functions, and whose functions have a different set of parameters and local variables than the original program had, yet the two programs perform the same operations in the same sequence.

We envision a two-phase effort. During the first 12-month phase we will develop a transformation tool that operates on source code. We will obtain and prepare vulnerable applications and exploits for testing. The applications may require some modification to put them in a form needed by the tool. They may also require some manual annotation to assist the tool. We will test and evaluate the effectiveness of the transformation tool at protecting software. We will develop a correctness and safety proof for the algorithm and its implementation. And we will develop a demonstration to show the tool and its effects. During the second 6-month phase we will extend the tool to transform executable code. We will test the new tool using the same applications and exploits from the first phase. And we will develop a demonstration to show the tool and its effectiveness.

Dream Team — LGS, to provide technical leadership; a university's compiler group; a commercial software vendor (e.g., Microsoft), to provide application code for experimentation; a software security company (e.g., Symantec), to provide attacks/exploits/malware for experimentation.

Proposal for National Cyber Leap Year - Request for Input (RFI)
Authentication Indemnification and Risk Transference
December 15, 2008

Who we are: We are a group of friends and business associates with government, business, and academic experience in the information security domain: Mike Hamilton (CISO, City of Seattle), Douglas Barbin (Product Manager, Verisign Managed Security Services), Fred Langston (Product Manager, Verisign Consulting Services), Ernie Albers (Independent Security Consultant), Mark Baenziger (Principal, Hexsaw Consulting), Randy Richey (VP, Professional Services/Sales at Govplace), David Matthews (Deputy CISO, City of Seattle).

Game changing dimension: Morph the game-board

Concept: Authentication indemnification and risk transference through the creation of an ecosystem between organizations who own the data, organizations who authenticate users, and back-end insurers where risk of loss is transferred in return for adherence to basic standards of good practice.

Vision: In the history of information security, true authentication of an individual has been one of the hardest challenges to address. While improvements in encryption and authentication technology have been significant, little has been done to address the human side of authentication (actually knowing who you are issuing a credential to in the first place.) Today, the owner of the data and issuer of the credentials has 100% of the responsibility, 100% of the liability, and typically very little expertise on how to manage the risks associated with identity and access management. Even with the increased role of third-party authentication services, the liability still remains with the organization to authorize the user to access and particular application.

Our vision begins by creating a legal means to shift some of the burden of risk from the authorizing organization that controls the data to the organization that provides the authentication services (a company or government agency) through indemnification. In turn for this indemnification, the authorizing organization agrees to adhere to a standard set of practices and data sharing requirements to ensure accountability and transparency. On the back end, the authentication company may offset some its risks through insurance in return for authentication company performing the necessary due diligence on the organizations that it protects.

With the above, our concept is designed for more than merely transferring risk. The ultimate goal of our concept is to create an ecosystem whereby you have an increased number of participants involved in the process of authentication, each with financial stake in the process, and such acting as checks, balances, and incentives for the others.

Method:

First, changes would need to be made to any relevant statutes to provide for the ability for an authentication services provider to indemnify the owner/authorizer of the data. We also believe that either mandating such risks be identified (whether indemnified or not) will drive an otherwise voluntary process. Second, the ecosystem would need to be created with standards of good practice for each of the participating groups. The below diagram is a high-level depiction of how this process would work.

Note: This proposal and any opinions are that of the contributors and not necessarily our employers.

Proposal for National Cyber Leap Year - Request for Input (RFI)
Authentication Indemnification and Risk Transference
December 15, 2008

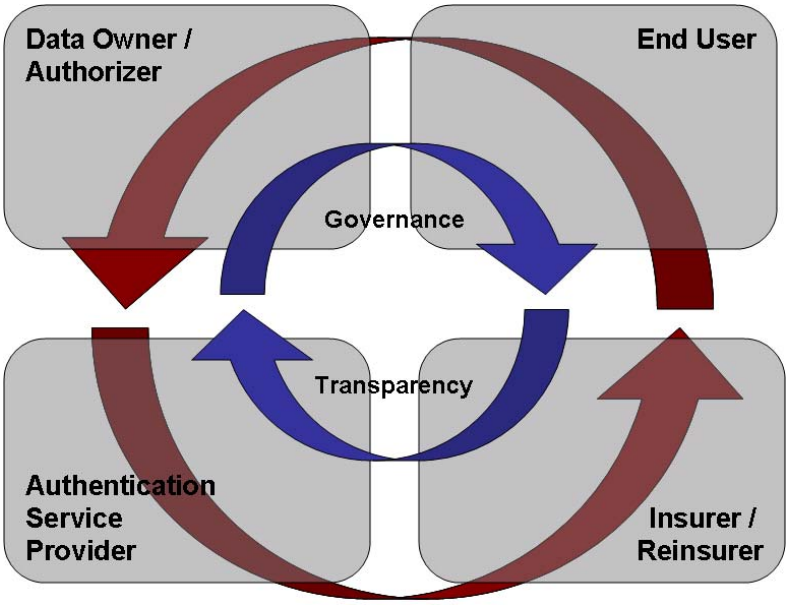

Indemnification / Risk Transference

The overlying goal is that when one group indemnifies another, the group which now bears the burden also has the responsibility for "underwriting" or performing the necessary due diligence of the other group. Thus, the authentication service provider would have the ability to audit the owner / authorizer to ensure these processes were being followed so that if a compromise occurred and the owner / authorizer had not followed these standard practices, the service provider would not be liable for the damages. The same would be consistent with the insurers who would be required to review and audit the practices of the authentication service providers to ensure that they are taking standard steps to minimize risk. This type of ongoing review, often referred to as underwriting, is common-place for any insurer or re-insurer.

The above framework is high-level and conceptual in nature. If selected as a leap-year concept, phase II would include the following:

- Defining draft standards for the ecosystem (to be shared/published for comment)
- Developing actuarial-type models of compromises and losses to determine the what the appropriate level (or limits) of indemnification at the individual account and/or organizational level
- Evaluating a potential ecosystem models through a series of interviews/discussions with government agencies and private companies such as financial institutions, health care providers, and authentication service providers.

Dream Team: Our dream team would include professionals from the following disciplines:

- Governance / Audit - The core of the group responsible for creating this would be persons with strong information security risk and governance backgrounds from both the public and private sector. These would be CISSPs, CISAs, CPAs, CFEs, etc.
- Legal - At least one attorney with expertise on the legal environment of authentication and/or breach litigation.
- Insurance / underwriting - Someone from the underwriting side of the insurance industry would be needed. Ideally this could be someone familiar with fraud or cybersecurity risk.

Note: This proposal and any opinions are that of the contributors and not necessarily our employers.

Proposal for National Cyber Leap Year - Request for Input (RFI)
Security Event Aggregation, Correlation and Alerting for Local Government
December 15, 2008

Who we are: We are a group of friends and business associates with government, business, and academic experience in the information security domain: Mike Hamilton (CISO, City of Seattle), Douglas Barbin (Product Manager, Verisign Managed Security Services), Fred Langston (Product Manager, Verisign Consulting Services), Ernie Albers (Independent Security Consultant), Mark Baenziger (Principal, Hexsaw Consulting), Randy Richey (VP, Professional Services/Sales at Govplace), David Matthews (Deputy CISO, City of Seattle).

Game-changing dimension – Raise the stakes, by:

- Making it more difficult to attack critical infrastructure by implementing currently non-existent detective controls in local government
- Disseminating real-time, actionable alerts to a broad base of customers
- Assisting in the collection and provision of intelligence information for anti-terrorism efforts
- Provide a venue and data set to test research and development projects for new attack detection technologies, thus accelerating the transition to widely-deployed defensive capability

Concept – Local government supports and maintains 85% of critical infrastructure. Transportation systems, public safety and in some cases utilities are managed by cities and counties. These are the systems that, when rendered unavailable or unreliable, have an adverse impact at the scale most affecting the quality of life, and life itself. Despite the criticality of the infrastructure and services maintained by these organizations, local governments cannot afford to attract and retain qualified security staff, or even implement the detective controls such as event aggregation, correlation and alerting that are common to organizations in the private sector. A non-profit managed service conducting these tasks would remove cost as a barrier to the capability, and ensure that critical infrastructure is being monitored for attack by organized crime, nation-states and terrorism actors. Further, the system would act as a test bed for research and development projects that are funded by DHS; new and "fringe" technologies could be deployed in an operational setting.

Events of interest, for example firewall and IDS logs, would be mapped geographically, so that consumers of the aggregate information would see events and trends in a regional context, using an at-a-glance "mashup" of attack taxonomy and geographic components. As most information is contributed, higher-value processing products may be delivered, for example alerts identifying internal systems that should be examined for botnet infection. Various classifications of alerts might be delivered to the public, data contributing organizations, local, regional, state and national information distribution channels such as NWWARN and WACIRC in the Pacific Northwest and US-CERT nationally. Note that this capability also provides compliance with several regulatory requirements.

Note: This proposal and any opinions are that of the contributors and not necessarily our employers

Proposal for National Cyber Leap Year - Request for Input (RFI)
Security Event Aggregation, Correlation and Alerting for Local Government
December 15, 2008

Another reporting vector is envisioned for events and trends that may have intelligence value. Should the determination be made that a specific event is worthy of further investigation, it would be communicated to the most local Fusion Center for examination by an analyst. In this way, seemingly targeted attacks may be investigated for their national security implications. This aligns the project with the goals of HSPD-7 for the protection of critical infrastructure from terrorist threat.

A tertiary vision is that of research facility. Information contributed by participating local governments may be adequately anonymized, such that its use in evaluating new technologies will not have privacy, intellectual property, or other concerns that prohibit deployment in the private sector. Research-grade systems may be "bolted on" for evaluation, and the value-added products communicated to the base of participants for feedback. This service will identify those technologies that may have wide applicability, and which should be developed as commercial products.

Vision – The vision is a non-profit corporation that is essentially identical to a managed security service provider (widely-adopted in the private sector). Event traffic will be aggregated locally and pre-processed per privacy policy, and transferred over an encrypted tunnel to the operating center. There, analysis would be performed by machine and with the use of analysts. Several classes of consumers would have access to various "views" of the data; the simplest representation would be a display of attack frequency and types over a geographic region, mapped using colors and drill-down capability. This would provide at-a-glance situational awareness on attack traffic, which is sufficient for most consumers. Data contributors might have access to internal events, or compared directly with geographic and logical "neighbors". Initially, local governments may contribute simple perimeter firewall logs, and subject to appropriate confidentiality and privacy agreements, include other data sources that indicate the frequency and taxonomy of events internal to networks such as on-access virus detections, trojans removed in weekly sweeps, etc.

The value proposition for local government participants is this:

- Subscription to the service will be at a cost far below commercial competitors
- The system will meet several compliance requirements
- The system will manage retention requirements for electronic records
- The system may provide e-discovery services
- Internal systems will be surgically identified for remediation, to minimize human resources expended on security tasking

Method – Combine Security Information Event and Information management technology with MSSP development experience, experience and name recognition in the local government sector, and a group of focused entrepreneurs.

Dream Team – The list of individuals contributing to this RFI response.

Note: This proposal and any opinions are that of the contributors and not necessarily our employers

Proposal for National Cyber Leap Year - Request for Input (RFI)
Inducing a Backbone Network Security Risk Marketplace
December 15, 2008

Who we are: We are a group of friends and business associates with government, business, and academic experience in the information security domain: Mike Hamilton (CISO, City of Seattle), Douglas Barbin (Product Manager, Verisign Managed Security Services), Fred Langston (Product Manager, Verisign Consulting Services), Ernie Albers (Independent Security Consultant), Mark Baenziger (Principal, Hexsaw Consulting), Randy Richey (VP, Professional Services/Sales at Govplace), David Matthews (Deputy CISO, City of Seattle).

Game changing dimension: Morph the game-board

Concept: Create strong incentives for Tier 1 backbone network providers to reduce "dirty" (heavy intrusion signature/botnet/spam) traffic, in order to create a network security risk market that drives reduction in such traffic throughout the cloud and eventually down to the endpoints.

Vision: A major challenge to improving security of the global network is the large number of insecure endpoints that continually create spam, botnet, worm, and other security attacks against the rest of the global network. The significant number of insecure machines means that all elements of the global network are heavily impacted by "dirty" traffic – making it challenging to identify higher threat attackers, and negatively impacting network bandwidth and availability.

Our vision is to introduce some form of market-based incentive/disincentive to influence Tier 1 backbone network providers to highly value network connections (peers) that have lower levels of the "dirty" traffic discussed above. If this market (dis)incentive is strong enough, the financial impact of connecting to the backbone will create similar incentives for Tier 2 and 3 ISP's to find ways to reduce the amount of connections that allow "dirty" traffic onto the Internet. This will further impact pricing and availability for more secure endpoint connections.

This in turn will improve the marketplace for systems that help improve security of the endpoints and ISP's, and will eventually result in a much cleaner network that will both reduce the impact of botnets, spam, malware, and other security attacks, and allow truly dangerous attacks to stand out more clearly.

Method: This proposal, given that it intends to influence Internet peering points and exchanges that fundamentally affect Internet stability and capacity, has potentially far-reaching and unpredictable consequences. It is important that it is examined deliberatively, with significant modeling and testing on a small scale prior to implementation at a larger level.

In order to ensure that different approaches to this are effectively evaluated, it makes sense to assemble multiple teams with different membership, with each team having representatives from the disciplines discussed in the "Dream Team" section. Each team should create models and initial testing plans, and conduct live tests for different approaches to this proposal. There are a large number of potential approaches to any solution that need to be independently evaluated.

Approaches to measuring how "dirty" traffic is:

- **Direct signature measurement**. Measure the relative amount of certain types of inappropriate signatures within traffic. (i.e., IDS/AV-based signatures).

- **Aggregate statistical measurement**. Measure gross traffic characteristics (distribution of IP sources/destinations, unusual TCP/UDP ports, ICMP unreachable counts, invalid HTML/flash/images) that reflect evidence of inappropriate traffic.
- **2nd order signature measurement**. Measure impact of IP addresses on the rest of the network (i.e., if an endpoint shows evidence of worm infection following a visit to a website, then count the website as a "dirty" endpoint)
- **Peer feedback**. Allow Internet members to provide feedback via existing (RBL, Dshield) or new (social media forums) methods on the security safety of endpoints.

Approaches to influencing Tier 1 backbone providers:

- **Direct financial incentive.** Provide direct funding to providers which maintain lower "dirty" traffic levels.
- **Indirect financial incentive.** Create tax incentives for providers which maintain lower "dirty" traffic levels.
- **Regulatory.** Fine or regulate companies which maintain higher "dirty" traffic levels.
- **Fund Competition.** Create competitors to existing Tier 1 backbone providers which provide low-cost (or no-cost) connections for peers with lower "dirty" traffic levels.
- **Cap and trade**. Regulate Tier 1 backbone providers' ability to accept "dirty" network traffic, and allows these providers to trade surpluses or deficits of "dirty" network traffic.

Key Points

- Determining how dirty a network connection is does not require comprehensive analysis of all traffic. Statistically valid sampling can determine the relative "security cost" of a network. This significantly simplifies determining the "security cost" of a connection.
- Measuring and creating security incentives/dis-incentives for network connections may not be feasible among Tier 1 providers, due to the need to maintain the settlement-free nature of the majority of Tier 1 peering points. Creating (dis)incentives for Tier 1-Tier 2, or Tier 2-Tier 3 connections based on "Security Cost" is a good alternative.
- Influencing global Internet backbone providers is within scope - failure to consider how to influence non-US Tier 1 providers will not create an effective marketplace.

The final approach implementing this marketplace may include any or all of the above approaches (or others defined by the team), and will have to be carefully designed following modeling, evaluation, and testing.

Dream Team:

- Network engineers and executives from Tier 1 "backbone" providers.
- Engineers from very high speed IDS or Tier 1 network analysis companies.
- Economists familiar with the introduction of market-based incentives to systems.
- Statisticians familiar with market analysis
- Security experts familiar with Internet-wide end-point security statistics
- Legal experts in US government tax-incentives/regulatory regimes.

ENABLING SECURITY MIGRATION FOR EMERGING CRITICAL INFRASTRUCTURES

Catherine Meadows
Naval Research Laboratory
Washington, DC 20375

Who I am: I am head of the formal methods section in the Center for High Assurance Computer Systems at the Naval Research Laboratory. I have been performing research in computer security and formal methods for over 25 years. One of the chief focuses of my research has been the application or formal methods to the analysis of cryptographic protocols. In particular, I have applied my techniques to a number IETF protocols, and have worked closely with the IETF assessing several of their security standards.

Game-Changing Dimensions: Change the Rules

Concept: We all know that it is difficult to retrofit security onto an infrastructure that was not designed with security in mind. But standards (both de facto and de jure) must usually be proposed and implemented before security needs are well understood. What is needed is a way of introducing standards that are flexible enough to incorporate security additions when needed, but do not try to lock in security solutions ahead of time. In other words, what we need for emerging infrastructures is not so much standards that provide security as standards that enable *security migration*.

Vision: The vision is to develop techniques for developing "securable" infrastructures, and to incorporate them into standards. These "securable" standards should not enforce any particular security solution, but should be flexible enough so that they security can be introduced later without a complete overhaul of the infrastructure. Consider, for example, cryptography. Many new communication technologies are introduced with minimal use of cryptography or authentication, if any. As the technology becomes more widely embedded into the infrastructure, the need for cryptography and authentication becomes more obvious. But by this time critical decisions may have been made about the standards that make this difficult. Another example is the ability to communicate with other communication infrastructures. It is common now, when a new infrastructure is introduced, to start adding capabilities for communication with other infrastructures that are already present: e.g. the Internet, GPS, Bluetooth, and so forth. But if one of these infrastructures is vulnerable to attack, then that could also put the new infrastructure at risk. Thus any new standard should make it possible to include the

capability for controlling or limiting communication with other infrastructures when necessary.

Finally, we need to take account of the fact that definitions of security can also change and evolve. For example, work on network security started out by focusing on needs such as confidentiality and integrity of data. However issues such as denial of service and privacy concerns have widened the scope of network security mechanisms. Of course, we should not expect to able to make standards so flexible that they can be extended to encompass countermeasures to fundamentally new and unforeseen types of security threats, but new security threats are often foreseen long before they become an actuality (e.g. the notion of "availability" as a security goal was widespread long before denial of service became a concrete threat).

Method: In order to be successful, this concept must be approached on a number of different fronts. First, a set of "best practices" for enabling security migration for infrastructures must be established. Much of this can be done by studies of the infrastructures that are already in place. The last few decades have seen a number of new infrastructures introduced: not only the Internet, but the cellular telephony infrastructure, Peer-to-Peer communities, wireless communication infrastructures, SCADA, and so forth. This should give us an ample source of lessons learned that we can use to compile a best practices list. The goal here will be to discover what early decisions impeded or facilitated the later adoption of security solutions, and how the lessons learned could be applied to other technologies.

The next task is harder. We want to make sure that any changes we make to the standardization process do not stifle innovation or impede the introduction of new technologies. This will require economic studies of what the impact, both positive and negative, of these new types of standards would be, as well as studies of the types of incentives that could be applied to promoting their adoption. Research on the economics of adoption of new technologies and on the economics of security already exists, and it is likely that much of this would be useful to us. However, new research may also be necessary.

Once the list of best practices has been developed and the economic implications are well understood, it will be necessary to put together a list of recommendations for new standards. This will require the input of standards experts from different technologies and possibly from regulatory agencies. Once the list of recommendations is complete, it may become apparent that some or all of them should become part of new regulations for security of emerging infrastructures.

Dream Team: As is clear from the above discussion, the dream team for this concept involves people from a number of different areas, including security experts, economists, members of standards bodies, and possibly regulatory agencies.

From: Mike Ford
Sent: Wednesday, October 22, 2008 11:28 AM
To: Leapyear
Subject: Submission of Idea

My idea is to create a system of IP Authorities similar to the system of Certificate Authority, which could be used to allow end users to create whitelists for IP traffic. The concept is simple. I as an end user get to decide if my PC will only talk to certain companies or entities. My operating system could be configured to default to only the OS vendor, therefore ensuring that a new installation would only be able to talk to the vendor for updates and patches until I decide otherwise. I then can add entities to my whitelist I deem trusted or contract with a vendor who could maintain a list of trusted sites. This could spawn an industry of companies that offer users options like "Whitelist companies that offer sporting goods sales" and I could choose to opt into allowing connections to these companies either permanently or for a temporary period of time. Software vendors could offer preconfigured whitelist entities for their products so that games and other highly connection oriented software would still work well.

For example a game might list “Game Spy” and “NVidia” as necessary whitelist companies to run their software. I would put these into my system which would connect to my IP Authority and give me the current list of IP addresses these companies use to communicate and add them to my whitelist. An internet based game might maintain an entity called for example “AOE - Current Online Game Hosts” which would be updated in real time so that home users could still connect to home users for gaming, but after the game session was over, the user would be removed from the whitelist and the user would remove the entity from his trusted list.

This setup would stop me as a home user from making silly mistakes such as connecting to a phishing site in China instead of my bank, and prevent hackers from scanning me with impunity. Also if I do get a Virus, it would severely hamper its ability to propagate to other computers by stopping my connections to them and by limiting the number of computers which will allow me to connect to them.

Thanks for the consideration,

Mike Ford
4402 Augusta Ave.
Richmond, VA 23230

National Cyber Leap Year (NCLY)
Request for Information (RFI)

Who we are – Southwest Research Institute® (SwRI®). Southwest Research Institute, headquartered in San Antonio, Texas, is one of the oldest and largest independent, nonprofit applied research and development (R&D) organizations in the United States. Founded in 1947, SwRI provides contract research and development services to industrial and government clients. SwRI consists of 11 technical divisions that offer multidisciplinary, problem-solving services in engineering and the physical sciences, with a staff of over 3,100. SwRI is currently conducting research in many areas of cyber security, including development of innovative methods for analyzing and improving the security of software applications.

Game-changing dimension – change the rules

Concept – Software vulnerabilities are the Achilles heel of our national information infrastructure. Many, if not most, attacks on information systems are enabled by exploiting flaws in the design and implementation of software code. Because of competitive pressures, software producers rush applications to market that have not been tested for security based on the assumption that they can "fix the holes later." Today's consumers have generally accepted the premise that there is no effective way to influence software developers to produce more secure applications. *What if we changed the rules* to require consumer product information labels with software security ratings, similar to government labeling requirements for food, automobiles, household appliances, drugs, clothing, and many other types of products?

Vision – The vision is a *standard rating and labeling system for the software "quality of security."* All software products display a standard label that provides the consumer with information about the product's rating on standard security vulnerability measures. Consumers use this information to make informed trade-offs between product cost, features, and security risk. Market-based economic incentives drive software manufacturers to produce home and business software with fewer vulnerabilities. As a result, fewer home computers are compromised and used in botnets. Government involvement is limited to regulating product rating and labeling, similar to programs for product safety labels by the U.S. Consumer Product Safety Commission (CPSC), food nutrition labels by the Food and Drug Administration (FDA), new car mileages ratings by the Environmental Protection Agency (EPA), and Energy Star ratings by the EPA and Department of Energy.

Method – Government funding agencies, partnering with NIST, solicit research on tools and metrics for rating software vulnerabilities. One agency such as NIST or the FTC is tasked with developing rating and labeling requirements. This agency sets up a commission that includes leading members of the scientific community, industry, and government to provide input on methods, tools, and procedures for producing ratings. The commission should also include lawyers familiar with government regulations and laws on product labeling. The commission should produce a report within the first year,

and the responsible agency should produce a draft version of the rating process and labeling requirements by the end of the second year. During this time, the research programs into tools and metrics should be producing results that can incorporated into the draft requirements.

Dream team – NIST, CERT, FTC, DHS, software security industry, major software development companies, universities and other research organizations

The Boeing Company
Intelligence and Security Systems

Providing Mission Assurance in the Face of Cyber Attacks and Other Disruptive Events

In Response to:
Request for Input (RFI) – National Cyber Leap Year
C/O National Coordination Office for Networking and Information Technology Research and Development

15 December 2008

Submitted To:
C/O National Coordination Office for Networking and Information Technology Research and Development
Suite II-405, 4201 Wilson Blvd. Arlington, VA 22230

Submitted By:
The Boeing Company, Intelligence & Security Systems
1330 Inverness Dr, Suite 330
Colorado Springs, CO 80910
Tel : 719-572-8188

Technical Contact: Ismael Rodriguez

Contracting Contact: Terri Ferrari

Title: **National Cyber Leap Year 2008 – Submission**

"**Cyber Sentinel Agents**"

Cyber Sentinel Agents

Who We Are

The Boeing Company, Intelligence and Security Systems, www.Boeing.com ,
Boeing Phantom Works and Integrated Defense Systems are at the forefront of research, development, and implementation of advance cyber solutions. Additionally, Boeing maintains one of the largest worldwide network infrastructures supporting all of its business operations with international partners.

Game – Changing Dimension: Morph the gameboard

This approach brings our network defenses closer together, working actively, stealthily, in real-time, and with minimal human intervention so even zero – day attacks can be prevented.

Concept

This approach uses multi – intelligent mobile agent architecture and expert systems as the cornerstones for dynamic, distributed, real – time cyber attack protection to include: distributed attack prevention, detection, correlation and reaction while providing self-healing and learning from new attacks.

Vision

Intelligent mobile agents have several desirable capabilities such as autonomy, rationality, reactivity, and inferential capability. These qualities lend themselves to dynamic improvements and even overcoming most of today's network protection limitations. For example, when an intrusion detection system needs to examine large amounts of data, it typically transfers the data to a node to be processed. Intelligent mobile agents can simply move to the node where the data resides and process it there. This reduces the transfer load on the network and preserves bandwidth. Our approach would manage a set of intelligent mobile agents with specialized functionality acting autonomously, in real – time and reporting back to a central cyber security management station per mission design. Their individual designs will rely heavily on expert systems technology. Most of these technologies already exist but are not very mature and have not been integrated to work as we envision. Our research will concentrate on; developing an effective architecture that includes all the essential elements of protection, investigating the strength of applicable expert systems technologies, looking at ways to eliminate false positives, and making these agents operate in stealth – mode while preserving immunity to unauthorized control.

Method

The process we used to formulate this idea follows the below logical steps:
1- Gather network event data from target networks during normal operations. Network event data shall be gathered spatially and temporally. This data will serve as the training set for various intelligent mobile agents and as formal representation of 'normal' user activities in the expert system.
2- Develop Expert Systems: The expert system shall have at least three major components: the formal representation of 'normal' activities; known attack signatures; and learning capabilities. This last component will enable the real – time expansion of the knowledgebase with newly discovered attack signature algorithms. Also, the expert system shall expand the knowledgebase

of normal user activities. As a precaution, a human Subject Matter Expert (SME) shall be consulted before the knowledgebase of normal user activity is expanded. Once an event has been deemed 'normal', intelligent agents can be retrained on newly discovered normal activities thus reducing the number of false positives detected. The third component will include any new disposition actions the expert system has been SME trained to take.

3- Develop Agents with Specialized Functions:

a. Special Collector Agents (SCA): These agents are interested in pre-defined sets of events categories. SCA patrol the network collecting event data of categories of interest to it. These categories would be defined and tested during our research.

b. Intelligent Correlator Agents (ICA): These agents are trained on a subset of 'normal' events. Collectively, ICA represent normal user activities. ICA receive data from SCA and correlates it against 'normal'. If the resulting correlation returns events outside 'normal', the events are tagged anomalous. The ICA could take actions like: (i) report anomalous events to Intelligent Analyzer Agents (described below); (ii) poll ICA peers external to its network to determine if the same subset of anomalous events are present; if so, the ICA peer could set an 'urgent' flag (meaning this set of anomalous events have been observed on another network) and send the anomalous event report to the local Intelligent Analyzer Agents.

c. Intelligent Analyzer Agents (IAA): We envision several types of IAA. There are IAA trained to recognize known attack signatures, IAA cognizant of each network's observed anomalous events, IAA that create intelligent mobile agent task forces to be dispatched to external networks looking for traces of anomalous events previously reported on other networks, and IAA with statistical models computing the likelihood that observed anomalous events could be a trace of a new attack signature or simply normal user activity. IAA work to detect, react, and decrease the likelihood of distributed coordinated attacks. IAA communicates observed activities to Intelligent Manager Agents (described below).

d. Intelligent Manager Agents (IMA): IMA receive inputs from IAA and alerts the expert systems to the types and locations of anomalous events and identified attack signatures. The expert systems decide the disposition action and communicates these actions back to the IMA. The disposition action is carried out by the IMA.

Dream Team

The Boeing Company

Boeing Phantom Works and Integrated Defense Systems have been researching, developing and integrating intelligent mobile agents for military agencies for more than twenty years.

Exsys, Inc.

Exsys, Inc. has twenty – five years of developing knowledge automation and expert systems technologies and consulting to businesses, government and military agencies.

Recursion Software, Inc.

Since 2001, Recursion Software, Inc. has been providing middleware for developing intelligent applications using mobile agent technology. Recursion Software, Inc. actively partners with various universities, research programs and industry partners.

The Boeing Company
Intelligence and Security Systems

Providing Mission Assurance in the Face of Cyber Attacks and Other Disruptive Events

In Response to:
Request for Input (RFI) – National Cyber Leap Year
C/O National Coordination Office for Networking and Information Technology Research and Development

15 December 2008

Submitted To:
C/O National Coordination Office for Networking and Information Technology Research and Development
Suite II-405, 4201 Wilson Blvd. Arlington, VA 22230

Submitted By:
The Boeing Company, Intelligence & Security Systems
1330 Inverness Dr, Suite 330
Colorado Springs, CO 80910
Tel : 719-572-8188

Technical Contact: Ismael Rodriguez

Contracting Contact: Terri Ferrari

Title: **National Cyber Leap Year 2008 – Submission**
"Embedded Packet – Level User Credentials"

Embedded Packet - Level User Credentials

Who We Are

The Boeing Company, Intelligence and Security Systems, www.Boeing.com , Boeing Phantom Works and Integrated Defense Systems are at the forefront of research, development, and implementation of advance cyber solutions. Additionally, Boeing maintains one of the largest worldwide network infrastructures supporting all of its business operations with international partners.

Game-changing dimension – Morph the gameboard

Embedding a unique user ID to each TCP – IP session – packet requesting access to any network resources (internal and external).

Concept

By embedding a small set of user credentials into each packet of a TCP – IP session and making at least one instance of these credentials survive during packet re - assembly, we could achieve consistent end-point authentication, authorization and auditing for every computer generated session. This method shall be similar to the way the TCP – IP protocol identifies packets for re – assembled to form a complete computer session. This unique Identification (ID) will identify a particular user working from a specific platform. Every ID shall be at a minimum, a part of the user's biometric signature and the hardware identifier from the platform/location the user is working. Therefore, the ID is not portable. Every platform operated from will validate the user's login information against a unique ID previously created and registered using previously specified user – credentials and part of a particular operating platform. If authorized, a user can still get network access from any platform within the organization, but each shall have a unique ID for every authorized user. As a common network resource, every computer platform will validate the user login credentials against a local network manager's Access Controlled List (ACL) were all registered users IDs reside. Once validated, from that point on the operating system (OS) will embed the unique ID into every session - packet created from that machine. This will help us eliminate or reduce unauthorized access to our government networks more than 98% of the time. It will also help us identify the originating source of any attacks without having to rely on complicated and legally challenged hack – back techniques.

Vision

Our vision is to encode and embed a very small set of user credentials (based on a subset of a combination of biometrics authentication, digital certificate, access location, hardware MAC ID, etc). These credentials will be unique for each end-point user and shall serve to ID the user and every move they make in a network at all times during a session. This ID shall be embedded in every packet and survive every re – assembled session until closed. This ID can be used to make access control decisions at every point during a session at a lower level of granularity and more persistently than current PKI or other access control systems. Access to any destination network resources should be dependent on matching, specific users credential IDs against the local ACL for a particular network resource. Each unique ID shall be encoded, embedded and

tamper protected. The definition, encoding, embedding and survivability techniques for this unique ID would be part of our proposed research in developing this idea.

Method

We reached this idea from reading and hearing daily news about the difficulties our government endures every day trying to determine where successful computer hacker intrusions, virus, worm, and malware infestations are originating. Our proposed idea would give us the ID of the perpetrators and show us undeniably every step taken to achieve their goals. Also, wide implementation of our idea would help us prevent unauthorized users from achieving even step one of any cyber attack. Insider attacks can be deterred from the knowledge that any malicious move can be detected and the source easily identified and stopped. Unsuccessful authentication and authorization within this networking paradigm will result in a blocked attempt to the resource, an audit record of the invalid access attempt recorded, and an access denied error returned to the originator. Network resources themselves (applications, printers, databases, etc.) do not need to determine or enforce access at this level. The modified networking implementation supporting this concept can operate at the kernel level of the OS on a given management device to govern access to resources – only successfully authorized attempts will go through, all other unauthorized attempts will be prevented from ever reaching the resources. Audit records of all attempts will be created at each decision and enforcement point. This will facilitate effective reconstruction of user activity and tracking of unauthorized attempts back to the origination point.

This method provides capabilities and facilities above and beyond those currently found in IPSec and IPv6 specifications, which support enhanced authentication as well as authorization, enforcement and auditing at the network level. IPSec can be exploited by internal (trusted) users with minimum detection. This new enhanced method ensures user identification is inherently attached to each packet so any unauthorized attempts can be easily stopped before it is too late. Each packet's origin should be identifiable and dealt with appropriately.

Another part of our research would be to study best methodology for creation, registration, encoding, embedding, distribution and management of these credential IDs and synchronization with local ACL's. We also need to look into making these functions an integral part of Network Management systems, such as Tivoli and HP Open View products. Making this part of the day-to-day management (updates, revocation, registration, auditing, etc.) and security configuration of our network resources will improve consistency and effectiveness.

Dream team

- Cisco, IBM, PKI and Biometrics vendors
- NIST, NSA, DISA, DoD, IATFF, & ITF
- Boeing – as Large Systems Integrator for developing this idea

Who you are – http://www.itl.nist.gov/ – We are the Nation's premiere measurement laboratory, devoted to improving quality and security in information systems.

Game-changing Dimension – Change the rules

Concept – Organizations worldwide are adopting open-source software as the basis for network-connected enterprise and infrastructure. A fledging industry is researching, developing and marketing practical tools for static and dynamic analysis of software. What if we fostered an environment that encouraged tool developers to analyze open-source code and to publish results from those analyses?

Vision – Users consult publicly available analyses to assess quality and reliability of key open-source software products. Widespread user acceptance encourages open-source projects to repair identified code weaknesses. Competition arises among developers of analysis tools. Identifying weaknesses in tools stimulates R&D investments to enhance capability. Improved capability expands the market for analysis tools. Overall quality of software improves measurably. Global enterprise and infrastructure becomes less vulnerable to attack.

Access and Use – NIST creates and evolves a program that solicits and selects key open-source software products (e.g., Apache web server, MySQL database server, BIND domain-name server and Firefox web browser) for static and dynamic code analyses. NIST also solicits open participation by tool developers (e.g., Mathworks, Green Hills Software, LDRA Software Technology and PRQA Programming Research). Working with stakeholders, NIST establishes methods for analyzing open-source software and reporting results. The methods and analysis results are published on a searchable Web site. After reaching a critical mass, NIST publicizes the site for users of open-source software and analysis tools. Users, open-source projects and tool developers are encouraged to expand the scope of coverage by the site. NIST uses data derived from the site to measure progress in improving the quality of open-source software and also analysis tools. R&D funders use the site to assess the current state-of-the-art in analysis tools and to identify key gaps and promising avenues for funding. Eventually, the process is spun off into an industry-run organization.

Feasibility – Current practice regarding software quality is to certify that code is developed and tested following specified processes. The Holy Grail is automatic generation of code using provably correct transformations from provably correct specifications. We advocate a pragmatic middle ground. As David Rice points out (Geekonomics) the software market suffers from "asymmetric information", where buyers cannot be sure what they are getting. Rice proposes a model where software is rated on a five-star scale using an objective measure. The problem, of course, is defining a useful measure. We propose a pragmatic approach that allows open-source projects (some mentioned above), purveyors of analysis tools (some mentioned above) and software consumers to cooperate and compete in an open environment. Rather than define an

objective measure, we exploit evidence from multiple competing analyses with different tools. Rather than define a rating scale, we let the evidence speak for itself. At the same time, using data that becomes available, we develop measures to assess the overall trajectory in quality for open-source software. Perhaps our measures might eventually be adopted by developers of proprietary software to assess their own progress toward improving software quality.

Next Steps – Achieving our vision requires creating a positive feedback loop: tool developers publish analyses; users rely on analyses to demand corrections in open-source software; open-source projects improve software quality; tool developers compete to improve analysis tools; R&D funders target identified needs; analysis capabilities improve; users require proprietary software developers to exploit analysis tools; overall software quality improves. To begin, we need to persuade a core subset of tool developers to work with us to provide high-quality analyses of selected open-source software. This may require some cost-sharing. Subsequently, we need to promote initial results in three ways: (1) convince open-source projects to address identified weaknesses, (2) engage tool developers in comparing results from analyses of open-source software and (3) attract interest from potential R&D funders. Sufficient analyses of open-source software can lead to creation of success stories for marketing to users of open-source software. After three years of analyses, NIST should have enough data to quantify quality improvements in subject open-source software and to characterize the state-of-the-art in analysis tools. Further, injection of R&D funds into tool developers should encourage expanded participation in publicly available analyses and tool evaluations, especially if such participation is a condition for funding. Finally, user demand can stimulate adoption of analysis tools for assessing proprietary software and, thus, expand the market for such tools.

Method – This idea was stimulated by a whitepaper from Green Hills Software, which applied their *DoubleCheck* static analysis tool to the Apache Web server software and identified several code weaknesses. Unfortunately, the Apache Web server still exhibits the identified weaknesses. We began thinking about missing feedback loops and then conceived a world that engaged tool developers and stakeholders to create mutually beneficial positive feedback. Further, we foresaw a path leading from existing analysis tools to substantial public value; augmented by a path leading from academic research into enhanced, pragmatic analysis tools. The missing ingredient is leadership to create an environment motivating tool developers to analyze open-source code and to publish results from those analyses. The ultimate payoff: improved software quality.

Dream Team – Initially, commercial developers of analysis tools (e.g., Mathworks, Green Hills, LDRA, PRQA, Coverity, Parasoft, Sun, Enerjy, Klocwork, Valgrind and IBM) and Google and managers of key open-source projects (e.g., Apache, BIND, MySQL, Firefox, Linux, and SendMail). Early involvement from relevant research funding agencies (e.g., DARPA, DoE, NASA, NSF and OSD) would also be productive as a lever to encourage participation by tool developers and, later, as a conduit to create connections between academic researchers of analysis techniques and commercial tool developers. And of course, NIST should participate.

Security Information Exchange

Internet Systems Consortium
950 Charter Street, Redwood City, CA 94063
https://sie.isc.org

December 15, 2008

Who you are – www.isc.org - Internet Systems Consortium, Inc. (ISC) is a nonprofit 501(c)(3) public benefit corporation dedicated to supporting the infrastructure of the universal connected self-organizing Internet — and the autonomy of its participants — by developing and maintaining core production quality software, protocols, and operations.

Game-changing dimension – Morph the Gameboard

Concept – The Security Information Exchange (SIE) has been created as a trusted, private framework for information sharing in the Internet Security field. Participants can operate real time sensors that upload and/or inject live data to SIE while other participants can subscribe to this data either on-site in real time, or by query access, or by limited and anonymized download.

Vision – Attempts to combat criminal Internet activity have been inefficient, thwarted or not even attempted because security professionals do not have access to adequate comprehensive real-time information. SIE has created common tools, services and communication infrastructure and is bringing together security researchers, Internet service providers, government agencies, abuse desks, universities, businesses, and law enforcement to share and analyze Internet Security information. The organizations utilize a common privacy and legal framework to protect the use of data within the common SIE infrastructure. Tight restrictions are placed on data entering and leaving SIE. The network effect from such collaboration with consolidated data will allow each participant to build correlations between disparate data sets in real time that would otherwise not be possible. Results from the correlations will build new real-time tools available to the participants, including law enforcement for the investigation and prosecution of Internet crime. Results from the collaboration will enable security researchers to prevent new attacks before they happen and close the window of opportunity for any new attacks.

Method – Led by Paul Vixie (ISC) and David Dagon (Georgia Tech), SIE was created by ISC in 2007 with development and operational support from ISC and supplemental funding from the US Army (hardware), NSF (programming & tools), and DHS (operational support). SIE privately requested several ISPs and services to donate real-time DNS sensor information for research purposes and as a result made correlations of DNS data toward identifying phishing and botnet activity. SIE currently maintains the largest collection of passively collected DNS information on the Internet and can provide insight into world-wide changes in DNS resolution patterns by malware and botnets or DNS cache poisoning. Anonymized information collected by SIE has been used to help correlate server information for actions by law enforcement. As a result of presentations and consultations, several commercial and public-benefit security organizations are joining or at least evaluating SIE. We built our first production broadcast facility in June 2008 and plan to build more across the world where data sharing is more efficient or laws prohibit the export of data across borders. We created and will continue to maintain and improve free and open source software that is able to efficiently collect, anonymize, store, and share any type of information across a specially designed network and server cluster. SIE plans to maintain a public registry of security data definitions for use with the software to help avoid duplication of effort between security researchers.

ISC is uniquely positioned to serve the Internet with the Security Information Exchange because:

1) we are a non-governmental organization which commercial and foreign interests can trust more than any individual government,
2) we are a not-for-profit organization that does not compete with commercial Internet security interests,
3) we have a core mission and exceptional competence toward supporting vital Internet infrastructure,
4) we have gained the respect and relationships within the Internet security community because of its public benefit support toward the Internet security community, and
5) we are already building the infrastructure and have operational experience in this area.

We plan to help Leap Year participants understand what information they have, understand what information others in SIE have or need, and build new sharing methods for their data ("stone soup").

Dream team – No single set of key organizations should be defined (and it's against our privacy policy to disclose current participants to outsiders without permission); instead, the dream team is *everyone* SIE can enable to work together to share Internet security information. We desire and request participation from everyone involved in Internet security including Internet service providers (commercial, educational, government), commercial Internet security companies, educational and government-sponsored security research teams, national law enforcement agencies, national and public benefit computer emergency response teams, non-profit public benefit security data collection projects, and others. Contributing participants will provide SIE real-time information about malware, phishing, DNS queries, domain registrations, network mapping, unauthorized network traffic ("darknet"), geolocation, unsolicited emails ("spam") and messaging, network flows, botnet command and control activity, web search data, domain and address-based reputation, and any other information which gains more value as the number of researchers who examine it increases. Any organization that can share security data with another should consider whether SIE can help them more efficiently share the data with everyone who needs it. Any data that can be requested of a network provider can instead be shared with SIE to ease the collection burden for the provider. As a result, law enforcement, network providers, and security companies will have more tools and real time information from the participants.

Cyber Highway Patrol

Who we are – James Horning, Ph.D., Erik Mettala, Ph.D., Stephen Barnett, David Balenson, Stephen Schwab, Howard Weiss, Andrea Colegrove, members of the SPARTA National Security Systems Sector, a group that among ourselves have 182 years of experience in computer and network security research, design, development, and operation.

Game-changing dimension – Change the rules

Concept – A key point in the development of civil societies is the replacement of private security forces by police as the primary maintainers of public order. At least ideally, police provide equal protection to all, and their cost is amortized over the society that benefits. *We propose the creation of a national Information Super Highway Patrol (ISHP) chartered and empowered to enforce relevant laws and regulations and to ensure order on public networks through the legitimized use of various kinds of force.*

Cyber crime is international. An important issue to be settled is whether to build a national force, planning to cooperate with other nations' forces as they emerge (sort of like Interpol), or whether to build an international force from the start—although that seems much harder to launch. An international force would not have to be fully global to be more effective than a purely national force.

Vision – Public networks are regularly monitored and "patrolled" by a publicly funded and accountable force that responds to disturbances, investigates violations of laws, regulations, and protocols, tracks malicious activity to its source(s), and takes action to eliminate or mitigate it. The ISHP is empowered to take actions that are not generally available to private parties (or to CERT), such as requiring ISPs, hosting services, core routers, DNS registrars, etc., to promptly disconnect service to serious violators.

The ISHP is a first responder to cyber-attacks of all sorts (Cyber-911). Distributed monitors measure traffic levels at key points on public networks to quickly identify DDoS and spam attacks and track them to their sources. ISHP also looks for known precursor event patterns of attacks, such as probing.

The ISHP is nearly invisible to ordinary users of public networks as long as they are operating within the rules, and don't encounter a situation where they need help defending against malefactors. But knowing that the networks are being patrolled gives users a level of comfort and confidence that was not possible in the pre-ISHP world.

It is no longer necessary for private vigilantes—who are "always outnumbered, always outgunned"—to bear the brunt of counteracting criminal gangs. Instead, they can report criminal activities to the ISHP, which is empowered to take effective action against them, including activities that are illegal for private parties (e.g., taking botnet zombies offline, examining their files, and/or disinfecting them; seizing control of botnets; directing DDoS against botnet controllers; arresting violators; confiscating equipment—all subject to Fourth Amendment restrictions). [Recall the significant, but temporary, global drop in spam in November 2008 that resulted from taking a single web hosting service off-line, achieved by unprecedented private-sector coordination and cooperation.]

The ISHP has an intelligence section that collects information from CERT, private sources, and the public, as well as by infiltrating and observing hacker networks, operating honeynets, etc. It

"follows the money" that motivates and funds malware development and deployment. It investigates the nature and sources of attacks, takes preemptive action to counter emerging threats, performs statistical analysis, and informs policymakers and the public of trends and emerging vulnerabilities. With strict oversight (to preserve Constitutional rights) it may sometimes analyze the content of network traffic (e.g., by deep packet inspection). It is a repository of expertise for detecting, tracking, mitigating, and punishing network crime.

The ISHP coordinates with police and regulatory agencies in other countries, enlisting their cooperation (to the extent possible) in acting against criminals, who increasingly operate across national borders. Failing such cooperation, in extreme cases it takes direct action against violators outside US jurisdiction.

Making it happen: Most of the technical tools that the ISHP needs are already in use by existing security organizations. Scaling up the tools, effectively dealing with false positives, and providing the mechanisms for oversight will surely raise new technical issues. Relations with existing local, national, and foreign organizations must be worked out. But the critical requirements are

- a clear mandate,
- a legal and regulatory basis for operation and oversight, and
- adequate staffing and funding.

Method – In response to the RFI, we had a pair of brainstorming sessions (followed by multiple email interactions) on game changing ideas. We started with a discussion of the fundamental reasons that adversaries have the advantage in today's cyber infrastructure. Ideas to counter these were collected into themes that gave rise to this and our other submissions.

Dream team – We would like a team that includes our proposal team and diverse experts from: Electronic Frontier Foundation (EFF), Computer Professionals for Social Responsibility (CPSR), ACLU, NAACP, American Bar Association, Center for Democracy and Technology, Department of Justice, Federal Communications Commission, Secret Service, CERT, Government Accountability Office, as well as other security researchers, e.g., IOActive, Kaspersky Lab, McAfee, Spamhaus, Symantec, Trend Micro, Websense Security Labs; Dan Boneh (Stanford), Ed Felten (Princeton), Peter Neumann (SRI), Avi Rubin (JHU), Bruce Schneier, Gene Spafford (Purdue), Dan Wallach (Rice).

"NewNet"

Who we are – Stephen Barnett, Erik Mettala, Ph.D., James Horning, Ph.D., and David Balenson, members of the SPARTA National Security Systems Sector, a group that among ourselves have 111 years of computer and network security policy formulation, research, design, development, and operation.

Game-changing dimension – Morph the gameboard

Concept – Create from scratch a cyber infrastructure providing all the capabilities of the current Internet, whose primary *added value* proposition is security, with the concomitant adoption of rules, rights, restrictions, responsibilities and technology whose absence from the Internet today gives advantages to adversaries. By changing the technical, regulatory, responsibility, accountability, and compliance enforcement foundations of cyber space, NewNet eliminates many of the conditions that enable adversaries to operate within cyber space successfully, and with impunity. NewNet changes the game by forcing adversaries to operate within constraints that decrease their chance of success and increase their risk of detection, removal, and punishment.

Vision – NewNet initially protects vital information and critical services, e.g., SCADA systems. It enables information sharing to support homeland defense and critical infrastructure protection, and ultimately grows to largely replace the Internet (as the Internet replaced Usenet). There is no anonymity. NewNet holds all users accountable for their actions, with strong attribution, authentication, and policing. NewNet uses active protection that prevents successful attacks by eliminating vulnerabilities, and active defense that takes the fight to adversaries, who are tracked down and eliminated. It is built with trustworthy components; all applications are designed and built to run on these components via standard interfaces to ensure that they run securely. NewNet protects information objects, not just the systems that store and process them. It actively and continually analyzes and monitors to ensure that all participants are meeting technical, operational, and behavioral standards. Those who are not are denied access to and use of NewNet, and may face further punitive action. Hardware and software suppliers, IT service providers, and system operators are contractually responsible for creating, maintaining, and as necessary updating their products, services and systems to provide the strongest level of protection possible. Failure to do so results in punitive action such as paying consequential damages to the victims, and removal of their products from NewNet until the problem is fixed (akin to today's safety recall programs). There is a legal and insurance framework that enforces NewNet's underlying principles and provides incentives to the cyber community (users, and technology and service providers) to transition to NewNet. (Perhaps the pending switch to digital broadcast TV serves as a model; adapt or lose service.)

Proper individual use of NewNet is the key to success. There is a process for training, testing, periodically retesting and indelibly identifying licensed users. To access and use NewNet, individuals will have to be trained and use their non-forgeable cyber identity. As with other specialists serving the public, only trained and "board certified" cyber engineers, operators, and maintenance personnel will design, build, operate, police, and maintain it.

Why is this possible? Individual, corporate and government users of the Internet face substantial losses of their information and services due to threats that currently exist on the Internet. A trustworthy infrastructure with well defined rules for using it safely, and with active policing to

ensure conformance to those rules, will be a more valuable asset for both users and service providers. Much of the fundamental technology is available.

For this to become real, we must research and develop a legal basis and policy doctrine to underlie NewNet. Standards must be established, including an architecture for information protection, authentication, attribution, access control, and active monitoring. A security technology gap analysis must identify which technologies exist and which need further R&D. Key issues will be technology scalability in deployment, operation, support, and management; incremental transition to NewNet; providing controlled interaction with the existing Internet; securely "wrapping" legacy COTS solutions; a secure way for the general public to attach enclaves to NewNet in a controlled manner, benefitting from some of its protections without entirely replacing legacy systems; promoting the transition until NewNet is self-sustaining; and creating an effective workforce training and certification program.

What exists and what is needed? The technology pieces to provide a trustworthy foundation for the new infrastructure and to protect information objects exist. What's still needed are: a scalable and manageable system for multifactor authentication; mechanisms to support attribution and provenance; additional tools and techniques to monitor and detect non-compliant actions in a dynamic, high volume information processing environment; trustworthy applications; improved techniques to defend against, or predict and disrupt, attacks; systems that are automatically in "secure mode," with an easy to understand and use interface for ordinary users; a strategy and funding for subscribers to transition to NewNet, while continuing to get services and information from the Internet without compromising NewNet's protections; and a secure interface device for attaching legacy enclaves to NewNet. The development of NewNet also provides the opportunity to design and implement a different protocol model that better supports the security functionality. There must be government and industry commitment and investment to making NewNet a success during its development and fielding.

Method – A group of our senior technical staff members brainstormed on game changing ideas in response to the RFI. We started with a discussion of the fundamental reasons that adversaries have the advantage in today's cyber infrastructure. The ideas were collected into themes which gave rise to the structure of this and our other submissions. We assume that individuals and organizations are willing to comply with rules and restrictions in order to gain better confidentiality, integrity, and availability. With adequate demand, quality cyber security products and services that are easy to use correctly can be developed, refined and distributed with little, if any, additional cost to the end user.

Dream team – This proposal involves the cultural, legal, financial, business, indemnity, and technical issues that prevent the Internet from becoming a safe place in which to operate. We would like a team that includes our proposal team and experts in the following areas: law, insurance/risk, security technology, industrial IT, finance, privacy, history of technology (paradigm shifting examples), and education. In security technology, the team should contain people such as: Dave Clark (MIT), Tim Gibson (DARPA), Charles Herzfeld (Potomac Institute), Tom Leighton (MIT), Larry Peterson (Princeton), John Wroclawski (USC/ISI), Ron Rivest (MIT), Larry Roberts (Anagran), Pradeep Sindhu (Juniper), Bob Taylor, Loren Thompson (The Lexington Institute) and Brian Witten (Symantec).

Trusted Application Suites

Who we are – Erik G. Mettala, Ph.D., James Horning, Ph.D., Stephen Barnett, David Balenson, Stephen Schwab, Hugh Harney and Howie Weiss, members of the SPARTA National Security Systems Sector, a group that among ourselves have 186 years of experience in computer and network security research, design, and development, and in the development and operation of secure computing systems and environments.

Game-changing dimension – Morph the gameboard

Concept – Institute a "Manhattan Project" for Network and Application Security. Design, develop, test, and release as open source a suite of trusted applications enabling complete shielding of users by cryptographic protection additions built on top of Security Enhanced (SE) Operating Systems (SE-Linux, SE-BSD, SE-Darwin (Mac OS-X)), married with hardware roots of trust, such as commercially available Trusted Platform Modules (TPM). The Trusted Application Suite would include SE-eMail, SE-OpenWord, SE-OpenPowerpoint, SE-Firefox, SE-MySQL, SE-Apache, SE-Lucene, and SE-WfMS.

Vision – A long-established Security Engineering principle is to use trust chains as a way of describing and maintaining the security between sender and receiver in networked communication. Trust management implies that authenticated trust is established among a sending user, his end-host hardware, operating system, and services, and a receiving user, his end-host hardware, operating system, and services; confidentiality, integrity and availability of all intervening systems that move the data are also part of the trust chain, through encryption of data in transit in a detailed and forensically auditable way to ensure the security of communication. Establishing trust chains between computers and networks, as it turns out, is far easier than establishing trust between an authenticated user and data that the user prepares, saves, or transmits.

To facilitate widespread adoption of end-to-end, user-to-user and application-to-application, cryptographically-secure communication, we must lower the pain of security integration for the end user. Our goal is to provide end-application hosted mechanisms that enable creation of documents, spreadsheets, presentations, e-mails, etc., where trust relationships and classification tokens are accessible by default to application users, and integrated at the time of document creation or modification. Applications create documents that in turn transfer and preserve the trust relationship from the user-application pair to the operating systems services, such as file saving and file attachment to e-mail, so the user identity and security properties can be trusted by receiving user-application pairs.

> "In anticipation of emerging encryption product capabilities as well as for device authentication, DOD Components shall ensure all new computer assets (e.g., server, desktop, laptop, and PDA) produced to support the DOD enterprise include a Trusted Platform Module version 1.2 or higher where such technology is available."
>
> ~ US DOD Memorandum, Article 4, July 2007

Why is this feasible now? For three reasons: First, low cost hardware roots of trust are now widely available in the form of Trusted Platform Modules built in accordance with the Trusted Computing Group's recent specifications. Second, Security Enhance Operating systems such as SE-Linux now have mature forms of the mechanisms for mandatory access control, audit, and security policy enforcement that are required to bridge the security gap between user applications and the underlying hardware.

Third, the open source software community has developed a number of extraordinary user applications, enabling free access to enterprise-grade end host applications, as well as world-class refactoring tools that would enable very rapid integration of trust relationship preservation properties into user applications.

The Trusted Platform Module (TPM) is a hardware-based security and cryptography chip built into virtually every enterprise-class desktop and laptop computer—PC or Mac—that ships today, as well as numerous consumer and small business configurations. Recent reports indicate more than 100 million computers have shipped with a TPM installed, and a number of RFPs from the Fortune 1000, as well as numerous government agencies, including the Department of Defense, explicitly require a TPM for all new computers. Even though the chip is widely available, and management tools ship with enterprise PCs, many organizations, and most individual users, have not yet put this valuable security tool to work.

Hardware Root of Trust: The TPM creates a hardware-based foundation of trust, enabling enterprises to implement, manage, and enforce a number of *trusted* cryptography, storage, integrity management, attestation, and other information security capabilities.

Security Enhanced Operating Systems: SE-Linux, SE-BSD and SE-Darwin (MAC OS-X) each have mechanisms for mandatory access control, separation kernels, object security, audit, and security policy suitable to enforce strong trust relationships between TPM identity registers, and applications to preserve trust relationships in nominal trust chains.

Security Enhanced End-Host Applications: It is now feasible to securely integrate trust relationship management and trust credentials into major applications, where user involvement in the selection to use security enhancement features is embedded in SE-eMail, SE-OpenWord, SE-OpenPowerpoint, SE-Firefox, SE-MySQL, SE-Apache, SE-Lucene, and, SE-WfMS. By integrating trust management into the normal usage patterns of applications, most, if not all difficulties in using strong trust will be eliminated, making the adoption and dissemination of strong trust the normal mode of operation, as opposed to the exceptional case.

Integration and Transition: Our recommended approach would be to build a dream team of open source developers and security developers to integrate strong trust through SE-operating systems up through end-host applications that are usually found to be the source of the documents, web pages, and e-mail that carry malware today. This approach would create secure applications that would enable secure information sharing among all federal, state, local and tribal government agencies. By providing open source applications that ease the interface to a strong trust model—through e-mail, web browsers, web servers, office applications and workflow—we create protected enclaves of trust suitable to protect government, business, and end user use that do not depend on their connection by secure networks. This leads to protecting both applications and the user's data by careful engineering of the "stack" including server processes, middleware, OS and trusted hardware TPMs.

Method – We had a pair of brainstorming sessions (followed by multiple email interactions) on game changing ideas in response to the RFI. We started with a discussion of the fundamental reasons that adversaries have the advantage in today's cyber infrastructure. The ideas were collected into themes which gave rise to the structure of this and our other submissions.

Dream team – We would like a team that includes our proposal team and experts from: The Trusted Computing Group (TCG), Apache Software Foundation (ASF), Free Software Foundation (FSF), Open Office Organization, NSA, anti-virus and other security researchers, CERT, and the Government Accountability Office.

Who you are – We are the Steganography Analysis and Research Center (SARC) (www.sarc-wv.com), a Center of Excellence within Backbone Security (www.backbonesecurity.com), a private sector company specializing in advanced computer security. Since June 2004, the SARC has continuously and exclusively focused on digital steganography (data hiding) research and the development of state-of-the-art steganalysis products to detect and extract digital steganography.

Game-changing dimension – Morph the game board and raise the stakes

Concept – Bad guys are using any of the 1,000+, and growing, easy to find and use digital steganography applications freely available on the Internet today to conceal evidence of criminal activity. Criminals, including terrorists, also communicate covertly with zero risk of detection because there are currently no commercially available network security appliances capable of detecting digital steganography.

For example, the adjacent photograph contains a 110-page, 37,0245 word, extract of a terrorist training manual, invisible to the naked eye. The photograph has room for an additional 72,094 characters! The potential use of steganography for covert terrorist communication and concealment of criminal activity represents a significant threat to national and civil security yet few are aware of the magnitude of this threat.

What if we morph the game board by deterring and pre-empting terrorist activity and concealment of other criminal activity? Criminals and terrorists are currently able to use digital steganography to conceal criminal activity and hide their messages without being detected. We can change that by exposing the use of digital steganography and making it much easier to extract hidden information. We can also raise the stakes by increasing the probability of detecting and extracting information hidden by the bad guys by making the technology available to those who need it most.

Vision – Detect the use of digital steganography to plan terrorist acts and conceal criminal activity. When use of steganography is detected, provide the capability to find and extract valuable evidence for prosecution.

We can prevent terrorist acts and prosecute criminal acts planned or concealed through the use of steganography by providing Federal, state, and local homeland defense and law enforcement authorities with the means to identify the illicit use of steganography. We can design and implement network security capabilities to detect bad guys downloading or using digital steganography as traffic passes in and out of government networks. These capabilities can be implemented on a dedicated platform or integrated with other network security appliances such as intrusion detection systems, intrusion prevention systems, network behavior analysis systems, data leak prevention systems, etc. The key to success is to design the systems and make them readily available to the organizations that work to prevent and prosecute terrorists and criminals.

Deploying a capability such as this would raise the chances of detecting bad guys attempting to download or use steganography. They would not be able to use digital steganography to hide

information without running the risk of being detected: thereby thwarting their plans and significantly increasing the probability of their successful prosecution.

Method – Our proposed method includes three phases.

Phase I: Design a system to integrate the SARC's evolving steganography detection and hidden information extraction capabilities into Einstein 2.

Einstein is the network monitoring tool developed by DHS' US-CERT to automatically monitor and analyze Internet traffic transmitted into and out of federal computer networks, filtering packets at the perimeter. Anomalies that may represent unauthorized access or other hacker activity are reported to US-CERT, the operational arm of NCSD, which then moves to quarantine and disinfect affected machines. Whereas the first version of Einstein was designed to passively monitor network traffic flow, Einstein 2 is being designed to actively analyze network traffic content. Application Program Interfaces need to be designed to interface Einstein 2 with the application and signature detection and automated extraction capability developed in the SARC

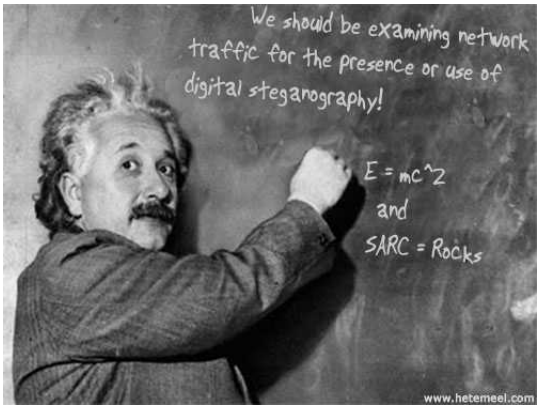

Phase II: Deploy the SARC's current steganography detection and hidden information extraction capabilities in Einstein 2 to all government agencies currently using Einstein.

Detecting steganography applications in network traffic is a very strong indication the applications will be, or have been, used to steal sensitive information. Automated detection notifications will be generated and sent to agency network security administrators to facilitate increased surveillance of the sending/receiving parties. Subsequent monitoring may yield carrier files from which hidden information hidden may be extracted. Expanding the functionality of Einstein 2 to detect the presence or use of digital steganography applications will raise the stakes by dramatically increasing the probability of detecting covert information flows which could reveal terrorists plans or criminal activities in progress.

Phase III: Expand steganography R&D to discover additional applications and signatures and develop additional automated extraction capability for subsequent versions of Einstein.

Expand research to find additional steganography applications and discover additional signatures. Expand development to develop, test, and integrate additional detection and automated extraction capability for incremental updates to Einstein 2.

Dream team – DHS, HSARPA, DoD, DARPA, DC3, DISA JTF-GNO, TOC, IRC, NSTC, NITRD, NCS, NSTAC, NSIE, network communications engineers and security specialists, IDS/IPS specialists, network and system administrators, database administrators, programmers.

Granular Caching Resolver: A Poison *Resistant* DNS Server

David Dagon[1] Manos Antonakakis[1] Wenke Lee[1]
[1]Georgia Institute of Technology Information Security Center,
801 Atlantic Dr. Atlanta, Georgia, USA 30332-0280

December 15, 2008

1 Who You Are

We are GTISC, the Georgia Institute of Technology Information Security. Our area of research has focused on botnets, spam, malware and DNS security. We have written numerous papers covering the topic of DNS security, DNS poisoning, and other threats to the resolution infrastructure.

2 Game-Changing Dimension

We must *change the stakes* in DNS poisoning. Such DNS attacks are currently an all-or-nothing game, where the successful attacker completely controls the victim cache contents.

We need to limit the impact of poisoning so that even successful attacks on a resolver's iterative phase causes limited damage to stubs. Ideally, DNS servers should be able to preen their caches and remove suspect or poisonous records. Periodic reviews of DNS caches (with selective purging or cache replacement) will greatly reduce the harm caused by poisoning.

3 Concept

At present, it is extremely difficult for IDS equipment to monitor DNS servers for poisoning attacks. For example, cache evictions performed by DNS servers are are not visible to outside sensors. Similarly, whether or not a DNS server drops or caches additional or authority records may not be clear, unless the DNS implementor has detailed their exact interpretation of RFC 2181. We need a standardized way to expose contents of DNS caches to IDS equipment. Further, we need a way for IDS sensors to ask DNS servers to purge selected items from their cache, based on additional information (e.g., TCP session data) gathered outside the DNS server.

Because DNS caches are currently black boxes, poisoning events are treated as a single fatal event: once the poisonous record is accepted by the resolver, the cache is effectively polluted forever (or for a lengthy `TTL`), and the DNS daemon must be restarted to recover.

We think DNS poisoning should be instead treated like a snake bite: the patient is indeed poisoned, but has a window of time to purge the poison. If DNS servers had an API that exposed their cache to external analysis, then poisonous records (and all subsequent records derived from it) could be identified and removed. Clean, correct cache entries could also be injected by trusted, ACL'd analysis nodes.

The benefits of creating a standardized DNS cache API include:

- ***Fewer victims.*** Right now poisoning is all or nothing. DNS servers should instead have the ability to heal themselves even after being poisoned. As a result, malicious records would be given out to fewer users.
- ***Improved forensics.*** By tracking poisonous records (e.g., `NS` replacement) and associated malicious records, we *significantly* facilitate the growth of new field: *DNS forensics*.
- ***IDS-DNS Hybrid***. By exposing the DNS cache to external inspection and cleaning, we create an opportunity for existing IDS equipment to perform more rigorous network analysis of DNS traffic, and independent verification/re-evaluation of cache entries. To the IDS, DNS servers and their cache are a blackbox. By exposing internal state of the caching resolver, we foster the development of an IDS-DNS hybrid.

4 Vision

We need transitional DNS security technology, since DNSSEC will likely require years to see wider adoption. At present, DNS resolvers lack fine-grained *policy filters*. Resolvers are engineered for high-availability, and beyond the RFC minimum (e.g., RFC 2181, RFC 1034), are not skeptical of the answers they receive.

But once the internal state of a DNS server's cache is exposed via an API, IDS and firewall technologies can review answers and purge poisonous records. This greatly limits the number of "wrong" answers a DNS server will provide to users. It also allows a DNS server to operate in a compromised state (still providing known good answers while questionable or poisonous answers are held under review). Further, external review of DNS cache lines can provide an early warning of attempted DNS poisonings, providing some measure of IDS functionality to DNS servers.

5 Method

We observe that the DNS vendor community is aggressively hiring security technologists. We also note that, while DNSSEC is seeing wider adoption, DNS technologists are focused on "interim solutions" such as SPR, DNS-0x20, and bailiwick logic. There is significant vendor interest in making the current resolver footprint of the Internet more resistant to forgery, pending wider acceptance of DNSSEC. At the same time, DNS poisoning attacks are increasing, particularly after the discovery of protocol-level vulnerabilities.

We note that at present, IDS technologies usually perform extremely limited inspection of DNS traffic, since the DNS cache is a blackbox. The architecture of some DNS servers is complex (BIND) or close (commercial offerings), making external cache examination difficult.

We propose the creation of an API to allow remote inspection and monitoring of DNS caches, so that IDS technology can provide deeper, bailiwick-level inspection of DNS traffic. A simple read/write cache API could allow a DNS server to broadcast DNS cache evictions, NS replacements, and other events. Further, the API could potentially leverage other cache management innovations (e.g., `memcached`).

6 Dream Team

Implementers would include: Georgia Tech's Information Security Center, the Internet Systems Consortium, NLnet labs, OARC, the major DNS vendors, and NIST. This group would develop a vendor-neutral cache-manipulation API, release patches for the open source DNS resolvers (e.g., BIND), and pursue appropriate standards through the IETF process. From the other end of the telescope, we'd also need to involve the IDS/Network Security community (Sourcefire, Trend Micro), to leverage the exposure of DNS cache lines.

RFI Name - National Cyber Leap Year - RFI-3

Title of Concept - Using a Positive Security Model with Smart Data for the *Security Triple Play* – Secure Data, Simplify Management with Attribution

RFI Focus - "Morph the Gameboard"

Submitter(s) - Scott Ruple, CEO/President
Tracy Crowe, Vice President
FireRock Technology Corporation
14835 E. Shea Blvd., Suite 103-286
Fountain Hills, Az 85268

Submitter Summary – FireRock Technology Corporation is a privately held Delaware corporation with headquarters in Fountain Hills, Arizona. The company was started in 2006 by a group of computer security, systems, storage, and networking professionals with a commitment to provide solutions for those who are not satisfied with the current point solutions for data security that are burdened with compromises.

Concept – The cyber playing field continues to grow and evolve at a relentless pace. In less than five years, the Internet has gone from primarily a "push" model to one that is more interactive; mobile phones have gone from voice-based communications to smart-phones that are ultra-portable PCs and storage devices small no bigger than a person's thumb with capacities over 50 GB. But, data security is still based on reactive, negative protection models. In other words, turn systems into "vaults" with strict gatekeepers, or worse yet, wait until a threat reveals itself, then build a defense against it. In the former, while secure, everyone's accessibility to information is severely impacted. In the latter, it's usually too late – the data/system has been compromised and the best that can be done is to use forensics to identify affected areas, vulnerabilities and hopefully the culprit. In these cases, the impact on business productivity, systems reliability and most importantly, information protection is tremendous. It should also be noted that these systems do not lend themselves very well to a global, highly interactive network where zero-latency, ease of data access and flexible data access technologies are key.

FireRock Technology is morphing the game-board with their development of an architecture that is based on a positive data security model using "smart data" to enable the *Security Triple Play*. The Triple Play means the content of a file is secure from anyone who isn't allowed to see it, the system is self-managing, based on a policy-based management system that uses business rules and ensures the who, what, where and how information about the files is always available for enhanced accountability – attribution. In this model, information protection is a function of identity, acknowledged escalation of privileges and this information is embedded in the file with the data, hence "smart data." The positive data security model with smart data ensures the

content of file is secure, accessible by only the appropriate individuals or groups without impacting its manageability.

Vision – FireRock's vision is for a positive security model that protects data in a truly ubiquitous manner that supports business objectives, flexibility and can function on a variety of clients and servers inside and outside of corporate boundaries – in other words, the over Internet. It is a system that is non-intrusive, ensures minimal interruptions to business processes, rules-based for simplicity; integrates seamlessly with existing security and data management solutions and includes facilities for improving performance through reduction of data access times.

By embedding critical security and other information about the contents of a file within the file, as smart data, applications, regardless of the platform they are running on, will only need to "filter" the intelligence off the file before processing it – similar to the way email clients strip and interpret MIME header information out of a message before presenting it on the screen. This will allow all levels of data security, from simple, lightweight encryption all the way through FIPS certification, to be enacted on the data. Furthermore, the system is designed to provide key information on the data like who, what, where and how it was used for enhanced accountability and management.

Method – With well over 100 combined years of experience in security, storage, systems, and networking, the FireRock team has gained broad knowledge of industry requirements along with an extensive understanding of technical and customer issues that provided the foundation for their vision. Upon the infusion of additional capital, FireRock will complete development of the first generation of their suite of products. Once completed, the team will move into extensive testing with industry partners and select customers before making a commercial version of their system available to the public.

Once released and well proven, FireRock will publish their specification for "smart data" and make key components, or examples thereof, publicly available via their website and through industry standards committees. Basically, FireRock will follow examples set before them by companies like PGP, Microsoft, and Sun Microsystems for standardizing technologies/interfaces available for free, while using proprietary enhancements for commercial growth.

Dream Team – Initially, the core team will be FireRock Technology but will grow to include security solution providers (e.g., Symantec, RSA, Thales, etc.), along with key system, software, and storage vendors (e.g., IBM, HP, Microsoft, VMWare, EMC, NetApp).

The final members of the team and perhaps the most important will be the NCO (NITRD), DARPA and even the U.S. Postal Service. Although doable with them, involvement by these organizations will help promote the standardization and ensuing adoption of the model ensuring a faster, more complete global adoption thus minimizing delays in deployment leaving the "game board" exposed. In fact, with the growth of social networking, the USPS could position itself as the ultimate, centralized certificate authority with a fee-based model for registration and help promote the U.S. as the center of the world's most cyber-safe community.

Who We Are: Secure Decisions is a leader in cyber security situational awareness, with expertise in computer science, computer security, mathematics and psychology. We take on hard problems such as building systems to understand mission impact of cyber attacks, and modeling complex dependencies. We offer to the Cyber Leap Year project our ability to make progress where others were stymied, manage teams of smart specialists to achieve a common goal, support testing of advanced security technology, and address the human aspect of the problem space.

Secure Decisions will collaborate with BBN, an organization that was around before the internet was universal and ubiquitous, capable of thinking of how things were then. They are cognizant of the multiple approaches and operating systems from the beginning of the internet to the present. Their world-renown expertise in information security, intelligent systems, networking, and adversary modeling will be needed in this very ambitious concept.

POC: Dr. Anita D'Amico, Director, Secure Decisions division of Applied Visions, Inc.

Game Changing Dimensions: We both morph the game board *and* change the rules. We envision an unprecedented level of protection by implementing what appears to attackers to be a hitherto unseen operating system. We envision such Disposable Software Frameworks implemented within participating sub-networks or in cloud configurations.

Concept – Disposable Software Framework: One would be hard-pressed to envision a bigger shift in the game than a Disposable Software Framework. We embrace the fundamental assump-tion that error-free, completely secure networks cannot be built because completely secure soft-ware is not feasible (some would say, impossible). The more time an adversary is given to devel-op an intrusion plan, lying in wait for a vulnerability to appear – whether due to human, software or physical security lapses – the greater the chance for success. Newly mounted attacks rarely target lesser-known operating systems. This fact led us to the insight that today's systems must operate in a less static, less universally inviting environment.

In our rework of the game board, the adversary will be confused, challenged and discouraged by what appears to be an unknown operating system. Attackers' occasional successes are much less likely to result in a single vulnerability that can be repurposed to attack thousands or millions of machines possessing the same flaw.

Vision: While the underlying concepts are not new,[1] the emergence of cloud computing has encouraged a fresh round of thinking about how to build systems. For example, the designers of SmartGRID envision "a fully decentralized grid scheduling framework supported by swarm in-telligence."[2] We extend this newly vital force in thinking about systems to network security. Key objectives in the Disposable Frameworks approach include:

- Robust perishable network topologies
  - Servers that "disappear" by adopting new "operating systems" in place
  - Hub-and-spoke server/workstation clusters that disassociate and regroup

[1] I. Foster, "Automatic Generation of Self-Scheduling Programs," IEEE Transactions on Parallel and Distributed Systems, vol. 2, no. 1, pp. 68-78, Jan., 1991

[2] Y. Huang, A. Brocco, P. Kuonen, M. Courant, and B. Hirsbrunner, "Smartgrid: A fully decentralized grid scheduling framework supported by swarm intelligence," Grid and Cooperative Computing, International Conference on, vol. 0, pp. 160-168, 2008. [Online]. Available: http://dx.doi.org/http://doi.ieeecomputersociety.org/10.1109/GCC.2008.24

    - Servers that are repurposed, or appear to be, seemingly instantaneously
    - A new morphing, perishable "layer" in the OSI stack
    - Remapping objects such as network ports to different, changing software abstractions
    - Destination-less network traffic, such as with P2P methods
- Multiple disposable "perishable" operating systems
    - Task Managers with differing methods for starting, monitoring and killing tasks
    - Integrated process migration, even to dissimilar architectures
    - File system schemes capable of remapping block-level data to alternative security access control lists (ACL) frameworks
    - Distributed re-authentication schemes; permissions migrate from one system to another, but require time-based re-authentication[3]
- Temporary, task-specific cloud-resident networks that dissociate
- Authentication and crypto schemes to support shortened OS lifecycle requirements
- Automatic software generation techniques to create new frameworks on the fly
- Self-migrating, decomposing, reconstituting systems
- Algorithms for OS layer substitution, migration, disappearance

We envision phased implementation of Disposable Frameworks using middleware proofs of concept. Earlier phases would reuse existing, well-understood frameworks.

Method: Based on prior work[4,5] we will identify the components required to build systematically disposable frameworks capable of communicating with existing platforms such as Windows servers, DOS, Linux or mobile networks. Such communicating could take the form of "hosting" (as with virtual machines) at the *operating system* level, or it could be at the *application server* level, such as hosting MySQL but with neither Linux nor Windows beneath it.

Automatic software generation plays an important role in the method we envision. It would not be possible to undergo the conventional build-test-deploy method that is widely used to deploy software. A Disposable Framework identifies its successor, learns how to migrate / restructure / transition / decompose itself. What was a single "server," could identify a successor operating system, decompose its current application responsibilities task by task, and migrate these to one or more servers. Disposable framework research is linked to work in digital preservation, which is concerned with migration of obsolete technological artifacts. Metadata preservation is one important element in this strategy. Related work in automatic service migration shows promise.[6]

Our method revives proven concepts in OS research that have been superseded by the dominance of Windows and Linux, and the complacency in the OS community to accept this dominance.

Dream Team: The dream team requires intellect in networking, distributed OS, automatic software, cyber security and project management of creative forces. Examples: BBN, IBM Research, Microsoft Labs, MIT CSAIL, Secure Decisions, SRI.

---

[3] E.g., as with a VPN key fob, authentication is based on a point-in-time synchronization.
[4] IEEE International Conference on Composition-Based Software Systems, http://www.iccbss.org/2008/.
[5] Milojičić, D. S., Douglis, F., Paindaveine, Y., Wheeler, R., and Zhou, S. 2000. Process migration. *ACM Comput. Surv.* 32, 3 (Sep. 2000), 241-299. DOI= http://doi.acm.org/10.1145/367701.367728
[6] Renata Bandelloni, Fabio Paterno, Zigor Salvador, "Dynamic Discovery and Monitoring in Migratory Interactive Services," percomw,pp.604-607, Fourth IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOMW'06), 2006

Who We Are: Secure Decisions is a small business best known for cyber security situational awareness. Our talents extend beyond that area: our staff spans the domains of computer science, information systems, computer security, mathematics, and psychology. We do not shy away from hard problems; we make progress even on technical challenges where others have been stymied, such as analyzing the mission impact of cyber attacks. In fact, that particular work inspired our interest in developing metrics of information value, and automatically determining the value of information content stored within information systems. Keeping the highest-value information furthest from the attackers is the underlying theme of our proposed concept.

POC: Dr. Anita D'Amico, Director, Secure Decisions division of Applied Visions Inc.

Game changing dimensions: Both morph the game-board, and change the rules.

Concept – "Spread Spectrum Data Store": The value of an information system to its owner or to an attacker derives largely from the *value of the information content* stored or transmitted by those systems. Yet we secure the information systems in a relatively uniform way, using perime-ter defense and access controls. Once attackers have gained access to a network they can access information content ranging from the mundane to mission-critical. In our zeal for standardized, centralized information storage to foster information sharing we have created static targets for attackers searching for specific content. The static nature of information "at rest" also fosters malware such as Banker-AJ,[1] which lies dormant for months, waiting for remote wakeup signals.

Attackers use this static landscape to reduce their effort to discover high-value information. High-value files and data tables remain in consistent, readily identified locations; authentication practices and challenge-response systems remain unchanged; methods for naming and locating high-value information are constant; typical database practices co-locate both high- and lesser-value information. Even information security practices that protect information based on value categories do not dynamically re-rate information based upon current mission and needs; information whose value has decayed often remains protected, using up valuable resources.

The proposed concept makes it more difficult for an attacker to obtain high-value information. We propose to automatically measure the value of information "on the fly," then break up the highest value information into pieces and distribute those pieces across distributed stores. Individually the pieces will reveal little; authorized users would have the "keys" to locate the constituents of the information and reassemble it on demand. To add yet another barrier to the attacker, we propose to periodically *move* the higher value information around the distributed stores, with the rate of movement correlated to the value of the information. Finally, we envision that some of the data stores would be semi-permanent, appearing and disappearing over time. Attackers should find it far less appealing to search for high-value information within a shifting, decentralized, disappearing data store that requires dynamic re-authentication.

Vision: To understand the difference between our proposed approach and the status quo, consider a hypothetical emergency responder registry that contains extensive information about responders, e.g. work schedule, physical abilities, specializations, immunizations, last-recorded GPS location, hire date, and contact information. If the information is stored in a traditional master table of a relational database, attackers with access to this table could use its contents to thwart a timely response to a terrorist attack.

[1] http://www.us-cert.gov/cas/bulletins/SB04-322.html#bankeraj

By contrast, in the "Spread Spectrum Data Stores" concept each piece of information about the emergency responders is automatically scored for its value. High-value information such as specializations (e.g. bomb disposal, language skills) would be decomposed and distributed, while lower-value information (e.g. last GPS location, which has a rapid rate of decay) would be left in a static location where resourceful attackers might find it but not use it effectively.

The same concept can be applied to a high-value *file*, such as a detailed response plan. Once identified as high-value, the contents of that file would be spread throughout the distributed data store. Only users with the proper keys to the distribution will be able to reconstitute the file.

We believe this concept is technically feasible. It exploits several recent trends: cheap, multi-core servers, distributed database technologies, cloud computing, and hybrid wired/wireless self-organizing networks.[2] Distributed data technologies such as Google's BigTable, which was motivated by scalability, have shown that data can be robustly distributed across thousands of machines; but they have not investigated the best ways of doing this for security purposes. Our proposed concept may also build upon Microsoft Lab's Leslie Graph[3] to enable dynamic network reconfiguration at the server level – literally "moving" the database to a different host.

Current research in the following areas would also need to be intensified and transitioned: migratory, decaying data stores; distributed re-authentication schemes; information valuation methods; mission-dependent valuation reassessment; data distribution; and data concealment.

Method: Our work on mission impact of cyber attacks highlights the fact that some attacks are far more important than others, because they affect critical operations or exploit sensitive information. Our own prior work,[4] which focused on methods for automatically identifying mission-critical information assets (e.g. devices, files, servers), and the work of the Air Force Institute of Technology (AFIT)[5] have inspired us to improve methods of measuring *information value*. As we started researching this area, we realized the potential approaches and benefits of being able to rapidly rate the value of information, including its use to create barriers to attacks.

We conducted a quick assessment of the state of the art and developed the proposed concept, in which value ratings can be assigned at the time data is committed to a distributed data store, and dynamically reassessed – and *relocated* – based on policy or other network intelligence. We also extended the concept so that data relocation schemes could be coupled with time-based authentication. Just as a "one time password" key fob is used to re-authenticate a connection, using a program wishing to obtain data that has been moved must first re-authenticate itself for that specific data. We also brainstormed how to incorporate information hiding and decoy services into this new concept. The ready availability of multi-core servers, deployed en masse, creates an opportunity for redundancy and incorporation of systematic decoy servers.

Dream Team: The dream team incorporates intellect in the areas described under our Vision, e.g. distributed DBs, distributed authentication, data concealment, information valuation, statistics, applied economics. Candidates are: AFIT, Google, Johns Hopkins Applied Physics Lab, Oracle, Microsoft Research, Red Hat, Secure Decisions, SRI, UC Berkeley, and Versant.

---

[2] http://www.nist.gov/public_affairs/factsheet/improving_first_repsonder_communication.htm [*sic*]
[3] http://research.microsoft.com/apps/pubs/default.aspx?id=64515
[4] "Visual Representation of Cyber Defense Situational Awareness Final Report: Automated Analysis of Mission-Relevant Content," Applied Visions, Inc., under Phase 2 SBIR Contract No. DAAH01-01-C-R044, December 2003.
[5] Hellesen, D. (2008) An Analysis of Information Asset Valuation (IAV) Quantification Methodology for Application with Cyber Information Mission Impact Assessment (CIMIA), Master's thesis, AFIT

Resilient Smart Grids for Power

Who you are: We are researchers at the Information Trust Institute, University of Illinois with expertise power systems, reliability and security. Himanshu Khurana (Principal Research Scientist, Distributed System Security, Bill Sanders (Professor of Electrical and Computer Engineering, Reliability and Security) and Pete Sauer (Professor of Electrical and Computer Engineering, Power Systems).

Game-changing Dimension: Morph the Gameboard

Concept

The North American power grid is an extensive cyber-physical infrastructure supporting life and essential services. Owing to deep penetration of computer monitoring and control technologies that were never designed with security in mind and the complex nature of this system of systems, this infrastructure has become vulnerable to cyber attacks as evidenced by recent events. These risks are likely to increase with deployment of a "smart" grid that employs advanced communication and computer systems. This white paper proposes the design and development of resilient smart grid tools and technologies that are built with security and availability in mind from the ground up. By doing so, a fundamentally stronger defense can be realized for this essential critical infrastructure. Furthermore, solutions are likely to have applicability to other critical infrastructures including those in energy and process control.

Vision

The vision is ensuring the resilience – that is, availability and security – of the smart power grid. To do so, research will focus securing the low-level devices, communications, and data systems that make up the power grid, to ensure resilient operation during normal conditions, cyber-attacks, and/or power emergencies. This will involve research in ways to (1) combine hardware, firmware, and software techniques to provide low-overhead, robust protection against both accidental and malicious faults, (2) ensure that both data protocols and communication systems that carry the data protocols are secure and available, and (3) model, simulate, emulate, and experiment with the various subsystems in the power grid to allow for adequate quantitative and qualitative validation of our research efforts.

At the device level research is needed to explore ways to combine hardware, firmware, and software techniques to provide low-overhead, robust protection against both accidental (non-malicious) and malicious faults, and hence to enhance the trustworthiness of the power grid. The major research themes include (1) the use of various types of hardware trust enforcement to provide adequate trust in smart grids, as well as (2) the demonstration of some of the developed/adapted techniques on large-scale applications in a realistic testbed setting.

At the communication and protocol level research is needed to explore ways to ensure that both data protocols and the communication systems that carry the data protocols are secure and trustworthy. Data protocols and communication systems include those that gather information from sensors, process it at intermediate levels, and take it all the way to authoritative centers and coordinators to ensure reliable power grid operations. Security and trust technologies include cryptographic techniques to protect data along with their associated key management infrastructures, adherence to real-time and quality-of-service requirements, and policy negotiation and management for data sharing and control.

The power grid is a complex system of systems that includes interconnected energy conversion devices, cyber-infrastructures, communication systems, and markets. Understanding

this complex system is crucial to developing resilient grids; furthermore, the ability to experiment with a complete system is crucial for validating the results of the research efforts. In this area, research is needed to explore means to model, simulate, emulate, and experiment with the various subsystems in the power grid to allow for adequate quantitative and qualitative validation of all research efforts.

Research efforts will need to focus on a range of representative smart grid applications that are stated to be the driving force behind the smart modernization of the grid. These include Plugin Hybrid Electric Vehicles (PHEV), Advanced Metering Infrastructure (AMI) and Phasors for Wide Area Measurement Systems (WAMS). By focusing on representative applications the developed tools and technologies will have a clear path for validation and product development.

Development of successful technologies will require an effective government-industry-academic partnership that is informed by regulatory and national security policies. The research effort will therefore involve workshops and other events [CSPCS] that involve these various stakeholders and promote discussion and action.

The Roadmap to Secure Control Systems in the Energy Sector [Eis06] makes it clear that much has to be done to provide a secure environment for energy control. Near-term actions to develop and integrate protective measures are limited to distributing consistent training materials on cyber and physical security for control systems. Long-term actions include development of a next-generation control system that can survive malicious attacks without loss of critical functions. The research proposed in this work will focus on such next-generation control systems where COTS devices, operating systems and software will be widely deployed and used. Furthermore, the proposed work on monitored, timely and secure communications will provide a complete systems context for the emerging smart grid.

Methods

The authors lead the Trustworthy Cyber Infrastructure for Power Grid (TCIP) [TCIP] research center funded by the National Science Foundation, Department of Energy, and the Department of Homeland Security focusing on security and trust for the Power Grid. 50 researchers from the University of Illinois, Cornell University, Dartmouth College and Washington State University along with 35 industry partners from the electricity sectors are working to achieve the center's goals. The topics proposed in this white paper are a result of many meetings, workshops and summer schools involving center participants including extensive interactions with the industry partners.

Dream Team

Researchers in the TCIP center (tcip.iti.uiuc.edu), Government agencies and laboratories (such as the Department of Energy National Laboratories), as well as key industry organizations (vendors, utilities and reliability operators) would be an ideal team to work together and realize the outlined vision.

References

[CSPSS] Cyber Security for Process Control Systems Summer School, June 16- 20, 2008, Wisconsin. http://www.iti.uiuc.edu/events/SummerSchool2008.html.

[Eis06] J. Eisenhauer, P. Donnelly, M. Ellis, and M. O'Brien, Eds., "Roadmap to Secure Control Systems in the Energy Sector", January 2006. http://www.controlsystemsroadmap.net.

[TCIP] Trustworthy Cyber Infrastructure for Power Grid Center. http://tcip.iti.uiuc.edu.

Who we are: NIATEC, is a consortium of academic, industry, and government organizations to improve the literacy, awareness, training and education standards in Information Assurance. As the federally designated cornerstone for essential education and training components of a strong Information Assurance initiative, the mission is to establish an effective Information Assurance infrastructure for academic, industry and government organizations. NIATEC is associated with Idaho State University Center of Academic Excellence. The Centers of Academic Excellence and NIATEC are components of a plan to establish a federal cyber-corps to defend against cyber-based disruption and attacks.

Game-Changing Dimension: Change the game board

Concept: The Internet, as designed, does not lend itself to security. Building security into this insecure architecture is difficult. We propose that a second "Internet" be created in parallel to the original that will be used for secure applications.

Vision: The vision is that the current TCP/IP protocol stack and infrastructure that is currently employed continue being used for day-to-day operations on the Internet. However, the option of using a second network that has had security designed into it from its inception would be available for purposes requiring higher levels of confidentiality and integrity. This would be transparent for end users, with the burden of configuration being placed on content providers. Capabilities in computing logic and power should be able to accommodate this new development without difficulty. Research and testing would need to be done in order to develop mathematically sound, demonstrably secure infrastructure and protocols. Infrastructure elements may require redesign or the second network may have to run over hardware that is completely independent of the current Internet. Methods to confirm compliance would also need to be developed.

Method: This idea was one of three generated via computer aided nominal group techniques. We encouraged idea submission from NIATEC alumni, as well as from current members.

Dream Team: Private industry, including computer scientists, mathematicians, security experts, software providers, hardware developers. Government's role in this process would likely be limited, considering the nature of the project, to that of champion – encouraging the development through funding and legislative support – and eventual user of the system.

Who I am

This submission is written by Jonathan King, who is an information assurance research analyst with the National Information Assurance Training and Education Center (NIATEC).

Game-changing Concept

People are the weakest link in the realm of computer security and information assurance. Research and practice demonstrates that technology alone does not suffice as a method for protecting information assets, since people often neglect to apply the proper security practices consistently or effectively. In order to address this issue it is vital to change the culture and mindset of how people think about information by integrating computer and information security into the education curriculum. A federal mandate requiring that important security concepts be taught in the education system at an early age is imperative. A solid and early educational foundation relating to computer and information security concepts will help prevent and deter the growing number of computer and information security incidents that occur everyday.

Vision

The Department of Education would play a fundamental role in establishing standardized policies and procedures for disseminating the curriculum throughout the nation. With a nation-wide curriculum in place, fundamental computer and information security concepts will be integrated into the culture at a young age creating widespread awareness. The goal is for computer and information security practices to become second-nature and ingrained. This concept is similar to the way health education has become an integral part of today's education curriculum.

Method

Research and idea generation was done separately and anonymously. Afterwards, a list was compiled and the top ten ideas were voted upon, ranked, and chosen by our group of analysts through the use of specialized software that kept the originator of the ideas and voters anonymous to remain unbiased. The top three choices were selected, and this concept of mine was picked as the first choice.

Dream Team

The team should be composed of the Department of Education, cyber security professionals and educators, NIST, the National Science Foundation, and cyber security related government agencies.

Response to RFI for National Cyber Leap Year

Who: Nicira Networks, a startup developing security technology for the DoD.

Dimension: Morph the gameboard, by making it easier to secure enterprise networks.

Vision: Create a secure "Cloud-Style" approach for government enterprise networks that can provide security from the edge, requiring only changes to virtualized hosts and/or edge switches.

Problem statement: Today it is hard to run an enterprise network flexibly and securely,[1] largely because current network management mechanisms (*e.g.*, ACLs and VLANs), require significant manual configuration to adapt to changing network conditions and security policies.

Approach: Cloud computing provides a lesson in how we can move beyond this rigid and insecure *status quo*. By "cloud computing" we mean running computation at a remote site operated by another entity; effective cloud computing requires that the computation be fully portable and the computing requirements be specified in terms of high-level abstractions (*e.g.*, number of VMs) and not in terms of low-level host characteristics. The cloud approach leads to a clean separation of concerns: the cloud operator maintains the cloud resources (*i.e.*, the hosts and network), while the cloud customer is responsible for their particular application.

We contend that future enterprise networks should embody these two characteristics: portability and well-separated concerns.

- *Portability*: Operators should control the network by declaring policies over high-level abstractions (*e.g.*, users, hosts, groups): the network (or its management system) should translate these high-level policies into the required low-level switch configurations. Whenever the network topology changes, or new users arrive, or VMs migrate, the network should automatically adjust its low-level configuration to ensure that the high-level policies are enforced. This makes network management portable in that these policies can be applied in any network, no matter the topology, brand of switches, or user-population.

- *Separation of concerns*: Network management portability translates directly into a clean separation of concerns. There are three relevant components, each overseen by a separate group of administrators:
  - Network infrastructure: this includes the hardware (switches) and the basic management software (see below). Local network operators are responsible for ensuring that the infrastructure is functioning correctly.
  - Security and management policies: these policies are dictated by security officers and IT managers, not by local network operators. These policies are portable, in that they are defined in terms of high-level abstractions, not network-specific characteristics.
  - Directories: policies can refer to user attributes and groups. This information is stored in a set of directories, and can be updated only by those with the proper authority.

  Each set of administrators can manage their associated piece of the infrastructure independently; for instance, upgrading someone's clearance in a directory would automatically lead

[1] One can run a network securely by locking it down, but then it becomes hard to get work done; similarly, one can run a flexible network, but security holes are apt to appear. The true challenge is providing both security and flexibility.

to a change that person's network access (if the policy so dictated); no change in network infrastructure or management policy is required.

Design: So far we have sketched a vision of a future network. We now talk about a particular design that realizes that vision. First, all switches should support the OpenFlow abstraction (http://www.openflowswitch.org/), which provides a uniform interface to all switches, regardless of vendor. Second, the network should have a centralized management system that can translate high-level policies into low-level configurations. We have proposed a *network operating system* called NOX (http://noxrepo.org/doc/nox-ccr-final.pdf) that has this capability. NOX also provides comprehensive visibility into the network, making it easier to operate and troubleshoot.

Deployment: At first glance, it might appear that this approach requires a complete "fork-lift" replacement of an enterprise's networking infrastructure, which is clearly not a viable path to progress. To the contrary, this approach can be incrementally deployed in the following manner.

1. Deploy a network management system (like NOX) that can control OpenFlow-enabled switches. This requires only a few servers connected to the network running the desired software.

2. Equip hosts with a virtualization hypervisor that supports OpenFlow.[2] This enables access control over all VM communications (even between VMs colocated on the same host). Note that at this point in the deployment, we have a "secure cloud", where all inter-VM communication is controlled according to high-level policy (in fact, the same policy that would control their communications in a non-virtualized setting). This secure cloud requires only the use of an OpenFlow-enabled hypervisor and a NOX-like central management system; the networking infrastructure itself can remain unchanged (and since it only needs to deliver packets, not enforce security policies, it becomes easier to manage).

3. For hosts that aren't running on an OpenFlow-enabled hypervisor, or when the level of trust in the hypervisor isn't sufficient, the access switch should be replaced with an OpenFlow-based access switch. These access switches are now extremely inexpensive (*e.g.*, $2500 for $48 \times$ 1GE), so the deployment barrier should be low.

4. Eventually, as more comprehensive control is desired, the rest of the switching infrastructure can migrate over to OpenFlow-enabled switches.

To summarize, this approach cleanly separates the issues of infrastructure management, security policies, and directories, allowing each group of administrators to pursue their mission independently. By imposing network-level access controls via the hypervisor and/or edge switches, control over access (and other network security measures) can be enforced, while leaving the rest of the networking infrastructure intact. This approach provides more comprehensive security and control than network security appliances, yet is far easier to deploy than a complete fork-lift upgrade of the switching infrastructure. As such, we recommend it as the universal "first-step" in improving enterprise network security and providing a clean separation of administrative concerns

Method: We have prototyped this system, and have been using it operationally for months.

Dream Team: Cloud and network operators, security specialists.

[2]We are about to release an OpenFlow-enabled software switch for Xen, and we hope to have others shortly.

Response to RFI for National Cyber Leap Year

Who: Nicira Networks, a startup developing security technology for the DoD.

Dimension: Morph the gameboard, by making is much harder for malware to exfiltrate data.

Vision: Change the paradigm of exfiltration-prevention from "is this a normal transfer?" to "did a human initiate (or approve) this transfer?".

Problem Statement: There is no question that one of the most pressing problems facing the cybersecurity community is the malware-initiated exfiltration of data from compromised hosts in enterprise networks. Historically, exfiltration has been detected by finding signs that an outgoing data transfer is "abnormal" (*e.g.*, using an unusual protocol). Unfortunately, today's malware uses the same set of widely adopted communication mechanisms and protocols that are used by standards applications, so this approach of looking for low-level anomalies is no longer viable. Moreover, with widespread adoption of encryption, there is no way to detect the presence of sensitive data in outgoing transmissions. Thus, we need a new paradigm for dealing with exfiltration.

A different approach: The first step in finding a new paradigm is to recognize that we are only dealing with malware-initiated exfiltration (which is a far more widespread problem than insider-based exfiltration). The main difference between malware-initiated exfiltration and normal user-communication lies not in the low-level mechanisms but in the simple fact that the human user isn't involved in the exfiltration process. Thus, to detect and/or prevent malware-initiated exfiltration, we can pose a simple question: *was this transfer initiated (or approved) by a human user?*

Assumptions: To build a system embodying this new paradigm, we make three assumptions:[1]

1. Users are honest and inappropriate exfiltration is being done by malware, not humans.
2. The adversary cannot snoop on traffic intended for other destinations so, for data to be exfiltrated, it must reach a computer controlled by the adversary.
3. Users do not intentionally contact sites that are under adversary control.

When these assumptions hold, data exfiltration attempts can be prevented by only allowing user-initiated (or at least user-approved) transfers to cross the network perimeter. Our approach is superficially similar to that taken in host-based systems like *little snitch* and *blackice*, and in proxy-based systems like *whitetrash*. However, our approach is distinguished by how it deals with the three-fold challenge necessary to make this approach viable; creating a system that is *understandable*, *tolerable*, and *general*.

- *Understandable:* When asking the user whether they approve the transfer, it must be done in a way that's meaningful to the user. A simple request to cnn.com elicits a flurry of auxiliary traffic to destinations such as Akamai proxies, hosted advertisements, and IP analytics services; most users don't realize that their accessing websites like CNN leads to these other transfers. To avoid this problem, we propose to reduce the stream of low-level network

[1] We realize that these assumptions don't always hold, and that other techniques will be needed to deal with malicious insiders or adversaries that can snoop on outgoing traffic.

"transfers" to a set of high-level application "transactions" (such as sending an e-mail or visiting a web-page). These transactions should be at a level understandable by users as resulting from their actions.

- *Tolerable:* The number of times the user is contacted for approval must be kept to a tolerable level. To do so, the system must allow the operator to employ a flexible set of policies about automatically approving visits to sites that have been previously visited by that user, or other trusted users.
- *General:* Malware can use any allowable protocol to transfer data offsite. Thus, in order to be sufficiently general, the system must support all commonly used protocols, such as HTTP/HTTPS, SMTP/SMTP+SSL, IMAP, POP3, FTP, SFTP, AIM, Jabber, SSH. Moreover, the platform should be extensible, allowing for the simple addition of new protocols.

Technical Details: Our approach requires one new physical component, a proxy, which intercepts and analyzes traffic across the network perimeter. By *proxy* we are referring to a network appliance that terminates TCP and thus has full access to application data. This proxy will *reduce* the many ongoing transfers into a smaller set of user-understandable transactions. It will also elicit *feedback* from the user as to whether these transfers are approved (and it must do so in a secure manner). The proxy will also apply *policies* about which transactions require user-approval.

Best Practices: We should make clear that the effectiveness of this approach relies on enterprises adopting a set of best practices, and that if these are ignored, the system could be circumvented by malware. Among the requirements are:

- *Limit or ban the use of peer-to-peer.* By construction, peer-to-peer protocols are excessively permissive and often require hosts to operate as content servers as well as clients. We don't believe it reasonable to expect a user to authorize peer-to-peer connectivity in a safe manner.
- *Allow interception for all encrypted sessions.* The ability to intercept and decrypt standard encryption protocols (*e.g.*, SSH or SSL) is standard in many commercial appliances. If interception is allowed, the approach described in this proposal can be applied to all traffic leaving the network. Further, it allows the interposition of application-level proxies which can ensure that a long-standing encrypted session is not being used to exfiltrate data.
- *Limit webmail.* Webmail users must approve all sizable uploads and use a local smtp server for outgoing mail.
- *Limit access to public posting sites.* Access to sites where users post content (such as social networking and auction sites) should be curtailed (or uploads prohibited).

Progress: We are building a prototype, and are looking for sites where it can be tested.

Method: We are basing this approach on the assumptions listed above, and plan to evaluate this approach through analysis of traces from sites and live user testing.

Dream Team: We would particularly welcome participants with expertise in malware exfitration (Chris Eagle is a member of our team, so we already have some experience here), and site management (particularly those with access to traces and/or a willingness to do trial evaluations).

Response to RFI for National Cyber Leap Year

Who: Nicira Networks, a startup developing security technology for the DoD.

Dimension: Morph the gameboard, by defining a new system architecture for security.

Vision: Using a secure hypervisor as the *thin-waist* of the security architecture.

Problem statement: Our security efforts have produced a patchwork-quilt of useful security mechanisms at all layers of the software stack, but we still have no overall security framework that lets us think clearly about where in the architecture various aspects of security functionality should be implemented. In addition, security vendors are constantly playing catchup with rapidly evolving application and operating system software and, as a result, new applications are often disabled or blocked by outdated or overly aggressive security applications.

As an instructive example of a system architecture that overcame these hurdles, we turn to the Internet. There, the adoption of a "thin-waist" (*i.e.*, the IP protocol) provided a framework that dictated what functionality resided in the Internet (simple end-to-end connectivity) and what belonged to the end hosts (everything else). This division of labor enabled rapid and radical innovations to occur both above and below the IP layer.

In this proposal we suggest moving away from the current security practice of reactive, ad-hoc patches and towards a similar "thin-waist" architecture that provides a framework for implementing security mechanisms while allowing innovation at all layers of the software stack.

Requirements: A "thin-waist" for security should have the following properties:

1. *Minimum interference:* The solution must be minimally invasive, integrating simply with legacy components (*e.g.*, current networks, hardware, applications and operating systems) and, at the same time, allowing innovation in all portions of the system architecture. Proposals to build security architectures around newly designed components such as secure computing hardware, secure operating systems, or secure programming libraries, fail this test, as they would render large portions of the installed base unusable.
2. *Effective with untrusted code bases:* Currently there is substantial security "fate sharing" between software components. If one component is compromised, so generally is the security of the full system. Because it is not possible to audit all software on a system, security measures must maintain effectiveness even when deployed with untrusted components.
3. *Useful security primitives:* Security has many facets, thus the design must be sufficiently general to enable a broad array of security functions. Based on recent trends, we believe this layer must provide the following "services" from which secure systems can be built.
    - Support for *trusted computing*. The system should allow software operating systems and applications to attest that they are running known and unmodified software. Trusted computing is the basis of many recent security initiatives.
    - Ability to *control information flow*. Information flow control is necessary to stop inadvertent and malicious transference of sensitive data off of computer systems. Effective solutions to this pressing issue appear untenable in the traditional software architecture.
    - Support for *network filtering and isolation*. Perhaps the most pervasive and effective security mechanism against remote attacks has been the ability to shield end-hosts.

Whither the Waist?: These requirements call for a security layer that provides significant support for security mechanisms while being "thin" enough to interoperate with the vast installed base of legacy software and support future innovations in hardware and software. Where can we find such a layer? *The central technical claim of this proposal is that advancements in virtualization technology enable the creation of a hypervisor that satisfies these requirements.*

Hypervisors were designed to be transparent to the rest of the system, while having complete control and visibility of the running software on the computer. Moreover, hypervisors can be implemented with a code base vastly smaller than current operating systems, so that one can foresee building provably secure hypervisors using modern programming techniques.

More specifically, hypervisors have been shown to support the three properties requisite for a thin security waist. First, hypervisors are designed to support unmodified (or slightly modified) operating systems and they require no change to existing applications, hardware, and networks. Second, hypervisors have been shown to provide strong security guarantees even when deployed with untrusted software [3].

Finally, and most fundamentally, hypervisors have sufficient control to be broadly useful to security mechanisms in all system levels. Hypervisors have been used for bootstrapping trusted computing [1]. A leap-year submission by Katti *et al.* describes a hypervisor-based approach to tracking the flow of sensitive information. Hypervisors can also be used to control network connectivity, as demonstrated in [2] and in other approaches using OpenFlow-based virtual switches in the hypervisor (see the leap-year submission by Nicira Networks on securing the cloud).

To summarize, we believe that a hypervisor-based approach to security can change the way system security is implemented in both the short and long terms. In the short term, this approach can, among other things, control network access in legacy networks, control information flow in legacy operating systems, and provide support for trusted computing. These are security properties that are desperately needed, and we believe that they can be made widely available in the near future through this hypervisor approach. Moreover, these and similar hypervisor-based advances will remain relevant in the long-term, as the hypervisor thin-waist allows for continued innovation in all areas of the system architecture. Thus, not only does this approach provide more comprehensive security coverage than is available today, it allows for that same coverage to extend to systems built tomorrow.

Method: Various pieces of this approach have been prototyped; what remains is unify these efforts and clearly articulate the role hypervisors should play in an overall security architecture.

Dream Team: Hypervisor vendors (*e.g.*, VMWare, Citrix, Microsoft), government security experts (*e.g.*, NSA), and various academic and industrial researchers.

References

[1] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and B. Boneh. Terra: a virtual machine-based platform for trusted computing. In SOSP 2003.

[2] N. Zeldovich, S. Boyd-Wickizer, and D. Mazieres. Securing Distributed Systems with Information Flow Control. In NSDI 2008.

[3] X. Chen, T. Garfinkel, E. C. Lewis, P. Subrahmanyam, C. Waldspurger, D. Boneh, J. Dwoskin, D. Ports. Overshadow: A Virtualization-Based Approach to Retrofitting Protection in Commodity Operating Systems. In ASPLOS 2008.

Changing the Board with the Policy Machine

Contact: David Ferraiolo, NIST

Who am I: David Ferraiolo; Acting Manager, Systems & Networking Security Group, NIST; 18 years of research into access control mechanisms, models, standards, and implementations; 25 peer reviewed papers and journal articles; best selling book on RBAC; DoC Gold metal and Federal Laboratory Consortium award for Technology Transfer.

Change the board: The ability to control access to sensitive data in accordance with policy is perhaps the most fundamental security requirement. Towards better meeting end-user needs, Role-based Access Control (RBAC) was proposed in the early 90s as an alternate to the existing Discretionary and Mandatory Access Control (DAC & MAC) standards of the day. Today, RBAC is the dominant model for managing user permissions with most network and operating system, and enterprise security management vendors implementing some form of RBAC. While RBAC has advanced privilege management, the ability to specify and faithfully enforce enterprise policy remains in a dismal state of affairs. This is because access controls as implemented are not comprehensive, typically do not offer control at the process/inter-process level, and/or lack expressive power. Today, for instance, a user with read access to data can typically make a copy of that data and paste its contents into an email message and send it to anyone else in the world, regardless of enterprise policy. RBAC and DAC do not prevent the leakage of data to unauthorized principals through malware, malicious or complacent user actions, or administrator error. Today we must all but assume that programs contain vulnerabilities that can be exploited or an attacker may install malicious programs that can easily thwart policy. Although MAC can control the leakage of data, its control is limited to multi-level security, and it is heavy handed (in a session a user is restricted in performing actions for which the user is ultimately authorized). Using DAC, MAC, and RBAC together it is impossible to determine who has access to what data.

Concept: NIST, with the early support of DHS, has designed and developed a reference implementation for a standards-driven, enterprise-wide security policy enforcement framework, referred to as the Policy Machine (PM). The PM is a logical "machine" comprised of a fixed set of data relations for expressing any access control policy, and a fixed set of functions for making access control decisions and enforcing policy based on that expression. In its simplest and most general form, the PM standard architecture is comprised of one or more *PM clients*, one or more *PM servers*, a *PM database*, and one or more resource servers. Its objective is to provide a unifying framework to support any attribute-based policy or policy combination through a single mechanism that requires changes only in its data configuration.

PM configurations specify capabilities that users and processes "can" perform (under permission relations), and "can not", and "can only" perform (under prohibition relations). In addition to these relations, the PM defines obligations that dynamically specify capabilities that users and processes "can now", "can no longer", or "now can only" perform. Permission, prohibition and obligation relations provide the basic ingredients for expressing a wide range of attribute-based policies. Policies that are enforced through applications need to be taken into consideration. Any application that affords services through access control such as email, or workflow can provide those services through PM configuration and enforcement. This suggests that the PM and access control could be more fundamental to computing than anyone expects. Furthermore, by defining objects as logical names that map to physical content, we are able to comprehensively and uniformly apply policy to any object regardless of its type (e.g., files,

messages, records, fields, work items, attachments, clipboard, ports…) or the physical location in which its content is stored. PM protection can be applied to inter-process communication, and process data exchange with the external world. If a process (e.g., a worm) reads from the network, the process can be prevented from writing back to the network.

Vision: User and vendor needs can be accommodated through the adoption of the PM. Rather than reacting to the policy du jour or the individual needs of their customers, the vendors only need to implement the appropriate standardized PM components once while guaranteeing interoperability with other vendor products. Developers of applications that provide services through access control can provide those services through adherence with a standardized API exposed by the PM client and a prescribed PM data configuration. Application vendors that provide services independent of access control simply need to abide by the APIs. The big winner would be the customer that gets to implement their individual and precise policy requirements through acquiring PM components and the translation of those requirements into a PM data configuration. To facilitate this translation, standard configurations for a variety of policies can be made available as a library of parameterized policy configurations. This reduces the burden on administrators in specifying and configuring policies.

To instill confidence in this vision consider our reference implementation. We can now demonstrate the enforcement of a diverse set of policies to include instances, combinations, and hybrids of DAC, MAC, RBAC, Chinese wall, ORCON, object-based SoD constraints, etc. Not only can we demonstrate the enforcement of these policies on files but we can demonstrate them within and across a rich user environment that includes the Open Office suite of applications, email, workflow management, and records and forms management. With the PM's open environment, we see opportunities for the development of innovative applications that provide services through PM access control. The use of the PM dramatically simplifies the application logic, increases operational assurance, and facilitates the secure sharing and interchange of data with other compliant applications. Under any policy (to include MAC) the PM does not limit the actions for which the user is ultimately authorized (e.g., TS users can r/w TS, S, and U data within a session, without compromising policy). A critical property of access control is that it is not by-passable. Our reference implementation currently performs enforcement through a kernel simulator. An important next step, in our PM research, is to implement the PM in a high assurance environment. Alternatives include enforcement within a real OS kernel, or an entire PM implementation within a virtualization infrastructure environment. We are currently engaged with INCITS in promoting the development of a suite of PM standards.

Method: The PM was developed through examining fundamental characteristics of existing models and requirements of numerous policies. In 2005 we developed our first PM specification and an accompanying reference implementation, under which we were able to observe and analyze the enforcement of a multitude of policies. In an iterative fashion, we were able to refine our specification and implementation, while supporting increasingly sophisticated policies and applications through fewer and simplified relations and functions, while at the same time making access control increasingly transparent to the user.

Dream Team: David Ferraiolo, Serban Gavrila (current PM implementer and architect), Lee Badger (assurance advisor) of NIST; Prof. Arif Gafoor and students (currently creating novel applications) of Purdue University; an OS or VM vendor; and a Stds Development Organization.

Massively Parallel Combinatorial Testing – contact: Rick Kuhn, NIST

Who we are – Researchers in the Math and Computer Security divisions of NIST (Raghu Kacker, Rick Kuhn), Computer Science Dept., Univ. of Texas, Arlington (Yu (Jeff) Lei), Pi Shurlok, a UK control instruments company (Mike Ellims).
Game-changing dimension – Rather than changing Board, Rules, or Stakes, change player.

Concept – Attackers often exploit obscure faults in software, rare combinations of inputs that developers and testers never thought to try. One approach that attempts to find these rare combinations is fuzz testing, but it typically finds only vulnerabilities that result in a system crash and, because fuzz testing is random, some relevant combinations may be missed.

This project will use new developments in algorithms and inexpensive cluster processors to test all *t*-way combinations of inputs (to a pre-defined level of *t*). Much like computer chess systems, a combination of new algorithms and raw computing power (e.g., executing10^9 combinatorial tests) can be used to out-play opponents by finding obscure faults before they do.

Vision – pseudo-exhaustive testing for software. Empirical research suggests that software faults involve relatively few variables interacting. (In our research, so far we have not encountered a failure involving more than 6-way interactions.) These results have important implications for testing. If all faults in a system can be triggered by a combination of *t* or fewer parameters, then testing all *t*-way combinations can provide high confidence.
***Generating high-strength covering arrays*:** Until recently, pairwise testing (*t*=2) was the only form of combinatorial testing used in practice, because good algorithms to generate higher-strength combinations were not available. Project members have developed algorithms that make it possible to generate high-strength (up to *t*=6) covering arrays, i.e., arrays that specify a set of tests covering all possible *t*-way combinations. Lei's IPOG algorithm has been implemented in a covering array generation tool that produces compact arrays in times that are in some cases orders of magnitude smaller than commercial tools, and we are now distributing this tool.
***Solving the oracle problem*:** For any covering array algorithm, the number of tests produced is proportional to $v^t \log n$, where v = number of values per variable, t = interaction strength, and n = number of variables. As with any test methodology, v must be kept small using techniques such as equivalence class and boundary value analysis. Taking advantage of combinatorial testing requires a large number of tests: for real-world software the number may exceed 10^7. With such a large number of tests, it is impractical for human developers to analyze each test case and determine the expected results. Thus even with efficient algorithms to produce covering arrays, the oracle problem remains, but advances in model checking and other areas make it possible to solve the oracle problem, using methods described below, for large-scale combinatorial testing, and we have demonstrated the integration of these methods with combinatorial testing in proof-of-concept projects (see: http://csrc.nist.gov/acts).

Embedded assertions within code ensure proper relationships between data, such as preconditions, postconditions, or input/output value checks. Sufficiently strong assertions can be used in proofs, but when coupled with *t*-way testing can provide strong assurance by showing that assertions pass for combinations up to the value of *t*. The embedded assertions serve as an executable form of the specification, thus providing a (partial) oracle for the testing phase. Executing all *t*-way tests on code with embedded assertions demonstrates that no assertions are violated for all *t*-way combinations of inputs, providing strong assurance.

Another approach, *model-checker based test generation* uses a mathematical model of the system under test (SUT) and a model checker to generate expected results for each input. A model checker is particularly valuable because it not only reports that a claim is false, but also provides a counterexample that includes a trace of parameter input values and states that will prove it is false. In effect, this is a complete test case—that is, a set of parameter values and the expected result. It's then simple to map these values into complete test cases in the syntax for the SUT.

***Proof of concept*:** To date, NIST has produced proof-of-concept demonstrations of these methods in access control, simulation, and avionics applications, and developed algorithms to produce combinatorial tests for much larger problems (hundreds of variables) than previously available methods. Because the tests are based on a formal model of the system (assertions or model checking) assurance is much stronger than simple crash testing. In addition because tests can be run independently of each other, the process is trivially parallelizable, and we run 100 tests simultaneously, with scaling to 1,000 or more parallel runs dependent only on resources.

***Scaling up*:** Cluster systems of 1,000 processors are within the reach of most large organizations, and will become larger and less expensive in the near future. With 10,000 processors, test suites of 10^9 tests are entirely practical for many applications (e.g., at 100 seconds per test, 10^9 tests can be run in less than 4 months). Computer science departments have been teaching formal methods for many years, so a core body of knowledge already exists which many practitioners understand (but may not use now). Microsoft Research has invested extensively in methods for integrating strong assertions, based on proof techniques, into code during development, and these efforts are already paying off for them. Model checkers capable of processing large, real-world specifications are available. Coupled with advanced covering array algorithms developed by our project and others, the pieces are already in place – and all are open source or freely available – to bring this vision to reality. Commercial firms can package and streamline the application of these methods with better test environments, however research is needed to understand the types of faults and interactions that occur in different application domains, and investigate the effectiveness of alternative approaches to combinatorial testing (i.e., assertion or model-based, as outlined above, integrations of both approaches, prioritization of tests).

***Measures of success*:** Testing is only one component of software assurance, but it is essential for all software. Methods described here are applicable to a wide range of applications. The current state of practice for ultra-high assurance software testing, as required for avionics software, is modified condition decision coverage (MCDC) testing. We expect to be able to demonstrate equivalent or better fault detection at significantly reduced cost as compared with MCDC testing as practiced today, and demonstrate similar results for more conventional software testing.

Method – The concept developed from our research into the number of variables interacting in faults in real-world software. Across a variety of application domains, a maximum of 4 to 6 variables were involved in failures, which suggests that testing combinations up to 6-way (or slightly higher) could be effective for assurance. While we do not claim that all failures are attributable to the interaction of 6 or fewer variables, the assumption derived from empirical research is that the interaction strength that must be tested is far below exhaustive testing.

Dream team – NIST and UTA staff, who are experts in combinatorial testing, Pi engineers who are using combinatorial testing in real world software development, and university researchers experienced in conducting and monitoring experiments in software testing. We are also discussing options for cooperative work with Accenture, a global management consulting, technology services and outsourcing company, whose customers include major industrial software developers in aerospace, banking, health care, and many other critical infrastructure industries.

Online Desktop for SOHO

NITRD - National Cyber Leap Year Proposal

Who you are

Christophe Veltsos (PhD, CISSP, CISA, CIPP, GCFA) is a faculty member at Minnesota State University, Mankato, President of the Mankato ISSA Chapter, and President of a security consulting company called Prudent Security, LLC. He is a member of ACM, HTCIA, IAPP, ISC2, ISACA, ISSA, Infragard. He regularly teaches Principles of Information Security and Information Warfare classes.

Game-changing dimension

Change the Board – Reduce the number of zombie machines by providing SOHO with a cloud-based secure desktop environment.

Concept – Online Desktop for SOHO

The idea behind Online Desktop for SOHO is to provide end-users and small-business users with cloud-hosted desktops that provide baked-in security for web surfing, email, word & spreadsheet processing. Users of this platform would get in-the-cloud, top-of-the-line protection against malware and rootkits in an easy-to-use window that looks like (and could replace) their desktop. This would reduce the number of bot-controlled machines by routing this class of users towards a highly secure and controlled environment with enterprise-class security tools and processes.

Vision

In the war on malware and botnets, end-users and small businesses (SOHO) are left behind. They have neither the time nor the resources to deal with the current threat environment. Further, with few exceptions, a majority of SOHO use the same set of basic applications, namely a browser, an email client, word processing, and spreadsheets.

What if we lived in a world where the SOHO's computing environment was subject to the same best practices for security, virtualization, and active monitoring as an enterprise-class environment? Why should SOHOs need to worry about file attachments, compromised web sites, cross-site scripting, cross-site request forgeries, cookie-hijacking, wireless security, etc?

As soon as she stepped in her office, Jane powers-on her computer. While her login prompt looks similar to a Windows XP/Vista prompt, it also sports a connectivity icon which lets her know her online desktop is ready for her. After entering a username and password, she is greeted with her online desktop, complete with icons representing application programs, email, web, and the documents she was working on yesterday. A few minutes later, while checking her email, Jane clicked on a PDF attachment which it turns out, contained malware. With a gentle beep, her screen fades, and a message appears to indicate that her online desktop has entered a self-healing phase. Ten seconds later, Jane's desktop is once again operational, and the offending email has been quarantined.

Bill was working from home today. After logging into his online desktop, he accessed his LinkedIn inbox and found a message from what seemed to be a potential business partner. Curious to find out more about that opportunity, Bill clicked on the link but the page that opens up has some strange characters instead. A message appears on the screen informing

Bill that his session is being routed to a fresh, healthy machine. After ten seconds, Bill's screen is operational and a message informs him that the link he had just clicked on contained malware and that Bill should simply discard the message containing the link.

What makes Online Desktop for SOHO possible is the cloud-based computing platform that allocates resources from the cloud to provide each SOHO user with an online desktop that feels and functions much like his/her computer desktop. However, the project combines malware and web-attack detection with a self-healing capability that simply redirects a compromised session onto a new, healthy, session in a matter of seconds.

In order for this project to become reality, we need a convergence of security technologies wrapped in a user-friendly (seamless) desktop environment. While medium and large enterprises often use virtualization and remote computing, combining the two into the cloud will allow for on-demand, always-on, computing that can free the user from many of the more complex security procedures that the current threat environment requires. In other words, while we currently have virtualization, cloud-computing, and remote computing, we do not yet have solutions that can replace the computer desktop with a safe, user-friendly, accessible-anywhere version.

Method

The impetus for this project was to be able to reduce the attack surface by reducing the number of machines that become zombies, i.e. under the control of some outside and malicious entity. Asking home users or small businesses to invest hundreds of hours or thousands of dollars into anti-malware products, business continutiy strategies, and security incident and event management is simply not feasible. This class of users needs something simple that requires little or no additional work on their part, namely a solution like the Online Desktop for SOHO.

A major assumption is that in order to reduce the number of bots, this service would be offered at minimal or no cost to SOHO users. A marketing campaign involving US Internet Service Providers would likely help reach a wider audience.
Platform support should initially include Microsoft Windows, and be extended to Linux as well as instant-on computing technologies.

In terms of dependencies, the project will require the successful integration of remote-desktop technologies such as Citrix and VNC with the SOHO computer. Also, the remote computer (or the application mimicking that behavior) will need the capability to detect when it is infected with malware and redirect the user's session to a healthy session.

Dream team

Profs. Doug Jacobson & Tom Daniels (Iowa State U.), information assurance faculty
Stephen Northcutt, Director of SANS Institute
Rich Mogull - Security researcher and blog author at Securosis.com
Jeremiah Grossman - CTO at WhiteHack Security
A representative from CERIAS (e.g. Gene Spafford)
A representative from Citrix
A representative from Microsoft or Google

Who you are – We are the Energy Certification Council, ECC, a recently formed New Mexico based 501c3 whose mission includes improving communication security of legacy control systems infrastructure. The managing director of the ECC, Robert Sill, has over seven years of experience in developing Control System security technology and witnessing the development (and subsequent failure) of many TCP/IP based products in the control system community.

Game-changing dimension – Morph the Game Board by developing an inherently secure high speed serial communication technology.

Concept – Create "Secure High-Speed Serial" communication infrastructure for wired and wireless applications. The project is to develop, test and proliferate a completely secured communication environment that operates using today's existing serial based equipment and at the same time gives the ability to create faster and more capable technology that can be installed when the existing technology reaches the end of its normal useful life. This approach will require expertise from control system developers and end-users, IT security experts, universities, and national laboratories to evaluate the technologies.

Vision – The current state of the existing control system serial communication technology is at a point where many decisions about its future are being made. This provides a unique opportunity to create a fundamentally new technology to radically change control systems communication without making existing technology obsolete.

The US energy grid infrastructure, petrochemical plants, manufacturing facilities, commercial buildings, transportation (trains, subways, airports, sea ports), and the military use all use commercial industrial control system equipment that utilize the same low speed serial communications.

The communications to and from the machines across these networks is open and vulnerable to hackers. Much effort is being expended to design new control systems that communicate across TCP/IP networks, however this is being resisted by the operators of serial control systems as the incumbent cost to replace them is prohibitive and without a business case. Furthermore existing equipment is still operational and well within its operational life. Most importantly, moving to TCP/IP communications makes the industrial control systems susceptible to the security issues inherent in TCP/IP.

This project is to utilize existing equipment by upgrading its communications capabilities to gigabyte speeds while providing inherent security across the network. In short, create the best of both worlds: Develop, test and proliferate a completely secured environment that operates using today's existing equipment and at the same time gives the ability to create faster and more capable technology that can be installed when the existing equipment reaches the end of its normal useful life.

Because of its high level of security, in addition to control system data, critical data from any source could be transported across the technology as well.

Method – To realize the concept, we will canvass the existing user base to develop accurate technical and functional specifications. Universities would be critical in gathering the data as well as to create the necessary functional specifications. We would give special consideration to the requirements of the Control Systems environment which are not currently being addressed by TCP/IP technology.

Developers would create the communication technology and testing in a laboratory setting with eventual proliferation in targeted markets would complete the testing. Security methods, such as those created at MIT would be integrated into the communication layer, thereby eliminating any security issues inherent in the TCP/IP stack.

Including Cisco/Juniper on the team would allow for wider proliferation of the technology by having them add a “Secure Serial” port or ports to standard TCP/IP routers.

Dream team –Department of Energy, Department of Defense, Department of Commerce, NIST, Applied Control Solutions, General Electric, Rockwell, Emerson, Honeywell, Siemens, Areva, ABB, McAfee, Cisco and/or Juniper, Pacific Northwest National Laboratory, Idaho National Laboratory, University of Illinois, Mississippi State University, MIT, End users in the Industrial community (Utilities, Military Branches, General Motors, Boeing, General Dynamics, etc)

CONTACT: ROBERT SILL, ECC Managing Director

National Cyber Leap Year Proposal
Game-Changer: Digital Situation Awareness
Don O'Neill
Center for National Software Studies

Who are you?
Don O'Neill
President (2005-2008)
Center for National Software Studies
http://www.CNsoftware.org

The Center for National Software Studies (CNSS) is a 501 3c not for profit organization dedicated to conducting objective studies of software issues of national importance. Following its National Software Summit in 2004, the Center for National Software Studies (CNSS) published the "Software 2015 Report" identifying four initiatives: Trustworthy Software Systems, Software Innovation, Software Workforce, and Software R&D.

Don O'Neill has authored the following articles published on the CERT *Build Security In* web site:

1. "Business Considerations and Foundations for Assuring Software Security: Business Case Models for Rational Action", Build Security In web site, February 2007, https://buildsecurityin.us-cert.gov/daisy/bsi/articles/knowledge/business/676-BSI.html
2. "Maturity Framework for Assuring Resiliency Under Stress", Build Security In web site, July 2008, https://buildsecurityin.us-cert.gov/daisy/bsi/articles/knowledge/business/1016-BSI.html
3. "Calculating Security Return on Investment", Build Security In web site, February 2007, https://buildsecurityin.us-cert.gov/daisy/bsi/articles/knowledge/business/677-BSI.html

Keywords
critical infrastructure protection, critical infrastructure resilience, digital situation awareness, distributed supervisory control, intelligent middlemen, operation sensing and monitoring,

Game-changing Dimension
The Critical Infrastructure Protection (CIP) model is insufficient to ensure continuity of operations for critical missions. In addition to CIP, a Critical Infrastructure Resilience (CIR) model needs to be adopted. We need to move from static lock and chain protection beyond the combination lock to the strategy of a chess board with moving parts capable of anticipating, avoiding, withstanding, minimizing, and recovering from the effects of adversity under all circumstances of use. We need to shift the locus of control for the nation's Cyber Security from protection to resilience.

Concept

The current paradigm for Cyber Security is based on protection. Protection depends on identifying vulnerabilities and applying countermeasures to neutralize their effects. These are complex human based activities whose results are uncertain and not capable of according 100% assurance. While used with some effect for components, applications, and stand-alone systems, the paradigm of protection is insufficient for assuring systems of systems, such as, the nation's critical infrastructure and DOD's Global Information Grid. For systems of systems, the paradigm for Cyber Security must be based on resiliency.

Resiliency is the ability to anticipate, avoid, withstand, minimize, and recover from the effects of adversity whether manmade or natural under all circumstances of use. The essential capabilities in composing, fielding, and operating resilient systems of systems are coordinated recovery time objectives, operation sensing and monitoring, distributed supervisory control, interoperability, and reconstitution of data and information.

Vision

The challenge lies in anticipating and avoiding the effects of adversity, and this depends on highly refined situation awareness. So it is in the area of operation sensing and monitoring that a game-changing innovation can be found. What is needed is to obtain digital situation awareness so as to anticipate cascade triggers in the critical infrastructure and deploy effective distributed supervisor control protocols that can avoid these triggers.

Method

Digital situation awareness can be derived from traffic flow and volume. The method envisioned to anticipate and avoid cascade triggers in the critical infrastructure is based on traffic flow and volume and is specified as follows:

1. Identify industry sectors of interest to cyber security resiliency
2. Identify each enterprise and organization in each industry sector of interest
3. Identify each computer system of interest in each enterprise and organization
4. Identify each I/O port on each machine of interest
5. Record traffic flow and volume on every port for every second of every day for up to twelve months
6. Using recorded traffic flow and volume, determine expected normal operation based on derived upper and lower control limits for varying time intervals
7. Using traffic flow and volume scenarios, derive operating protocols, such as, shutdown, switch to backup, and switch to a designated alternate mode, for use by intelligent middlemen charged with distributed supervisory control of critical infrastructure operations

Dream Team

1. Government: NSF, NIST, NSA, DOD, DHS
2. Industry: IBM, Cisco
3. Academia: Carnegie Mellon University, George Mason University
4. Critical Infrastructure Sectors: Finance and Banking, Electrical, Telecommunications, Transportation, and Medical

Review considerations

Would it change the game?

1. Fielding digital situation awareness would impact the “go along, get along” culture that exists within the public/private partnership. Influenced by the concerns of moral hazard, industry executives resist knowing too much about crosscutting impacts currently neglected.
2. A by-product of this game-changer would be to improve the coordination of recovery time objectives and the derivation of distributed supervisory control protocols.

How clear is the way forward?

1. Recording, archiving, analyzing, and retrieving digital situation awareness through traffic flow and volume at every port is technically achievable within the state of practice.

What heights are the hurdles that may be found in the way forward?

1. Deploying resiliency and the digital situation awareness game-changer throughout the critical infrastructure is a public policy challenge.
2. A variety of public policy measures are available to assist the deployment of resiliency maturity. These are assessed in "Public Policy Strategy for Deploying Resiliency in the Critical Infrastructure", The Competitor Vol. 11 No. 6, July 2008, http://members.aol.com/ONeillDon2/competitor11-6.html
3. The self-help remedy with indemnification as the incentive appears to be very promising as a means to lowering the height of this hurdle.

RFI Name – RFI–2 – National Cyber Leap Year

Title of Concept – Block Watch – Cyber Attack Alerting/Coordination system

RFI Focus Area – Morph the game board

Submitter's Contact Information – Shane Macaulay, Security Objectives Corp, 240 FORSYTH ST, BOCA RATON, FL, 33487

Summary of who you are – Predominantly, Security Objectives is a for profit corporation. Our team's core member's each are 10+ year veterans of information security. Shane Macaulay (K2), in 2001, pioneered static signature evasion methods (ADMmutate), cited and studied in over 33 books (Including Peter Szor), hundreds of IEEE/ACM journals and countless online publications. Our team has 4 members, but our long time presence in the industry affords us a large substantiative network of peers.

Concept – A Rapid/Coordinated Malware/Cyber attack response system. During and after a cyber compromise, the individuals responsible for resolving the attack are largely working in the dark. There exists almost no ability for disparate groups of security professionals to co-operatively react and aid or otherwise facilitate resolution. Essentially, the same tasks are forced to be repeated countless times. Providing a means for cyber attack victims to coordinate rapidly will drastically alter the existing concept that attackers are afforded infinite time.

Our system also enables a database of "known-good" applications; the closest current system is provisioned by NIST, the NSRL. We are pioneering a comprehensive model which, is network based, self-organizing (groups may maintain private or semi-private registry's) and maintains usefulness across files, memory, network or any other application.

Current thinking is of the opinion that there is no way to apply cryptographically secure hashing beyond file analysis.

We have found that we can apply this to memory/networks/any unstructured data source.

Vision – No virus, hacker, malware or any other form of unwanted program/code would run on any computing system.

Users may simply download and begin to use immediately. Over time, users may build or grow personally specific extensions, to suit their needs.

A user could be a desktop user, being alerted *before* unknown code is executed in memory on their computer, or a network administrator charged to oversee the protection of intellectual property violations or theft.

Recent advances in storage media, performance and magnitude have made some aspects of this feasible, along with our recent prototype's success.

We have also devised protocols for use in P2P scenarios which ensure integrity even from un-trusted hosts (or at least the identification of a malicious host).

Method –

[redacted]

Dream team – The applicability of this type of data indexing/retrieval/identification/coordination/search could involve any number of agencies concerned with information assurance.

RFI Name – RFI–2 – National Cyber Leap Year

Title of Concept – RADE – Automated vulnerability discovery system

RFI Focus Area – Morph the game board

Submitter's Contact Information – Shane Macaulay, Security Objectives Corp, 240 FORSYTH ST, BOCA RATON, FL, 33487

Summary of who you are – Predominantly, Security Objectives is a for profit corporation. Our team's core member's each are 10+ year veterans of information security. Shane Macaulay (K2), in 2001, pioneered static signature evasion methods (ADMmutate), cited and studied in over 33 books (Including Peter Szor), hundreds of IEEE/ACM journals and countless online publications. Our team has 4 members, but our long time presence in the industry affords us a large substantiative network of peers.

Concept – There are currently no solutions for automatically discovering security vulnerabilities in binary code, without false positives. This goal has long been regarded as unattainable; however using a system of systems approach it is possible to overcome the challenges. The capability to discover all of the vulnerabilities in a set of software would enable network defenders to implement impenetrable attack countermeasures such as patching all useful attack vectors or deploying perfect IDS rule sets. Attackers with the same ability would seem virtually omnipotent to frustrated defenders.

All of our systems running in concert provide a capability that changes the game board. When combined with significant computing resources, our system provides a decisive advantage in defending friendly infrastructure and controlling or disrupting infrastructure belonging to an adversary. To defend against an adversary who can exhaustively discover technical flaws in software by brute force would require either equivalent capability or strict avoidance of commercial off-the-shelf software. It is simply not possible develop a modern, functional IT infrastructure without COTS.

Vision – Our vision is of a vast datacenter full of servers, each methodically testing different parts of the same target application, until every possible code path has been tested. Software would be significantly more reliable and users would suffer less application crashing. Clear and objective metrics would be established for software robustness, creating true market incentives for vendors to invest upfront in building secure products.

We invented and implemented all of the individual systems for our solution. (Currently supports Microsoft Windows, and can be ported to other platforms) The final step is integration into a fully automated system of systems, and deployment to a high performance computing environment.

Method –

[redacted]

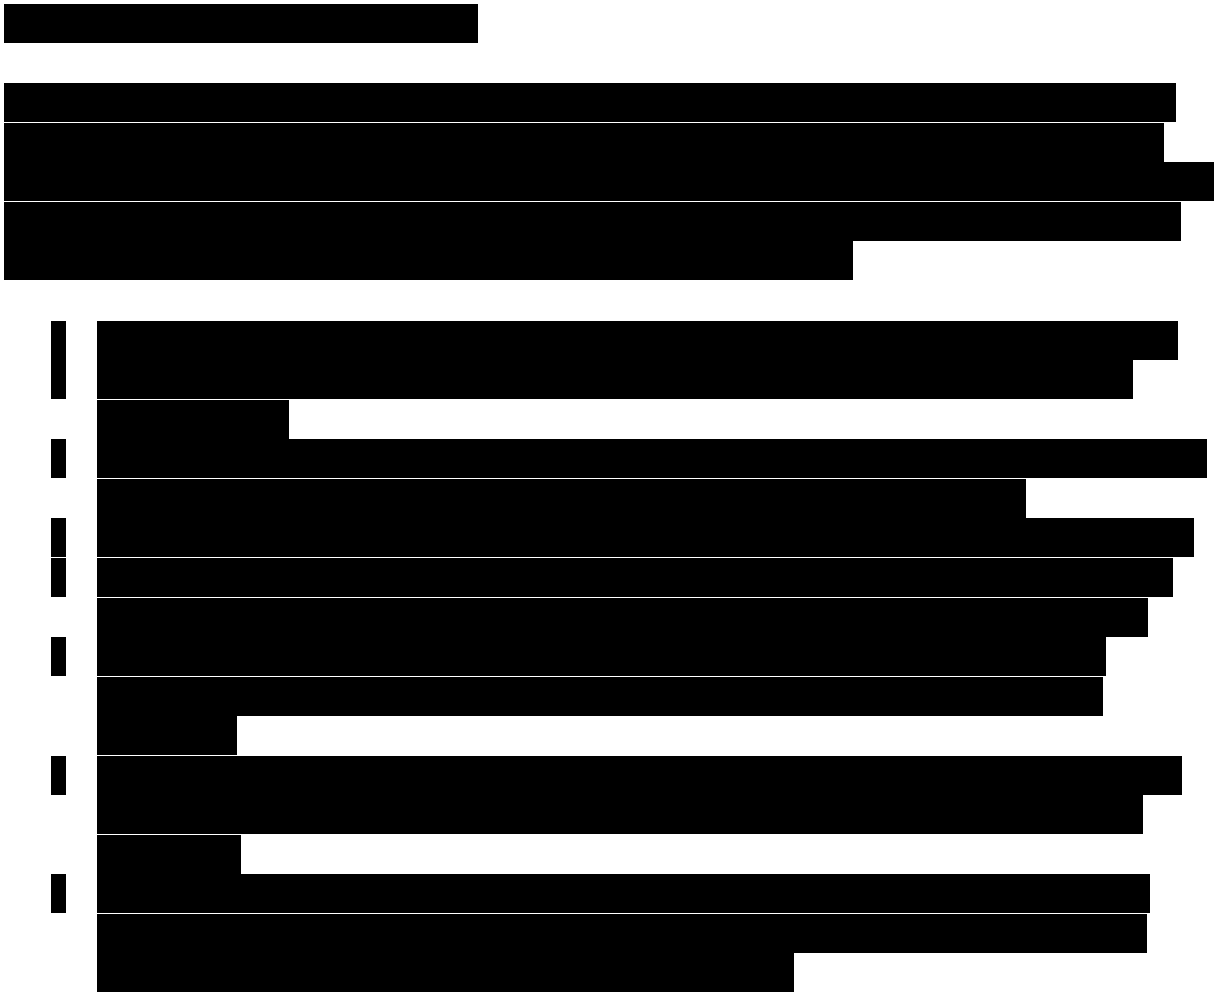

[redacted]

Dream team – The applicability of this type of automated binary analysis, reverse engineering, and automated vulnerability discovery/development could involve any number of agencies concerned with software security assurance.

Keychain Computing

Who you are - BBN Technologies is a 700 employee business that customers turn to for solutions to complex problems that provide safety and security using breakthrough thinking to bring new innovations to a competitive market. We have been actively engaged in the development, deployment and security of the Internet from its roots in the ARPAnet, and continue to drive evolving protocols and standards.

Game-changing dimension – Changing the game board

Concept – Composable untethered secure computing: Imagine a computing platform that is comprised of small cheap processors (with storage) that are plugged into a chassis which provides user I/O, processor interconnect, and access to external resources. One instance might look like a smart phone with multiple SDIO card slots. Another might be a laptop. In short, any form factor that provides power and I/O to the cards would suffice.

Vision – Today we are accustomed to the idea of installing one or more virtual machines on a separate high-powered physical machine. A number of groups are active in this area including all the VM vendors, but also Intel and AMD with their VT and Pacifica architectures. Another underappreciated player is the One Laptop Per Child (OLPC) program and the BitFrost security model. OLPC uses a small and cheap processor (AMD Geode) to run a Linux kernel. The BitFrost innovation is that OLPC gives each application its own virtual machine and consequently its own view of available resources. Keychain computing takes the OLPC model to the extreme by putting virtual machines and applications onto removable "chip" media like SDIO cards, thereby making the base platform little more than an I/O provider.

Method – SDIO cards are processors. The SDIO specification is intended for control of content and intellectual property. Some, like the Eye-Fi photo card, not only have 2GB of memory, but also carry a 400Mhz MIPS processor and an 802.11 transceiver. Such a device is more than capable of serving both a communications and processing function. It could easily "run" an email or web browser program in addition to providing networking functions. What the card lacks is power and a keyboard, screen, and mouse. Another SDIO card of note is the Spyrus Rosetta card with a processor and approval for storage of Top Secret data. Other cards have Bluetooth and GPS capability.

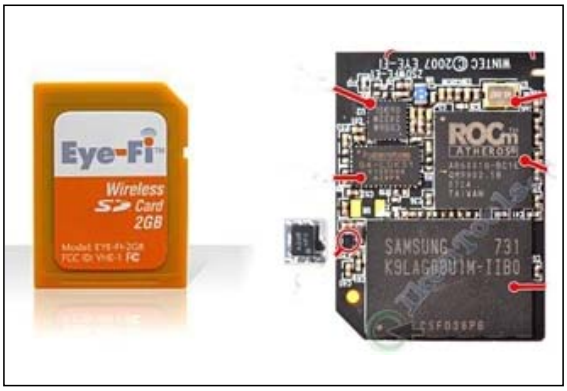

The further development of these cards as well as their composition into novel computing platforms is easy to envision. Security is the major concern and here several independent developments could be extended and integrated. The first is the development of the Trusted Computing Group standards for trusted BIOS and Hypervisor loading based on the Trusted Platform Module (TPM). These already exist and are part of the Intel VT chips and TPMs shipping today. Together, these form one trust anchor that allows a user to confidently insert their SD chip into an unknown platform. The second is the processing capability of SD chips themselves. The Spyrus Rosetta card is evidence that strong cryptography combined with FIPS 140-2 level 4 tamper protections is achievable. Such capability can be used not only to challenge and affirm the configuration of the base I/O platform, but the configuration and integrity of any other SDIO chips sharing the environment. Finally, the BitFrost security model allows these cards to carry not only virtual machines, but applications, significantly reducing the complexity and power needed for effective processing.

Several scenarios can be imagined for how this might be used in the real world.

- Special Forces or Marines could carry a base platform device like a smart phone, PDA, or Netbook that could provide nothing more that power and I/O. The configuration of the device would depend on the SDIO cards presented by each soldier. One could bring a GPS, secure storage, mapping application, and RF card to perform mapping and planning. In a serial fashion, another soldier could later use the base platform device to create log entries and gather photographic and audio evidence. The weight reduction in base platforms and batteries is significant. The base platform and cards are paired so that zeroizing the base platform destroys the ability to extract information from the cards. Breaking (snapping) the cards violates the tamper boundary and makes them similarly unusable.
- A grandmother receives a Netbook-like base platform for Christmas. Some time later she wants to create a photo album and goes to the store to buy an SDIO card with the appropriate application and storage. The camera she uses is able to write pictures to the card without overwriting the application that permits editing and display. Later she goes to visit her friends and conveniently has her SDIO card on her keychain. This makes it easy to view and share pictures on the friend's SmartPhone, including editing the photos using the application already present on the SDIO card.

There are a wide variety of base platform candidates. Basic I/O, power, and card slots are all that are needed from a hardware point of view. Note that for enhanced security, a system with VT or Pacifica combined with a TPM and appropriate hypervisor would be required. The BitFrost security model is already defined although it isn't clear that BitFrost is reusable for this purpose. However, this should be a relatively straight-forward design and code activity. Building SDIO cards is not terribly difficult. The standards are available and joining the consortium is a matter of a small fee. The hardware for the cards is not difficult to produce and prototypes and experimental versions are well within the capability of universities and small companies.

Dream team – The dream team would consist of: an experienced SDIO developer, an experienced base-platform developer, security experts and developers, and a suitable test team.

Institute a Network of Fairness

Who you are - BBN Technologies is a 700-employee company that specializes in innovative and effective solutions to complex problems. We have been actively engaged in the development, deployment, and security of the Internet beginning with its roots in the ARPAnet, and continue to drive evolving protocols and standards.

Game-changing dimension - Changing the Rules

Concept - This idea would institute a network of fairness by using a quality of service approach. Currently, high bandwidth users get a disproportionate share of the network bandwidth, with no penalty for being piggish. One could implement a new form of enforcement via SLA-type mechanisms, with additional mechanisms in the network core whereby users agree to abide by certain policies in return for an assured high quality of service.

Vision – Distributed denial of service attacks (DDoS) in the Internet can be viewed as a failure of existing quality of service (QoS) and fairness mechanisms in Internet routers, and, to a lesser extent, in host operating systems. It can be argued that any DDoS attack is clearly unfair to the users being locked out. A truly fair enforcement mechanism wouldn't allow DDoS attacks to succeed, or would at least require 10-100X more attacking nodes in order to successfully appear to be normal user behavior, and thus not unfair.

The focus is on DDoS attacks in which the target system is overloaded by normal-appearing traffic, and not partial connections (e.g., SYN-flooding) or other protocol-related DDoS tricks. Traffic-overload DDoS attacks cannot be stopped by simple inspection of individual packets, such as by a firewall. In addition to direct attacks from botnets, DDoS attacks can include indirect social mechanisms, such as the "slashdotting" attack - posting a target URL to a widely-viewed location on the Internet with the intent that ordinary users will follow the link in excessive numbers, causing a traffic overload. Another related phenomenon is the "flash crowd", which is identical to the slashdotting attack, except that the coordinated traffic overload occurs due to natural (non-malicious) causes, such as a widely-followed news report.

Unfortunately, existing Internet router behavior under congestion generally rewards the largest incoming flows and flow aggregates with the most outgoing bandwidth when congestion occurs. Various strategies have been developed to combat this problem, such as stochastic fair queueing and Diffserv three-color marking, but these address only a small subset of the problem space, and are generally ineffective against DDoS attack techniques.

Specific resources, such as servers, can be protected against unfair traffic overloads by specialized network devices, such as load distributors and traffic flow controllers that track excessive usage by specific remote peers. Similar technology is now being used by some access ISPs (notably Comcast) with the intent of limiting the (unintentional) impact of a few excessive bandwidth users on normal user traffic. However, there is no system-

wide solution for the Internet in general - only individual sites or access systems. Protecting the end system is insufficient when the attacker can move the DoS point back into a network link leading to the target. So ultimately, the Internet core itself has to be part of the solution.

Method – BBN believes that many of these problems can be successfully addressed by "changing the game" - re-examining current approaches to network flow control, congestion management, and QoS enforcement and identifying new mechanisms that would make the network and end systems robust against distributed traffic load attacks.

Possible research topics include novel (and probably disruptive) QoS schemes, strategies based on economic behaviors (finding a way to make the cost to the attacker much higher than for legitimate users), or credential-based access to network and end-system resources. One caveat here is that these types of mechanisms have, in the past, sometimes offered new avenues to deny service, so any proposed mechanism must address the fact that its own enforcement mechanisms might be leveraged as an attack avenue.

One insight is that simple changes to the queuing mechanisms can have a dramatic effect on network fairness. Currently, queuing mechanisms reorder packets in attempt to differentiate service between flows. An alternative is to concentrate on how packets are dropped. Instead of marking packet flows, each packet could be marked with a drop precedence or drop cost. This would allow the edge-systems and internal routers to drop the least important part of a flow based on the user's own determinations.

The new fairness mechanisms must be simple to configure and manage, or the ISPs will not adopt them. The implementation must be simple to implement, or the Router vendors will not adopt them. The end users must see an advantage for the special treatment, or they will fall back to best-effort service. That is, the ISPs need to make money by offering fair services and the end users get a predictable service.

Dream team – In order for the new fairness services to be fielded, there must be agreement between ISPs, router vendors, server farms, and the end users. While these players have many conflicting special interests, they can all agree that a network infrastructure that gives unfair advantage to unscrupulous users hurt all of players. Getting key players from all these groups to focus on a few simple changes to improve network fairness will have a profound effect on perceived and actual network robustness.

Input to NITRD National Cyber Leap Year RFI

Who we are

The National Institute of Standards and Technology (NIST),
Information Technology Laboratory (ITL)

Founded in 1901, NIST is a non-regulatory federal agency within the U.S. Department of Commerce. NIST's mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life. NIST employs about 2,900 scientists, engineers, technicians, and administrative personnel at its headquarters in Gaithersburg, Maryland, and its laboratory in Boulder, Colorado.

ITL conducts research and develops test methods and standards for emerging and rapidly-changing information technologies. ITL focuses on technologies to improve the usability, reliability and security of computers and computer networks for work and home. ITL employs 329 professional and support staff and 147 guest researchers.

This information and more can be found at http://www.nist.gov.

Game-changing dimension – Change the rules

Concept

The traditional security authentication model is described as using one or more of the following approaches: something you know (such as a PIN), something you have (such as an ATM card), and something you are (such as a fingerprint). Despite known weaknesses and breaches, the prevailing model in use today is single factor authentication, username and password. Multifactor authentication methods that are interoperable, usable, and cost-effective are needed to securely operate in untrusted environments.

Vision

Satisfying the goal of this initiative will enable secure, automated access to Federal services by multiple Federal agencies, state and local governments, industry, and individual private citizens. This is absolutely necessary to meeting National goals for increasing the efficiency and transparency of Federal operations and services. It will also serve to protect tomorrow's needs in a highly-connected, ubiquitous computing environment upon which the Nation's security and economy is becoming increasingly reliant, as well as increasingly vulnerable to attack and misuse.

Method

NIST will develop a framework and implementation plan for interoperable tokens that contain biometric and cryptographic credentials to support logical access control on a multi-platform and multi-operating environment basis. This will eliminate the difficulties with achieving interoperability that have impeded the widespread use of multifactor authentication technologies, which are more secure than single factor, e.g., password, security systems. For example, while interoperable biometric templates do now exist, operating system and applications interfaces necessary for enterprise-wide reliance on biometric authentication are still rare. Standardization of the protocols, interfaces and data structures as well as usability will be undertaken in this initiative in coordination with vendors as well as government agencies and departments, to enable interoperability.

A leap-ahead will be based on projecting tomorrow's needs and addressing the major challenge of securely and robustly authenticating identities on a global scale, by :

- Utilization of new types of information, such as detecting geo-spatial information or measuring a neuro-response to a challenge;
- Sustainability of biometrics identifiers through developing measurement and evaluation of revocability and liveness algorithms;
- Portability of the factors, by matching on a token (e.g. fingerprints on a USB flash drive or voice and fingerprint on a phone) or not relying on hardware; and
- Fusion of factors for high confidence in authentication mechansims.

Dream team

NIST ITL has two divisions – the Computer Security Division and the Information Access Division – and a Program in Identity Management Systems, which include expert scientists and engineers working in areas relevant to this problem, including standards development and metrology in biometrics, key management, and credentialing.

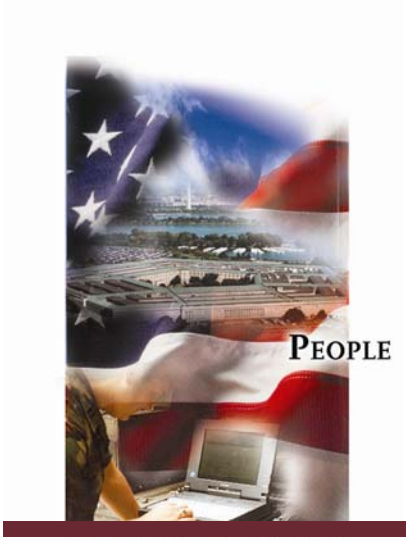

Networking and Information Technology Research and Development

National Cyber Leap Year

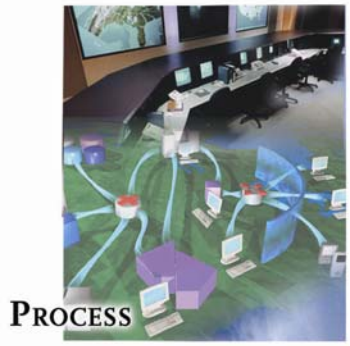

NCI Information Systems, Inc.

Request for Information

National Coordination Office for
Networking and Information Technology
Research and Development (NITRD)

December 15, 2008

Prepared by:
NCI Information Systems, Inc.
11730 Plaza America Drive
Reston, VA 20190
www.nciinc.com

This Page Intentionally Left Blank

Introduction

NCI Cybersecurity Prowess

- Our support of the cyber threat analysis and fusion center has earned US-TRANSCOM (supported by NCI) top National Security Agency (NSA) recognition: the Frank B. Rowlett Award
- Developed detailed data mining and complex mathematical modeling applications on high-volume data streams for our classified customers' counterterrorism missions.
- Designed and implemented grid infrastructures and other emergent technologies for large-scale, high-performance information sharing between disparate organizations.

NCI Information Systems, Inc. (NASDAQ:NCIT), a leading information technology and cyber security implementation firm, is pleased to provide this response to the National Coordination Office for Networking and Information Technology Research and Development (NITRD) Request for Information (RFI). In this response, we will share our insights, knowledge, and ideas for successful implementation of the National Cyber Leap Year program. NCI has served Federal customers with the highest quality of innovation and technical solutions for 20 years. We have a history supporting leaders of diverse and complex organizations to achieve, sustain, and surpass mission goals and objectives. NCI brings the highest caliber talent to projects that make a difference. There is no project as foundational as the Comprehensive National Cybersecurity Initiative (CNCI).

Since its inception in 1989, NCI has earned clients' trust by understanding their dynamic technical environments and responding to their needs with technical innovation and implementation excellence. NCI has more than 2,400 employees, nearly 100 locations worldwide, and more than 70 percent of our staff maintains active Government clearances. A record of repeat business with customers such as Army National Guard, U.S. Transportation Command, and Department of Energy, and the Government Accountability Office demonstrates the confidence that customers place in NCI. Our support is guided by our clients' requirements and expectations, the talents of our staff, and industry's most advanced technologies to deliver innovative, cost-effective, long-term solutions. For NITRD, we propose to leverage our Intelligence Community (IC) expertise into a game changing technical platform that can be morphed immediately to gain momentum to address threats stemming from asymmetric attacks in cyberspace.

Game Changing Dimension

We propose to *change the rules* of the game by leveraging our information dissemination and robust analytical processing expertise. NCI will use its experience in applying massively scalable and automated Knowledge Management/Information Extraction techniques to accelerate deployment (and utilization) of counter terrorism-focused capabilities across the IC.

Concept

Numerous intelligence and data resources across the Federal Government, including the Department of Homeland Security, Department of Defense (DoD), and IC, support their own unique missions. Because no one organization has the sole responsibility for cybersecurity, the organizations involved in this effort need a comprehensive approach. A common unified knowledge management and collaboration solution would enable these organizations to dynamically share facts, information, and analysis leading to full knowledge of the situation as, or even *before*, it unfolds. This solution would change the game by providing the capability to quickly and effectively respond to impending and ongoing cyber attacks. Similar to the attackers' ability to share

information about targets, exploits, tools and defenses, this knowledge-based environment will *deliver decision advantage* by aiding in discovery, prediction, and ultimately interdiction of those seeking to do harm through cyber space.

Vision

NCI envisions a proven method of acquiring and analyzing massive amounts of data across multiple intelligence domains to address global threats and attacks. NCI will exploit this data, using tested and scalable knowledge management/information extraction tools and techniques at unrivaled velocity within massive processing and storage environments. The result will be rapid decision making abilities that in turn lead to actionable results. These results can be used to initiate rapid responses to perceived and predicted malicious activities from adversaries. Using these techniques will enable the customer to identify individual as well as networks of cyber attackers, isolate them and their means of creating cyber attacks, and subsequently defeat their purpose.

Method

NCI has partnered with its IC customers to successfully deploy mission-centric analytics as integrated components on a Service Oriented Architecture. We have integrated technologies for deploying extraction and identification methods that uncover "interesting" information and relationships within and across unprecedented volumes of previously uncorrelated data. NCI is using these techniques today, leveraging an architecture that can scale horizontally and vertically as processing requirements demand. NCI's has successfully implemented knowledge management/information extraction approaches within massive amounts and flows of data. We have found that resource allocation, information sharing, and collaboration in the IC share common usage patterns needed to move from an implementation framework of "need to know" to "need to share." Emerging distributed computing technologies, (e.g., grid computing) provide an ideal solution for managing massive data movement by taking advantage of the latest large-scale cluster computing, reliable messaging, and entity-based security. NCI uses emerging distributed computing technologies intelligent monitoring features that are fully aligned with high-availability and fault-tolerant requirements. NCI has leveraged grid and globally scalable technologies to cost effectively enable dynamic and coordinated data sharing and aggregation of information and services. This successful approach is based on Google's industry leading cloud computing, search, and advanced visualization capabilities. NCI gladly extends an open invitation to NITRD to attend a demonstration of this technology in its secure facility in Maryland.

Dream Team

NCI has the business skills and relationships to form and lead a team with expertise in grid architecture, reference ontology, data enrichment, advanced information extraction techniques, geospatial visualization and cyberspace exploitation skills. NCI would leverage the best minds assembled from our partner companies, including Google, Microsoft, Cisco, LexisNexis, and Teradata. These companies are leaders in commercial industry as well as Federal contracting, and have expertise honed by serving the IC. In addition, our team would be further supported by cross-Government participation including NSA, Central Intelligence Agency, Defense Intelligence Agency, DoD, Department of State, Department of Justice and the EOP OS&T. NCI will form and manage a team that possesses the tools, knowledge, frameworks, and expertise to provide the highest level support to NITRD and the CNCI with the National Cyber Leap Year.

Sufficient Cyber Attack Attribution

Who you are — We are the Institute for Information Infrastructure Protection (the I3P): a consortium of 27 leading universities, national laboratories and nonprofit institutions dedicated to strengthening the cyber infrastructure of the United States. Managed by Dartmouth College, the I3P fosters a multi-disciplinary collaborative approach to R&D related to cyber security.

Game-changing dimension — Raise the stakes for malicious actors to use the Internet.

Concept — Our legal and policy frameworks for responding to cyber attacks cannot work unless we have adequate attribution; these frameworks remain incomplete because we lack the basis (sufficient attribution) to actually use them. Without the fear of being caught, convicted, and punished, individuals, organizations, and nations will continue to use the Internet to conduct malicious activities. We need attribution to create a system of deterrence. In such a world, malicious activity on the Internet would start to look much like that on the telephone system – bad things may happen, but much of the low level malicious activity goes away or is dealt with in a straightforward manner. Better attribution would bring reductions in cyber crimes and cyber terrorism, and it would lay the groundwork for an international doctrine for cyber security.

Vision — The vision is a set of policies, supporting technologies, and incentives to provide sufficient attribution on the Internet.

In the short to medium term perfect attribution of the source of traffic across the Internet will not be achieved, but we certainly can improve on today's capability. If there is a future transition to a new networking technology (and enhanced attribution is built in) we may approach perfect attribution – akin to what exists, and is socially accepted, in telephony.

However, near-perfect attribution will not be necessary to achieve many socially beneficial goals. The degree of attribution required will vary by situation, so that the degree of attribution assuredness required to issue the equivalent of speeding ticket for low level attacks would be lower than that required for a military/law enforcement response to acts of cyber terrorism or war.

The technical challenges to improving attribution capabilities may be less important than the policy/legal/social issues which attribution raises. For example, which cyber activities need attribution? Which cyber activities need non-attribution? There is an issue of Internet governance here: just as a court decides when a wiretap is appropriate, someone must decide when to use attribution, assuming that it is not automatically employed, and someone must decide when to act upon the attribution. Attribution may require the cooperation of multiple entities, and that cooperation may be conditioned upon having some say in the eventual outcome.

This endeavor will include collaboration efforts across three key areas:

- Establishing mechanisms for inter-jurisdictional cooperation, at policy, operational, and technical levels, to trace back or otherwise attribute network traffic to its source machine, and to marry source machine identification with social/human factors techniques to link corporeal individuals to a particular machine;

- Creating incentives, thereby making attribution valuable to those who engage in legitimate business transactions while keeping non-attribution possible in the realm of idea exchange. The incentive structure should also penalize those who engage in malicious activities or retribution for idea exchange. The goal of such an incentive structure is to develop social acceptance for the concept of attribution and to guide technological development toward appropriate and sufficient attribution without destroying non-attribution;

- Building legal/policy/doctrine frameworks that would address key questions: what is adequate attribution for a particular event and contemplated response? How are conflicts between different jurisdictions in desired attribution goals (e.g., what traffic should be non-attributable) resolved? Who pays, and how, for the costs of an attribution system? How are non-cooperating jurisdictions to be dealt with in attempting to achieve attribution? While some of these questions are economic, many will require new criminal procedures, additions to the law of war, and new military doctrines for dealing with cyber attack.

It will be challenging to build many of the elements of a system of improved (but not perfect) attribution. More challenging, however is the need to build multi-jurisdictional, including international, cooperation at both a legal/policy level as well as a real time attribution capabilities level. Some technological evolution may also be required.

We believe that a move toward sufficient attribution is possible now because: (1) it is increasingly apparent that Internet security is worse today than a decade ago; (2) technology alone has not provided a silver bullet for security: and (3) if we cannot reliably attribute the source of cyber attacks, we will continue to make little progress fighting cyber crime, war, and terrorism.

Method — The I3P funds "headlights" white papers authored by its members. These papers are to describe a future security problem that the author feels will pose a serious threat to the global information infrastructure within the next five to ten years. One of the very best proposals, and subsequent white papers, was one on the topic of attribution, which has led to this submission.

The main dependencies of this idea are on the carrier's and companies' willingness to participate while assuring privacy advocates of the idea's safety.

Dream Team — FCC; One or two major telecomm carriers; one or two major ISP's (US and Canada at least, EU a plus); DoD; DoJ; DHS-NCSD; NIST; privacy professional, (e.g., Lisa Sotto of Hunton & Williams); Internet technology companies (e.g., Jon Stewart of CISCO).

RFI Name: RFI-2 – National Cyber Leap Year
Title of Concept: Resilient Coordination of Autonomous Offensive Software Agents
RFI Focus Area: Change the Rules
Submitter's Contact Information: Richard Murphy
Noblis
3150 Fairview Park Drive South
Falls Church, VA 22042 Tel: 703-610-1635 Fax: 703-610-1699
Summary of who you are: Noblis, a public interest nonprofit organization, is often called upon to act as a "trusted partner" in the assessment of cyber-security products, programs, and services. Our conflict-of-interest free structure permits us to provide independent objective assessments, assistance, and support services for government agencies that are often operating within a complex environment of competing commercial interests. Noblis' accomplished and experienced technical staff is comprised of over 550 engineers, scientists, analysts, researchers, specialists, and management experts. The majority of these staff members have served in positions of trust and responsibility in private industry, academia, and the government**.**

Concept: The key innovation to the research proposed in this white paper is the development of a resilient, reliable, and adaptable C^2 infrastructure which is used to operate tamper-resistant agents. Formalizing these ideas into a reusable framework will potentially create an infrastructure that can make real gains in system resilience.

Vision: Cyber Defense depends upon situational awareness. Systems can defend themselves by detecting hostile activity and deflecting those attacks by blocking the network hosts involved in the malicious activity. Centralized sensor systems such as intrusion detection and intrusion prevention systems operate close to the systems being defended (often on the same physical network), which provides them the ability to quickly respond to attacks. However, there are disadvantages to this approach. First, the information available to such a sensor is limited to whatever traffic is directed toward the systems which it is defending. Given a highly distributed attack network (as is easily constructed using compromised home Internet-connected hostss) and a large number of systems which must be defended, it becomes clear that a distributed sensor network is necessary to have a full picture of offensive activity. Without such a broad-based view, attackers can simply spread their attacks over thousands of sources such that the attacks are "lost in the noise" at a given network endpoint.

A similar argument can be made in the reverse direction. A system component which reacts to attacks by launching counterattacks is easily thwarted as the attacker can simply block counterattacks from the target network. A counterattack network therefore should be widely distributed if it is going to be effective at eliminating the sources of detected attacks. A counterattack network would ideally be agile (moving around the network periodically) as this would more effectively avoid detection.

The use of widely distributed networks for defense and counterattacks leads to a new problem, that of attacks against the command and control (C^2) systems. An adversary would quickly notice that reactive attacks were taking place in response to their activities. It is logical to assume that they would attempt to subvert the distributed sensors and counterattack agents so the

C^2 networks could then be penetrated. Simply blocking communications between agent systems and the C^2 network would be effective in deflecting counterattacks from the distributed network.

It can be seen that the problems in this area are similar to the problems found on the Internet. Specifically, as hosts have adapted to spam by improving their filters, spammers reacted by changing their content to avoid those filters. Observing the evolution of these malware instances can provide ideas that can be applied to the cyber defense problem.

Network-centric applications have embraced cryptographic protections in the area of widely-distributed networks of compromised systems which are called botnets. Botnet C^2 has evolved from highly centralized systems that were easily disrupted to use of so-called "fast-flux" networks of systems built upon peer-to-peer networking, load balancing, and highly redundant systems. When defenders have been able to infiltrate botnet C^2, the botnets have adapted by incorporating additional protections that allow the networks to avoid infiltration. Another example of a basis for C^2 resilience is tamper resistant software such as that used by the Skype voice over IP program, which uses cryptography to ensure privacy as well as continuous monitoring for tamper attempts with active response to debugging attempts, etc. We propose using these ideas to change the rules for resilient C^2 systems.

Method: During the research envisioned in this paper, Noblis will investigate existing botnet C^2 mechanisms to determine which can be adapted to protect existing network services. From this work we envision construction of a C^2 toolkit that can be adapted to control existing network servers and daemons. A distributed service such as the domain name service (DNS) will be used to demonstrate the ability to adapt the framework to existing services while providing secure C^2 of those servers. DNS is a useful target for secure C^2 given that the DNS databases are changing frequently and given the fact that highly distributed DNS servers are helpful in ensuring system resilience and availability. Continued work during this phase will involve adapting the C^2 infrastructure to support distributed network sensors and network counterattack agents.

Building this C^2 infrastructure as a library that can be easily repurposed for other services will be a significant goal of this research. After the initial implementation, other critical services (such as service-oriented architecture components) will be chosen for demonstration of the viability of the approach.

For the second phase of this research, Noblis proposes abstracting the self-protection mechanisms used in products such as Skype as referenced earlier, and investigate methods by which such techniques can be widely used. We propose modifying an open source compiler suite to generate binaries with imbedded integrity checking, tamper resistance, and debugger detection so that arbitrary applications can be built with such protections imbedded.

This phase of the research expects to make real improvements in the integrity of applications by detecting tampering attempts. Detection of attempts to manipulate servers using a debugger can improve protections for cryptographic data such as keying material, helping a system to resist attempts to extract private keys from running programs.

RFI Name: RFI-2 – National Cyber Leap Year
Title of Concept: Designing IA Systems to Degrade Gracefully
RFI Focus Area: Change the Rules
Submitter's Contact Information: Richard Murphy
Noblis
3150 Fairview Park Drive South
Falls Church, VA 22042 Tel: 703-610-1635 Fax: 703-610-1699
Summary of who you are: Noblis, a public interest nonprofit organization, is often called upon to act as a "trusted partner" in the assessment of cyber-security products, programs, and services. Our conflict-of-interest free structure permits us to provide independent objective assessments, assistance, and support services for government agencies that are often operating within a complex environment of competing commercial interests. Noblis' accomplished and experienced technical staff is comprised of over 550 engineers, scientists, analysts, researchers, specialists, and management experts. The majority of these staff members have served in positions of trust and responsibility in private industry, academia, and the government**.**

Concept: An information-system-based attack can be expected to take place outside the scope of reasonable human response. Either the attack will be too swift (three minutes to peak attack for Slammer), too broad (massive Denial of Service, multi-point attacks), or too subtle (single-packet attacks over weeks or months). Rules and heuristics are required for automated response and notification. These rules and heuristics can also filter and prioritize events so that operators are directed to the most important information first, and not overwhelmed. Quality of Information Assurance (QoIA) is intended to describe the level of Information Assurance (IA) services available to a system. To make practical use of QoIA, common metrics are required. This white paper proposes the research necessary for defining a uniform QoIA standard which engages and unifies currently disparate efforts, defines how QoIA interacts with Quality of Service (QoS), and creates a methodology for establishing QoIA Protection Profiles.

Vision: There are currently no commonly accepted QoIA metrics or uniform methods of scoring QoIA Metric importance. Even where independent system designers and developers produce their own metrics, differences prevent interoperability without costly translation interface design. This ad-hoc approach increases system complexity and costs, as well as reduces effectiveness and efficiency. A uniform standard for Quality of Information Assurance (QoIA) is needed to aid the DoD in designing information systems that adapt and/or gracefully degrade when unexpected events occur. Uniform QoIA standards will assist application developers in managing the critical requirement for adaptation and graceful degradation.

Method: A quantitative set of QoIA metrics will allow different approaches to IA system degradation to be compared and evaluated, and used as a basis for developing a uniform standard protocol which will be the basis for designing IA systems to exchange information about their state in a much more fine-grained fashion. We propose the following multi – phased effort:

Phase I: Define a uniform set of QoIA metrics and their relationships
Step 1 - Leverage the following initial samples of common IA space QoIA Metrics of Confidentiality, Integrity, Availability, and Non-Repudiation (CIAN) to develop a

comprehensive uniform standard set. Because of a lack of coherence among current QoIA efforts these are subject to revision and addition based on the results of research and testing.

- Confidentiality:
  - Ability to encrypt data in motion or at rest
  - Ability to authenticate endpoints in a data transmission
- Integrity:
  - Integrity of stored, sent or received log data
  - Integrity of data in motion or at rest
    - Example: RAID 5 single-unit failure threatens data integrity and availability, but allows operations to continue.
    - Example: Hardware encryption accelerator failure decreases performance (availability) but allows processing to continue using software-only encryption.
- Availability:
  - QoIA metrics for Availability are generally related to availability of IA services
    - Antivirus (signatures out of date, etc.)
    - Host/Network IDS (signatures, bandwidth/CPU)
    - Authentication (LDAP/Federated Credential connectivity)
- Non-Repudiation:
  - Ability to digitally sign transactions, or identify parties in data transmissions
    - Example: Impending signing-key expiry at end of mission defines when non-repudiation will become unavailable. This may have greater or lesser impact depending on the value assigned to non-repudiation by mission planners. A GPS-guided munition that can still accept and authenticate instructions may still be of value to the mission while unable to respond in an authenticated manner.

Each of these Metrics will be assigned one or more levels for acceptable performance, and actions to be performed at those levels. Once the levels have been defined, interactions between Metrics will be mapped to determine what the impact of a specific event will be. The result will be similar to a state-table, which will lend itself to simple implementation with very high performance. Weighting may be applied to the state-table to accommodate varying operational requirements.

QoIA Metric scoring will be mapped to existing IA "levels," such as Basic, Medium and High Robustness to facilitate integration into current systems.

Step 2 - Identify levels of functionality from full, through degraded, to unavailable for each standard QoIA Metric

Phase II: Develop a prioritization of IA services affected by each QoIA Metric to determine what losses to tolerate and when to disable services
Phase III: Develop a CC Protection Profile Template and a Protection Profile for Firewall QoIA.
Phase IV: Use the Firewall QoIA PP to instrument an open-source firewall using the metrics, and test that the system responds appropriately to simulated attacks. This Firewall QoIA PP can be optionally sponsored for submission to the U.S. Common Criteria Scheme for evaluation and certification.

Dream Team: Noblis, NSA/NIAP, DISA

National Cyber Leap Year - Cyber City

Agnes Hui Chan, Professor and Associate Dean of Graduate School, College of Computer and Information Science,
Northeastern University,
Boston, MA 02115

Professor Chan received her Ph.D. in mathematics in 1975 specializing in the area of combinatorics. She joined the Northeastern University faculty in 1977 and is currently the Associate Dean and Graduate Director in the College of Computer and Information Science. She led the effort in establishing an interdisciplinary research Institute of Information Assurance with the Department of Electrical and Computer Engineering in 2005 and is instrumental in getting the University designated by NSA and Homeland Security as a Center of Academic Excellence in Information Assurance Education and Research. She designed and launched the interdisciplinary graduate program, MS in Information Assurance, in collaboration with the College of Criminal Justice.

Professor Chan focuses on cryptography and communication security. Her research considers coding schemes that are easy to implement, but difficult for others to eavesdrop. Professor Chan holds two patents, one on ultrafast pseudorandom sequence generator and one on software based stream ciphers. She has also published widely in IEEE conferences and journals, as well as in Crypto and Eurocrypt. Her research has been funded by NSA, NSF, DARPA and telecommunication industries.

Game-changing dimension: redesigning the game by introducing cyber cities with cyber players in electronic form. The design will rely on the interdependence and integration of three major dimensions of the game (a) the "game board", (b) rules of the game, and (c) payoff on the game.

Concept:

If we can have police cars catching speeding vehicles on the highway, why can't we have patrolling agents catching cyber highway offenders, such as those carrying viruses? If our transportation infrastructure can be partitioned into air, train and highway systems, why can't we subdivide our cyber connections into substructures for easier management and protection? The concept is to study how a cyber infrastructure can be re-organized and re-designed as a cyber city for management and governing purpose, so the security of the cyber city can be efficiently provided and guaranteed.

Today's internet has grown from a combination of different, independently developed networks, each of which may serve a different population for different purpose. For example, the ARPA net was set up for defense research work, the NSF net was set up for academic collaborations and network experiments, other commercial networks are set up for various banking purposes. The cyberspace resulted from this ad hoc formation can easily generate traffic congestion, bottlenecks and points of vulnerability for attackers. The game board (the architecture of cyberspace) has to be redesigned based on user behavior, purpose served, and other factors.

Currently, internet users are provided access to internet resources with little or no cost incurred. This easy access to "free" resources has emboldened users, legitimate or otherwise, to use these resources for personal gains and satisfaction at the expense of other users.

Vision: The vision is to model and to govern cyberspace as we have done with the physical world, except via *electronic personnel*. We will draw on parallel behavior and vulnerabilities observed in our physical world, design and introduce electronic professionals in order to maintain a manageable and secure cyber society.

Method:

A new design for a secure cyber society requires a thorough understanding of current behavior and capabilities of cyber space, as well as those of physical society. The process can be divided into two major stages: (1) examination of current status and (2) design of a cyber city.

1. Examination of current status: a series of workshops to study the parallelism between physical and cyber cities. For example, a workshop on *transportation* where routing architecture and rules of traffic on the electronic infrastructure can be compared to our transportation infrastructure. A workshop on *criminology* where white-collar crimes, terrorism and other criminal violence are studied and parallelism are drawn from cyber cities. Other workshops such as *vulnerabilities* and *protection* will be held. While technical and implementation aspects of the design remain important in the study, other factors such as economic incentives/deterrent, user behavior and/or judiciary systems need to play an important role of the study.
2. Design of the new game: Based on the results of the workshop studies, rules of the new game should be designed accordingly. We will concentrate on two main areas:

(a) Redesign the Architecture of a Cyber City. The network routes should be partitioned and arranged in a hierarchy, providing "express" service through routes analogous to "air transportation". The lowest hierarchy should represent local roads where mail and other chats can be delivered. By partitioning the routes into hierarchies, we can better manage and protect the critical cyber points. Economics of managing network routes and service delivery should play an important role in the design of the architecture.

(b) Creation of Electronic Citizens. Each electronic citizen is a software agent living in the cyber city, performing professional duties that are assigned. The initial challenge here lies in the identification of electronic professionals and their associated duties. Once identified, the challenge will be the implementation of the software agents "living" in the cyber city without adversely degrading the performance of cyber space expected by users.

Dream Team: To ensure that the game is indeed a new game, we need to rely on interdisciplinary teams of experts representing (a) social networks (L. Barabasi of Northeastern University), (b) policy and legal rights (D. Lazer of Kennedy School at Harvard), (c) criminal justice, (d) communication networks, and (e) economists. The workshop attendants should consist of representatives from academia, industry, government, law-enforcement, and business organizations.

Harnessing Ambiguity

Provided in response to
NATIONAL SCIENCE FOUNDATION
Request for Input No. 2 (RFI-2)
National Cyber Leap Year

February 16, 2009

Prepared For:
The National Science Foundation
The National Coordinating Office (NCO) for Networking Information Technology Research and Development (NITRD)

Prepared by:
Northrop Grumman Information Systems,
Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238

RFI Name: RFI-2—National Cyber Leap Year
Harnessing Ambiguity

Game-changing dimension: "Morph the gameboard"

Contact Information:

Chris Valentino
Northrop Grumman Information Systems Essex Operation
8666 Veterans Highway
Millersville, MD 21108
(410) 923-8415

Peter C. Canestaro
Northrop Grumman Information Systems Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238
(210) 706-4712

Company Description:

Northrop Grumman Information Systems – Essex Operation has been involved in Cyber Security and Information Assurance for over a decade, and has been recently awarded the DARPA National Cyber Range Phase 1 contract.

Concept

Ambiguity is a needed element in creative enterprises such as literature, humor and the Arts. As multiple and hidden meanings approach, a form of fission occurs, where new ideas are sparked, cascading into entirely new trees of thought. In the software engineering world, however, ambiguity is frowned upon. An unexpected interpretation by an optimizing compiler can cause loss of sleep; an unexpected race condition can threaten a project and one's livelihood. So, in software and computer engineering, ambiguity is hunted down and eliminated without mercy.

However, there is a general approach by which ambiguity can be used as the key element in increasing security, flexibility and in decreasing resource constraints. The basic idea is that in recording or transmitting information, the actual information is not manipulated, but rather a set of clues with which to reconstruct the information is transmitted or recorded. These clues would be orders of magnitude more compact than the actual information, and would be sufficient for later reconstruction of the original information, given a context. This differs from current encryption methods, in that there may be (uncountably) many contexts for which these clues may be decoded into coherent information. Someone intercepting a message may decode the clues by guessing a context, but he would have no confidence that he had reconstructed the original message. In fact, it would be possible to deceive an adversary blending multiple messages in a single set of clues, and "leaking" the wrong context to the adversary so that he would be fed coherent, but false, information.

A similar technique could be used for executables and DLLs, but for purposes of flexibility and security. Rather than the explicit executable code, clues and context information would be recorded. A reconstructed message corresponds to a variable-length piece of code, so it could map to a method call, or to a programming construct. Patching or updating code could be done with much more precision than replacing the entire binary file. This also allows for quick customization of software features needed for various security levels.

With respect to security, it would be possible to pre-identify code fragments corresponding to specific content cells (P/Q locations) as malicious. If a program attempted to use one of these cells, the operating system could raise an exception and stop execution. Similarly, combinations

of cells, executed in a certain order, may be designated as malicious and their execution prevented. The malware developer would have ever-increasing hurdles to overcome. (Note that for messaging, the content landscape can be changed periodically by prior agreement, but for executables, it would be fixed.)

Vision

The approach described above can begin immediately, implemented as a software layer in the protocol stack, and as part of a hardware abstraction layer, for executables. Operating system developers (probably Linux kernel developers, initially) would write monitor modules for preventing execution of malicious cells. For more effective and widespread use, the basic algorithms should migrate to hardware. The net effect of widespread adoption of this type of technique would be a more secure computing environment and more secure and efficient communications.

Method

The basic method employed in formulating this idea was to investigate uses of ambiguity to increase bandwidth. Security considerations were a result of a fruitful fission of ideas.

Dream Team

A “dream team” would consist of computer scientists, operating system developers and chip designers.

Operational Empathy

Provided in response to
NATIONAL SCIENCE FOUNDATION
Request for Input No. 2 (RFI-2)
National Cyber Leap Year

February 16, 2009

Prepared For:
The National Science Foundation
The National Coordinating Office (NCO) for Networking Information Technology Research and Development (NITRD)

Prepared by:
Northrop Grumman Information Systems,
Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238

RFI Name: RFI-2—National Cyber Leap Year
Operational Empathy

Game-changing dimension: "Change the rules"

Contact Information:

Chris Valentino
Northrop Grumman Information Systems
Essex Operation
8666 Veterans Highway
Millersville, MD 21108
(410) 923-8415

Peter C. Canestaro
Northrop Grumman Information
Systems Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238
(210) 706-4712

Company Description:

Northrop Grumman Information Systems – Essex Operation has been involved in Cyber Security and Information Assurance for over a decade, and has been recently awarded the DARPA National Cyber Range Phase 1 contract.

Concept

In our daily lives, it is simple enough to tell if a friend or colleague is feeling well or if he is under stress, or ill, or simply "not himself". We are constantly sending subliminal health-status messages to those around us, who in turn process these messages with no conscious effort. It is quite obvious when a normally cheerful person seems preoccupied, or when a typically serious person says something (intentionally) silly. Not only are we processing the content of communications, but the manner and style with which they are presented. When there is deviation from this expected manner, we notice and adjust so as not to put undue stress on the individual. Depending on the severity of the deviation from the norm, we may become concerned about this person, or we may even alert authorities.

There is currently no "operational empathy" in the computing world. Communication between computers simply transfers content or deals with protocol formalities. There are no health-status indications in the protocol, other than time-outs and packet re-sends. There is no indication as to whether a machine is operating normally, or is under duress. However, this situation is easily remedied (in theory, at least). If some minimal host operating characteristics (such as CPU and memory utilization, number of processes running and idle, etc.) were included as part of an expanded IP header, a number of possibilities unfold.

- **Enhanced Cyber-Community Security**
  All communicating nodes are now part of a "neighborhood watch" community, monitoring the health of all. It would be a simple matter for a host to profile the normal behavior of all nodes typically communicating with it. When the behavior deviated beyond a certain threshold, the community would be alerted. Appropriate action could be taken. Depending on the situation, it may be appropriate to quarantine the machine and search for malware.

- **Increased Difficulty in IP Spoofing**
  It would also be a simple matter to notice when IP spoofing was taking place, as the health-status information would not necessarily match current values. It may be

appropriate to automatically reject attempted connections whose health-status values do not match currently active sessions' values.

- **Increased Difficulty in Writing Effective Malware**
  Similarly, when malware resident on a host, attempts to contact or infect another host, leveraging a trusted relationship between the two, it would be required to craft packets that include accurate health-status information. Failure to do so may cause the connection to be rejected (see above). However, the activity of the malware itself may cause accurate health-status information to already be deviating from normal. The malware must somehow determine what normal values would be without its presence. (This is the dilemma of a "Heisenbug".)

- **More Efficient Server Utilization**
  From the initial TCP handshake, an application may be cognizant that the requested server may be less responsive than normal. The application could immediately attempt connections to mirror servers or duplicate servers on a pre-defined list. The net effect (pun not unintended) would be to decrease delay time and automatically balance load as appropriate.

Vision

If the IP protocol were augmented as described, and in wide use, the internet (and local networks and sub-nets) would become more secure, flexible and fault-tolerant. There are two main hurdles to overcome in realizing the vision.

- **Industry Refinement and Adoption**
  IPv6 has a number of unassigned header designations (134 – 254) that could be used; IPv4 has unassigned Option types. Standards bodies must agree on allocation of one of these unassigned designations for this proposed use. In addition, protocol stack software developers must have access to host resource consumption metrics available through the host OS, and must transfer this information to outbound packets. Application developers would readily make use of this information, if provided in each packet.

- **Privacy Concerns**
  There may be privacy concerns initially, that will be allayed, when it is noted that only summary information of the host resource consumption is being disseminated, and not specific application information.

Method

The basic method employed in formulating this idea was to investigate how computer networks may emulate beneficial aspects of human networks.

Dream Team

A "dream team" would consist of protocol standards committee members, protocol stack software developers, OS developers and major application developers.

Right-Brain Computer

Provided in response to
NATIONAL SCIENCE FOUNDATION
Request for Input No. 2 (RFI-2)
National Cyber Leap Year

February 16, 2009

Prepared For:
The National Science Foundation
The National Coordinating Office (NCO) for Networking Information Technology Research and Development (NITRD)

Prepared by:
Northrop Grumman Information Systems,
Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238

RFI Name: RFI-2—National Cyber Leap Year
Right-Brain Computer

Game-changing dimension: "Morph the gameboard"

Contact Information:

Chris Valentino
Northrop Grumman Information Systems Essex Operation
8666 Veterans Highway
Millersville, MD 21108

Peter C. Canestaro
Northrop Grumman Information Systems Essex Operation
6100 Bandera Road, Suite 505
San Antonio, TX 78238

Company Description:

Northrop Grumman Information Systems – Essex Operation has been involved in Cyber Security and Information Assurance for over a decade, and has been recently awarded the DARPA National Cyber Range Phase 1 contract.

Concept

The common (and simplified) view of the Human brain, with its left and right hemisphere functional distinctions, can serve as both a metaphor and roadmap for research related to an alternate computing platform. As a metaphor, it allows us to place current computing in a larger context; as a roadmap, it suggests a few components that do not yet exist. The current digital computer can be likened to left-brain functionality, in that it can deal with ordered sequences of logic and symbols; it can perform precise numerical calculations and apply rules. However, we do not currently posses computing machinery that can be likened to the right brain, where geometric and spatial manipulations take place naturally, and where a holistic or intuitive view of a problem can exist and be manipulated. And of course, we do not have an analog for the connection between the two hemispheres, the corpus callosum, allowing these two halves to communicate. Lastly, if the two halves are to collaborate for some purpose, an executive entity must exist to coordinate this collaboration, modulating the effect one side has on the other, towards that purpose. These "missing components" will be explored in terms of their possible form and function in an alternate computing platform.

This alternate computing platform would be more capable in a number of ways, including enabling intuitive programming solutions to visually intuitive problems. It would also have immediate implications for cyber-security:

- Constant introspection and health cross-monitoring
- Instant visual recognition of malicious code patterns in digital memory
- Difficulty of next-generation malware to disable all components simultaneously
- Increased fault tolerance and possible real-time work-around for component failure

Vision

There are a number of hardware and software components that do not currently exist, but for which requirements and functionality may be envisioned:

- **Orthogonal Memory View/Manipulate/Write (Corpus Callosum)**

The proposed functionality is predicated on an orthogonal view of the memory of a traditional digital computer. This orthogonal view allows immediate visualization of the current state of memory in a single event or machine instruction. From this view, normal patterns of activity may be learned, and malicious patterns may be discerned. For example, a long string of zeroes in a process stack space can indicate an attempted buffer overflow attack "landing zone". Such a pattern would be obvious as an unusual dark vertical line in the visual rectangle of system memory. There could be simultaneous visual representations of memory, at various resolutions and dimensions (1D – 4D) revealing different patterns and aspects.

This view must also allow write access to the memory. Then, when a malicious situation is noted, it can be directly corrected in memory and/or designated memory cells in the traditional computer can be set to inform it of the situation. In normal operation, this mechanism would be used to communicate back to the digital computer the results of "right-brain" computations.

There are several technical problems to overcome. This mechanism could not be implemented with current display and camera technologies. The memory access times are about 10^7 times faster that display refresh times, rendering obsolete any memory image so displayed and captured. It may be possible to have some visual property of memory cells (reflectivity, specific wavelength absorption, etc.) change with respect to current value. Then the problem would shift to capturing the image in a timely manner. A solution to these problems would make the rest of the endeavor possible.

- **Hardware-Based Spatial Computations (Right-Brain Circuitry)**
  Light-sensitive conditional circuitry must be implemented, enabling "immediate" recognition of lines, arcs, connected subsets, closed loops and nested entities.
- **Transition Machine (Right-Brain)**
  The right-brain component itself would be a "transition machine", rather than a state machine. A transition machine is an abstract machine whose chief characteristic is that it is always in transition. If it ceases (even for a moment) to be in transition, it is not a transition machine. Any snapshot of its activity could be called a "state", but these states are not very useful in describing or quantifying the behavior of the machine. Examples of transition machines are lava lamps, the weather, and time itself.

  This transition machine would be event driven, with various visual machine instructions overlapping others on different regions of the board or chip. The machine would receive memory images from the digital computer, and generate events for digital memory for itself and for the Executive Advisor.
- **Executive Advisor**
  The executive advisor would focus attention of either side by masking the orthogonal memory view. Where this software module should reside is an open question.

Method

The motivation for these ideas arose from frustrations encountered while approaching problems with intuitively simple visual solutions, but the coding of which was surprisingly complex, such as the "convex hull" problem.

Dream Team

A "dream team" would consist of materials scientists chip designers and computer scientists.

Managing Routing Trust

Paul Syverson
Naval Research Laboratory
Washington, DC 20375

Who I am: I am a mathematician in the formal methods section in the Center for High Assurance Computer Systems at the U.S. Naval Research Laboratory, where I have been researching, designing, and analyzing security and privacy systems for two decades. My work includes the design and application of logics for cryptographic protocol analysis, formalization of multilevel security in probabilistic systems, and traffic analysis resistance and anonymous communication. I am inventor of several technologies including onion routing and am designer of Tor, which is used to protect the communications of hundreds of thousands of people worldwide.

Game-Changing Dimensions: Raise the stakes

Concept: Manage the information about the history and path of documents and data to manage accuracy and confidentiality for forensics, fraud prevention, accountability, countering insider threats, and facilitating information sharing.

Vision: Attempts to cope with evolving accountability and visibility of government and private communications are increasingly emerging. One example is the quasi-official email messages of Governor Palin via a Yahoo account that was then exposed by a simple guessing of the sort of frontdoor personal information that allows one to reset a password. Another is the viewing of Barack Obama's cell phone records by Verizon employees or the viewing of his aunt's immigration records at INS. Lots of attention has been paid to confidentiality of data and authenticity of data. Less has been paid to confidentiality and accuracy of data flow. Knowing who said what to whom is often far more sensitive and significant than the actual contents communicated, indeed, sometimes to the point that this is significant even if nothing is known about what was communicated. Recently there has been an increasing awareness of data provenance: which includes metadata indicating who saw a document when and who altered it, when they did so, and how. Sometimes, however, the provenance information may be more sensitive, or more highly classified, than the underlying data. Very recently there has been awareness that confidentiality of provenance information is an issue (e.g., Braun, Shinnar, and Seltzer's "Securing Provenance" in *Proceedings of the 3rd USENIX Workshop on Hot Topics in Security (HotSec)*, San Jose, CA, July 2008), but theories and mechanisms have not yet developed. Data flows that are permissible and desirable may be infeasible if they are only possible with full provenance. Examples include

- anonymous tips for crimes or terrorist activities,
- non-attributable public disclosures of important information, and
- cross-agency information sharing that manages trust of shared information concerning sources and internal procedures.

Thus, as auditing of flows becomes more complete and accurate, it must also be managed. A general theory that incorporates all aspects of traffic flow and provenance trust is needed together with the mechanisms and systems to implement it.

Method: The proposal impacts almost every sector of information technology. Some examples are software development, SCADA systems, intelligence gathering and sharing, law enforcement, medical and healthcare systems, scientific research, banking information, business and governmental administration, personnel and communication records, etc. It would therefore be useful to start by having experts from at least some of these areas meet to describe perceived and anticipated emerging problems in data accountability and provenance, both where not enough information is available and where there is risk because too much is available.

Next, emerging provenance systems for scientific data, file systems, etc. should be examined so as to determine the open problems from the various application areas that they do not address. Some examples of questions that we already know they do not answer are

- How can we attach provenance data in a way that selected parties on a path can read some or all of it but others cannot or can read only other parts?
- Can we attach provenance information such that attribution of source, modification, or reading can only be made to a group or an office but not an individual?

(Answering these questions could require innovative use of cryptographic accountability techniques such as group signatures or blind signatures. It may also be necessary or useful to employ anonymous routing techniques so that a document's trajectory can be managed from the source or so that routing provenance is provided but obscured from unauthorized viewing or for controlled non-attribution.)

- How can provenance information be controlled as multilevel secure independent of how the data to which it is attached is classified?

(Formalizations from the access control and trust management literature are likely to be instructive starting points here.)

Once goals have been identified, mechanisms such as those from areas described above can be developed and prototype systems designed. Since provenance in general is starting to be recognized, but managing confidentiality of metadata is barely on the horizon, now is the time to develop theory and systems to manage these, rather than looking for covert channels and unintended flows after systems are in place.

Dream Team: Experts in

- Data handling and policy for specific application areas
- Provenance
- Traffic Analysis
- Trust management
- Routing
- Cryptography, e.g., blind and/or group signatures

Novel Transition of Security Technologies

Paul Syverson
Naval Research Laboratory
Washington, DC 20375

Who I am: I am a mathematician in the formal methods section in the Center for High Assurance Computer Systems at the U.S. Naval Research Laboratory, where I have been researching, designing, and analyzing security and privacy systems for two decades. My work includes the design and application of logics for cryptographic protocol analysis, formalization of multilevel security in probabilistic systems, and traffic analysis resistance and anonymous communication. I am inventor of several technologies including onion routing and am designer of Tor, which is used to protect the communications of hundreds of thousands of people worldwide.

Game-Changing Dimensions: Change the Rules

Concept: Many security technologies have been created and prototyped by government scientists or by government funded researchers at university and private labs. But most researchers have neither the skill nor the interest to be entrepeneurs and government customers have neither the interest nor the procurement path to pay for transition to a finished product. Thus, even if there is government technology and a government need, there is usually no development pathway from the technology to fulfilling the need without a detour through the private sector. Sometimes this works adequately, but good opportunities can also be wasted for reasons irrelevant to their viability, cost, or potential usefulness.

Vision: The vision is to develop means to transition government-created technologies to fill government needs regardless of whether or not someone has a plan for how to monetize the technology's use. The motivation to use COTS technology is based on satisfying a government need while saving money through the use of commercially developed products. While this works to some extent, there are numerous examples where government needs are different from the commercial sector's. This is true for many areas of technology but is especially so for security. Adapting commercial products to government needs often inadequately meets the actual government need, and it is not clear that this less-than-adequate adaptation of technology is even cheaper in the long run than developing appropriate technology in the first place. It is also risky for companies to develop technologies for general government use. The changes to the evaluated products process after the early nineties addressed these risks to some extent, but it remains the case that risk for businesses developing government security technology is not free-market-driven, and thus development of government security technology should minimize market dependence.

Method: This is a significant departure from existing practice (but is meant to supplement, not replace existing practice). So the first step would be to assemble appropriate people to research (brainstorm and then evaluate) potential issues and approaches. Some examples are

- Should there be a government-run development office (or agency or community specific offices) whose employees develop security products from prototypes or should it be handled in another way?
- What relationship is most effective and efficient between the group handling product maintenance and that responsible for development?
- Existing tech transfer offices are at research entities and responsible for finding private-sector developers for their research, while customers typically look for procurement-ready or in-development products from those developers. What sort of transition offices should exist to connect customers to prototype-level researchers directly? And where: at the developer, at the push or pull side, or are all needed?
- Matching is likely to be limited to research that is already in prototype or that could be prototyped in less than a year. Would it ever make sense to go after even earlier stage research, or is that already adequately covered by existing funding mechanisms and initiatives?

Once such questions have been identified and answered as best they can initially, likely prototype technologies and customers should be identified for test cases, and development mechanisms established. Likely test cases will be ones where there is no clear private sector developer but a clear match of technology and need. If the approach makes it to that stage, test cases should be evaluated for success of the transition. They should also be evaluated for cost of the transition. A challenge will be to improve on existing cost evaluation methods since many costs of transitions via the private sector are effectively hidden by the current process, making them appear cheaper than they are. Alternatively, if existing practices preclude comparison, evaluation might show the new approach to be more cost transparent.

Dream Team: Program managers from existing agencies that fund secure IT research. Doug Maughan at DHS has probably done the most within the existing framework to move in this direction and would be a natural. Also helpful would be tech transition specialists and procurement specialists with knowledge of security, who must recognize that the goal is to set up mechanisms that obviate their usual job of matching up with commercial developers. Finally, it would be useful to include government researchers and government-funded researchers who have experience with transitioning (or attempting to transition) their security technologies.

Generating Contextual Data for Security and Privacy

Who you are – This is a joint submission by PARC [7] and SRI International [10]. PARC is an independent research business and a wholly owned subsidiary of Xerox Corporation. PARC has contributed to the creation of more than 30 companies and is celebrated for such innovations as laser printing, distributed computing and Ethernet, the graphical user interface (GUI), object-oriented programming, and ubiquitous computing. SRI International is a large nonprofit research institute based in Menlo Park, CA, with 60 years of historic innovations in computing, business, education, materials, and biosciences.

Game-changing dimension – Morph the gameboard.

Concept – Mobile devices and wireless sensor networks are becoming more prevalent and powerful, and computer systems are becoming more adept at anticipating user needs through intense analysis of various context data. Because of these trends, the US information infrastructure is already seeing the addition of huge amounts of ubiquitously recorded context data. Such data includes location, colocation, calendar, electronic communication, "twitter" status updates, and web history. For individuals, this data increases productivity and enhances social relationships. For corporations, this data is mined to provide improved services, for instance targeted advertising and location-based services, e.g. [8, 6].

On the flip side this data gives rise to new levels of privacy and security risk. At an extreme, knowledge of a person's real-time location, can facilitate kidnapping or assassination. In addition, the very data mining tools that enable improved services can be used to entrap users through convincing phishing and fraud attacks that are personalized to the target and consequently harder to detect.

Defenses for these data driven-attacks often consist of methods for perturbing the data. However, in the case of location data, modifying the data to protect a person's identity may make the data so imprecise that it becomes useless [5]. Indeed, for an arbitrary data set (location-based or otherwise), preserving privacy even when data access is limited to statistical queries requires such a large amount of perturbation to have serious utility consequences [3].

Another approach to this problem is to rely on security architectures to limit access to this data. For example, many consumers trust Google to keep this data, believing that Google will implement appropriate policies and security measures to protect the data. This approach, however, is unrealistic. Given the size, pervasiveness, and intelligence value of this data, Google cannot be expected to identify all the potential avenues for privacy breaches.

Perhaps the best known precedent is web search history. Search engines build profiles of users based on their web search histories, but web searches can be potentially quite private, as highlighted by the AOL incident in August 2006 [1]. In response to these privacy concerns, the major search engines agreed to limit the retention period of IP addresses and cookie data associated with search data. (Of course, as demonstrated by the AOL incident, a search history can be identifying even without an identifier attached to it.) One effort that is in the direction of this paper is TrackMeNot [11], a browser plug-in that obscures a user's actual search queries by automatically generating a multitude of other queries.

We propose that this faking of context can be generalized to other kinds of data, and that the widespread ability to generate convincing fake contextual data would be a game-changer. Users would in general give out both real and fake data but also have the ability to authorize and enable selected service providers to be able to distinguish the two. The difference with today is that giving this authorization would be explicit instead of implicit. This authorization may be as simple as not giving fake data to a service provider. When giving authorization, the user would have the power to weigh improved service versus reduction in privacy. The user would be able to fall back on generic service but full privacy, unlike today.

Note that it is critical that the fake data be convincing. If fakes can be detected through various algorithmic or probabilistic methods, then the fake data offer no protection to the real data. Also, the various kinds of context data need to be convincing in concert; a location trace showing movement would not be consistent with an accelerometer trace showing no movement.

Vision – In our vision, users take advantage of contextual services and are potentially free of privacy worries. Their devices generate fake contextual data (perhaps modeling fake users, aka "Sybils") and the other players (e.g. service

providers, fraudsters, etc.) are unable to distinguish the fake from the true without explicit authorization. The user might authorize a provider to see actual context data if there are real benefits, such as letting friends know where you are. Note that giving the user indirect incentives (e.g., a cash payment) to provide actual data may not work so well - the user may simply take the incentives and provide fake data.

To realize this vision, one needs first of all the ability to generate realistic contextual data. Most of the previous work in this area has concentrated on search engine contextual data. TrackMeNot has gone through several revisions to make their fake queries more realistic. There is still more work to be done on their solution; for instance, an approach to the presence of specific, identifying terms in search queries needs to be found. A more academic study of the effects of injecting noise to protect search privacy is [12].

For location data, PARC has been examining techniques for generating fake location traces. A red/black team is in place, with the red team attempting to distinguish fake location traces out of a set of combined real and fake traces. PARC has found that the generation of fake traces is non-trivial because of path generation and the noise model, but we are optimistic that convincing fake location traces can be generated.

Building on the aforementioned work in faking location traces, there is a strong need for tools that generate contextual data for more interactive settings. In particular, to the best of our knowledge there are no tools for generating fake activity data, social data or detectable audio events.

Finally, to make this vision actually useful, there needs to be a UI in which a user can authorize distinguishing fake from real. The difficulty here is that the end user needs to understand the consequences of giving actual data to a provider. Some of these consequences are the inferences possible with real data, and this is a hard problem. Nevertheless, there has been a lot of work in this area of human-computer interaction and end-user privacy. A good survey is [4].

The benefits to the government in sponsoring research in this area are twofold. First, the government will have the ability to generate convincing fake data to protect its operatives and citizens. And secondly, it will understand better the various techniques used to generate fake data in order to detect fake data generated by enemies.

Method – PARC and SRI's cybersecurity groups collected ideas internally and then met as a group. The group voted on the top two ideas, which were then fleshed out into Leap Year RFI submissions.

Dream team –

1. PARC - Expertise in ubiquitous computing, data mining/statistics, usable security, and user-centered design.
2. SRI - Expertise in computer, network, and database security; artificial intelligence.
3. Company recording and using contextual data, for example Google or Yahoo!
4. Sensor experts
5. User privacy advocate, for example EPIC [2] or Privacy International [9]

References

[1] http://en.wikipedia.org/wiki/AOL_search_data_scandal

[2] http://epic.org

[3] Dinur, I. and Nissim, K. Revealing information while preserving privacy. In Twenty-Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, pages 202-210, 2003.

[4] Iachello, G., Hong, J. (2007). End-User Privacy in Human-Computer Interaction.

[5] Krumm, J. Inference attacks on location tracks. Pervasive Computing 2007.

[6] http://www.nytimes.com/2008/11/30/business/30privacy.html

[7] http://www.parc.com

[8] http://www.nytimes.com/2008/06/22/technology/22proto.html

[9] http://www.privacyinternational.org

[10] http://www.sri.com

[11] http://mrl.nyu.edu/~dhowe/trackmenot

[12] http://www.cs.ucdavis.edu/research/tech-reports/2008/CSE-2008-10.pdf

Turning the Tables: Using Behavioral Models to Foil Cyber Attacks

Who you are – Pacific Northwest National Laboratory is a multi-program Department of Energy laboratory with broad expertise deploying operational cyber analytics technology. We have formed a partnership with Stanford University to envision broadly deployable future network protection techniques. This effort is being lead by William Pike, Ph.D. at PNNL and John Gerth at Stanford, with industry partnerships being formed.

Game-changing dimension – Morph the gameboard

Concept – Cyber attackers are increasingly successful today not only because their malware can be modified more quickly than detection signatures can be distributed, but also because the individual communications events that malicious applications (such as bots and trojans) have with their command-and-control networks or peers can be made to look like routine traffic. Network defenders must devote ever-increasing computational resources to tracking this well-hidden activity. Defenders must also update signatures constantly as attackers modify the details of the attack and use extensive botnets to disperse malicious traffic. What if the tables were turned? What if defenders knew exactly what kind of communications each machine on their network should be having, with whom it should connect, and when? On such a gameboard, it becomes extremely expensive, and thus difficult, for attackers to hide their activities.

Vision – In our vision, the future gameboard will inherently favor defenders – not attackers, as the case is today. The burden of resource intensive data processing will be on attackers (to hide their traffic) rather than on defenders (to detect it). In the future, a combination of opt-in systems including network-wide passive monitoring and active instrumentation on individual hosts will permit monitoring activity and communication patterns to create **time-varying models of each host's approved behavior**. In an enterprise, models could be shared such that new machines inherit the models of others until they have exhibited enough activity of their own. Each host-specific model is used as a baseline of acceptable behavior against which current activity is compared; activity that does not fit this model is immediately flagged as suspicious. We propose that organizations and individuals can choose to deploy this modeling system as a means of protecting themselves and can choose whether or not to share their models to help protect others.

An attacker who gains a foothold on a machine has gained a resource which is only valuable if it can be communicated with reliably in the future. To avoid detection attackers need to be able to hide the activity their code generates. In our future vision, defenders have **perfect per-host knowledge of what a legitimate activity profile looks like**. In order to avoid detection attackers will need to be able to match that profile exactly, requiring target-specific knowledge – for every potential target. For instance, command-and-control traffic for bots would need to go through servers or peers plausibly used by the potential victim.

The only way for attackers to gain this knowledge is to replicate the model. To do this they could attempt passive network monitoring to try to build their own model of acceptable behavior; this is computationally intensive, but more importantly, requires attackers to obtain a **comprehensive vantage point** since they need to analyze all traffic from each target in order to build the model. This is a difficult task which defenders can make almost arbitrarily more difficult by simply diversifying their traffic routes. Alternatively, an attacker could try to monitor the activity of a

compromised machine and exfiltrate the activity data to process it into a model externally. This presupposes that the attacker also controls a machine that the intended victim communicates with and can fit the exfiltrated data within its communication profile. Attackers, not defenders, will be forced into the position of collecting and processing huge amounts of network traffic – which in itself will arouse suspicion. In essence, attackers will have **"nowhere to hide"**.

Method – White lists that permit one machine to engage in communication with another are common, but coarse, ways to approve certain network communications today. What we propose instead is a "white model" that describes the kinds of *behavior* in which a machine should engage. Here a behavior is envisioned to be the sequence of activities necessary to accomplish a larger task, for example, printing a file. Attackers would need to be able to discern this white model, for each behavior and every machine, in order to be able to avoid suspicion. Unlike existing behavioral analysis work, this method does not model the distribution of aggregate flows, but compounds sequences of activity over time into larger behaviors.

Existing research suggests that it is possible to create the kind of user- or host-level activity models we envision. To create them, we need to refine statistical techniques for time-based modeling (what activities are typical at what times of day, for what users). We anticipate that processing could be done locally on each host, but it is also possible to centralize the processing to a server on the enterprise network (or even outsource the processing to a third-party vendor).

The behavioral modeling system we envision could be deployed as a client application on each host. Organizations will be able to choose to install this client on their machines, and the client will, with the permission of the end user, collect application and transaction records and transmit them to a modeling system elsewhere which will combine the reports with passive network monitoring to create the model. We envision that the client could also request models dynamically from a server, such as when multiple users share a machine. A behavioral model will encompass both activities that users initiate (such as web surfing or emailing) as well as those that occur automatically by the applications they use (such as beaconing for software updates). Knowledge of the full space of network activities in which a potential victim engages is only available to the network defender who controls this client; this forces the attacker to do extensive reconnaissance and continually update massive databases on every possible target.

The **business case** for this concept centers on 1) the modeling systems being deployable as managed services, providing long-term revenue streams to service providers; and 2) it benefiting from economies of scale as more organizations opt-in to sharing anonymized activity models thereby increasing the shared protective benefit. Obstacles to broad deployment include the need to encourage sharing of activity models across users and organizations, ensuring that models can be applied with a low false positive rate, and tuning the modeling approach to work within providers' existing deployment models. This concept also assumes that, for the foreseeable future, game changing technologies that simply preclude all malicious use are not feasible.

Dream team – Commercial vendors of host security software will be involved, as they have the market penetration to deliver our solution to millions of desktops. We have begun working with industry to form these partnerships. We also need statisticians, psychologists, and ethnographers to help refine the modeling strategies that these tools will use to correctly represent behaviors.

Submission: Networking and Information Technology Research and Develop (NITRD)/National Cyber Leap Year

Who we are--We are a team consisting of QinetiQ North America (QNA), OPNET, IOMAXIS, Carnegie-Mellon, Johns Hopkins (Applied Physics Laboratory), and MIT. QNA is a company of 6,000 employees with experience in information technology infrastructure, integration, and cybersecurity. Our team reflects strengths in cyber security, advanced high fidelity modeling and simulation, and next generation technology R&D for the both public sector and government markets.

Game-changing dimension—Morph the board; change the rules.

Concept—Bots (web robots) can be used for many malicious purposes often without corrupted computers or users being aware of it. We propose to integrate three technical concepts to target BotNet command and control, reduce infection response time (indications and warning), and recognize patterns/anomalies.

Vision--What's the pattern? Activity correlation across the command-and-control of the Bot (see graphic), infection, and traffic patterns can yield important, real-time information on the character and intent of a BotNet. BotNet command and control generates unusual anomalous traffic (such as very small packets) that can be detected – especially in a fashion that is matched to the character of BotNet traffic types can provide early warning of potential attacks. This would support tactical I&W.

- Approach: The primary thrust of this research is the estimation of threat environment and early warning based on the characterization of the three main types of BotNet traffic at key internal locations, like an early warning system. Research will be performed to establish static and dynamic signatures for detecting activity that raises the threat profile of a network under protection. This includes the challenge of establishing statistical classifiers in a multi-sensor, dynamic environment. Utilizing this research, coupled with key partnerships that can provide signature development, signature mining, and visualization can lead to a powerful new methodology for detecting and tracking BotNets in networks.

Vision--Clever communications. Command-and-control traffic is key to the operation of a BotNet. This traffic is typically low rate and obfuscated in some manner—making detection difficult. If we can detect and understand this traffic, we can neutralize the BotNet infection.

- Approach: This work is a focused look at low-rate covert channel traffic analysis. These channels can take various forms ranging from single packet 'harmless' transmissions, to embedded communication channels in innocuous traffic, to full–up encrypted anonymized tunnels. The challenge here is detecting these channels and isolating their signature/fingerprint/features, so that an infected network receives no commands.

Vision--Shifting the time advantage. Often the most serious Internet attacks are those that exploit a weakness that had been previously unknown. The detection of this (0-day) exploit is usually performed by observing the large-scale impact that it creates. However, in some cases, we can provide an early warning of unusual activity based on some key metrics. The development of a capability to statistically characterize Internet traffic is important in determining what is

Mission Solutions Group
2677 Prosperity Ave
Suite 400
Fairfax, VA 22031
Tel 703.852.4000
Fax 703.852.2200
www.QinetiQ-NA.com

normal and what is unusual. This characterization is a difficult problem in that the traffic patterns on the Internet are non-stationary and burst at all timescales. Research has been performed over the past 15 years to understand this character, and we believe that we can leverage that research to provide a real-time picture of the character of traffic on a network – specifically identifying new and unknown behaviors. This would support strategic I&W.

- Approach: Research will be performed to establish a parametric, statistical model of the non-stationary process that Internet traffic exhibits. This is a critical piece in developing a classifier to automatically detect threat behavior. This model can feed early warning correlation engines and can provide a trajectory of behavior that can be used to observe previously undetectable attacks. We would use realistic loads and scale to develop a traffic parameter estimator and a corresponding classifier for detection of large and small scale unusual behaviors.

This strategy would achieve better situation awareness that yields the opportunity to act with a range of decisions in the right time frame. It would yield autonomous response, pattern recognitions algorithms, and new metrics (what is the right thing to measure and provide better confidence levels) that provide insight into parameters for better cybersecurity risk management.

Method—We researched NITRD R&D priority documents, Department of Homeland Security high priority technology needs, critical infrastructure R&D requirements (under the IT sector work plan, and National Infrastructure Protection Plan, the Cyber Security and Information Assurance/Interagency Working Group cyber sector plan, and other needs drivers); held discussions with senior scientists and stakeholder counterparts; and performed literature searches for science, technology, issue/trend data. We then examined the desired end state (lower overall risk), then refined these into enduring and/or emerging hard problem sets. We needed to include emerging problem sets as an extended R&D time domain might have us solving one problem trend while missing a new one. The integration of the underlying technologies in our submission provides a more fundamental approach for addressing the larger cybersecurity problem while dividing it into its enabling technology parts. They are interdependent and rely on understanding such mundane enterprises as what is normal and what is abnormal (even before a problem is manifest or other symptoms appear). The capability to detect other types of low-level traffic can provide early indications and warning. In our approach, we assume availability of data sets and representative traffic for analysis. In addition, we will depend upon identifying detectable patterns in covert channels—a very difficult challenge as this is often embedded within normal communications. Moreover, the successful technology outcomes of this approach also depend upon (an incentive for) users to be aware of what comes out of their own network—as such, there is a policy dimension to this vision. In summary, what we are fundamentally proposing is to look for a very small needle (covert channels) in a very large haystack (large-scale traffic analysis). This is inherently daunting, but we think that is where many of the solutions lie.

Dream Team—National Institute for Science and Technology, Cooperative Association for Internet Data Analysis (CAIDA), Cyber Defense Technology Experimental Research Testbed (DETER), University of Michigan, and other R&D universities or companies working on cyber output/outflow communications.

Contact—Keith Rhodes, Chief Technology Officer, QinetiQ North America, (703) 852-1384

QinetiQ North America

Resolving the Bot problem requires a holistic approach.

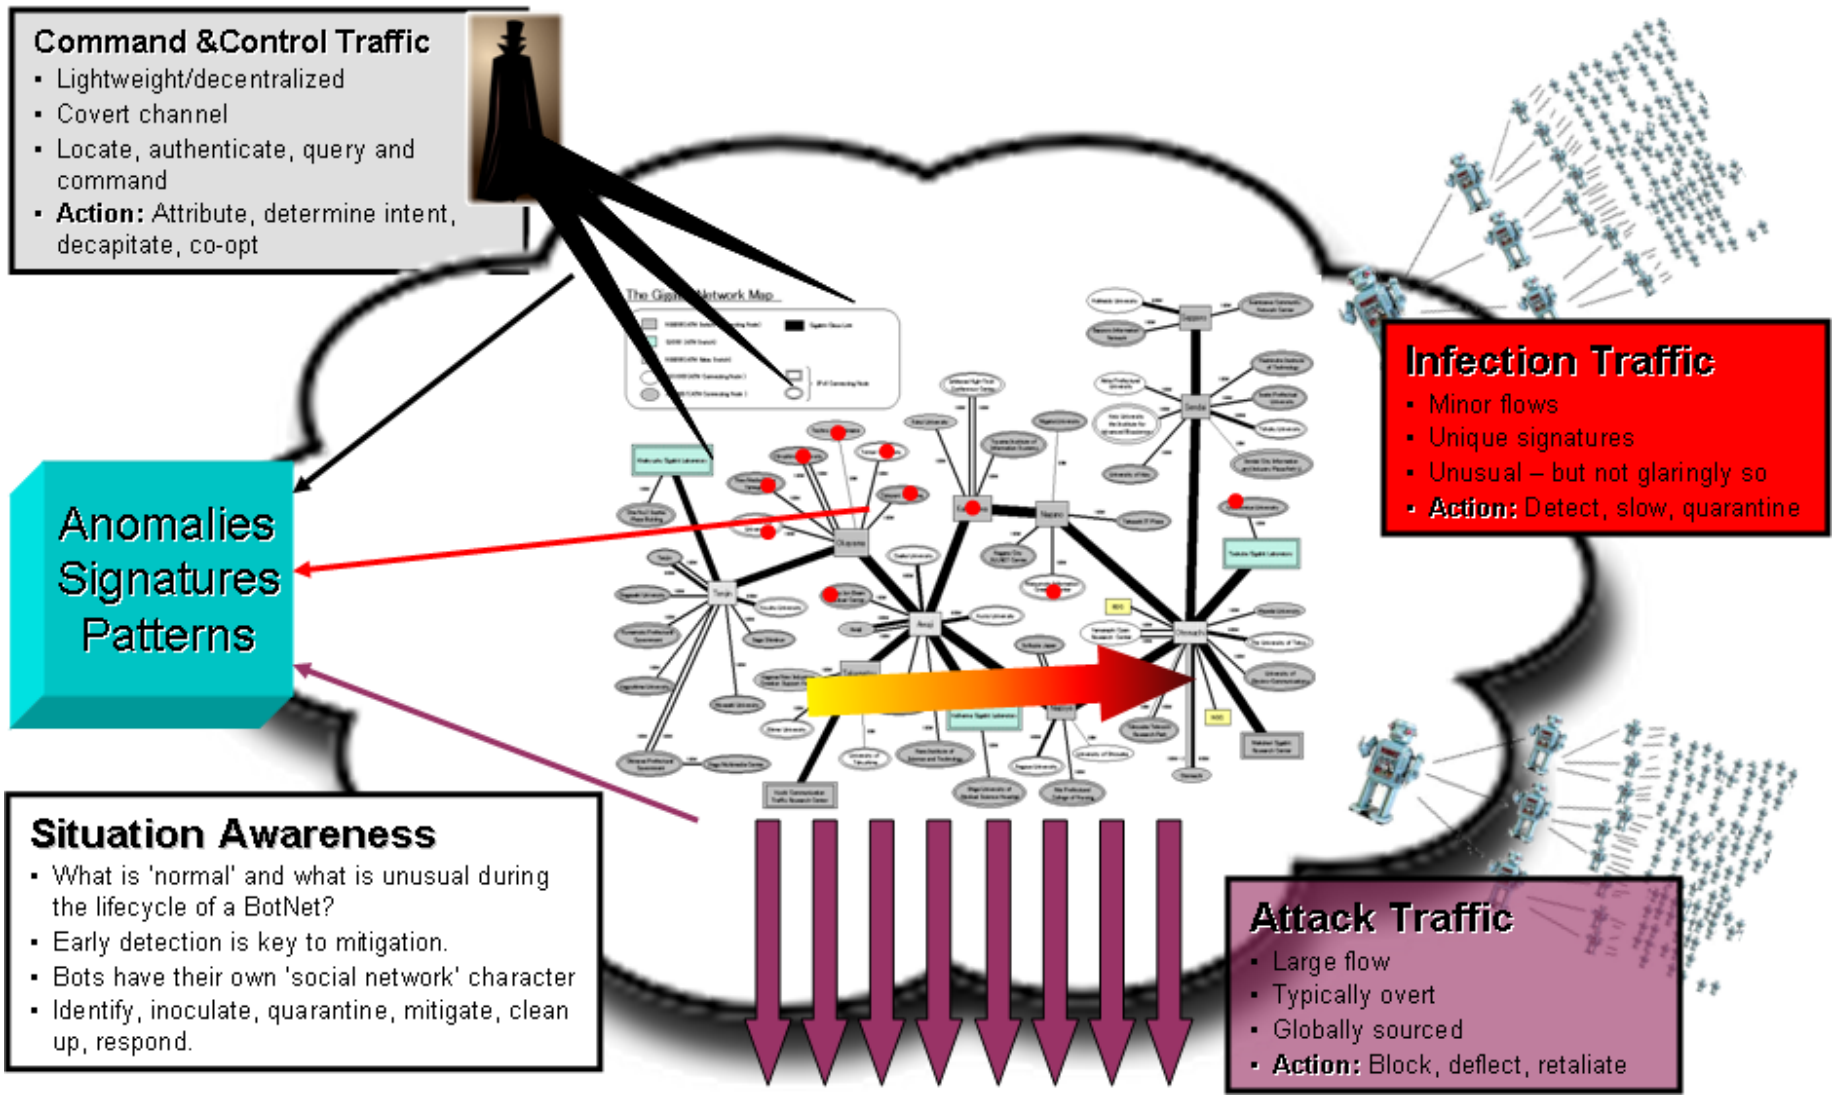

Mission Solutions Group
2677 Prosperity Ave
Suite 400
Fairfax, VA 22031

Tel 703.852.4000
Fax 703.852.2200
www.QinetiQ-NA.com

RFI on Security Risk Analysis of Computer Networks

Anoop Singhal
Computer Security Division
NIST

Who You are: We are members of the Computer Security Division of the National Institute of Standards and Technology (NIST). The URL is http://csrc.nist.gov

Game Changing Dimension: Security Metrics, change the rules.

Concept: Currently, it is difficult to answer important questions such as "are we more secure than yesterday" or "how should we invest our limited dollars on improving security." In this research we plan to develop models and tools to evaluate and mitigate security risk in enterprise networks.

Vision: At present, computer networks constitute the core component of information technology infrastructures in areas such as power grids, financial data systems and emergency communication systems. Protection of these networks from malicious intrusions is critical to the economy and security of our nation. To improve the security of these information systems, it is necessary to measure the amount of security provided by different network configurations. The objective of this research is to develop models and prototype tools for security risk analysis of computer networks. Standard models for security analysis will enable us to answer questions such as "are we more secure than yesterday" or "how does the security of one network configuration compare with another one". Also, these models will bring together users, vendors and researchers to evaluate methodologies and products for network security.

An essential type of security risk analysis is to determine the level of compromise possible for important hosts in a network from a given starting location. This is a complex task as it depends on the network topology, security policy in the network as determined by the placement of firewalls, routers and switches and on vulnerabilities in hosts and communication protocols. Traditionally, this type of analysis is performed by a red team of computer security professionals who actively test the network by running exploits that compromise the system. Red team exercises are effective, however they are labor intensive and time consuming. There is a need for alternate approaches for security risk analysis. The models and metrics we plan to develop will allow one to objectively assess the effect of applying patches and security controls on the overall security of a network. We expect that these models can also be used to suggest configuration changes that mitigate the threats to an enterprise.

The long term objectives of our research is to devise and validate a set of models for representing, comparing and measuring different aspects of network security. The models and metrics form a theoretical foundation from which practical issues such as network hardening, Return on Investment (ROI) and attack prevention responses can be formulated as an optimization problem. One example of an analysis is a "what if scenario", if a new vulnerability is discovered in a software system on one of the hosts, how does it impact one of the assets in the network. Another example is to determine the impact of patching a vulnerability on the ability of the attacker to compromise a certain asset (e.g. become a super user on a database server machine).

The following important problems will be addressed in this research:

- Devise a formal framework for identifying the high level modeling components and requirements.
- Propose numerical models of vulnerability information to integrate causal relationships encoded in graphs[1] with measurements of individual vulnerabilities provided by existing scoring systems. Recent deployment of Vulnerability Databases such as the National Vulnerability Database (http://nvd.nist.gov) and Symantec DeepSight (https://tms.symantec.com/Default.aspx) make this analysis and integration feasible.
- Apply the proposed models to derive quantitative methods for network hardening, Return on Investment and attack responses. Evaluate these methods with extensive experiments.

This research will result in a network security management tool that can assist a human user to manage the security of an enterprise network. By simulating incremental network penetration, and propagating attack likelihoods, we plan to measure and manage the overall security of a networked system. To help administrators in defending their network, the tool can automatically generate recommended actions to improve network security. The Graphical User Interface will allow the users to explore these recommendations (a list of vulnerabilities that must be patched) and their impact on the combined asset value.

There are several challenges for this research. Firstly, the model used in the analysis must be able to automatically integrate formal vulnerability specifications from the software community. Secondly, the analysis must be able to scale to networks with thousands of machines running several applications containing vulnerabilities.

Dream Team: NIST, Universities, Symantec or other Computer Security Companies

References

1, A. Jaquith, Security Metrics: Replacing Fear, Uncertainty, and Doubt, Addison Wesley, 2007.

2. L. Wang, T. Islam, T. Long, A. Singhal and S. Jajodia, "An Attack Graph Based Probablistic Security Metrics", In Proceedings of 22nd IFIP WG 11.3 Working Conference on Data and Application Security (DBSEC 2008), London, UK, July 2008.

3. M. Frigault, L. Wang, A. Singhal and S. Jajodia, "Measuring Network Security Using Dynamic Bayesian Networks", ACM Workshop on Quality of Protection, October 27th 2008.

[1] A *graph* is a collection of nodes connected by links. It is an abstract representation of a network.

Game changer: Object capabilities

Who we are. We are GROC (Group Researching Object Capabilities), an informal consortium of open source developers. We work with the object-capability model of secure cooperation, which we have applied to domains as diverse as operating system microkernels, applications, and the Service-Oriented Architecture. Affiliations are for identification purposes only and do not imply official endorsement or organizational participation.

David Wagner (UC Berkeley) Charles Landau (Strawberry Development Group)
Carl Hewitt (MIT emeritus) Alan Karp, Marc Stiegler, Tyler Close (HP Labs)
Doug Crockford (Yahoo!) Chip Morningstar (WeMade Entertainment USA)
Dean Tribble (Microsoft) Norm Hardy, Bill Frantz, David-Sarah Hopwood (consultants)
Mark Miller (Google) Matej Kosik (STU Bratislava) Jonathan Shapiro (EROS Group)

Game-changing dimension: Change the terrain. Our metaphor is a game we'll call *Night of the Living Zombies*. The players are zombies under control of unseen forces and survivors who have not yet been taken over. Zombies win the game by biting all the survivors; survivors win by staying unbitten until the zombies starve. The game begins with the survivors gathered in an isolated farmhouse as zombies approach from all directions. In a conventional defense, the survivors board up the windows and doors and defend them as zombies find weaknesses. They may also barricade the doors to various rooms as a form of defense in depth. However, once a breach is discovered, the zombies come pouring through and overwhelm the defenders.

Our game-changing approach limits the amount of space the zombies gain when they break through a barrier. While the survivors board up the windows and doors, they recognize that there are unknown weak points. Rather than futilely guessing where to strengthen the defenses, they use their resources to turn the inside of the house into a maze of little rooms. When the zombies break through, they only get a small amount of space, which is unlikely to have a survivor. The zombies need to find another weakness to exploit, but most of the ones they do find take them back outside, into spaces where they've already been, or to dead ends. The zombies' problem has now become far harder: they must find many weaknesses that line up just right, forming a complete path to a survivor. Most of the survivors go undiscovered while the zombies starve.

Concept. In today's cyber war, attackers have a compelling advantage over the defenders. Defenders need to plug all holes, while attackers need to find only a single flaw they can exploit. Defense in depth helps because the attacker must find a flaw in each layer. However, any flaw in a layer compromises the entire layer. Our approach changes the terrain so that an attacker must find exactly the right combination of flaws in order to cause noticeable harm.

Vision. The essential feature of our solution is the application of the Principle of Least Privilege (LP) at the finest available granularity. While today's systems make some attempt to apply LP, they do it at too coarse a granularity. The operating system is a monolithic whole. Like the farmhouse, a single breach exposes everything. Users are also at risk because each program they run has all their privileges, so attackers need only find a single exploitable flaw in any application to compromise all of the user's data. In our vision, a single flaw grants the attacker such a small set of rights that the attacker will typically be unable to achieve his goals.

The end state is better than just being more secure. It also enables more cooperation by limiting the exposure the parties have to each other. It simplifies management by replacing Federated

Identity Management (FIdM), with all its inherent complexity, with a system of Federated Access Management (FAccM), which makes distributed policy management a tractable problem.

Since our vision requires rewriting operating systems and applications, we're initially targeting new environments. These include the Services Oriented Architecture, browser-based applications, and software for mobile phones. As the success of our approach becomes apparent, we expect other software to evolve to use our approach. We have built prototypes at each of the following levels of granularity, giving some assurance of success.

At the coarsest level each program the user runs is granted the least set of privileges it needs to do the job the user wants done. This approach dramatically limits the damage a malicious or erroneous program can do. Still, a single breach in the program can compromise all of the rights the program has, as if the farmhouse maze consisted of relatively large rooms.

We reduce the possible damage further by modularizing applications, with a privilege-separated architecture. Consider a mail program, which has a network component, send and receive functions, a contact list, and a renderer. The renderer has no need to access the contact list, the receive function has no need to send, etc. Applying LP to the individual components means that a successful attack against the renderer doesn't let the attacker send messages to people in the contact list. The attacker must find multiple exploits that can be combined to achieve the desired behavior.

It is possible to do even better by using object-capability principles, which primarily involves restricting the code to obey good object-oriented design. With this approach, LP is applied at the level of individual objects. An exploited flaw in an object can only abuse references the object holds to other objects. Good design leads to objects that have only the references they need. To do harm, the attacker must either find a flaw in one of a very small number of objects that control critical resources, or find multiple flaws in a number of less powerful objects with resources that can be combined to the attacker's advantage.

We see object capabilities as an exciting research agenda that could bring together a broad variety of participants from industry, open source developers, and researchers from multiple areas, including computer security, programming languages, operating systems, and usability, and we would be delighted to grow the community of people working on these issues. We see an opportunity to make an impact in multiple areas: more secure operating systems; new programming languages and libraries designed for security, and secure subsets of existing languages; compilers and refactoring tools for programs written in existing languages; extensions to next-generation hardware architectures that support capabilities; frameworks for building secure and useful applications; and better tools for building RESTful and XML-based web services securely.

Method. We meet regularly and use open mailing lists to coordinate our work in these areas; this submission is based on our combined experience. Our past work (CapDesk) showed how, in our approach, security follows naturally from the normal actions that users already take as they interact with applications, rarely requiring extra clicks for security. We have experience building operating systems (EROS, KeyKOS, CapROS), verification tools (Joe-E), infrastructure (Waterken server), distributed services (Zebra Copy), and are well-positioned to apply these ideas broadly.

Dream Team. OS: developers of KeyKOS and its derivative CapROS. Applications: the Virus Safe Computing group at HP. Services: SOA developers familiar with ZBAC. Tool developers: Joe-E team from UC Berkeley. Ideally, developers from industry and interested researchers from the programming languages, security, and OS communities.

Envisioning a Larger Role for Internet Service Providers

WHO YOU ARE

Our team is composed of three primary researchers:

- ***Mr. Brent Rowe***: Mr. Rowe is an economist at RTI International with over five years leading studies for DHS and NIST. He has conducted multiple studies that identified costs and benefits of cyber security techniques and technologies. Mr. Rowe recently published a book entitled *Cyber Security: Economic Strategies and Public Policy Alternatives*. Mr. Rowe received a BS degree in electrical engineering and an MA in economics from N.C. State University.
- ***Dr. Douglas Reeves***: Dr. Reeves is a Professor of Computer Science and Electrical and Computer Engineering at N.C. State University. His research focuses on network security and peer-to-peer computing, with current funding from NSF. He was the general chair of ICICS06 and P2P2006. Dr. Reeves received his PhD in computer science from Penn State.
- ***Dr. Michael Gallaher***: Dr. Gallaher is a Program Director at RTI with over 10 years of experience leading projects for DHS, NIST, NTIA, EPA, and NSF modeling the economic impact of new technologies. He has conducted studies involving many industries including automotive, aerospace, computer hardware and software, chemical, and construction. Dr. Gallaher received his PhD in economics from Boston College.

GAME-CHANGING DIMENSION

We propose a "morph the gameboard" type concept. Our proposed solution would change the security framework in which attackers conceive of and execute their attack plans and would reduce the cyber security costs to Internet users as a whole.

CONCEPT

We propose a bold goal: to break the stranglehold that botnets have on the Internet community by motivating Internet Service Providers (ISPs) to increase the security measures they provide to customers. Because of the distributed nature of the Internet, no central organization or stakeholder group has the proper incentives to ensure security for all Internet users, contributing to the widespread vulnerabilities of the current system. Past and current solutions have proven to be insufficient; thus, we propose to investigate a range of ISP-led initiatives aimed at improving both the efficiency and the effectiveness of the Internet.

If ISPs provide more security, home users and small businesses' networks could be made significantly more secure. Similar to a neighborhood entrance security checkpoint that provides a measure of security to all houses, small-scale Internet users would be much better protected by their ISP. Users would still be wise to buy secure products, similar to the way houses in secure neighborhood still usually lock their doors. However, hackers would no longer be able to easily compromise hundreds of thoughts of individual computers to conduct illicit activities (e.g., botnets); thus, medium and large businesses and government agencies would be more secure.

VISION

We propose that ISPs take an increased role in providing security to home and small business Internet users. We believe this to be a technically feasible and cost-effective means of increasing overall Internet security. In this new world, home and small business users would no longer concern themselves with purchasing security products or services. Instead, ISPs would provide extensive scanning and filtering of incoming and outgoing traffic, supplementing existing PC-level security that comes "built in" to software products. PC-level security would thus become the last line of defense, as opposed to the first and last, as is often the case today.

In order to make this a reality, we need both technical knowledge and economic knowledge that does not exist today in the public domain. A technical analysis of potential ISP security options is needed. Many ISP security services are widely known, and some are used today; however, other potential services are either not well known or not yet conceived. A focused, objective effort to identify and analyze potential ISP security services is needed.

We also propose an economic evaluation of the costs of each service and the willingness to pay of various stakeholder groups. We need to study who can and will bear the cost of these new services. Are home users and small businesses willing to pay more for additional security from their ISP, as well as potentially give up some performance? Would large or medium businesses be willing to help subsidize the additional costs to secure small users? Alternate business models could include (1) higher fees to all subscribers, (2) higher fees to large and/or small and medium businesses, or (3) government subsidization to cover ISPs' costs.

METHOD

We will use a novel approach to develop new data and investigate revolutionary approaches to improve the current ineffective security paradigm. Similar to the NSF-funded INDEX Project, this study aims to develop the economic and technical information needed to improve Internet security by developing more robust security service offerings by ISPs. In order to conduct the required research, we have partnered with several ISPs, including Internet Texoma Inc. and the North Carolina Research and Education Network (NCREN), to measure the penetration of their networks by botnets and to assess the technical success of and economic factors related to using various security solutions. Data will be collected through honeynets and passive network traffic monitoring (using, for instance, Bothunter), and alternate technologies will be tested for feasibility, costs, and benefit characteristics.

The team will also estimate both individual and business Internet users' willingness to pay for increased security. A robust stated preference conjoint-style survey approach will be used to discern home users' desire and willingness to pay for increased security of varying levels. The results of this survey effort will be a set of data more robust than any other that is publicly available today on individual Internet users' demand for security.

DREAM TEAM

In addition to our research team, our "dream team" would include computer security experts Wenke Lee at Georgia Tech University and Guofei Gu at Texas A&M University; Google Chief Economist Hal Varian, a technology economist formerly at UC-Berkeley; and Ed Amoroso, Chief Security Officer at AT&T.

Submitters: Dr. Jeffrey Voas, SAIC
Dr. Phil Laplante, Penn State U.

Game-changing dimension: Change the **board**

Concept: The concept is to research what is a *true operational environment* for software. It has been a serious mistake to consider the traditional definition of an operational profile as the assumed proper definition of an environment. There is an explosive number of "invisible/phantom users" affecting a system during operation. These users are not people, rather underlying processes that affect resource allocation and background functionality. They are also opportunistic from the standpoint that they can mask malicious behavior. Understanding who they might be and where they might come from is a terribly important issue in assessing system trust.

Vision: A great deal of research energy has been expended over the years in creating means of reducing software to its functional components for purposes of re-composition. Formal methods and functional languages have emerged to fuse functionally aligned code. Unfortunately, far less research interest extends to dynamically composing the environments in which software operates. As these environments define the code, its behavior, its utility and its surety over time, understanding their dynamics as they affect code is essential to any successful functional composition activity. Regrettably, these environments are seldom static for long periods of time.

Time is an enemy of software and system behavior. This does not necessarily imply that software grows old or decays, even though code obsolescence is a reality, but rather that the software environment changes so dramatically with the passage of time that the static functionality of the software becomes a secondary concern. Policies governing software must not be static. Information Technology (IT) governance practices must also be malleable. The ways that users employ software evolves over time sometimes in ways unintended by the designers. And finally, new threats emerge to expose new software, hardware or human vulnerabilities.

Interconnected applications mature and gain added nodes that change their behavior and the result is a mutated environment. Operational systems and systems software morph and require differing interfaces. Value propositions migrate. User expectations remain fickle. Each of these environmental factors influences the viability of any given set of code in terms of usability, compatibility, assurance and robustness. Computational environments, affected by the impact of time, certainly influence, if not perhaps even inhibit effective information processing.

Environments are the most difficult to bound of all of the key ingredients of system trust. They tend to be discordant, chaotic, strewn with seeming garbage and loaded with extraneous noise. Such environments are almost always defined in highly unstructured ways. In fact, a workable metadata set to define an entire dynamic software environment is difficult to even conceive, much less implement. Reducing software environments to the "ilities", such as reliability, maintainability, sustainability, survivability, etc, falls short of adequately defining a given software environment. The 'ilities' themselves tend to morph the very software environments they intend to describe. They represent engineering entities to be traded across the software life-cycle either by design, intentional practice or, more often than not, happenstance.

Nonetheless, if one can effectively define the software environments at play in federated systems of systems, families of systems or any set of interactive systems, their combined impact may be evaluated. In essence, their ultimate fusion to an intersection of requisite functions to satisfy a combined and well understood operational environment would serve to increase resulting trustworthiness. The only missing variable is how to define the combined operating environment as an intersection of the various environments comprising the constituent software components. That is a game changing research challenge.

We propose that NSF tackle the "never before resolved" problem of understanding and modeling what is the true operational environment of software, considering the interconnectedness of systems of systems, environments within environments, and threats within threats. While we do not know if this is feasible, we argue that it would be game changing. We recommend a research agenda that encourages a combination of mathematical techniques, such as those traditionally used in formal models of software, as well as new ones such as Category Theory to tackle this problem. We also recommend exploring non-empirical techniques such as ethnographic analysis to help understand the context of rapidly changing software environments and threats, with the intent to be able to understand the "known unknowns" and to help anticipate the "unknown unknowns" that create the environments that reduce software systems' integrity. We consider this to be basic systems theory research, but with a software-centric and security-centric focus. We also consider this to be a 5-10 year objective.

Method: The method to do this research begins outside of the application layer. It begins in the registry, associated data bases, calls to and from the operating system, and other middle-layer functions, many of which will be off-the-shelf functionality. It involves DLLs, access to wireless communication as well as network communication. The actual approach is based on the concept of observability, but in this case, observability of the invisible/phantom users. That opens up a huge range of possibilities to thwart malicious channels that can affect behavior. It also, for the first time, opens the opportunity for true software product certification, such as a "software" Underwriter's Laboratory).

Dream team: Dr. J. Voas (SAIC), Dr. Phil Laplante (Penn State University), Dr. Keith Miller (U. of Illinois, Springfield), Joe Jarzombek (DHS), Dr. Chris Michael (Cigital), Dr. Bret Michael (Naval Postgraduate School), Dr. Anup Ghosh (George Mason U.), Dr. James Whittaker (Microsoft), Dr. Joseph Williams (Microsoft)

Attribution of Network Transactions via Resilient Authentication

RFI Name: RFI-3 – National Cyber Leap Year

Category: Attribution

Title of Concept: Attribution of Network Transactions via Resilient Authentication

RFI Focus Area: Raise the Stakes

Submitter's Contact Information:

Lyndon Pierson
Sandia National Laboratories
MS 1072
1515 Eubank SE
Albuquerque, NM 87123

Who we are – Sandia National Laboratories' network operations and cyber security R&D groups. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Lyndon Pierson, idea lead; Ben Cook, coordinating POC for Cyber Leap Year RFI)

Concept – The current architecture of the Internet supports anonymous transactions well but lacks fundamental infrastructure for robust attribution of location of origin and path of network transactions. One method of achieving attribution would involve insertion of authentication tokens or watermarking information into network traffic at or near the source, to be validated at the destination. The remote authentication information insertion techniques can be attacked 1) via subversion in the early (creation) part of the life cycle of components intended to perform signing or inject watermarking information, or 2) via extraction of private keys in spite of anti-reverse-engineering protections employed, etc. and using these keys to spoof the attribution of network transactions. If we develop authentication techniques that do not involve protection of small secrets (private keys) that will be immune to extraction of secrets, then we can focus efforts on preventing, detecting, and deterring subversion to achieve high assurance in these systems.

Vision – We envision the evolution of an attribution service that could be offered by internet service providers alongside the legacy "anonymity" service largely embedded in the current design of the Internet. This attribution service that may be optionally required by a server would operate conceptually similar to our current commerce servers that require HTTPS to protect credit card or other financial transactions rather than unencrypted HTTP sessions. Establishing a commercially viable attribution service for network transactions will enable the incremental development of infrastructure within the Internet to support an "attribution service"

based on strong authentication of the location and path of network transactions. The attribution framework required to enable this service will likely involve changes in network architecture and implementation that must be supported by commercial interests.

It is also important to assure the fundamental operation of anonymous network transactions for those applications that need it. The missing infrastructure targeted by this proposal is essential to the development and successful use of "classes of service" (analogous to first class U.S. Mail vs "junk mail", for example) enabled by robust traffic admission policy enforcement. By separating network transport services into classes of service that are valuable enough to "pay extra to use", higher classes of service will become naturally more immune to common denial of service attacks by traffic flooding, etc.

There are two parts of this vision; 1) the development of a commercially viable "network transaction attribution service" that will justify the economics of making the required adjustments to internet infrastructure, and 2) development of robust hardware/software that might be infused into our infrastructure to support this service.

Method – Implementation of the first part of this vision will require collaboration between multiple commercial interests to specify and standardize the optional attribution service and to assure that it has economic value. Implementation of the second part of this vision may require development of extremely robust hardware to inject watermarking or other authenticating information into packet streams at remote locations. The securing of this hardware against subversion or penetration by a sophisticated adversary is an extremely difficult challenge, but may be accomplished by novel authentication techniques that do not depend on the maintenance of secrets in the trusted hardware that would be subject to extraction by an adversary, or by protecting these secrets using novel "resilient" techniques that place additional entropy in the attack path, or by finding cryptographic means of moving these secrets to another portion of the life cycle where they can be better protected with guards and guns, etc.

Dream team – Collaboration with standards bodies such as the Optical Internetworking Forum (of which Sandia is a member), Internet Service Providers (ISPs), Web Service Providers, and with manufacturers of telecommunication equipment such as CISCO, Nortel, Fujitsu, etc.

Deterrence of Integrated Circuit Subversion via Rapid Attribution of Manufacture

RFI Name: RFI-3 – National Cyber Leap Year

Category: Attribution

Title of Concept: Deterrence of Integrated Circuit Subversion via Rapid Attribution of Manufacture

RFI Focus Area: Raise the Stakes

Submitter's Contact Information:

Lyndon Pierson
Sandia National Laboratories
MS 1072
1515 Eubank SE
Albuquerque, NM 87123

Who we are – Sandia National Laboratories' network operations and cyber security R&D groups. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Lyndon Pierson, Todd Bauer, idea leads; Ben Cook, coordinating POC for Cyber Leap Year RFI)

Concept – Of the three categories of computer attacks; Inadvertent Disclosure, Penetrations, and Subversions[1], Subversions are the most difficult and most dangerous because they are so difficult to detect if done well[2]. Subversions also separate sophisticated and well-resourced Nation-State level adversaries from lesser adversaries[3]. Protection of the IC supply chain for critical military functions is now getting some attention[4][5]. But these adversaries also target unclassified and non-government owned systems which comprise the majority[6] of our critical infrastructure. The risk is a cyberspace entirely of potential zombie nodes on loan to commercial and civilian applications that can be re-directed to nefarious purposes whenever the most crafty nation-state adversary chooses to wage cyber-war.

Vision – We envision a method of deterring the insertion of hard to detect subversions of ICs into COTS infrastructure by enabling low cost means of identifying individual ICs at various points in the life cycle. We further envision that COTS manufacturers would be economically attracted to implement this means in order to minimize loss of profit due to hard-to-detect insertion of counterfeit parts into the supply chain. The ability to identify a specific Integrated Circuit at various times in the life cycle would form a deterrent against wholesale substitution of the component with a counterfeit and/or subverted part, due to the improved forensic

---

[1] Subversions are defined as the clandestine insertion of an artifice in one part of the life cycle to be used to advantage by an adversary in a later part of the life cycle.
[2] Meyers, Philip A., Subversion: the Neglected Aspect of Computer Security, June 1980 Masters Thesis, http://csrc.nist.gov/publications/history/myer80.pdf
[3] J. Gosler, "Vaults, Mirrors, Masks, page 182"
[4] "Report of the Defense Science Board on High Performance Microchip Supply", http://www.acq.osd.mil/dsb/reports/2005-02-HPMS_Report_Final.pdf
[5] DMEA Trusted IC Supplier Accreditation Program, http://www.dmea.osd.mil/trustedic.html
[6] "Critical Infrastructure: The National Asset Database" http://www.fas.org/sgp/crs/homesec/RL33648.pdf

ability to trace discovered subversions to their source and the perpetrator's aversion to discovery and attribution.

We envision a method of positively and non-destructively recognizing an integrated circuit component (die) in an assembly or system as the same component as identified in an earlier part of the life cycle. The means of IC identification must be made extremely low cost and attractive to COTS IC manufacturers for detection of counterfeit parts. This vision has two parts, a) developing a nondestructive, near real-time method of tracking ICs through a wide span of the product life cycle, and b) developing the non-real-time forensic tools to attribute that portion of the ICs that are not easily associated with known foundries via method (a). The combination of cost-effective tracking of IC die with the forensics required to reliably attribute non-tracked die will deter insertion of subversions due to fear of correct attribution, mitigation of discovered subversions, and/or retribution. A separate RFI-3 response discusses the non-real-time forensics (see RFI-3 submission entitled "Attribution of Counterfeit and Subverted Integrated Circuits").

Method – By utilizing "Physically Unclonable Functions" (PUFs), the measurement of small manufacturing variations from one die to another and the use of these measurements as a "chip fingerprint" [7] can be used to form a non-destructive, robust chip identification in near "real-time". When combined with the right infrastructure and high integrity design and implementation, small, evaluatable circuits can be designed and implemented in a small portion of each IC to reliably measure and communicate this fingerprint information, and to present a "challenge-response" protocol for even more sophisticated identification that would reveal no information that could be used in a fingerprint playback attack.

Pieces of this concept have been under development over the past few years in the open literature[8]. This proposed work is to accomplish the refinement of these concepts and integration into a "JTAG-accessible" function block that can be universally inserted into almost any digital and/or mixed signal IC (using industry standard interfaces requiring no additional I/O pads or pins). These concepts are to be prototyped in FPGAs and examined together with industrial partners who must assess a business case for inclusion in their products. Follow-on work would include the incorporation of industry requirements and subsequent standardization required to inject this technology into a broad number of COTS products. Independently of this work, programs to evaluate a wide sampling of ICs for subversion will augment the deterrent enabled by the forensic capability targeted by this method.

Dream team – Collaboration between Sandia's Information Assurance experts and Sandia's Microsystems and Engineering Sciences Applications (MESA) Lab, with both "Fabless" and "Pure-Play" Integrated Circuit Manufacturers to standardize and broadly implement this method when mature (e.g., NVIDIA, XILINX, National Semiconductor, IBM, etc. and industry consortium such as the Global Semiconductor Alliance (GSA).

---

[7] G. E. Suh and S. Devadas, *Physical unclonable functions for device authentication and secret key generation*. In *Proceedings of the 44th Annual Conference on Design Automation* (San Diego, California, June 04 - 08, 2007). DAC '07. ACM, New York, NY, pp. 9-14. 2007.

[8] Ying Su, Jeremy Holleman, Brian Otis, "A Digital 1.6pJ/bit Chip Identification Circuit Using Process Variations," *IEEE J. Solid-State Circuits* , Jan. 2008.

Communication Policy Enforcement in Trusted Hardware

RFI Name: RFI-3 – National Cyber Leap Year
Category: Policy-based Configuration/Implementation
Title of Concept: Communication Policy Enforcement in Trusted Hardware
RFI Focus Area: Morph the gameboard
Submitter's Contact Information:

Lyndon Pierson
Sandia National Laboratories
MS 1072
1515 Eubank SE
Albuquerque, NM 87123

Who we are – Sandia National Laboratories' network operations and cyber security R&D groups. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Lyndon Pierson, Rob Armstrong, idea leads; Ben Cook, coordinating POC for Cyber Leap Year RFI,)

Vision – This effort is to develop information protection systems robust against high exposure to nation-state level adversaries while implementing policy-based sharing of this information among dynamically specified "communities of interest" and well-authenticated individuals with valid "need-to-know" for the information. We envision an architecture and an implementation that would eventually enable email and/or xml transactions between high level and low level systems (JWICS addresses and NIPRNET addresses, or between widely disparate environments for command and control nodes, or for coordination of emergency services and emergency warnings between different agencies with different jurisdictions, for example). An implementation of this architecture would assure robust authentication and attribution of originators, recipients, and intermediate processes, and of transformations allowed by policy, such as decryption for delivery on a protected distribution system, encryption for delivery on an unprotected distribution system, and even message handling policies, such as "originator controlled (ORCON)", and handling procedures for protection against increased sensitivity due to aggregation of information, etc. This Trusted Object-Oriented Hardware Architecture (TOHA) would involve high assurance methods to process data objects that include specification of how to communicate the data via especially trusted hardware. Within this hardware lies the means of identifying and authenticating a trustworthy hardware platform to policy servers and to users/clients, securely configuring the hardware with the cryptovariables required to open the data objects received and to act on the enclosed instructions with extremely high integrity. Also associated with this architecture is a "policy checker" to assure that a collective policy distributed into the trusted hardware nodes achieves the intended results even in an evolving environment of new policy, information regarding compromised nodes, new attacks, etc. This policy checker would operate both to assure correctness of global policy with respect to global objectives, and also to qualify local policy modifications as consonant with the global policy before implementing a proposed local policy specification.

To apply this vision to high consequence systems, this work would employ and build on methods of detecting and deterring the insertion or substitution of hard to detect subversions of Integrated Circuits (ICs) proposed separately, and on methods of countering unknown residual vulnerabilities using compositions of diverse implementations in fault and vulnerability-tolerant systems also proposed separately.

Concept – The TOHA would use a data-centric processing flow model, whereby context, status, and security information are kept with the message object itself. This removes unnecessary steps and simplifies the work done at particular process nodes in the system, thus, eliminating excess traffic. The TOHA elements can wrap and unwrap context, provide security, and reduce processor overhead and traffic between nodes. Communication between TOHA nodes would be performed using IP (Internet Protocol) so the TOHA nodes are connected between the jurisdiction's enterprise[1] and the Internet, and the TOHA node itself forms a secure divider from the Internet.

Unique benefits of this work include:

- Nonrepudiated Message Admission. This ensures the message was actually sent by the entity who claimed to send it.
- Transfer of Trust. As the message flows, each TOHA element can be trusted as the delegated information carrier.
- State Encapsulation in the Message Data Object. Additional routing, security, and policy information is carried with the message.
- Enforced policy for message admission and for message delivery.

Method – This proposed work is to accomplish the refinement and integration of multiple concepts. Object processing and policy checking would be developed in conjunction with multiple collaborators. High integrity IC concepts might be prototyped in FPGAs and examined together with industrial partners (IC manufacturers) who must assess a business case for inclusion in their products.

This architecture and its implementation must have a concrete plan to protect against subversions, penetrations, and inadvertent disclosures[2]: 1) to counter hardware and software subversion (by attribution techniques proposed separately); second, to counter penetration (by composing multiple diverse implementations in a unknown-vulnerability-tolerant system), and third, to counter inadvertent disclosures (by implementing policy that can be enforced by automated high assurance means). To assure robust implementation of protection measures, some level of "adversary emulation" will be required (multiple teams on design, subversion, detection of subversion, iterating to improve protections).

Dream team – Collaboration between Sandia's Information Assurance experts and Sandia's Microsystems and Engineering Sciences Applications (MESA) Lab, with both "Fabless" and "Pure-Play" Integrated Circuit Manufacturers and with members of standards organizations such as OASIS to standardize and implement this method when mature.

---

[1] Here we try to make abstract the notion of various security "enclaves" in differing environments, in order to apply these architectural concepts to a wide variety of government and even civilian applications.

[2] Meyers, Philip A., Subversion: the Neglected Aspect of Computer Security, June 1980 Masters Thesis, http://csrc.nist.gov/publications/history/myer80.pdf

RFI Name: RFI–3 – National Cyber Leap Year
Title of Concept: Changing the Cyber Playing Field: Rebooting the Critical Infrastructure
RFI Focus Area (Learning the Game to **Raise the stakes**): What will it take: Preparing by Practicing
Submitter's Contact Information –
George S. Davidson, Sandia National Laboratories, NM 87185-1319

Who we are – Sandia National Laboratories' infrastructure and information systems analysis, network operations and cyber security R&D groups, plus our high-performance computing and informatics research organization. Sandia is a multiprogram laboratory operated for the U.S. DOE fulfilling a national defense and national security mission, with a rich history of applying science and technology to protect some of the nation's most critical information systems. (George Davidson idea lead; Ben Cook, coordinating POC for Cyber Leap Year RFI)
Concept – A cyber infrastructure that robustly reboots to a known and trusted state would raise the stakes for adversaries, a game changing move, by lessening the significance of an adversaries attack. Creating & demonstrating such a response will involve political/industrial coordination, reliable models of the infrastructures, and a way to practice (as in war games) recoveries.

While purely technological solutions excite our interest, the human element remains the essential component of a robust recovery strategy. We suggest research involving human-led cyber-war gaming supported by cyber-disaster & recovery simulations and optimizations to explore (1) how bad things could become, (2) how to train the leaders and the operational work force directly involved in the recovery, (3) how feedbacks can create waves of social problems impeding recovery, (4) how these same social conditions could instead be harnessed for calmer, quicker recovery, and (5) how to identify and temporarily suspend those policies and regulations that might become an obstacle to recovery. This approach is designed to manage wide-spread crisis and lead to full operational recovery.
Vision – Preparedness is the essential tool for recovery of 'failed cyber nations'; and no country is immune from such a fate. Cyber systems are so interconnected that recovery is likely to require preparedness from the local community up through the national levels. However, the ability to eventually restore the flow of bits may not itself fix coupled infrastructure failures (e.g., the electric power grid, the gas distribution system, and the banking system). Policy planners and those responsible for recovery will need sophisticated tools and models of human and technical systems for war gaming to evaluate and practice their responses.
Method – We begin with the assumption that catastrophic cyber attacks are likely to include elements of 'hot' war; specifically the loss of critical facilities, personnel, and expertise. Insider threats are likely to be activated, all of which suggest the possibility that even the most perfect of technical/engineering protections will not guarantee the survival of the US infrastructures. The capacity to 'reboot' these infrastructures as quickly as possible could determine the fate of such conflicts, and provide for the security of our population and its economy.

We further note that the problem is not just technical, but will involve people. Certainly, longer delays in restoring minimal infrastructure functionality are going to be associated with larger social disruptions, and a greater chance for adversaries to exploit that chaos, too. These social effects are likely to be the most critical of the post event problems, and may be the least likely to have been considered if our starting stance is to approach a cyber attack as just a technical problem, which only requires clean, sharp engineering responses.

In considering how to change the game significantly, we were moved by successes in other fields. For instance, war fighters, police, and pilots are all trained under realistic stress conditions while practicing their responses. We conclude that cyber recovery requires a similar preparedness, together with regular practice; again at all levels.

We propose that changing the game would involve activities of three kinds. First, much of the actual infrastructure and discovered weaknesses must be protected at the classified level. Further, many policies and recovery plans will also have classified components. As such, these elements will be included only in simulations at secure facilities where classified planning and gaming events can be conducted.

Second, the best minds in the world must be brought to bear against the difficulties. Research must be encouraged with the objective of a fuller understanding of the issues (especially, those involving crowds and perhaps splinter groups). Better system-wide simulation techniques would inform improved engineering approaches and as well as the risk management strategies associated with the as-built infrastructure. Thus, academic research will be encouraged to explore the problem space, and to train the required work force and future leaders.

Third, we envision giving the American people the education, tools, and responsibility to plan and to participate in our collective defense; which is the best kind of prevention. This population wide awareness could begin through open gaming environments (think cyber-risk, or cyber-sim-city) possibly in stand-alone configurations or freely embedded in other games. The better the population has explored the implications of these cyber attacks, the more aware and responsive they are likely to be in real life. One challenge will be how to present realistic, open simulations without disclosing actual vulnerabilities and other classified aspects.

Attention to historic events must be cross-cutting theme to ensure simulations & scenarios reflect actual behaviors under relevant conditions gleaned from history. Regional and smaller historical events will be important because there are more small events extreme ones such as the 1919 flu pandemic or the Black Death, circa 1347. We, also, expect it to be important to understand the potential for violence and looting following a cyber attack (consider the looting in Iraqi or of war stores in Japan within hours of the Emperor's radio message ending WW II).

Dream team – Expertise is available throughout the DOE laboratories; the DHS National Infrastructure Simulation and Analysis Center; the USC-led National Center for Risk and Economic Analysis of Terrorism; and the CNCI Initiative Centers.; and the anticipated, DARPA National Cyber Range. The private telecommunications/ISP sector (e.g., AT&T) will be critical to any cyber recovery and should be involved from the beginning. Legal aspects of cyber events and recovery strategies must necessarily be addressed by experts, such as Lawrence Lessig (Center for Internet and Society). The Electronic Frontier Foundation and the Center for Unconventional Security Affairs at UCI can provide insight into public perception and unconventional security challenges associated with cyber attacks.

The modeling of cyber disruptions can build on SNL-developed scalable-simulation methods for war gaming WMD attacks, and modeling of cross-border vehicular traffic, and the economic consequences of disruption to critical infrastructures. Relevant, on-going & proposed work includes operation of the DOE National SCADA Testbed; simulation the functioning of cyber assets and the consequence of their failure; work on DARPA's Cyber Test Range BAA; large scale emulation of networks and hosts, scalable modeling (using parallel high performance computers) of economic interdependencies between critical infrastructures (undertaken as part of DHS NISAC); scalable solution of large integer programming problems (used for solving resource allocation problems); and scalable modeling of populations in large urban areas.

Transformation of Cyber Response through a Federated Defense

Who we are – Sandia National Laboratories' network operations and cyber security R&D groups. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Declan Rieb and Bridget Rogers, idea leads; Ben Cook, coordinating POC for Cyber Leap Year RFI)

Game-Changing dimension – Change the stakes for attackers and defenders.

Concept – A cyber-response plan needs to be developed for large-scale attacks on critical infrastructures to ensure our national security. Federated intrusion response is an idea to help bridge the gap between the top-down structure of CNCI Initiative 5 ("Connecting the Centers") and the more bottoms-up requirements of the implementers of intrusion detection to communicate with each other.

Vision – Our vision is of effective coordinated response that thwarts sophisticated attacks, enabled by defenders from different enterprises seamlessly communicating with one another, as if they were part of one large virtual organization. This vision requires incident sharing standards / technology; organizational transformation to overcome reluctance to share; in short, collaborative tools and the inter-organizational mores to help support a defense strategy for the federation.

Before an operational plan like this can be developed a smaller scaled focused plan needs to be developed and vetted. This plan needs to contain considerable detail and define successive stages of action and information sharing, and in principle should rely on an overarching cyber strategy. Today, typically little coordinated action or information sharing exists among sites let alone between industry and government.

Sandia and the DOE complex are a large target for attackers. We have been the target of these kinds of attacks for many years. In response we have informally set up a system of defenders who communicate and try to warn other DOE sites when we see something awry. Capabilities include SIN (Sandia Incident Notebook), and the DOE-sponsored CIRC/CIAC "SILC" server, and newer efforts such as G8 and Headsup. DOE/NNSA have proposed VCCIRT (virtual collaborative cyber incident response team) that can become the prototype for a national federated defense. There are many heterogeneous systems in place throughout the complex, all of which aid in mitigating attacks. The DOE complex would be a great exemplar to develop and test a focused plan and determine the necessary details the nation could use if it were under attack. This model will promote the growth of a dispersed community of network security analysts

interacting collectively to achieve a more secure information infrastructure for DOE/NNSA. These can be built up from the existing capabilities as well as informed by and support results from CNCI efforts.

Method – This concept has been under development by Sandia's network operations group in consultation with others in the DOE complex and internal research collaborators. The concept was summarized for this RFI in coordination with the larger team of Sandia researchers that is responsible for two other submitted Sandia concepts, advanced informatics and super resiliency. The technologies being developed for VCCIRT will support a flexible data-sharing relationship and accommodate precise and sophisticated controls for access to shared resources. The technologies being considered are based on a grid architecture and federated databases utilizing identity and access management systems.

Dream team – DOE complex (including weapons, science and energy laboratories and plants) and other government agency partners. Cooperation with the CNCI Initiative 5 "centers."

Transformational Cyber Event Discovery and Attribution through Highly-Scalable Advanced Informatics

Who we are – Sandia National Laboratories' high-performance computing and informatics research organization. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Suzanne Rountree, idea POC; Ben Cook, coordinating POC for Cyber Leap Year RFI)

Game-Changing dimension – "New game" employing transformational research (that also morphs the game with adaptive technology, changes the rules, and increases the risk for attackers).

Concept – Our local and national cyber infrastructure increasingly is vulnerable to stealthy intrusions that occur over time, coordinated intrusions to multiple sites/systems, and self-morphing malware – all of which must be detectable within massive amounts of streaming cyber data. What if we forced the intruders into a new game that coordinates multi-level, multi-timeframe event detection and attribution by integrating rapid response with scalable time-based event discovery and prediction?

Vision – The vision is for transformational cyber event detection (and the "discovery" of potentially related intrusions over multiple computing sites/systems), prediction of future intrusions, anticipation of consequences, deterrence and/or recovery, and attribution to the source – that integrates a multi-level approach combining triage of streaming data events, adaptive learning to better inform the triage criteria, and scalable informatics for a more static and detailed intrusion analysis, combined with innovative high-performance computer architectures for advanced data analysis.

This transformational informatics and visual analytics research can be integrated with a multi-tiered cyber strategy that includes resilient and coordinated response through federated defense (described in other Sandia proposals).

We intend to transform cyber security and cyber analysis by integrating a multi-level approach that includes:

- Fast heuristic approaches and algorithms for rapid initial processing to triage streaming cyber data in real-time;
- Scalable and adaptive learning methodologies to improve the heuristics and algorithms and more effectively triage the streaming cyber data;
- Event/relationship-based cyber analysis that focuses on a more detailed analysis of large-scale, time-based collected data using scalable analysis algorithms, information visualization, and large-scale data systems; and

- New computing architectures for enabling informatics in high-performance computing to effect orders-of-magnitude improvement in advanced informatics analysis.

Sandia has ongoing informatics research in “large data problems,” addressing not only new analysis algorithms R&D and R&D into alternative hardware architectures, but also research into scalable systems for storing, retrieving, processing, and managing large quantities of data from numerous sources – along with associated security, multi-classification handling, data uncertainty, and data privacy issues. Integrating novel data machines, along with specialized alternative architectures for advanced analysis, into an overall systems architecture is a challenge that (if met) will provide an end-to-end capability from raw data storage and processing to fast analysis and visual presentation of results to analysts. We have demonstrated a prototype capability focused on cyber security. As we gain better understanding of scalable capabilities for analyses through our current research, the next challenge is to move into scalable real-time analyses to detect, predict, and deter unwelcome cyber situations and anomalies as they occur.

Method – This concept was developed by a Sandia team of informatics researchers working in partnership with Sandia’s network operations and cyber security R&D groups. The informatics team has worked for over a year to understand the opportunities and challenges of extending scalable informatics to perform revolutionary cyber analysis. This concept was refined in consultation with the larger team of Sandia researchers that is responsible for two other submitted Sandia concepts, super resiliency and federated defense.

Dream team – Industrial partners (in computer architectures, in large-scale and real-time information search-and-retrieval systems, in scalable high-performance data/database systems), Federal R&D agencies with national security and cyber challenges, university partners (in statistics, in math- and graph-based algorithms, in scalable data analysis), national and international academics.

Transformation of the Cyber Infrastructure through a Super Resiliency Cyber Architecture

Who we are – Sandia National Laboratories' network operations and cyber security R&D groups. Sandia is a multiprogram laboratory operated for the U.S. DOE with national defense and national security programs. Sandia has a rich history of applying science and technology to protect some of the nation's most critical information systems. (Julie Perich, idea lead; Ben Cook, coordinating POC for Cyber Leap Year RFI)

Game-changing dimension – A morphing of the game board to permanently and adaptively change the defensive terrain of cyber.

Concept – A super resiliency cyber architecture will provide the foundation for implementing hardened, survivable systems that can deflect, morph and continue to operate through attacks. What if we changed the game board, and kept changing the game board so that it really didn't matter if adversaries attacked?

Vision – The vision of a "Super Resiliency Cyber Architecture" is one of transformation through informed cyber science and engineering. A super resiliency architecture (SRA) will provide a strong foundation for implementing a secure information and computing environment. SRA would result in a network architecture where system security is a design criterion just as important as functionality, reliability and speed. The architecture will consist of a series of elements that would respond to current and evolving cyber threats through a single operational environment. Development of such architecture would use emulative computing techniques and adversarial assessment (red teaming) to evaluate and refine performance under stress.

This type of stratified system design would ensure confidentiality of information through trusted pathways, while addressing integrity and authenticity of user, application, operating system, hardware, protocols, etc. It is a way to ensure independent elements of a system work together by each having their own security criteria *and* each element working with the other to monitor access and protect information.

Key SRA elements that we have considered in our notional architecture include but are not limited to the following:

1. Trusted Anchors: trusted out-of-band monitoring and management systems that remove the ability of adversaries to circumvent or disable security.
2. Ubiquitous Type-1 Encryption for Data at Rest: extension of Type-1 encryption to large-scale storage systems for sensitive information.

3. Increased System Inertness: a comprehensive response that inhibits the illegitimate use of extraneous functionality in COTS hardware and software.
4. Deception Layer: a proven network layer that detects and redirects attacks to a defender-controlled environment.
5. “Opt-in” Attribution: tools and techniques, including a new protocol, to provide inherent, verifiable network traceability.
6. Sophisticated models, methods, and tools that mimic real-world adversaries: development and application of the models, methods, and tools to test, validate, and improve SDA systems.

We recognize that SRA will depend upon and indeed must leverage advances in the scientific rigor of cyber security, e.g., developing the necessary knowledge that will allow us to design and engineer key components to be analyzable for security flaws (so that both hardware and software can be inspected thoroughly) to ensure trustworthiness. We also recognize that elements of SRA as outlined here have been proposed in the past and in some cases may exist today in various stages of maturity. However, no comprehensive, systems-oriented design and engineering has been undertaken to realize SRA, and this must be done for us to move forward and restore confidence in our systems.

Our concept for SRA goes even further than what is sketched above, encompassing a multi-tiered cyber strategy that transforms the underlying information infrastructure to be more resilient through advancements in the science and engineering of cyber security, increases our broader understanding of the cyber threat environment through information integration (described in Sandia’s second idea for event discovery and attribution through scalable informatics), and allows us to be effective in response through coordination (described in Sandia’s third idea for a federated defense).

Method –
Recognizing the escalating importance of protecting the nation’s cyber infrastructure, a group of Sandia operational and research staff have spent considerable time over the last year exploring new, holistic architectural approaches to securing critical information systems. This concept has been abstracted from that group’s recommendations and then coordinated with the larger team of Sandia researchers that is responsible for two other submitted Sandia concepts, scalable informatics and federated defense.

Dream team – Industrial partners in computer hardware and software, federal R&D agencies with national security and cyber challenges, academic partners conducting cutting-edge research in cyber security, machine learning, predictability theory, and related areas.

Who you are -- We are computer science faculty and security researchers (Stefan Savage, Vern Paxson and Geoff Voelker) at the University of California. For the last four years we have run the Collaborative Center for Internet Epidemiology and Defenses (an NSF CyberTrust center project) and we have direct research experience with the technology of large-scale cyber-attackers, the underlying threats they target and the economic framework in which they operate.

Game-changing dimension – Change the rules *and* raise the stakes.

Concept – Our concept is very broad but is decidedly game-changing: rather than defend our systems only on the *technical fronts* defined by our adversaries (i.e., updating our defenses in response to each new attack) we argue that we should also focus on the *economic fronts* defined by our attacker's underlying value-chain. To put it another way, compromising any particular computer is usually only a small (and relatively low-value) part of the attacker's overall enterprise. Instead of just trying to defend all machines, we should focus on the high-value points within the attacker's value-chain; the points that most **undermine their ability to profit.** For example, rather than simply creating ever better defenses to prevent the sending or reception of spam e-mails, we should instead focus on quantifying the attacker's return on investment and then take the most efficient actions to undermine their profitability; (e.g. by reducing the revenue produced by responders, increasing the overhead of extracting that revenue or increasing the costs of hosting spam-advertised pages).

Vision – The structure of today's cybersecurity game has long revolved around a small set of *technical battlefronts* (e.g., anti-virus, anti-spam, intrusion detection/prevention, software defect analysis, firewalls, authentication/encryption products, etc.) with comparatively little effort focused on understanding the adversaries themselves. However, the playing field has changed dramatically over the last decade and the sophistication and scale of cyberattacks now are driven by an underlying *economic value-chain* that is largely orthogonal to the particular technical modes of attack. We believe that investigating and addressing the vulnerabilities in the attacker's economic value-chain is likely to offer a broader and more successful security capability than simply playing catch-up with the software vulnerability flavor-of-the-month. By analogy, our military does not train against a hypothetical adversary with hypothetical resources, strategies and interests, but instead we design our forces and stratagems based on what is known about real and potential adversaries. Our vision has four parts:

1) *Characterization*: We need to first measure and analyze the underlying value proposition in cybercriminal actions, typically focused on monetizing a service or stolen information. For example, for "carding" scams (the monetization of stolen debit/credit account info) this involves analyzing the costs for stealing financial information directly (e.g., via phishing, skimmers, banking Trojans, e-commerce compromise, etc.), the commissions on third-party monetization (e.g., via mules, resale to underground markets, etc.) and the rate of return as a function of these methods and ecosystem.

2) *Bottleneck identification*: Typically we focus our security efforts on the end host by installing host (e.g., antivirus) or network (e.g., firewalls) defenses. However, this is arguably the most expensive asset for us to defend (hundreds of millions of hosts and users, thousands of security vulnerabilities, etc.) while it incurs the least impact to the attacker's value chain since the attacker can typically choose to evade the defenses or just target those hosts without them. We believe that our security efforts must also include identifying the economic bottlenecks that most significantly impact the value delivered to our attackers. Continuing the previous example, there is strong anecdotal evidence that "carding" is limited by the ability to monetize stolen accounts at scale; there are far more compromised financial accounts than there appears to be ready capacity to liquidate those accounts quickly.
3) *Economic disruption*: Analyzing the economic bottlenecks provides the guidance to focus on the appropriate technical defenses and offensive/enforcement actions to undermine the attacker's value proposition. For example, in the context of carding again, it is probably difficult to prevent phishing, spyware infestation, or e-commerce server compromises, but making more extensive use of one-time credit card credentials could minimize the value in stealing these assets to begin with. Similarly, infiltrating and undermining the markets in which these accounts are bought and sold could make it more expensive to monetize accounts at scale.
4) *Success metrics*: In general, *measuring security* in any comprehensive way has long been a grand challenge research problem. However, by focusing on the attacker's point of view this challenge can be neatly side-stepped. In our model, the efficacy of any intervention can be directly measured based on its impact on the attacker's return-on-investment. Thus the same measurements used to characterize our adversar's value proposition can be used to measure the effectiveness of our intercessions.

Method – It is widely observed that the effective monetization of cybercrime over the last five years has become the engine driving dramatic advances in the scale and sophistication of Internet-based cyber attacks. However, while it is understood that these attacks only take place because they result in profit, it is only recently that there have been significant efforts to characterize and quantify these value chains. Our proposal is informed by this emerging sub-field, at the interaction of computer security and economics, that includes both our own team's work on the dynamics of underground markets (Franklin et al, CCS07) and the return-on-investment for spam campaigns (Kanich et al, CCS08), as well as that of other researchers investigating topics including spam-based securities manipulation (Boehm et al, WEIS06), phishing (Moore and Clayton, eCrime07,eCrime08), and AV scams (Stewart, SecureWorks08). This field is only in its infancy, but we believe that these initial results demonstrate that it is possible to quantify significant portions of the economic value-chain behind complex attacks.

Dream team – Computer security researchers in both academia (e.g., ourselves, Tyler Moore, Nick Feamster, etc.) and industry (e.g., Joe Stewart, Jose Nazario, Dan Hubbard, Rob Thomas, etc.), Economists, electronic payments industry representatives (e.g., Paypal, MC/Visa, Banks, Merchant associations), legal scholars, FTC and Law Enforcement agencies.

Who I am – Martin Schulman, twenty year Internet veteran and former Chief Technologist for Juniper Federal Systems.

Game-changing dimension – Morph the gameboard, or more colloquially: "add optional gutter guards for novice users of Internet bowling alleys".

Concept – Define an optional, backward-compatible, policy enforceable URL checksum that stops minor typographical and memory errors from enabling typo-squatters[1].

Vision –Any frequent computer user has misremembered or mistyped a URL or email address. While it usually results in "Page Not Found" error on the web, returned message, or other innocuous result, it can sometimes lead a person to annoying ads, malware, or content inappropriate for younger surfers posted by typo-squatters at common misspellings of popular web sites. For a while, innocent visitors who meant to visit "whitehouse.gov" but instead entered "whitehouse.com" were greeted with pornography.

DNSSEC and increased use of server certificates or encrypted email will not address this problem, and may over time lull users into believing their legitimacy. Users might see a site like "www.redxcross.org" uses DNSSEC and SSL certificates and assume they are the well-known international relief organization when in fact they are a scam – credit card information entered to make donations will simply be sold to the highest bidder.

Misspellings could be lessened if domain names were required to be sufficiently distinct, but it's too late for the registrars to require now, and even if it wasn't they couldn't reasonable police what is legitimate. Left to defend against typo-squatters, some organizations resort to preemptive registration of obvious domain name variants or lawsuits, but both can be cost prohibitive for smaller entities and ineffective when spanning international boundaries. The introduction of new TLDs will increase both the potential for squatting and the cost of this defense.

This proposal is for an optional extension to URLs of a few characters algorithmically derived so that minor omissions, insertions, and transpositions are unlikely to produce the same result. Suppose for example the check consisted of 3 numerals appended to the usual URL. The Whitehouse would offers its web site address as:

www.whitehouse.gov.562

Users who enable this anti-typo feature who enter "www.whitehouse.*com*.562", "www.*whithouse*.gov.562", "www.whiteh*uo*se.gov.562", or any other mistake would be warned by their browser that the URL is invalid and they should check the spelling.

Similar error correction is required in other common identifiers like credit card numbers and barcodes, but it's important to emphasize this check must be optional - both to

[1] http://en.wikipedia.org/wiki/Typosquatting

facilitate incremental deployment and to keep experienced users and embedded systems from requiring change. As deployment becomes widespread, system security policies may be configured to require the checksum for all manually entered URLs. Businesses, people who administer PCs for their parents or children, and other users may require it, just as they might require anti-virus software or limit software installation.

This example also illustrates one way backward compatibility could be achieved by making the checksum a numeric extension appended to the FQDN, where it cannot be confused with a (currently legitimate) TLD. Other approaches are possible – for example, using three numerals in place of "www".

This does not completely solve the problem, as people may still typo-squat and game search engines to redirect unsuspecting users to their site. Hopefully ranking pages by traffic will minimize their effectiveness.

Method – McAfee's "The State of Typosquatting 2007"[2], "Cyber-Fraud is One Typo Away"[3], and informal anecdotes, along with familiarity with error control coding prompted this suggestion.

Dream Team – Would consist of subject matter experts in the following areas:

- Error Control Coding – to identify algorithms that ensure different checksums for strings with low Levenshtein distance
- CHI – to asses what and how many characters to use and where to place them
- Browsers – to assess integration into the major HTML rendering engines
- DNS – to ensure compatibility with existing records and emerging protocols
- Operating System providers – to verify compatibility with and discuss applicability to other applications

Martin Schulman

[2] http://us.mcafee.com/root/identitytheft.asp?id=safe_typo&cid=38296

[3] http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&arnumber=4509853&isnumber=4509595

Who You Are
A PhD mathematician (Scott Guthery) and a PhD university professor (Mary Cronin) in Boston, Massachusetts

Game-Changing Dimension
Analog computations to secure digital information

Concept
Safeguard bits by irrevocably binding them to analog properties of physical objects

Vision
Things are easier to keep track of than bits.

If bits are irrevocably bound to a thing then …

1) if you know where the thing is, then you know where the bits are
2) if you alter the thing then it is a new thing and the bits are lost forever

Method

Merge the computational capabilities of the analog properties of a physical object with the bits to be safeguarded so that the computation must complete correctly in order to achieve an error-free recovery of the bits.

Predecessor work along this line is called a Physical Unclonable Function (PUF) [1]. This RFI proposes to actually compute with the analog properties of physical objects rather than simply use these properties for unique identification as is the case with PUFs.

Shine a laser light through a bowl of Jell-O. The Jell-O performs a particular computational transformation on the light. If the Jell-O is altered however slightly, the computation will change. Furthermore, it is virtually impossible to create a bowl of Jell-O that performed exactly the original transformation.

In the proposed method, the bits in hand correspond to the laser light, the bowl of Jell-O is the physical object, and the light coming out of the bowl are the safeguarded bits. The safeguarded bits are irrevocably bound to a particular bowl of Jell-O. The Jell-O performs an analog computation on the bits at hand to create the safeguarded bits. [2]

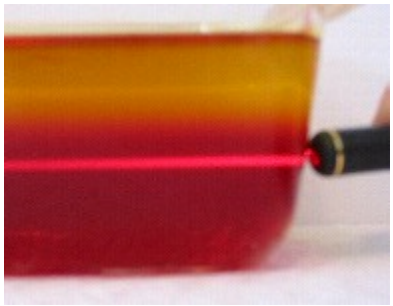

A finite state machine is another familiar example of transforming bit streams. It has been shown that state transitions in finite state machines can be bound to analog phenomena [3, 4, 5]. This is but one of many ways that bits can be irrevocably tied to things.

In general, analog computing [6, 7] is currently an overlooked source of cheap and efficient computing power. By providing examples and applications of analog computing to digital data security, this project will demonstrate that analog computation can make effective and economically compelling contributions in today's digital computing networks.

Dream Team

The people on the submitting team can do the mathematics and envision the use cases and applications. The dream team would need to include a person with deep understanding of each analog property to be harnessed. Examples of such a person would be an applied physicist, a chemist, a material scientist, a biologist, or a sonic engineer.

[1] Pappu, Ravikanth, Ben Recht, Jason Taylor, Neil Gershenfeld, "Physical One-Way Functions," *Science*, Vol. 297, September 20, 2002, pp. 2026-2030.

[2] Mills, J., N. Miller and J. Nakamura. "The Jell-O® Brand Gelatin Processor: A Prototype Colloidal Computer," B644 VLSI Design. Indiana University, 2004

[3] Brockett, R.W., "Dynamical Systems that Sort Lists, Diagonalize Matrices and Solve Linear Programming Problems," Proc. 27th IEEE Conf. Dec. and Control, Austin, TX, pp. 799-803, Dec. 1988

[4] Herrera, Ruben and Rahul Sarpeshkar, "Spike-Triggered Asynchronous Finite State Machine," US 6292023, September 18, 2001

[5] Gy. Turan, F. Vatan, "On the computation of Boolean functions by analog circuits of bounded fan-in," *Journal of Computer and System Sciences* 54, pp. 199-212, 1997

[6] Shannon, Claude, Mathematical Theory of the Differential Analyzer. *J. Math. and Physics*, 20 337–354, 1941

[7] Rubel, L. The Extended Analog Computer. *Advances in Applied Mathematics*, 14, pp. 39-50, 1993

Function Extraction Technology:
Computing Software Behavior for Software Assurance and Malware Analysis

RFI: RFI-3--National Cyber Leap Year

Focus Area: Change the rules

Submitter: Richard C. Linger, Principal Investigator

Contact: CERT, Software Engineering Institute, Carnegie Mellon University
4500 5th Avenue,
Pittsburgh, PA 15213

Credentials:
The CERT organization of the Software Engineering Institute at Carnegie Mellon University is a major contributor to technologies and methods for software assurance and malware analysis. Richard Linger is Technical Manager of the Survivable Systems Engineering group in CERT, where he directs the Function Extraction (FX) project. He is an expert in the function-theoretic foundations of software behavior computation, having contributed to much of the requisite technology over the past 20 years of research and development. While at IBM, he focused on rigorous software engineering technology for developing and certifying high reliability systems.

Concept:
The national economy and defense are dependent on software systems of ever-increasing scope and complexity. Yet these systems continue to experience errors and vulnerabilities despite best efforts, and present costs and risks that are difficult to control. It is a reality of present-day software engineering that programmers have no practical means to determine the full functional behavior of software. This technology gap has been the source of many software problems experienced by DoD for decades. It is vital for effective engineering and management to know all of the behavior of software, whether intended or unintended, benign or malicious. While current tools can help analyze specific properties of software, what is needed is an "all cases of behavior" understanding of what software does.

CERT is closing the gap by developing the leap-ahead technology of Function Extraction (FX) to calculate the behavior of software with mathematical precision to the maximum extent possible. Computing the behavior of software requires deriving its net functional effect, that is, how it transforms inputs into outputs in all circumstances of use, without heuristics or approximations, to define "all cases of behavior." The objective is to move from slow and fallible human analysis of software behavior to fast and correct computation of behavior at machine speeds. Controlled experiments showed that users of an FX prototype were several orders of magnitude faster than users of manual methods in determining software functionality, and that programmers were about 15 times more productive in analyzing software when using computed behavior.

Vision:
As a first application, CERT is developing a Function Extraction system that computes the behavior of programs written in or compiled into Intel assembly language, with a principal focus on improving the speed and accuracy of malware analysis. FX technology can be applied to

other languages and processors, and can provide automated support for many software engineering activities. The behavior databases produced by FX can provide input to other value-add applications. The following examples illustrate application of FX across the life cycle:

Software development
- Programmers get immediate feedback on functionality as code is written.
- Errors, vulnerabilities, cost, time, and risk are reduced.

Software verification
- All behavior is computed -- intended, unintended, benign, and malicious.
- Checking behavior against requirements reduces testing and improves quality.

Malicious code analysis
- Fast determination of malicious intent enables fast countermeasures.
- Current asymmetric advantage to intruders swings to the good guys.

Legacy system evolution
- Computed behavior reveals legacy functionality as a basis for evolution.

Anti-tamper analysis
- Defense systems can be checked for malicious content and corrupted functionality.

Pedigree and provenance
- Computed behavior lets the software speak for itself, no matter its origin or history.

Software quality
- Software without computed behavior becomes suspect, too expensive and risky to use.

Method:
FX technology is based on the mathematics of denotational semantics expressed in terms of sets and functions. It is implemented in terms of extensive function composition augmented by function-equivalent abstraction of computed behavior for fast human understanding and analysis. A structure theorem defines a constructive process for transforming complex, spaghetti logic into algebraic structured form, and a function theorem defines the starting point for transforming procedural logic into non-procedural as-built specifications of behavior. Research breakthroughs include methods for definition of the functional semantics of the 1100+ op codes on the Intel chip, creation of a behavior expression language with only a single statement form, major advances in control flow determination and structuring theory, and mathematical foundations for computing the behavior of loops that makes the effects of theoretical limitations arbitrarily small.

These advances characterize a new science base and technology to help transform software engineering into a computational discipline for the 21st century. In recognition of this potential, the FX project received SEI's highest award for creativity.

Dream Team:
The FX team at CERT is composed of six senior researchers with proven R&D track records and in-depth knowledge of the function-theoretic mathematics that are the basis for behavior computation. Development of FX theory and implementation is well along, but is not completed. The team will collaborate with sponsors to build out FX technology for their specific applications. Sponsorship is required to complete the FX research and its implementation.

Title: Virtualizing the Cyber Warrior

Our Company

Sentar, Inc.
4900 University Square, Suite 8
Huntsville, AL 35816

CAGE #: 0A168
DUNS #: 17-426-5736

Sentar, Inc. is a women-owned small business that specializes in innovating, building and securing network centric systems, employing experts in the field of information assurance and cyber security to provide thought leadership and service to military, government and commercial customers. To complement our exemplary staff, Sentar has several patented and trademarked technologies. Dr. Leigh A. Flagg, the Director of Research and Development, is an Information Systems Security Professional with 13 years of experience in heterogeneous systems integration as well as a New York University trained screen writer, film director and film producer who has co-written, co-directed and co-produced a feature length film.

Game-changing dimension

Our concept centers on leveraging non-traditional cyber warrior role and employment paradigms to enact cyberspace courses of action. Rather than pit military against military in cyber space, we propose empowering a "cyber posse comitatus" to police and enforce desired behavior. Civilian cyber warriors, guided by military cyber warriors, would be motivated to monitor and secure cyber space from within a virtual world. By changing the environment, roles, responsibilities, economics, reality, and social image of cyber warriors, we will *change the rules* of the cyber security game.

Concept

Our concept is to integrate a 3D virtual space into the infrastructure of the internet and the GIG. This would be akin to layering Second Life on top of a networked systems infrastructure. Avatars of military cyber warriors would act in the virtual world, defending people and places and neutralizing threats, whose effects would be carried out in the actual systems. The cyber warriors would include civilian cyber warriors to assist in policing and in defensive/offensive engagements. The employment of civilians would be modeled as a "job posting" paradigm such as those used in immersive, role-playing simulations or job auctioning sites, such as RentACoder.com, where people can bid on programming tasks. Civilian cyber warriors would be offered jobs by the military cyber warriors based on their experience, past performance, reputation, history, etc. in the virtual community. This will allow civilians to build trust with their "employer" much the same way they would in the real world. The result would be a population of civilian cyber warriors "crowdsourced" and guided by military cyber warriors, offering broad and large-scale defensive and offensive capabilities in cyber space.

In order to influence the behavior of civilian cyber warriors, payment and rewards such as real money, virtual promotions, virtual property, virtual business, virtual good and services, etc. would be offered in exchange for the work they perform. Furthermore, intrinsic rewards or punishments would be used to deter "bad" behavior. These could include a forced change in appearance of their avatar, an adjustment of virtual credit, a reduction in the trust status afforded their avatar, or turning other civilian cyber warriors against the offender. These economic and social rewards and punishments are crucial because they entice the populous to help while it applies measures to dissuade wrong doing. Thus, it influences behavior in a way that is desirable to the military cyber warriors.

Vision

Our vision is to re-direct some of the enormous amount of energy currently expressed in virtual world gaming toward the security of cyber space. The vision integrates the entertainment and cyber security worlds, enlists the talents of hackers, geeks and hobbyists currently doing both good and harm, and establishes a new and potentially very appealing career path for young military recruits. With this concept, the role and image of the cyber warrior can be changed. This includes their self-image, the image they display to others, and their role in the defense of our national interests. There is a vast group of untapped resources who up until this point have been considered too inaccessible if not downright nefarious to participate in cyber security. Our vision is to make these vital resources visible and useable such that attacks can be predicted, anticipated, charted, halted, prevented, and countered in real time.

Method

Invention process. The vision described here is the product of our extended research and development in knowledge-based systems for information assurance. It is also a result of our research and development in applying modeling, simulation, and training technologies to security. This includes extensive development of a set of advanced products for certifying and defending network systems as well as providing advanced, game-based command and control experimentation.

Assumptions and dependencies. The most significant assumption underlying this concept is that military cyber warfare activities would interface their systems and networks to a virtual environment and economy. Another significant assumption is that an appropriate virtual world can be chosen, implemented, or otherwise generated that would serve our needs. Furthermore, the idea assumes that the virtual world will assume enough intelligence and self-security to automate some of the job recruiting, hiring, tasking, and compensating responsibilities. This is necessary in order to allow for real-time cyber defensive and offensive action.

Dream team

David Tillman, Director of Advanced Training Solutions, Miltec. His team leverages the latest advances in commercial video game technology, applying these technologies to advanced learning solutions.

Anita D'Amico, Ph.D., Director of Secure Decisions, Applied Visions. Dr. D'Amico is both a human factors psychologist and an information security specialist. Her area of expertise is cyber defense situational awareness, particularly improving cyber defense decision-making through visualization.

Artificial Humans Group of the Intelligent Systems Division of the USC/Information Sciences Institute (ISI) at the University of Southern California. Major current thrust at CARTE is "socially aware learning environments," which emulate aspects of human social interaction in engaging with learners and which support social learning processes. They have a long-standing interest in pedagogical agents, and in developing "drama-based learning environments," environments that exploit dramatic structures and techniques in order to make learning experiences more understandable and engaging.

Linden Labs, the creators of Second Life™

Title: Fuzzying Reality in Cyber Warfare

Our Company

Sentar, Inc.
4900 University Square, Suite 8
Huntsville, AL 35816

CAGE #: 0A168
DUNS #: 17-426-5736

Sentar, Inc. is a women-owned small business that specializes in innovating, building and securing network centric systems, employing experts in the field of information assurance and cyber security to provide thought leadership and service to military, government and commercial customers. To complement our exemplary staff, Sentar has several patented and trademarked technologies. Dr. Leigh A. Flagg, the Director of Research and Development, is an Information Systems Security Professional with 13 years of experience in heterogeneous systems integration as well as five years experience developing game-based command and control experimentation aids and cyber security training technologies.

Game-changing dimension

Our innovation is to *morph the game board* by implementing multiple, competing virtual copies of critical systems to which we can move the cyber war and thus eradicate an attackers first mover advantage. This innovation will change the game of cyber warfare by proliferating the equivalent of mirror (virtual) images of our real mission critical systems and networks allowing us complete control of the environment including even the rules of the game. Here we can learn our adversaries tactics and directions. We can understand in real-time their tools and methods, but, most importantly, we can manipulate their reality. We can change how the world looks to them and behaves (e.g. systems and network characteristics can be manipulated). We can build a variety of controls into our virtual system/network decoys to enable us to accomplish changing the rules. In this environment, of mixed real and virtual copies of systems and networks, reality becomes Fuzzy for our adversaries. They can no longer be certain of the effects of their actions or even know if they are fighting in the right war.

Concept

Our concept of Fuzzying Reality incorporates development of simulated/virtual replicas of critical networks and systems that will run in tandem to their real counterparts, to which we can, unbeknownst to them, re-route and shuffle cyber adversaries during an engagement and manipulate their reality. Along with the process of keeping/moving attackers in the virtual systems we will also be hiding the real systems from them (e.g. playing a dynamic shell game). This will be accomplished by building on, and combining, the elemental concepts of: Honey Nets, simulated network-centric warfare, network-centric warfare effects and controls, and network attack deception methods modeled on sleight of hand.

An additional component of our Fuzzy Reality concept is the protection of our real mission critical system and network assets. This is akin to the use of Anti-Tamper methods to protect Critical Protection Items/Critical Technologies (CPI/CT). Once the virtual space consisting of the copies of our assets has been developed and integrated with the real one, we can employ a variety of sleight of hand type methods to continuously hide the real systems from our adversaries. We can enhance the integrated space by putting layers of protections in both virtual and real components. We can also modify the protection of the real assets as we learn what is happening in the virtual copies. A final advantage is that this whole integrated space can be used for training our cyber forces

Vision

As things stand today, the attacker in the cyber world has the first mover advantage. Due to the speed of events (seconds, minutes at most) in this space and the rapid evolution of methods and tools, the first mover advantage is overwhelming. It is very difficult (if not impossible) for the defender to identify, attribute, negate and counter a relatively sophisticated adversary (of whom there are many). Thus the case is often one where an attack is discovered long after the damage is done. In mission critical systems (e.g. real time weapon systems, national infrastructure, etc.) these scenarios pose grave and unacceptable risks. Our vision is to re-direct cyber war to a space where our adversaries cannot tell the real systems from virtual ones that we control and we, the good guys, have the upper hand.

Methods

Invention process. The vision described here is the product of our extended research and development in knowledge-based systems for information/knowledge management and its application to cyber security along with a lengthy history of engineering and development of military Command and Control Systems. It is also a result of our research and development in applying modeling, simulation, and training technologies to security. This includes extensive development of a set of advanced products for defending and certifying networked systems as well as providing advanced, game-based command and control experimentation. The technologies and development expertise will be augmented by more then 100 man-years of real world experience with Command and Control in both conventional and cyber warfare.

Assumptions and dependencies. The most significant assumption underlying this concept is the ability to scale up the integrated Real/Virtual meta-system. Other significant assumptions include the ability to manage this meta-system and the methods needed to implement the obfuscations at this scale.

Dream team

David M. Nichol, Ph.D., Professor of Computer and Electrical Engineering, University of Illinois-Urbana, and Theme Leader, Critical Infrastructures and Homeland Defense, Information Trust Institute. Professor Nicol has been studying the security properties of large-scale systems, examining large-scale system behavior and developing a modeling and simulation methodology that supports evaluation of that behavior.

Dave Gursky, MSIA, CISM, CISSP, Sr. Principal IA Engineer, Raytheon Company, Intelligence and Information Systems. Mr. Gursky is an Information Assurance Manager and researcher at Raytheon Intelligence and Information Systems in Crystal City VA. Mr. Gursky is considered a SME on Behavioral Based Intrusion Detection Systems, Multi-Level Security, Attribute Based Access Control and other IA technology.

David Tillman, Director of Advanced Training Solutions, Miltec. His team leverages the latest advances in commercial video game technology, applying these technologies to advanced learning solutions.

Title: Self-Managing Ontologies

Company:
Sentar, Inc.
4900 University Square, Suite 8
Huntsville, AL 35816

CAGE #: 0A168
DUNS #: 17-426-5736

Sentar, Inc. is a women-owned business that specializes in innovating, building, and securing network centric systems, employing experts in the field of information assurance and cyber security to provide thought leadership and service to military, government and commercial customers. To complement our exemplary staff, Sentar has several patented and trademarked technologies. Dr. Andrew Potter, Chief Scientist of Sentar, has twenty-five years experience in research and development, including principal investigator for numerous projects in semantic systems for cyber defense.

Game changing dimension

The Self-Managing Ontology (SMO) is a breakthrough technology that will enable the provision of self-sustaining dynamic ontological processes for use in assuring sovereign options in cyberspace. Changing the balance of power in asymmetric cyber attack will require evolvable reasoning systems. These systems will implement a dynamic information infrastructure for identifying, analyzing, and managing vulnerability risks within a rapidly evolving adversarial environment, for detecting attacks on at-risk assets, and for responding to attacks effectively. By fundamentally transforming the nature of semantic technologies from static structural encodings into a comprehensive constellation of dynamic adaptive processes, Self-Managing Ontology (SMO) will *morph the cyber security game board* sharply in our favor.

Concept

Evolvable reasoning systems require evolvable ontologies. However, current ontologies are static, inert structures, and this makes them difficult to evolve. SMO revisits the basic notion of what an ontology is—rather than treating ontology as a *structure*, we define it as a *process*. Viewed as processes, ontologies not only represent knowledge structures, they *are* knowledge processes, and the same mechanisms used to reason with ontologies *about* domains can be used to reason about ontologies *as* domains. By positioning ontologies thusly we can begin to address the problems of automated evolvable knowledge representation. By providing a technology that fully addresses this need, we can realize the long-held promise of this technology and apply it to an area of critical need, namely, reducing vulnerability to cyber attack.

Vision

Cyber warfare is constantly changing. This demands that our protective information infrastructure constantly evolve. That is why the self-managing ontologies are so important to reducing our vulnerability and enhancing our response to asymmetric cyber attack. Self-managing ontologies can be realized by taking an end-to-end approach to ontological functionality, construction, and utilization. This requires that ontologies know how to reason about themselves and about their relationships with other ontologies. This can best be understood within the context of the holy grails of ontology research: *adapt*, *repurpose*, *merge*, and *cross-domain information sharing*:

- *Adaptive and repurposing ontologies* are ontologies that can automatically update themselves as their environments change. These changes could entail changes in goals, in operating environment, or in domain knowledge structures. The adaptation process proceeds through several stages. First, the need for change is detected as a discrepancy between the ontology and its operating environment. The self-managing ontology responds by determining the scope of the change, identifying the affected concepts, relations, and other structures, defining metrics for delta transitions, and redefining goals as necessary. The ontology is then in a position to identify and evaluate options, using metrics and predicted outcome. Based on these values, the ontology then may reformulate, test, and assess; depending on the outcome of the test, the ontology may now redeploy itself.
- Self-managing ontologies *merge* by first identifying merge candidates and determining the need for a merge. This is followed by a process of identifying nominal consistencies and inconsistencies through semantic analysis. The identified consistencies can then be merged, and the inconsistencies resolved prior to merging. Finally complementary information contained in each ontology can be incorporated, and the merged ontology is then ready for assessment and redeployment.
- Self-managing ontologies accomplish *cross-domain information sharing*, first by identifying and normalizing their mutual interface points, then identifying their interface structures, and selecting and deselecting interfaces as needed to achieve their goals.

Method

Invention process. The vision described here is the product of extended research and development in knowledge-based systems for information assurance. This has included extensive development of a set of advanced products for certifying and defending network systems.

Assumptions and dependencies. The major assumption is that the ontological self-awareness necessary for an end-to-end self-managing solution can be achieved. This assumption will be tackled as an initial research question for the project. A second assumption is that self-managing ontologies can themselves be made secure. We see this as a bootstrap process, whereby the SMO technology is used to assure itself. A major dependency for the project would be that it will require non-trivial teamwork to achieve the shared vision necessary for success.

Dream team

Deborah L. McGuinness, Rensselaer Polytechnic Institute, works in the fields of knowledge representation and reasoning, description logics, the semantic web, explanation, and trust. She is also well known in the field for co-authoring the World Wide Web Consortium (W3C)'s recommendation for an Ontology Web Language - OWL.

Erik Mettala is Chief Scientist at Sparta, Inc. Prior to joining Sparta, he was Vice President and Director of McAfee Research.

Bruce Porter, University of Texas, Austin, has performed research leading to development of a new class of knowledge systems that are designed to be task independent and able to answer a broad range of questions within a domain.

Who you are Dr. Steve Dawson, Dr. Drew Dean, Dr. Hassen Saidi, Computer Science Laboratory, SRI International; Dr. Markus Jakobsson, Dr. Kurt Partridge, Dr. Ignacio Solis, Dr. Jessica Staddon, Computer Science Laboratory, Palo Alto Research Center (PARC). SRI International is a large nonprofit research institute based in Menlo Park, CA, with 60 years of historic innovations in computing, business, eduction, materials, and biosciences. PARC is an independent research business and a wholly owned subsidiary of Xerox Corporation. PARC has contributed to the creation of more than 30 companies and is celebrated for such innovations as laser printing, distributed computing and Ethernet, the graphical user interface (GUI), object-oriented programming, and ubiquitous computing.

Game-changing dimension Change the rules

Concept Unsolicited email (hereafter "spam") has been a problem for over fifteen years now with no solution in sight. The scope of the problem is increasing; current estimates have spam accounting for over 90% of all email crossing the Internet today[1]. More troubling is that spam has become a significant vector for the distribution of malware: while Dr. Dean's inbox is protected by 3 separate mail filters (including a commercial anti-virus product), a significant fraction of the spam that gets through the filters contains a malicious attachment. Indeed, the propagation of malware via email is estimated to have increased more than 250% in 2007 (IronPort's Threat Operations Centre).

The *practical* response of the Internet community so far has been to build increasingly complicated email filters (see the discussion under **Method** below for other research proposals). Organizations such as Spamhaus exist to maintain blacklists of bad email senders, updated and distributed in real time. However, given the impossibility of deciding what is spam (some email should be categorized as spam, except for a pre-existing relationship that the spam filter doesn't know about), there will always be the trade off between tuning a filter for false positives vs. false negatives. As the risk of false positives is very high, the filters must be tuned to minimize the false positive rate, which implies that some spam (false negatives) will always get through. The only lasting value of spam filters has been to provide social cover: "I'm so sorry, my spam filter must have eaten your email" is a completely plausible excuse for any unacknowledged message.

It has always been assumed that making fundamental changes to the email infrastructure is too hard, because of the $O(N^2)$ network effect. However, we argue that the dynamic nature of communication modes dramatically lessens the network effect. Consider the following facts:

1. The usage of FAX machines is declining, as many exchanges formerly done via FAX have moved to email[2].

2. USENET has seen extensive decline, overtaken by Web-based discussion forums[3].

[1] See, e.g. `http://sentra.ischool.utexas.edu/~adillon/blog/archives/127`
[2] `http://www.gartner.com/DisplayDocument?doc_cd=123711`
[3] `http://www.pcmag.com/article2/0,2817,2326849,00.asp`

3. Social networks (e.g., MySpace, Facebook, Bebo, Orkut, etc.), instant messaging, and cell phone SMS (texting) have started to displace email among teenagers and college students[4].

All of these technologies show that new communication networks rise and fall over time, even though the existing communications media have $O(N^2)$ network effects working against the new entrant. Given the above worked examples, the conventional wisdom that replacing SMTP or adding authentication will inevitably fail bears re-examining[5]. The conventional wisdom has resulted in an over-constrained problem space that unsurprisingly offers no solution. This is not acceptable.

Vision We propose redesigning email with a fresh set of eyes. The technical challenge to building a spam-resistant messaging architecture may well be the easy part of the problem: getting it deployed on a large scale is an even more formidable challenge. We need input from both the policy community, to make sure that the technical architecture has the proper value system embedded in it, and the economics community, to make sure that the architecture is incentive compatible with the needs of early adopters.

The era of "flag-day" changeovers on the Internet passed 25 years ago (with the transition to TCP/IP in 1983). Clearly a new messaging architecture will have to co-exist with SMTP-based email for a long time. However, we believe that if we can get a new architecture successfully bootstrapped with 1000 organizations, its value will rapidly become apparent and drive further adoption.

Method Many partial solutions to the spam problem have been proposed in the last 15 years, and deserve re-examination to see if they would fit into a new email architecture. We may or may not choose to replace SMTP; all options are on the table. In particular, authentication of email senders, while not without its own problems, may offer one avenue for progress. Authentication in and of itself is by no means a panacea, but it may enhance the capability of spam filters by giving them difficult to forge attributes to filter on. The modern (i.e. post-1990) version of *client puzzles* were reinvented to address the spam problem, but have not gained any traction. Other ideas that have merit have similarly not gained widespread adoption.

Dream team Computer scientists from SRI, PARC, Microsoft Research (Cynthia Dwork), and academia (e.g., Steve Bellovin, Ed Felten, Dan Wallach); anthropologists and ethnographers from PARC and other research labs; policy experts and economists, as well as computer scientists active in the economics of information security (see the Workshop on Economics of Information Security); product managers from large Web mail providers (Yahoo!, Microsoft, Google, et al.) and the consumer ISPs (AT&T, Verizon, Comcast, et al.)

[4] `http://slate.com/id/2177969/pagenum/all/#page_start`

[5] See the conventional wisdom at `http://www.rhyolite.com/anti-spam/you-might-be.html`.

Sun Microsystems Federal, Inc.
7900 Westpark Drive, Suite A110
McLean, VA 22102
703 204-4100 MAIN
703 280-3945FAX

THE NETWORK IS THE COMPUTER™

April 15, 2009

Dear Mr. Vagoun,

Thank you for accepting the Sun Microsystems, Inc. and Sun Microsystems Federal, Inc. RFI-3 -- National Cyber Leap Year submissions to the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity. We were very pleased to recognize that our responses to RFI-1 were used to create the prospective cyber security categories listed in RFI-3.

Summary of our concept in alignment with cyber security category follows.

The network sensors envisioned in the CAEWS invention may dramatically increase the overall manageability of networks and provide actionable intelligence of cyber threat information. Ideally, the network access control will evolve based on improved quality and resolution of sensors that may be used to model attack vectors and subsequently provide analytic as a feedback loop in the end-to-end network management system.

We look forward to providing our Leap Ahead concepts and working with the SSG during the workshops that are part of phase 2.

Should you have any questions regarding this letter, please contact Barry Sheldon, HLS Business Development,.

Respectfully,

Jean Edwards
Director, Federal Business Development

Name – www.sun.com/federal - We are a wholly owned subsidiary of Sun Microsystems, Inc. called Sun Microsystems Federal, Inc. also known as "Sun Federal".

Game-changing dimension – Morph the gameboard – Cyber Attack Early Warning System (CAEWS)

We are looking to change both the offensive and defensive terrain, as it has often been said; "the best defense is a strong offense" so, by making it easier for the Intelligence Analyst to monitor, analyze and squelch an attacker, we will also make it much harder for would be attackers and terrorists to exploit their cyber targets and achieve their goals.

Concept – Attackers, hackers, terrorists, cyber crooks and the like are all using cyberspace to coordinate their attacks, test their exploits and eventually compromise US assets. So, let's morph the gameboard to undergo transformation from a reactive defensive posture into that of a proactive offensive stance where it is the cyber equivalent of the Ballistic Missile Early Warning System (BMEWS). Let's call it the Cyber Attack Early Warning System or CAEWS, pronounced "say-waz" for short.

Vision – The vision is a sensor-to-analyst-to-ground cyber early warning system. Just as the BMEWS was the first operational ballistic missile detection radar and could provide long-range warning of a ballistic missile attack over the polar region of the northern hemisphere and also provided satellite tracking data, so will the CAEWS. The CAEWS can provide the "cyber radar" equivalent for early warning and protection of US and US interest cyber and physical assets as well as satellite tracking and UAV feeds directly to the attackers physical location where law enforcement and legates around the globe can apprehend the suspects.

Picture a radar screen circling the enterprise (where the enterprise is that comprised of assets of interest) at the installation sites, much like SITE I, II and III of the BMEWS, where cyber devices that are hardened virtual decoys are sitting in the open to be attacked, all the while the attack patterns being promiscuously recorded. Now, these sites would be equipped with the equivalent of a more modern "phased array radar" intrusion sensors as opposed to today's conventional Host based and Network based Intrusion Detection Systems (HIDS, NIDS, IDS) used for this purpose.

The Information received from the CAEWS sensors would be forwarded to the Cyber Attack Station, much the same "Cheyenne Mountain Air Station" equivalent of the BMEWS, where it is coordinated with data from other sensors at the CAEWS sites. The analyst would then be presented this information via a "Hybrid Air Station" desktop which has multilevel security access to the appropriate security enclaves, hi-resolution graphics capability for streaming video viewing and 3D application rendering purposes where it would display the origin of the attack source via geography on a world map.

This type of system has the ability to secure data and correlate and present the information to an analyst in such a way that will give different scenarios for the analyst to select from and the associated and required actions for that selection. This system will also give re-mediation actions to take to guard the systems from such an attack scenario as well as go on the offensive and negate the source of the attack.

Traditional efforts have primarily focused on external attacks, while overlooking the vulnerabilities of internal attacks which can be, if not more so a greater threat. Through role based policies and systems; devices and data will also be monitored real-time to capture when access is attempted, how, by whom, and why, whether through external or internal sources.

The stakes are now raised by making it much more costly for an attacker to find a hole to exploit and the analysts do not have to spend as much time on developing attack scenarios as the system does this for them in an automated as well as real-time mode fashion.

Method – Sun Federal has had a lot of past experience in the area of Host, Network and Anomaly based Intrusion Detections Systems. One of those experiences being the 'snort' open source Network based intrusion detection system project of which Sun Federal was an open source contributor and a now has turned into the commercial entity known as Sourcefire. Another was the 'Honeynet project' founded in 1999, by a Sun Microsystems, Inc. (Sun) employee as an international, non-profit (501c3) research organization dedicated to improving the security of the Internet at no cost to the public. Another, is Tripwire which is part of Sun Federal's iForce Integrated Security Solution, providing full scale data integrity and configuration auditing capabilities . The list of projects, partners and experience in this area can go on and on, however what is important to our method are the "lessons learned" from these current and past efforts and our carrying of those forward into this initiative.

We socialized the idea to various business development groups inside of our company and solicited government customer input from the Defense and Intelligence communities. We also socialized this concept with some of the companies and agencies of the dream team below whom we intend to partner. We plan to have brainstorming sessions with the interested parties of the Defense and Intelligence communities. We also plan to work with our Sun Federal Board of Directors to help develop, formulate and refine this idea into an even stronger and more sound concept. The Sun Federal Board will ensure that all the aspects described in the Leap Year RFI are addressed in our final submission.

Dream team – This will require a coordinated approach across many areas and agencies of Government and the private sector, our team would include: CIA, DARPA, DIA, DNI, FBI, iSight Partners, LGB & Associates, Inc., NGA, NIST, NRO, NSA, NSF, Network Solutions LLC, ObjectFX Corporation, Sun Microsystems, Inc., and the US Department of Commerce's InterNIC service and support from ICANN.

Name – www.sun.com/federal - We are a wholly owned subsidiary of Sun Microsystems, Inc. called Sun Microsystems Federal, Inc. also known as "Sun Federal".

Game-changing dimension – Morph the gameboard – Cyber Attack Early Warning System (CAEWS)

We are looking to change both the offensive and defensive terrain, as it has often been said; "the best defense is a strong offense" so, by making it easier for the Intelligence Analyst to monitor, analyze and squelch an attacker, we will also make it much harder for would be attackers and terrorists to exploit their cyber targets and achieve their goals.

Concept – Attackers, hackers, terrorists, cyber crooks and the like are all using cyberspace to coordinate their attacks, test their exploits and eventually compromise US assets. So, let's morph the gameboard to undergo transformation from a reactive defensive posture into that of a proactive offensive stance where it is the cyber equivalent of the Ballistic Missile Early Warning System (BMEWS). Let's call it the Cyber Attack Early Warning System or CAEWS, pronounced "say-waz" for short.

Vision – The vision is a sensor-to-analyst-to-ground cyber early warning system. Just as the BMEWS was the first operational ballistic missile detection radar and could provide long-range warning of a ballistic missile attack over the polar region of the northern hemisphere and also provided satellite tracking data, so will the CAEWS. The CAEWS can provide the "cyber radar" equivalent for early warning and protection of US and US interest cyber and physical assets as well as satellite tracking and UAV feeds directly to the attackers physical location where law enforcement and legates around the globe can apprehend the suspects.

Picture a radar screen circling the enterprise (where the enterprise is that comprised of assets of interest) at the installation sites, much like SITE I, II and III of the BMEWS, where cyber devices that are hardened virtual decoys are sitting in the open to be attacked, all the while the attack patterns being promiscuously recorded. Now, these sites would be equipped with the equivalent of a more modern "phased array radar" intrusion sensors as opposed to today's conventional Host based and Network based Intrusion Detection Systems (HIDS, NIDS, IDS) used for this purpose.

The Information received from the CAEWS sensors would be forwarded to the Cyber Attack Station, much the same "Cheyenne Mountain Air Station" equivalent of the BMEWS, where it is coordinated with data from other sensors at the CAEWS sites. The analyst would then be presented this information via a "Hybrid Air Station" desktop which has multilevel security access to the appropriate security enclaves, hi-resolution graphics capability for streaming video viewing and 3D application rendering purposes where it would display the origin of the attack source via geography on a world map.

This type of system has the ability to secure data and correlate and present the information to an analyst in such a way that will give different scenarios for the analyst to select from and the associated and required actions for that selection. This system will also give re-mediation actions

to take to guard the systems from such an attack scenario as well as go on the offensive and negate the source of the attack.

Traditional efforts have primarily focused on external attacks, while overlooking the vulnerabilities of internal attacks which can be, if not more so a greater threat. Through role based policies and systems; devices and data will also be monitored real-time to capture when access is attempted, how, by whom, and why, whether through external or internal sources.

The stakes are now raised by making it much more costly for an attacker to find a hole to exploit and the analysts do not have to spend as much time on developing attack scenarios as the system does this for them in an automated as well as real-time mode fashion.

Method – Sun Federal has had a lot of past experience in the area of Host, Network and Anomaly based Intrusion Detections Systems. One of those experiences being the 'snort' open source Network based intrusion detection system project of which Sun Federal was an open source contributor and a now has turned into the commercial entity known as Sourcefire. Another was the 'Honeynet project' founded in 1999, by a Sun Microsystems, Inc. (Sun) employee as an international, non-profit (501c3) research organization dedicated to improving the security of the Internet at no cost to the public. Another, is Tripwire which is part of Sun Federal's iForce Integrated Security Solution, providing full scale data integrity and configuration auditing capabilities . The list of projects, partners and experience in this area can go on and on, however what is important to our method are the "lessons learned" from these current and past efforts and our carrying of those forward into this initiative.

We socialized the idea to various business development groups inside of our company and solicited government customer input from the Defense and Intelligence communities. We also socialized this concept with some of the companies and agencies of the dream team below whom we intend to partner. We plan to have brainstorming sessions with the interested parties of the Defense and Intelligence communities. We also plan to work with our Sun Federal Board of Directors to help develop, formulate and refine this idea into an even stronger and more sound concept. The Sun Federal Board will ensure that all the aspects described in the Leap Year RFI are addressed in our final submission.

Dream team – This will require a coordinated approach across many areas and agencies of Government and the private sector, our team would include: CIA, DARPA, DIA, DNI, FBI, iSight Partners, LGB & Associates, Inc., NGA, NIST, NRO, NSA, NSF, Network Solutions LLC, ObjectFX Corporation, Sun Microsystems, Inc., and the US Department of Commerce's InterNIC service and support from ICANN.

Name: Radia Perlman

Credentials: Sun Fellow, Phd from MIT in computer science, 95 issued patents, author/coauthor of two textbooks in network protocols/network security widely used at top universities at the undergraduate and graduate level, adjunct professor at University of Washington.

Game-changing dimension – Morph the gameboard – Byzantine Robustness Assurance

---

We will eliminate the necessity for separate networks, and solve data containment in a much more scalable way. We will also build networks that are far more robust than today's networks, in that they will guarantee correct data delivery, with a fair share of resources for every authorized conversation, even if some of the components of the infrastructure have become compromised and are malicious.

Although the network protocols we are proposing are nonstandard today, the technology to build this new type of network can be implemented on today's hardware, so this design is practical right now.

Concept – The first observation we make is that it is impractical to solve the problem of data containment through separate networks. First it is incredibly expensive, second there is no practical limit to the number of networks that would be required, since it is not just levels of security, but the number of levels must be multiplied by the number of distinct coalitions. Second, it is not possible to have separate networks really separate, because some data (e.g., email, searching for unclassified information) must be allowed to move between the "networks".

So, we will solve the data containment problem in a way that is much less expensive and tractable, namely, through end-to-end authentication and encryption across the network, and enforcing data containment policies in the servers.

The second observation that there is one reason to have separate networks; for survivability. If some components in a network fail, or if there is some sort of process in one network that goes haywire and uses up all the network resources, it might have been nice to have had separate networks.

We propose solving that problem by using a network architecture built upon Perlman's PhD thesis, and recently extended to work in hierarchical networks.

A "Byzantine failure" is when a trusted component doesn't just halt, but instead becomes malicious. For instance, a router might lie about routing information, flood the network with garbage traffic, or perform the routing protocol correctly but then fail to forward data properly.

The network architecture we propose guarantees that as long as at least one non-faulty path connects nodes A and B, they can communicate, with a fair share of bandwidth, even if all the components not on the path were arbitrarily malicious.

The design works in a hierarchical way. So for instance, if it was desired to have 7 separate networks, instead the routers could be configured to allow guaranteed resources for each of 7 logical networks all sharing the same single infrastructure, and then allocating resources within each logical network, to guarantee fair resources for each authorized conversation, within the resources allocated for that logical network.

Vision - You don't need multiple networks.

I can explain construct of how to do this to any company who will listen. If we can get the government to help us get in front of the right companies, we can help them design and deliver the changes needed to support this construct.

Must change all of the routers – no just a few – same hardware can work – not that big a change to the software. With minor modifications can be get this done now and stay in front of the bad guys. It will refresh the infrastructure, reduce cost and eliminate separate networks.

Method - Data containment is enforced by using end-to-end authentication and encryption from client machine to server, and either having separate servers for each community (much more scalable than separate networks), or having servers that enforce data containment rules within the multiple communities using that server. Routing resistant to Byzantine failures is accomplished by:

a) implement robust flooding as per Perlman's thesis, that enables each source to flood to all the routers (in the domain) the most recently generated message from that source.

b) Use that flooding to disseminate link state routing protocol messages, in a link state protocol such as IS-IS or OSPF.

c) Do hierarchical resource allocation, in a way that only requires each router to keep state proportional to each level of the hierarchy.

This is not that different from what routers do today. Link state protocols already have a flooding mechanism. We just need to add a digital signature to link state routing messages. We can also use the robust flooding to disseminate public keys of all the routers, so that it is not necessary to pre-configure public keys for all routers in all routers.

The next step is for sources to use the link state database to calculate a path, and send a digitally signed message to the routers along the path about the desired path. The reason this is necessary is for source S to avoid a malicious router R that might be behaving properly for everyone else, but is dropping S's packets. With a link state database, S has enough information to find, perhaps through trial-and-error, a path that works for the conversation between S and D.

To enable this to work in a hierarchical network requires careful buffer allocation, and creation of guaranteed reliable virtual links between honest routers, so that the source need not compute the complete path, but rather, can hopscotch between honest routers.

Dream Team - Andy Bechtolsheim (Sun), CISCO/Juniper/other router companies, top universities, and DISA's Richard Hale, IA CTO

Willing to create a white paper that explains all of the details. Very open to public presentation and discussions. This leap-ahead technology concept has not been fostered outside of Sun to other technology companies. Reach into academia is very doable to align the right candidates for this dream team.

Name: Radia Perlman

Credentials: Sun Fellow, Phd from MIT in computer science, 95 issued patents, author/coauthor of two textbooks in network protocols/network security widely used at top universities at the undergraduate and graduate level, adjunct professor at University of Washington.

Game-changing dimension – Change the Rules – Ephemeral Key Management

a. Refuse access to a compromised laptop by change the rules of having low or high quality secret keys stored on a server and not the laptop.
b. Leverage a controlled timer to set actions in motion (decrypt, discard, backup or destroy) for private keys validation and termination.

Concept – Common wisdom says that data at rest is protected if it is encrypted. However, there are vulnerabilities beyond encryption. We propose some measures beyond straightforward encryption to further protect data at rest:

a) Allowing data to expire - In order not to lose data prematurely, it is necessary to make copies of data, and for cost reasons, one cannot assume that every location with copies will have guards to prevent theft of the copies. Encryption alone does not guarantee that after the data expires it become unreadable, because the data can still be read off the backup copies.

Sun has a new approach involving a service called an "ephemerizer". The ephemerizer publishes public keys along with expiration dates, decrypts with the corresponding private keys until a key expires, at which point the ephemerizer discards the private key. Backed-up data due to expire is encrypted with traditional methods, but then the ephemerizer public key is used as an extra key. If ephemerized data is needed to be recovered from backup, the ephemerizer must unlock its lock (use its private key). The ephemerizer does not make copies of its private keys; instead, robustness is achieved by using multiple independent ephemerizers (with independent public keys), perhaps in a quorum scheme (so that k out of n ephemerizers need to be available, and know the relevant private key, in order to unlock the backed-up data).

b) Data on mobile clients (laptops). It is not uncommon to lock laptops with a password, or to encrypt data on a laptop using a password. The problem with passwords is that it is often possible for a thief to guess the password, and it is also common for the real user to forget the password. Therefore, there must be a way of bypassing the password in order to recover data for the real owner.

Sun has a new approach for this that allows laptops to store data encryption keys on a server, along with some method of doing mutual authentication with the server. The laptop discards the data encryption key if it cannot communicate with the server, but easily recover it once it can again reach the server.

c) Thin client model: do not store data on the laptop. Sun has technology that supports virtual desktop (stateless devices that has no local cpu/memory/disk) which stores no data on the client. However, sometimes it is important to be able to work on some of the data offline, or for performance reasons, to be able to work on some data locally, that perhaps is not as security sensitive. We propose to organize a client so that some folders are only stored remotely.

d) The same technology that allows expiring data can also be used to have a remote high quality secret that can unlock all the data. In the case of a spy ship scenario, if the ship is captured, a small amount of

local storage can be destroyed, rendering the data on the ship impossible to recover without help from the remote location. So the data does not need to be destroyed; but it cannot be recovered through any information on the ship.

Vision – Data could be irrecoverably destroyed on a schedule. Also, data on laptops could be much more secure, without an all-or-nothing policy; some folders might never be stored even in encrypted form, others might have the laptop deleting the encryption key if it can't talk to the server every few seconds; other data might even be stored unencrypted.

Method – Current use of encryption practices and compromised agency/user data lost on stolen laptops.

Dream team – Sun, Symentac, CTO of PGP, Inc., and Johns Hopkins/Stanford/MIT (other top universities). This leap-ahead technology concept has not been fostered outside of Sun to other technology companies. Reach into academia is very doable to align the right candidates for this dream team.

Name: Sun Microsystems, Inc. (Sun)

Game-changing Dimension: Morph the gameboard

Concept: Mandatory Attack Resistant Security-focused High Assurance Language (MARSHAL), a new programming language aimed at high assurance and reliability.

Vision: Our aim is to make secure and reliable programming not merely possible, but attractive and convenient. We understand the need for more advanced system assurance programming languages. The underlying technical ideas are to leverage what we have learned from Java and Fortress to produce a new programming language and a development environment (tools) that have features to automatically optimize for security -- with ongoing feedback loops based on newly discovered attack vectors (or new taxonomies) that are part of the development environment.

Sun is a leader in the development of programming languages that are safe and secure in the face of both programming errors and hostile attacks. Sun created the Java programming language to allow safe and secure execution of downloaded code. Java has been deployed commercially for over a decade, is supported by numerous vendors, and has made the Internet significantly safer compared the code created in C, C++, and related languages. JavaFX has brought similar benefits to multi-platform, multimedia scripting. In the last five years, Sun has also created the Fortress programming language, designed to make programmers significantly more productive in the arena of high performance computing, was funded in part by the DARPA High Productivity Computing Systems Program. Fortress aims to make scientific programmers more productive through the following combination of design features: (1) A rich, parameterized, multiple-inheritance object-oriented type system for expressing detailed behavioral interface contracts that can be rigorously enforced. (2) Design-by-contract features for double-checking system state on entry to, and exit from, every function and method. (3) Features for automated unit testing. Every source code file that contains library or application code can also contain test code that verifies the intended behavior of the library for application code. Test code can be invoked before main program execution or as a separate verification step. (4) Invariant relationships among multiple data objects can be expressed as assertions. If test data is provided, testing code is automatically generated to verify the stated invariants. (5) Emphasis on making the language design modular and grow-able: Fortress is not just a specific language, but a framework for language design. Language syntax is malleable can be extended by libraries without altering the compiler. While the original design for Fortress supports the use of traditional mathematical operators for scientific computing, these operators are defined by libraries written in Fortress. The core language mechanisms used to define the mathematical syntax also support the creation of other domain-specific programming languages. (6) Fortress supports multi-threaded parallelism and uses it to implement basic language mechanisms such as "for loops". Automatic work-stealing balances the load among multiple processors or processor cores. Threads are synchronized by non-blocking transactional memory mechanism rather than locks.

We propose to use the Fortress infrastructure as a framework for developing MARSHAL as a domain-specific language for high security and reliability. Mathematical syntax may not be a requirement for this application, but multi-core execution is surely relevant. Key reasons why Fortress is a superior basis for this new effort: (a) A safe language with a type system that cannot be compromised. This should go without saying, but you never know. Buffer overflows and the de-referencing of null or dangling pointers (these are language design or implementation defects frequently exploited by malicious software viruses and worms) just shouldn't be an issue anymore; (b) Execution order is not over-specified. The ubiquitous

use of parallelism makes the precise order of execution less predictable, which has the virtue of making code harder to attack. Indeed, because execution order is not over-specified in Fortress, it creates the possibility of protective perturbations to execution without breaking the rules of the language. Consider the use of randomized heaps in C programs, consider Paul Kocher's timing and power attacks on RSA encryption in smart cards, consider banging-on-the-walls techniques for sending information out of secure compartments. Because the specification gives the implementation great freedom in determining these orders, and even the freedom to randomize them, it is possible to enhance an application's security at the platform level, without requiring changes to the application itself; (c) In Fortress, every object interaction goes through an abstract interface (conceptually, all accesses to data are intermediated by method calls). The design of Fortress enables, and indeed encourages, multiple implementations of the same interface. Code security can be enhanced by using implementations that have been hardened in various ways. As a trivial example, type String might be lightly perturbed in its stored form, for example, by XORing each character with a separate value, randomly chosen for each string; this would impose a small overhead every time a character is accessed (an extra XOR to recover the original character), but this would thwart the common malicious technique of encoding instruction sequences as string data. Such hardening techniques can be done without any changes to the application code.; and (d) Fortress is grow-able and domain-extensible. A common hole in servers is that SQL queries are constructed by concatenating strings, some of which are obtained from untrustworthy sources; the error lies in assuming that the untrustworthy strings are "well-formed". A rich type system such as that in Fortress can be used to distinguish strings from untrustworthy sources and ensure that such strings are first fed to a method that will vet their contents or insert appropriate escape sequences. There is also the issue of convenience: such bugs exist because plain old string concatenation is easier to use than libraries that create structurally correct SQL queries. With a grow-able, domain-extensible language that can support convenient rather than clunky syntax, the "right thing" can also be "the easy thing".

Method: The optimization will need to be better than what is currently done with source code analyzers and defensive programming "best practices" that are in use today. We want to include security features in an easy to use way that is transparent but context sensitive (cf. JavaFX), so that the programmer is not required to have math and computer science experience in order to produce verifiable code. We hope to borrow concepts from other information assurance research projects and make it an integral part of the software development life cycle and verification techniques (e.g. ACL2 or PVS).

Dream Team: Our Sun dream team experts include:
James Gosling, Vice President and Sun Fellow, Sun Microsystems, Inc.; Dr. Guy Steele, Sun Fellow -- Programming Language Research Group Project, Sun Microsystems Laboratories; and Dr. Whit Diffie Vice President and Sun Fellow -- Chief Security Officer, Security, Cryptography, and Policy, Sun Microsystems Laboratories

The MARSHAL Dream Team also includes, but is not limited to:
Formal methods centers of excellence such as University of Texas, NASA, and SRI; High assurance computing groups and security metrology research, such as NSA IAD, NIST CSRC, NRL CHACS, DHS BSI, and MITRE CWE; and Higher education institutes such as CMU SEI, UC Davis SECLAB, NPS CISR, and Purdue CERIAS

In order for MARSHAL to gain wide-spread adoption and general use in all system assurance development, we would need strong participation on the Dream Team from the vendors such as Microsoft, Intel, IBM Research, Nortel Government, and Cisco, etc.

Name: John Weeks

Credentials: Senior Staff Engineer, Sun Microsystems Federal, Inc., 30 years of industry experience, developed multilevel web services concepts including demonstration and presentation at JavaOne, project co-lead of OpenSolaris.org Flexible Mandatory Access Control open source project with NSA.

Game-changing dimension – Morph the gameboard – Virtual World Cybersecurity Gaming

Concept – We need more young minds to be exposed to the concepts of cyber security threats via video game development and game-play.

Vision – Today's young minds will bring fresh ideas to software security and exposing them to the intrigue of cyber security both as the attacker and defender will help them develop their skills before entering higher levels of education. The gamers will be able to apply skills learned during game-play to their real life computing experience making them aware of methods to prevent cyber security threats. Game developers will also find the environment challenging which could lead to a completely new form of gaming.

Method – Work with game console developers and network based virtual communities (e.g., Second Life) to promote security aware gaming. Bounty based games would provide rewards for successful hacking and deflecting of hacks. Players would be penalized for unsuccessful hacks and the lack of proper defense of their own resources. Virtual reality monetary systems could be incorporated (e.g., Linden dollars) to further stimulate competition between gamers and possibly allowing for the creation of custom methods and devices that could be sold/purchased as part of the game.

Cyber security gaming would go beyond current shooter or role playing since the actions of the players would dynamically change the game as it is played. If exploited methods are revealed by the attacker or defender, the reward could be higher, but the method might not work again since it would be known to others and in some cases could be used against the player that was tempted by the higher reward.

Sensor technology could be incorporated to demonstrate the benefits of early detection as well as methods of evasion. It might even be possible to develop new security related methods in a network-based virtual game before such devices could be built for the real world.

Providing games that allow for user generated content and customization such as Garry's Mod or LittleBigPlanet would allow gamers to develop and share their own modifications that could lead to greater adoption.

Dream team - Sun Microsystems, Inc., Sony, Nintendo, EA, Lucas Arts, Second Life and the open source community.

Name: – John Weeks

Credentials: Senior Staff Engineer, Sun Microsystems Federal, Inc., 30 years of industry experience, developed multilevel web services concepts including demonstration and presentation at JavaOne, project co-lead of OpenSolaris.org Flexible Mandatory Access Control open source project with NSA.

Game-changing dimension – Morph the gameboard – Standardized and Flexible MAC/DAC Implementations

Concept – We need to quickly move beyond bolt-on security fixes like virus protection packages and apply real security to the operating platform and application frameworks. The longer we allow relatively unprotected systems to connect to the Internet, the greater the risk to government, industry, and citizens from cyber attacks.

Simplification of security related development, concepts, and administration will be necessary to foster greater adoption and help drive demand for "security inside" products. Many security methods and practices exist today, but it will be across-the-board security unification that will make more solutions with built-in security generally available.

Vision – Whether it is a home PC, cellular phone, or a corporate server, the device will contain built-in safeguards from first use to prevent cybersecurity attacks. Today's lax computing environments are an open invitation to those looking for an exploitable attack surface.

Building more ridged Mandatory Access Control (MAC) based security mechanisms into generally available operating systems and application frameworks will make cyber attacks more difficult to accomplish. Recent Fedora statistics are showing that more than 70% of systems tracked are now running with SELinux enabled and in enforcing mode. This is solid indication that MAC-based systems can be deployed on a large scale thus providing greater protection. The greater adoption of secure systems/solutions will also help promote a greater pool of knowledgeable users, developers, and support staff.

It is time to start developing community-based holistic approaches to higher levels of assurance in the complete software stack while providing additional simplification for solution life-cycle and use. A uniform software security framework will allow sophisticated safeguards to be developed that will offer more stringent security boundaries.

Method – Continue to develop and advance reference monitor technologies such as Flask and type enforcement as open source projects making them available to researchers, community, and industry for inclusion in operating systems and software development environments (e.g., Java, web stack components). Work toward finer integration of Discretionary Access Control (DAC) and MAC to unify these types of controls into a more consistent mechanism that presents a singular security goal that could be applied from embedded devices to enterprise class server

solutions. Such combined solutions could offer both authoritative and restrictive controls allowing for a more flexible reference monitor implementation.

Further promotion of security API's should be provided at all levels of platform software to reduce development cost and provide consistency for security feature development. We really need to be thinking about unified security from the bottom-up and making a developer friendly framework that applies to all levels of software platform. This should not prevent competition between solutions providers since more of the value will shift to security component interaction rather than low-level security hooks.

The operating environment should offer security boundary isolation of unprotected operations and controlled data flows to validate information integrity before moving between environments. Some work has been done in this area such as Google's Chrome browser that associates each browser tab with a separate process. Methods such as these could be more tightly integrated into the platform and as well as creating standard patterns for developers. Furthermore, the extension of uniform security API's and reference monitor technologies would provide closer integration of virtualization methods by closing potential gaps between disparate security implementations.

A more unified security framework could allow for devices/systems that would prevent booting to a network-ready state unless all security boundaries are verified including, passwords for all users, remote session authentication, and encrypted data resources. Features like this would be mandatory and could not be disabled.

Presenting security related messages in a form that are more easily understood by the user/administrator with suggestions for corrective action and/or countermeasures would help safeguard systems from potential attacks. Simplification in the area of policy development and configuration will also be necessary to reduce the cost and increase adoption. Applying visualization to security related tools might reduce complexity and make security concepts more comprehensible to a larger user base. One approach might be a Grockker map style interface to show all subject-object permission relationships, or to use a similar visualization technique with an IDE such as NetBeans or Eclipse that would have the capability of detecting policy errors.

Standards have been a critical cornerstone of higher assurance computing solutions for some time. Bringing similar security standards to the consumer/commercial space would raise the bar for a larger audience and offer greater incentives for producer conformance. Just like we have network standards today (e.g., 802.11), security standards could be just as relevant to a larger base including consumer devices and help drive the adoption of products that meet such standards.

Dream team - Sun Microsystems, Inc., Red Hat, Microsoft, the NSA, and the open source community.

Sun Microsystems Federal, Inc.
7900 Westpark Drive, Suite A110
McLean, VA 22102
703 204-4100 MAIN
703 280-3945FAX

THE NETWORK IS THE COMPUTER™

April 15, 2009

Dear Mr. Vagoun,

Thank you for accepting the Sun Microsystems, Inc. and Sun Microsystems Federal, Inc. RFI-3 -- National Cyber Leap Year submissions to the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity. We were very pleased to recognize that our responses to RFI-1 were used to create the prospective cyber security categories listed in RFI-3.

Summary of our concept in alignment with cyber security category follows.

The FMAC project will deliver consistency and interoperability of fine- grained access controls to standardize implementations across diverse operating systems and application frameworks. Many of the security goals may be further advanced by making the policy configuration as complete, yet simple, to maximize the coverage of the policy servers enforcement and also minimize the complexity of managing the entire system. The benefits of leveraging a flexible framework should be considered a significant cyber security advantage that can be widely adopted as open standards, open systems, and open source software gain wider acceptance in the government open technology development plans.

We look forward to providing our Leap Ahead concepts and working with the SSG during the workshops that are part of phase 2.

Should you have any questions regarding this letter, please contact Rose Mucci, Program Manager

Respectfully,

Jean Edwards
Director, Federal Business Development

Name: – John Weeks

Credentials: Senior Staff Engineer, Sun Microsystems Federal, Inc., 30 years of industry experience, developed multilevel web services concepts including demonstration and presentation at JavaOne, project co-lead of OpenSolaris.org Flexible Mandatory Access Control open source project with NSA.

Game-changing dimension – Morph the gameboard – Standardized and Flexible MAC/DAC Implementations

Concept – We need to quickly move beyond bolt-on security fixes like virus protection packages and apply real security to the operating platform and application frameworks. The longer we allow relatively unprotected systems to connect to the Internet, the greater the risk to government, industry, and citizens from cyber attacks.

Simplification of security related development, concepts, and administration will be necessary to foster greater adoption and help drive demand for "security inside" products. Many security methods and practices exist today, but it will be across-the-board security unification that will make more solutions with built-in security generally available.

Vision – Whether it is a home PC, cellular phone, or a corporate server, the device will contain built-in safeguards from first use to prevent cybersecurity attacks. Today's lax computing environments are an open invitation to those looking for an exploitable attack surface.

Building more ridged Mandatory Access Control (MAC) based security mechanisms into generally available operating systems and application frameworks will make cyber attacks more difficult to accomplish. Recent Fedora statistics are showing that more than 70% of systems tracked are now running with SELinux enabled and in enforcing mode. This is solid indication that MAC-based systems can be deployed on a large scale thus providing greater protection. The greater adoption of secure systems/solutions will also help promote a greater pool of knowledgeable users, developers, and support staff.

It is time to start developing community-based holistic approaches to higher levels of assurance in the complete software stack while providing additional simplification for solution life-cycle and use. A uniform software security framework will allow sophisticated safeguards to be developed that will offer more stringent security boundaries.

Method – Continue to develop and advance reference monitor technologies such as Flask and type enforcement as open source projects making them available to researchers, community, and industry for inclusion in operating systems and software development environments (e.g., Java, web stack components). Work toward finer integration of Discretionary Access Control (DAC) and MAC to unify these types of controls into a more consistent mechanism that presents a singular security goal that could be applied from embedded devices to enterprise class server solutions. Such combined solutions could offer both authoritative and restrictive controls allowing for a more flexible reference monitor implementation.

Further promotion of security API's should be provided at all levels of platform software to reduce development cost and provide consistency for security feature development. We really need to be thinking about unified security from the bottom-up and making a developer friendly framework that applies to all levels of software platform. This should not prevent competition between solutions providers since more of the value will shift to security component interaction rather than low-level security hooks.

The operating environment should offer security boundary isolation of unprotected operations and controlled data flows to validate information integrity before moving between environments. Some work has been done in this area such as Google's Chrome browser that associates each browser tab with a separate process. Methods such as these could be more tightly integrated into the platform and as well as creating standard patterns for developers. Furthermore, the extension of uniform security API's and reference monitor technologies would provide closer integration of virtualization methods by closing potential gaps between disparate security implementations.

A more unified security framework could allow for devices/systems that would prevent booting to a network-ready state unless all security boundaries are verified including, passwords for all users, remote session authentication, and encrypted data resources. Features like this would be mandatory and could not be disabled.

Presenting security related messages in a form that are more easily understood by the user/administrator with suggestions for corrective action and/or countermeasures would help safeguard systems from potential attacks. Simplification in the area of policy development and configuration will also be necessary to reduce the cost and increase adoption. Applying visualization to security related tools might reduce complexity and make security concepts more comprehensible to a larger user base. One approach might be a Grockker map style interface to show all subject-object permission relationships, or to use a similar visualization technique with an IDE such as NetBeans or Eclipse that would have the capability of detecting policy errors.

Standards have been a critical cornerstone of higher assurance computing solutions for some time. Bringing similar security standards to the consumer/commercial space would raise the bar for a larger audience and offer greater incentives for producer conformance. Just like we have network standards today (e.g., 802.11), security standards could be just as relevant to a larger base including consumer devices and help drive the adoption of products that meet such standards.

Dream team - Sun Microsystems, Inc., Red Hat, Microsoft, the NSA, and the open source community.

Sun Microsystems Federal, Inc.
7900 Westpark Drive, Suite A110
McLean, VA 22102
703 204-4100 MAIN
703 280-3945FAX

THE NETWORK IS THE COMPUTER™

April 15, 2009

Dear Mr. Vagoun,

Thank you for accepting the Sun Microsystems, Inc. and Sun Microsystems Federal, Inc. RFI-3 -- National Cyber Leap Year submissions to the Networking and Information Technology Research and Development (NITRD) Program Senior Steering Group (SSG) for Cybersecurity. We were very pleased to recognize that our responses to RFI-1 were used to create the prospective cyber security categories listed in RFI-3.

Summary of our concept in alignment with cyber security category follows.

Our concepts for software assurance should be considered to be a forward thinking approach for a new programming language and software development tools to change the way code is constructed, risks managed while vulnerabilities and weaknesses are prevented. The idea is for MARSHAL to automatically build security-in the code itself and runtime execution environments, make it easy to develop technically sound architectures and robust software, desirable for everyone to adopt, and freely available to all.

We look forward to providing our Leap Ahead concepts and working with the SSG during the workshops that are part of phase 2.

Should you have any questions regarding this letter, please contact Rose Mucci, Program Manager.

Respectfully,

Jean Edwards
Director, Federal Business Development

Name: Sun Microsystems, Inc. (Sun)

Game-changing Dimension: Morph the gameboard

Concept: Mandatory Attack Resistant Security-focused High Assurance Language (MARSHAL), a new programming language aimed at high assurance and reliability.

Vision: Our aim is to make secure and reliable programming not merely possible, but attractive and convenient. We understand the need for more advanced system assurance programming languages. The underlying technical ideas are to leverage what we have learned from Java and Fortress to produce a new programming language and a development environment (tools) that have features to automatically optimize for security -- with ongoing feedback loops based on newly discovered attack vectors (or new taxonomies) that are part of the development environment.

Sun is a leader in the development of programming languages that are safe and secure in the face of both programming errors and hostile attacks. Sun created the Java programming language to allow safe and secure execution of downloaded code. Java has been deployed commercially for over a decade, is supported by numerous vendors, and has made the Internet significantly safer compared the code created in C, C++, and related languages. JavaFX has brought similar benefits to multi-platform, multimedia scripting. In the last five years, Sun has also created the Fortress programming language, designed to make programmers significantly more productive in the arena of high performance computing, was funded in part by the DARPA High Productivity Computing Systems Program. Fortress aims to make scientific programmers more productive through the following combination of design features: (1) A rich, parameterized, multiple-inheritance object-oriented type system for expressing detailed behavioral interface contracts that can be rigorously enforced. (2) Design-by-contract features for double-checking system state on entry to, and exit from, every function and method. (3) Features for automated unit testing. Every source code file that contains library or application code can also contain test code that verifies the intended behavior of the library for application code. Test code can be invoked before main program execution or as a separate verification step. (4) Invariant relationships among multiple data objects can be expressed as assertions. If test data is provided, testing code is automatically generated to verify the stated invariants. (5) Emphasis on making the language design modular and grow-able: Fortress is not just a specific language, but a framework for language design. Language syntax is malleable can be extended by libraries without altering the compiler. While the original design for Fortress supports the use of traditional mathematical operators for scientific computing, these operators are defined by libraries written in Fortress. The core language mechanisms used to define the mathematical syntax also support the creation of other domain-specific programming languages. (6) Fortress supports multi-threaded parallelism and uses it to implement basic language mechanisms such as "for loops". Automatic work-stealing balances the load among multiple processors or processor cores. Threads are synchronized by non-blocking transactional memory mechanism rather than locks.

We propose to use the Fortress infrastructure as a framework for developing MARSHAL as a domain-specific language for high security and reliability. Mathematical syntax may not be a requirement for this application, but multi-core execution is surely relevant. Key reasons why Fortress is a superior basis for this new effort: (a) A safe language with a type system that cannot be compromised. This should go without saying, but you never know. Buffer overflows and the de-referencing of null or dangling pointers (these are language design or implementation defects frequently exploited by malicious software viruses and worms) just shouldn't be an issue anymore; (b) Execution order is not over-specified. The ubiquitous use of parallelism makes the precise order of execution less predictable, which has the virtue of making code harder to attack. Indeed, because execution order is not over-specified in Fortress, it creates the

possibility of protective perturbations to execution without breaking the rules of the language. Consider the use of randomized heaps in C programs, consider Paul Kocher's timing and power attacks on RSA encryption in smart cards, consider banging-on-the-walls techniques for sending information out of secure compartments. Because the specification gives the implementation great freedom in determining these orders, and even the freedom to randomize them, it is possible to enhance an application's security at the platform level, without requiring changes to the application itself; (c) In Fortress, every object interaction goes through an abstract interface (conceptually, all accesses to data are intermediated by method calls). The design of Fortress enables, and indeed encourages, multiple implementations of the same interface. Code security can be enhanced by using implementations that have been hardened in various ways. As a trivial example, type String might be lightly perturbed in its stored form, for example, by XORing each character with a separate value, randomly chosen for each string; this would impose a small overhead every time a character is accessed (an extra XOR to recover the original character), but this would thwart the common malicious technique of encoding instruction sequences as string data. Such hardening techniques can be done without any changes to the application code.; and (d) Fortress is grow-able and domain-extensible. A common hole in servers is that SQL queries are constructed by concatenating strings, some of which are obtained from untrustworthy sources; the error lies in assuming that the untrustworthy strings are "well-formed". A rich type system such as that in Fortress can be used to distinguish strings from untrustworthy sources and ensure that such strings are first fed to a method that will vet their contents or insert appropriate escape sequences. There is also the issue of convenience: such bugs exist because plain old string concatenation is easier to use than libraries that create structurally correct SQL queries. With a grow-able, domain-extensible language that can support convenient rather than clunky syntax, the "right thing" can also be "the easy thing".

Method: The optimization will need to be better than what is currently done with source code analyzers and defensive programming "best practices" that are in use today. We want to include security features in an easy to use way that is transparent but context sensitive (cf. JavaFX), so that the programmer is not required to have math and computer science experience in order to produce verifiable code. We hope to borrow concepts from other information assurance research projects and make it an integral part of the software development life cycle and verification techniques (e.g. ACL2 or PVS).

Dream Team: Our Sun dream team experts include:
James Gosling, Vice President and Sun Fellow, Sun Microsystems, Inc.; Dr. Guy Steele, Sun Fellow -- Programming Language Research Group Project, Sun Microsystems Laboratories; and Dr. Whit Diffie Vice President and Sun Fellow -- Chief Security Officer, Security, Cryptography, and Policy, Sun Microsystems Laboratories

The MARSHAL Dream Team also includes, but is not limited to:
Formal methods centers of excellence such as University of Texas, NASA, and SRI; High assurance computing groups and security metrology research, such as NSA IAD, NIST CSRC, NRL CHACS, DHS BSI, and MITRE CWE; and Higher education institutes such as CMU SEI, UC Davis SECLAB, NPS CISR, and Purdue CERIAS

In order for MARSHAL to gain wide-spread adoption and general use in all system assurance development, we would need strong participation on the Dream Team from the vendors such as Microsoft, Intel, IBM Research, Nortel Government, and Cisco, etc.

Who you are – Symantec is a world leader in security software, focused on helping customers protect their infrastructures, their information, and their interactions. Symantec provides enterprise security solutions for all network tiers: the gateway, the server, and the client level, including PCs, laptops, and handhelds.

Game-changing dimension – Raise the stakes

Concept – Today we rely on commercial security software to differentiate *good* programs from *bad*. Unfortunately, these terms are often extremely blurry and what users really want to know is whether a program does *what it is supposed to do* and nothing unsuspected before the program is run. If we were to change the question from "is this program bad?" to "does this program behave as it claims?" then we could make the job of defenders much easier. By shifting the question we shift responsibility from the defender to the attacker. Instead of every recipient having to work hard to determine whether a program is acceptable, the creator will have to work a little harder to define what it does. This effort is easily amortized by commercial software, but could be prohibitively expensive for malware distributors.

Vision – The vision is a future in which users can state their expectations regarding program behavior simply, programs that meet those expectations – providing the desired functionality and nothing else – can be found, and systems can enforce the appropriate behavior in a "no surprises" manner. Before any new program can run, it will have to state up front what it will do, and then it will be constrained to behave in that manner. This changes the vision of security from differentiating "good" programs from "bad" to verifying and enforcing that a program does what the user wants, and only what the user expects.

Existing approaches to differentiate a "good" program from a "bad" program include the following:

> *Blacklisting*. This is the classic anti-virus approach. It's a great safety net. But we want to be able to protect *everyone*, not just those who come after the first wave of users have discovered new malware the hard way. And blacklisting is not scalable.
>
> *Behavior Blocking*. This is promising, but it suffers from problems with false positives. Further, behavior blocking is often too reactionary; i.e., by the time the malicious behavior is identified, it's already too late because the program has been executing.
>
> *Whitelisting*. Rather than trying to recognize "bad" programs and allowing anything else, whitelisting specifies the programs that are "good" or allowed and disallows everything else. Whitelisting can suffer from usability issues in general computing environments.

> *Provenance*. Another technique for recognizing "good" programs is via attribution of authorship. Techniques such as code signing fall into this category. While this is useful, in that there is some degree of accountability for the software, it does not work in practice. For example, a signed ActiveX control is only "safe" in some contexts. In addition, software vendors can't really say for sure that their own programs are safe – companies worry about backdoors in their code, programmers use libraries that they can't (or at least shouldn't) trust fully, etc. The fact that we know where a piece of software came from has not, and will not, prevent that software from doing unexpected or undesirable things.

All of these existing solutions address today's problems and remain a part of the overall, defense-in-depth strategy for providing security. But the real question we as a community need to learn how to answer is "*Does this program do what it says it will – and nothing else?"*

Put another way, today we are trying to find an approximate answer the question "does this program do anything unexpected" without knowing what the user expects. Programs need to have clear statements of what they do, and we need mechanisms for translating, verifying, and enforcing them.

Note that we are not talking about duplicating the Verification Grand Challenge, "an ambitious, international, long-term research program" that seeks to enable and promote formally verified software and development practices to produce such software. If successful, the Verification Grand Challenge could address a part of this problem. But in their current form, formal methods will not enable ordinary end users to make decisions about software written by ordinary developers.

Method – Achieving this vision of computing will require solving many hard problems and will incorporate many approaches. We will need good reasonable sets of policies. We will need ways of expressing policies that mere mortals can understand, and that won't allow devious misuses of the language to confuse them. We will need mechanisms for mapping of human-understandable policies to machine-verifiable policies in order to perform enforcement. We need static analysis and runtime containment that can enforce policies or recover from any damage caused by violations.

Dream team – Symantec; academic experts on policy manipulation; usability experts; Government stakeholders

The Fully Distributed Enterprise

Who we are – Symantec is a world leader in security software, focused on helping customers protect their infrastructures, their information, and their interactions. Symantec provides enterprise security solutions for all network tiers: the gateway, the server, and the client level, including PCs, laptops, and handhelds.

Changing dimension – Morph the gameboard

Concept – In our current environment, we all too often depend on scenarios of good insiders and evil outsiders while trying to draw uneven circles around changing sets of network endpoints. This fosters an increasingly futile effort to build "fences" that accurately corral the "good" while keeping out the "bad." These assumptions have been repeatedly proven false, and yet we continue to design systems and software that implicitly depend on them. Once attackers find a way through the "crunchy shell" they can operate with impunity in the "chewy center." We need to fundamentally change the playing field to eliminate the notion of a "safe" or "internal" network and instead shift toward *transient trusted communities of interest.*

We are using a growing variety of devices (e.g. laptops, mobile handsets) and connected services to do work once constrained to a few types of endpoints such as servers and personal computers. Simultaneously, our network infrastructure has been remained largely static. This has carried over into our perspective on system security, forcing our security practitioners to continue to fight "the last war" even as the terrain and tactics change around them.

What if we could replace implicit trust relationships between our increasingly diverse endpoints and replace them with explicit *arbitrary interlocking subsets of trust*?

Vision – We envision an environment where we have the ability to dynamically create purpose-oriented, mutually authorized, private connectivity between network assets for the *duration, services and endpoint communities required*. These relationships may be dissolved when no longer needed and re-established as necessary. Any given endpoint (or specific application) may be a party to multiple endpoint/service relationships as needed for the connectivity and security of any given session.

We must rethink our network, operating system, and application architectures, taking advantage of network infrastructures that are ubiquitous yet potentially untrustworthy.

We envision a system where connections can only occur after a system has successfully identified itself to some trusted authority and received verifiable authorization and configuration information necessary to connect itself to other network assets. A newly-connected endpoint, absent the "ignition key" would have no visibility into other traffic on the network to which it is not privy or subscribed, and potentially does not even possess a "promiscuous mode" through which it can see the presence of traffic not relevant to it.

If every network session operates within an encrypted ad-hoc community, we eliminate the all-or-nothing danger of current VPNs. Today, a machine that gains access to the enterprise VPN

gets total access to every system on the enterprise network. In the fully distributed enterprise, an endpoint will have access to only those systems and services that it would have had inside the "fenced" network.

From the perspective of an attacker, this morphs the gameboard. Target assets are now members of transient constantly changing closed communities. In addition to frustrating mapping activities absent additional knowledge, this model also scales to the level of applications and services. A system should not hold itself out as accessible for a given network access unless and until it is necessary to use or perform services.

Method – To effectively build such an environment we will need mechanisms for:

Secure autoconfiguration. Endpoints should be able to find a "broker" with information on credentials, policies, and perhaps neighbors that share needs for services or groups of users.

Dynamic Service Location. The specific locations of online services may change over time, so an endpoint must be able to locate relevant services within the confines of its policies and authorization.

End-to-end reachability. Ideally, endpoints should be able to reach each other in order to negotiate secure (encrypted) connectivity between themselves. An assumption of IPv6 connectivity is not inappropriate for this scenario. Autoconfiguration or service location functions may also provide means of NAT traversal, but it should be emphasized that if implemented completely, other systems on the underlying transport network would not be able to see or mix traffic without the proper credentials.

Pervasive encryption & better key management. The pervasive use of encryption underpins this model of connectivity. There will likely be services that hold themselves out for access by all comers, but it is not unreasonable to assume that even large portions of paths to them may be subject to opportunistic encryption for basic content protection, if not the entire path. We will require breakthroughs in the management of network encryption, as well as methods to ensure that keys are readily accessible without being stolen.

Explicit metadata on application connections. Given that each connection from each application may be separately negotiated, authenticated, and encrypted, it will be necessary to have specific knowledge of how an application will use the network and encode this for purposes of validation.

Education of the user base. Currently much faith is placed in the "inherent" security of wired networks or the inside of an enterprise perimeter. Such assumptions about the nature of network composition are often not grounded in reality.

Dream Team - This will require a concerted effort by software developers, hardware developers, network device vendors, appliance vendors, ISPs, computer science departments, PC vendors, and government stakeholders.

RFI Name: RFI–2–National Cyber Leap Year
Title: *A Post Quantum Secure, End-to-End, Universal Security Framework (PQSUSF)*

WHO YOU ARE – synaptic-labs.com/contact-us.html – A privately held company dedicated to the R&D of high assurance (HA), long term (100 year) post quantum secure infrastructure. **GAME-CHANGING DIMENSION –** Change the rules **CONCEPT** – The world's first 100 year PQSUSF for smart cards, desktops and servers over any network. Most networks and computer operating systems are, in themselves, fundamentally insecure. Even computer systems that employ US NIST and IETF security standards are at risk (or there is global resistance to acceptance) due to single points of security failure in their design, inadequate support for complex trust relationships in distributed (decentralised) systems, inadequate protection for all stake holder interests and a lack of assurance that they are fit for use across domains in our interconnected and interdependent Information Society. The British Government Technology Strategy Board adds to this list: "*The current way which organisations approach security can be recognised as an underlying market failure which consists of fire fighting security problems, silo'd implementation of technologies, uncontrolled application development practices and a failure to address **systemic** problems. Organisations tend to deal with one problem at a time that results in the deployment of point solutions to treat singular problems*." (2008). Wireless sensors and the billions of MCU RFID tags that form part of the Internet of Things and Ambient Intelligence visions often offer little to no security. Yet this is not the worst of it. In 2006, at the 30th Anniversary Public Key Cryptography Conference, Brian Snow [former Technical Director of the Information Assurance Directorate of the USA NSA] voiced the security industry consensus that RSA, D&H and ECC would be "***flat-lined***" (abrupt and complete security failure) by a quantum computing attack, and that finding an alternative for key exchanges was "***an open problem, an aching problem!***" Information protected using these 'at risk' crypto systems is trivially recorded, archived and decrypted at will in the future. In 2008 Professor Seth Lloyd of MIT advised that large code-breaking quantum computers could arrive after 2018. It will take 12 + years to migrate important security systems!

What if the global community could address the quantum threat *and simultaneously upgrade* our security primitives and protocols so that commodity MCU based smart cards could achieve 100 year security ratings with HA, be vulnerable to fewer single points of failure, offer increased transparency and accountability for all stake holders while protecting against inappropriate third party tracking? *What if* this could be confidently achieved using known and trusted cryptographic and networking techniques, with security levels approaching that of Quantum Key Distribution without the use of quantum physics and run efficiently over ANY data network? *What if* the PQSUSF could wrap around most existing security systems with minimal negative impact to provide rapid security improvements?
VISION – **Synaptic's vision of a PQSUSF** is a unified and interoperable security platform (SDK, end user products, systems, services) that enables 100 year secure e-Passports (ICAO MRTD), wireless credit cards (EMVCo-like), e-Commerce, eGovernment (and so on) that will deliver on all the "*What if*" points above. The architecture of the PQSUSF will be realised by completing the study on the security requirements of archetypical security application contexts/scenarios and normalising these requirements to create a single conservative interoperable protocol/framework. In this vision Internet e-Commerce transactions will feature the privacy characteristics required by RFID applications, the assurance levels on identity required by international e-Passports, and run at speeds comparable to SSL/TLS. This system will be engineered with the HA methods used in the aerospace industry. Similar to the Microsoft agenda to improve enterprise security through the use of smart cards, PQSUSF HA software will be installed into smart cards to provide increased protection of the desktop, servers and card holders identity and sensitive data. The security of the **global network** of computer systems (personal, corporate, government, critical infrastructure) will approach a uniformly high level of assurance**.**

RFI Name: RFI–2–National Cyber Leap Year
Title: *A Post Quantum Secure, End-to-End, Universal Security Framework (PQSUSF)*

Synaptic has been working for <u>5 years</u> to design a suite of HA 100 year secure cryptographic primitives required to realise our PQSUSF vision. The PQSUSFbuilding blocks include the use of US NIST standards plus Synaptic's complete range of PQS cryptographic primitives including key exchanges (between billions of international users), digital signatures (Lamport-Diffie-Merkle), data privacy, message integrity and cryptographic hash operations. The security of all Synaptic's primitives is firmly based on conservative pseudo random permutations that can credibly claim a full 100 year security rating against classical and quantum computers with HA on smart cards.[1] Synaptic's key exchanges can be implemented using SHA-2 or Synaptic's own hash function. [redacted] Synaptic's multi function PQS data privacy, message integrity and cryptographic hash functions called Post Quantum Secure DES and PQSAES show how to win 100 year security in a conservative construction that employs the full DES-56 or AES-128 cipher as the substitution operation within each round function invocation. The choice of unmodified NIST block ciphers will facilitate rapid cryptanalysis of these PQS primitives. The PQSDES cipher has been designed to overcome known security weaknesses in DES-56 and to provide 100 year PQS efficiently on billions of low-cost devices already in production such as smart cards (with existing hardware DES circuitry) and low-end desktop computers. No hardware retooling, no redundancy, no need to abandon the research and investments into DES-56! Synaptic is exploring HA processes and tools (such as Galois' Cryptol) for formal description, specification and implementation of the cryptographic primitives and the PQSUSF as a whole. A 10 person development team can complete high assurance implementations of the suite of data privacy, message integrity and hash primitives and one of the point-to-point key exchange technologies in 24 months. The remaining components of the full PQSUSF could be seq. implemented with the entire system available within 48 months.
<u>METHOD</u> – Synaptic builds on and completes the proven post quantum security strategies of IBM, Hitachi and the Technical University of Darmstadt, Germany. Synaptic has consulted widely with world leading experts in quantum computing to identify their future capabilities and with world leading post quantum security expert teams in the France, Germany and Japan to validate Synaptic's cryptographic assumptions and to independently review our key exchange technologies. Our NDAs prohibit the naming of Synaptic's world class collaborators. We can deliver on the published recommendations of Brian Snow who for years has called for HA cryptographic systems in the commercial and civilian environments. Synaptic draws on its extensive object orientated analysis and design experience and has already begun to identify and normalise generic cross domain security requirements published in academic papers and embedded within the security standards such as but not limited to EMVco, ICAO MRTD and the Transatlantic Secure Collaboration Program (BAES, Boeing, EADS, …) into a singular coherent security framework interoperable across a wide range of applications.
<u>DREAM TEAM</u> – (In no specific order) – United Nations, NIST, U.S. GSA, ARDA, ENISA, NSA Information Assurance Directorate and Brian Snow, Electronic Freedom Foundation, FP7 THINK TRUST, EU STORK, Galois, ClearSy, Hitachi, Technischen Universität Darmstadt, Laboratoire PRiSM: Université de Versailles, Government and University of Malta, Thales, AU DoD, ST Microelectronics, Scott Aaronson, and other international organisations and experts of equivalent stature and capabilities from across the globe representing a wide spectrum of socio/political interests.

---

[1] http://synaptic-labs.com/business/uvp-100-year-security.html

RFI Name: RFI–2 – National Cyber Leap Year
Title: Universal distributed decentralised federated post quantum secure eID (PQSeID)

WHO YOU ARE – synaptic-labs.com/contact-us.html – A privately held company dedicated to the R&D of high assurance (HA), long term (100 year) post quantum secure infrastructure. **GAME-CHANGING DIMENSION –** Change the rules **CONCEPT** – The world's first 100 year PQSeID acceptable to competitors, adversaries and social libertarians. The PQSeID is an extension of Synaptic's PQSUSF described in our first submission. The disparate eID policies and technologies implemented within and across different nation states, Corporate, Government and international identity systems have a profound impact on us all. This is a global problem. To quote the EU SecurIST Advisory Board's "Recommendations for a Security and Dependability Research Framework (2007)" (RSDRF): "*The security of the digital world has become a fundamental stake for the citizen with respect to his individual freedom and protection of his* ***computerized identity*** *and privacy, for the company with respect to the protection of its computerized industrial assets, the security of its* ***business transactions and the trust level of its information networks****, and for the state with respect to the reliability of operations and the reduction in the vulnerability of large and critical infrastructures: electricity and water distribution systems, communication methods and means, and information and communication systems pertaining to these infrastructures*." In addition Business and Government organisations can become legally exposed in many different ways if there is an incorrect eID or more than one eID assigned to a person. A globally acceptable eID framework is essential to **protect** against the many implications of identity fraud and increase efficiency in our interconnected and interdependent Information Societies. eID management issues are complex and subtle. For example, in our globally interconnected world of sovereign nations, who has the authority to make assertions over particular identities and how can this be regulated and controlled to protect against malice? In the context of the Internet there are over 50 root certificate authorities (CA) pre-loaded in the popular Firefox Web browser that are simultaneously trusted to provide digital certificates for website addresses. In the paper "MD5 considered harmful today - Creating a rogue CA certificate", Dec 2008: "*We have identified a vulnerability in the Internet Public Key Infrastructure (PKI) used to issue digital certificates for secure websites. As a proof of concept we* ***executed a practical attack*** *scenario and successfully created a rogue Certification Authority (CA) certificate trusted by all common web browsers.* ***This certificate allows us to impersonate any website on the Internet, including banking and e-commerce sites secured using the HTTPS protocol***." The security implications of this specific single point of security failure in the PKI architecture to national, regional and global economies is staggering. Modern federated ID systems such as the US PKI Bridge, and TSCP secure email also suffer from single-points of trust failure.

Increasingly, around the world, Government and Corporate policy is focussing on eID. Quoting the RSDRF: "*The EU citizen's requirements, therefore, are mainly focused around an individual, personal perception of security and dependability and all its related implications. Individual, personal, democratic, self-determined control is much more important to citizens than the traditional, historic, government controlled central approach to security and dependability. In the EU Information Society, security and dependability concepts must take into account not only central control requirements but also the individual need for security and dependability mechanisms that protect the citizens' privacy and identity.*" … Identity "*requirements can be illustrated by the following questions:* ***• The uniqueness of the identity • The ability to decide*** *– What can I choose? and What can you choose on my behalf?* ***• The privacy of personal knowledge and history*** *– What do I know? and What do you know about me?* ***• The ability to act*** *– What can I do that is right? and What can you do wrong?* ***• The ability to control*** *– What can I do to protect myself from risk? and How can I manage this risk?*" These issues are common to all countries and are not adequately address by existing eID systems. For example, it is now common place for applications that identify people such

as National ID Cards, ePassports and Credit Cards to use RFID technologies that promiscuously identify themselves. An extensive case study is published by Thomas S. Heydt-Benjamin, et al in "Vulnerabilities in First-Generation RFID-enabled Credit Cards", 2008: "*Using samples from a variety of RFID-enabled credit cards, our study observes that (1)* ***the cardholder's name and often credit card number and expiration are leaked in plaintext to unauthenticated readers****, […] (4) RFID-enabled credit cards are susceptible in various degrees to a range of other traditional RFID attacks such as skimming and relaying*." Similar technical problems have been published concerning ICAO e-Passports.

Unfortunately all eID systems based on standards based PKC face an even bigger security challenge. All digital certificates/signatures based on mainstream PKC, such as RSA, and ECC, are at risk of abrupt and catastrophic security failure from quantum computer (QC) attacks. Professor Seth Lloyd of MIT advises in 2008 that large code breaking QC could arrive after 2018, within the 10 year lifetime of e-Passports. It can take 12+ years to migrate important security systems!

What if the global community could comprehensively address the quantum threat *and simultaneously upgrade* our PKI architectures with a global eID system that mitigates against a global single points of security failure [redacted]? A system capable to support billions of international users across adversarial nation states and that enables digital signatures with 10-100 year life spans. A system where individuals through to Governments gain increased control over ID's they are responsible for and where Governments and large organisations can proactively participate to increase their confidence in identity assertions originating outside of their realms of authority. ***What if*** next generation e-Commerce authentication operations on the Internet have the privacy characteristics required by RFID applications, the assurance levels on identity required by international e-Passports, and run at speeds comparable to SSL/TLS? [redacted]

METHOD – eID policies are directly enabled, shaped and limited by technology. Digital signature technologies enabled the design of PKI and national e-ID schemes. The creation of large QC will, according to Brian Snow, "flat-line" the PKI technologies that enabled these systems. Synaptic has identified design patterns [redacted] that can be adapted to build more trustworthy international eID infrastructure suitable for Internet, Credit Card, National-ID and e-Passport applications and solve many of the technical problems described above. For example our designs will enable the creation of an application independent, interoperable eID platform [redacted] Synaptic envisages that our approach to eID and our PQSUSF combined with the experiences of organisations working in large trust regimes, and the legal expertise of eID policy makers, and the sensitivity of civil libertarians collaborating together towards a common application and country/regional neutral platform, will enable a revolutionary global eID infrastructure that satisfies the legitimate security needs of all stakeholders in a globally conscious, balanced and high assurance manner. **DREAM TEAM –** A collaborative joint industry and Government regime including the United Nations, NIST, U.S. GSA, ENISA, ICAO, TSCP, Northrop Grumman, Lockheed Martin, Thales, TSCP, NSA IAD, Electronic Freedom Foundation, FP7 THINK TRUST, EU STORK, STM, Infineon, Hitachi, Ross Anderson, Government and Uni of Malta, PRiSM and other organisations and experts of equivalent stature and capabilities from across the globe representing a wide spectrum of socio/political interests.

RFI Name: RFI–2 – National Cyber Leap Year,
Title: A Post Quantum Secure Universal Network Carrier Network of the Future (PQSUNC)

WHO YOU ARE – synaptic-labs.com/contact-us.html – A privately held company dedicated to the R&D of high assurance (HA), long term (100 year) post quantum secure infrastructure. **GAME-CHANGING DIMENSION –** Change the rules **CONCEPT** – A 100 year secure, including post quantum secure, Universal Network Carrier (PQSUNC) Infrastructure that delivers Network of the Future (Super Internet) capabilities. The Government of Japan states that it hopes to deploy a NoF by 2020. NoF research is a high priority in the EC Framework Platform 7. It is NOT clear that the lessons learned from the insecurity of our existing networks are being applied to these research agendas. We live in a globally interconnected and interdependent Information Society enabled by insecure purpose built isochronous, cell and packet based network protocols. Often these network protocol have been targeted to a specific application domain and evolved independently in response to their specific market forces. Example network protocols include AMBA, HyperTransport, PCI, Inifiband, USB, Firewire, Bluetooth, RFID, Ethernet, IP, ATM, POTS, ISDN, and SS7. Taken collectively these network protocols face three common challenges: Scalability, Convergence, and Security. Synaptic addresses these challenges in a combined UNC and NoF model!

Network scalability concerns several axis including the number of users/devices, number of routers, interconnectivity between routers, number of network sessions, bandwidth for each network session, network latency, the cost of hardware, the cost of carrying traffic, quality of service (QoS), privacy, integrity, availability and interoperability with legacy network nodes. In wide area networks the requirement to support legacy network nodes has resulted in protocol imposed limitations in network scalability that has proven difficult to reconcile.

Network convergence has been driven by the market to increase connectivity between devices and to improve performance while lowering cost. At one end of the convergence scale different classes of network are layered one on top of the other, such as with isochronous voice traffic over IP packet networks. At the other end a network protocol may be designed to replace multiple protocols such as with Infiniband's vision of a "System Area Network". Both approaches have serious drawbacks, for example the former suffers from a lack of Quality of Service guarantees requiring the insertion of new routers between users, and the later creates a new isolated network address space requiring new purpose built devices.

Security is conspicuously absent in practically all low-level network protocols. None of the network protocols mentioned above were **designed from the onset** to ensure continuous data protection, high assurance, high integrity, and high availability against malicious attacks from outside and inside the network. It is generally not possible to upgrade the security of a network protocol AFTER it has been deployed. Even when it is technically possible to increase end-to-end security, retroactively deploying upgrades uniformly across the network is extremely difficult. Dr Lawrence G. Roberts, one of the fathers of the Internet and a strong proponent of secure networks argues in his paper "The Top Five Lessons Learned from the ARPANET Applicable to IPv6" that "*Security does not sell, it must be mandated: Throughout history, security improvements have not been created by commercial demand, since one buyer cannot change the others. It has always required government direction or mandate to institute better security; this is true from police to networks*." The three factors limiting cryptographic networks have been the cost of encryption, the absence of long-term secure cryptography and international Government policies discouraging their design and deployment. Synaptic has systematically worked to resolve the first two problems while the later limitation is beginning to shift with the increasing international recognition of the asymmetric threats against critical infrastructures such as the Internet, Government, Corporate and personal networks. *What if* the three great challenges of Scalability, Security and Convergence could be comprehensively addressed in a singular wide area network protocol to deliver NoF/Super Internet capabilities?

RFI Name: RFI–2 – National Cyber Leap Year,
Title: A Post Quantum Secure Universal Network Carrier Network of the Future (PQSUNC)

VISION – Synaptic has been working for 10 years to design a robust distributed decentralised IT/ICT infrastructure. Over this time Synaptic has made significant progress in developing the technologies required to realise our vision of a PQS, high assurance, telco-grade, UNC/NoF. The Synaptic PQSUNC is designed to be backward and forward compatible with all existing communications infrastructure, running over its own native infrastructure or leverage existing infrastructure where available. [redacted]
The PQSUNC includes new congestion management and routing techniques that are designed to enable a global mesh network capable of sustaining up to 1 terabit/s bidirectional flows with full 1 second round trip latencies. The PQSUNC objective is to enable communications between any two homes located on different continents to sustain a 1 gigabit/s bidirectional flow. Every access point on the global network will be multi-homed (active nodes) to guarantee availability. New routing approaches have been designed to ensure the management of routing paths are automatically configured and maintained without humans.

Reducing the cost of PQSUNC routers has been a primary focus of our design. The PQSUNC network protocol has been optimised for efficient low latency execution exclusively on FPGA/ASIC chips. [redacted]
The cost of PQSUNC network access devices will vary based on the required host protocols. The mainstream access port may support telephony quality ISDN ports, ethernet ports to support distributed decentralised virtual private networks, and a dedicated ethernet port to support global Internet access. Other protocols can be incrementally supported as desired.

Synaptic is confident that the vision of an PQSUNC can be rapidly realised. Synaptic has worked for 10 years outside the conventional constraints of commercial organisations to design the network and systematically solve open problems that may have prevented its realisation. One published technology is the highly efficient hardware dedicated VEST cipher capable of 10 gigabit/s single pass authenticated encryption with 256-bit post quantum security ratings that are approximately 6 times more power efficient than NIST standards based authenticated encryption using AES-256. Another technology is the world's first scalable post quantum secure many-to-many key exchange technology suitable to replace at risk RSA, D&H and ECC technologies which has received favourable independent expert evaluation.

METHOD – Synaptic began its research by extensively studying the requirements for building a new distributed decentralised computing architecture that could run over existing wide area networks. From this analysis low-level requirements for a new communications infrastructure emerged and evolved into the PQSUNC project. Synaptic has achieved outstanding progress in the design of our PQSUNC by a) allowing the PQSUNC network protocols to be designed without the limitations of any given existing network protocol, topology or deployment, b) designing the system to host all mainstream network protocols, c) avoid dictating it's own protocol on new devices, d) investor and management support to perform radical architectural research and design without the requirement for incremental product releases.

DREAM TEAM – UN, Anagran, DARPA's Ultraperformance Nanophotonic Intrachip Communications program, Kotura, NIST ATP Terabit Photonic Integrated Circuits, Global Environment for Network Innovations, ENISA, Government and Uni of Malta, NIST, NSA IAD, EFF, FP7 THINK TRUST and other organisations and experts of equivalent stature and capabilities from across the globe representing a wide spectrum of socio/political interests.

Who You Are-- TechGuard Security, LLC, techguardsecurity.com, is a trusted and leading cyber defense company, and presents the Letters of Marque Leap Year Project. TechGuard Security was founded in February 2000 to address National Cyber Defense initiatives and US Critical Infrastructure Security. TechGuard provides trusted and award-winning Cyber Security Solutions through innovative research and development, consulting services and training for the DoD, Intelligence, DHS, Federal, Financial and Healthcare communities. TechGuardians ™ address the current challenges of cybersecurity and privacy, specifically the problems of information management, network vulnerabilities, firewall integrity and network security concerns created by e-commerce initiatives, global Internet connections and cyberterrorism. TechGuard Security, LLC, an 8a, small disadvantaged women-owned business enterprise, provides Trusted, Award-winning and Security Mission-focused Networking solutions addressing US Critical Infrastructure protection, offensive and defensive Cyber initiatives.
Additionally, TechGuard develops, tailors, and manages the Great Walls of Fire ® network security product line including the Poliwall™ with HIPPIE™ filter, firewalls, intrusion detection systems (IDS), and Spam filters. TechGuard performs research into offensive and defensive network attack solutions at their Centers for Adaptive Technology Security (CATS)™ labs, using their patented artificial intelligence, heuristics, micro and nano-technologies.
TechGuard provides expert leadership in various programs sponsored by the DoD, FBI, NSC, and Homeland Security. TechGuard holds the rights to the Heuristic Firewall Patent number 6519703, other filtering patents pending, and has expertise in Virtual World Security with an expert who has patents pending in Virtual World Security. TechGuard has held leadership positions in the FBI's InfraGard program, has contributed to President Bush's Strategy to Secure Cyber Space; was part of the national board of directors of the Cyber Security Industry Alliance; TechGuard President/CEO selected for the CxO Committee of the Information Technology Association of America; and was founded in direct response to Presidential Decision Directive 63 under the Clinton Administration. TechGuard has successfully teamed with the University of Missouri Nano-technology program under DoD cyber defense R&D contracts.

Game-Changing Dimension – Raise the Stakes

Concept – Letters of Marque allow the Congress to issue limited authorization for group(s) of private experts to engage and control piracy on the high seas. This same concept can be instantly applied to the Internet, which is the "high sea" of the Information Age.

Vision - Cyber crime is not a new invention specific to the Information Age alone but rather the natural morphing of ancient criminal behavior dating back at least 300 years and likely longer. The Internet of today has more in common with the oceans of the world today and some 300 years ago than at first is readily apparent. Both are used to convey the goods and services of the world's international markets. In today's Information Age the conveyance of information is the basis for the flow and trade of the global economy as well as the basis for national economies. The answer to piracy on the high seas was so well known it was included in the US constitution itself. Why forsake tried and trusted solutions? There exist today world class security experts who are more than capable of tracking, capturing, and controlling the cyber crime running rampant on the Internet today but they are hampered by the laws designed to protect us all.

Cyber criminals honor no such laws but find in them a shelter and safe habor the laws were not intended to produce.

Method – A consortium of IT companies will be created and provided the appropriate Letters of Marque that authorize them to seek out and engage cyber criminals of the Internet. These teams would be authorized to counter-hack, and use many of the same techniques cyber criminals are using against those very same criminals. This activity will be monitored and managed by the FISA courts to ensure that the companies are maintaining the highest standards of professionalism and adhering to all applicable law.

Dream Team - FISA Court, Techguard Security, Intel Experts, Immunity Sec., other individual hackers selected for their experience and dedication to upholding the law.

How will this change the game? It will place the United States finally on the same technical footing as the computer criminals using tried and trusted legal methods centuries old.

How clear is the way forward? The technology and expertise exists in the US right now, but is held down by restrictive laws. Just as in ancient high seas piracy the pirates are already long gone by the time search warrants are issued, and in many cases they can't ever be issued because the pirates are in another country. Using Letters of Marque which are already part of the US constitution, allows our world class experts to move just as fast against the international computer criminals as those very criminals move against US government and companies.

What heights are the hurdles that may be found on the way forward? Culturally the US must be reminded that it already has the legal precedent and legal structure to deal with computer crime. Strict monitoring must be used to ensure that the law is upheld by all involved.

Who You Are-- TechGuard Security, LLC, techguardsecurity.com, is a trusted and leading cyber defense company, and is teamed with the University of Missouri Nanotechnology Lab to present the Virtual World Security Leap Year project. TechGuard Security was founded in February 2000 to address National Cyber Defense initiatives and US Critical Infrastructure Security. TechGuard provides trusted and award-winning Cyber Security Solutions through innovative research and development, consulting services and training for the DoD, Intelligence, DHS, Federal, Financial and Healthcare communities. TechGuardians ™ address the current challenges of cybersecurity and privacy, specifically the problems of information management, network vulnerabilities, firewall integrity and network security concerns created by e-commerce initiatives, global Internet connections and cyberterrorism. TechGuard Security, LLC, an 8a, small disadvantaged women-owned business enterprise, provides Trusted, Award-winning and Security Mission-focused Networking solutions addressing US Critical Infrastructure protection, offensive and defensive Cyber initiatives.

Additionally, TechGuard develops, tailors, and manages the Great Walls of Fire ® network security product line including the Poliwall™ with HIPPIE™ filter, firewalls, intrusion detection systems (IDS), and Spam filters. TechGuard performs research into offensive and defensive network attack solutions at their Centers for Adaptive Technology Security (CATS)™ labs, using their patented artificial intelligence, heuristics, micro and nano-technologies.

TechGuard provides expert leadership in various programs sponsored by the DoD, FBI, NSC, and Homeland Security. TechGuard holds the rights to the Heuristic Firewall Patent number 6519703, other filtering patents pending, and has expertise in Virtual World Security with an expert who has patents pending in Virtual World Security. TechGuard has held leadership positions in the FBI's InfraGard program, has contributed to President Bush's Strategy to Secure Cyber Space; was part of the national board of directors of the Cyber Security Industry Alliance; TechGuard President/CEO selected for the CxO Committee of the Information Technology Association of America; and was founded in direct response to Presidential Decision Directive 63 under the Clinton Administration. TechGuard has successfully teamed with the University of Missouri Nano-technology program under DoD cyber defense R&D contracts.

Game-changing dimension— This approach changes the game board from the traditional Internet, to the Virtual world where malicious parties can move about undetected and unfiltered.

Concept – Virtual Worlds are graphic user interfaces for the Internet and typically require the user to use an "avatar". Virtual Worlds bypass firewalls by design, and the avatar can and often does use scripted objects whose trustworthiness is unverifiable. A new Firewall needs to be created for avatars to protect the avatar from the scripts the avatar is required to use as well as to protect the user's computer from the client software.

Vision – Firewalls protect networks from outside intrusion. This is the model at play in virtual worlds, except the "network" is the virtual world account and the user's computer. Protection is very difficult. The avatar is a section of memory on a remote computer's CPU, and the client occupies a section of memory in the user's local computer. Port/protocol/source and destination filtering are required for protection, and also artificial intelligence and heuristics to detect potential malicious activity and filter against it.

Method – TechGuard's experience in creating innovative AI firewalls will position it to provide leadership in the creation an avatar firewall. The method will be two-fold, to create a client-side software firewall using a "bump in the wire" device that lays between the client and the Internet. This device will have a key pad and smart card reader. The user will be required to insert a card and supply a PIN to the device. The device will then encrypt all traffic to and from the client side computer to a central server which would then decrypt and inspect the traffic using the AI firewall developed by TechGuard. Safe traffic is then re-encrypted and routed to the Virtual World server where it will undergo another inspection by the AI Firewall on the Avatar. In this inspection it will be decrypted and content-verified. Then it will be submitted to the Virtual World server for processing. Traffic bound for the user will undergo a similar process but in reverse. The AI Firewall on the Avatar will inspect traffic using the AI Firewall developed by the TechGuard team, encrypt traffic and submit it to the firewall central server. It will then be decrypted and deep analysis will occur using the AI firewall hosted there. Safe traffic will be encrypted and sent to the client user's Bump in the Wire device where it will be decrypted and content verified before sending to the user's computer.

Dream Team – National Defense University who heads up the Federal Consortium of Virtual Worlds; the trusted, small, innovative business TechGuard Security; existing relationships with the University of Missouri Nanotechnology Lab; Linden Labs who runs Second Life.

How will it change the game? Currently there are no safeguards in the virtual worlds and yet these systems are defacto Graphical User Interfaces for the entire Internet itself. They require holes in corporate and government firewalls and Virtual Worlds are quickly becoming ubiquitous. Recently the CIOs for several government entities within the DoD met at a conference regarding Virtual Worlds and the consensus is that the Virtual World environments are the future of the Internet and security tools do not exist for them. Such tools must be made now to stay ahead of enemy initiatives in the Virtual World.

How clear is the way forward? TechGuard has particular expertise in Virtual World security, whose original concepts have been captured in patents-pending. TechGuard is also a Firewall and filtering company specialist and have already experienced success in the research and development of heuristic solutions and holding the rights to the Heuristic Firewall Patent, which could be employed to address this emerging threat. TechGuard has outlined a way forward and has a proven success rate with partner the University of Missouri.

What heights are the hurdles that may be found in the way forward? Malicious activity in the Virtual World may fund cyber dominance and cyber terrorism/crime thus the importance of developing Virtual World security devices is essential to keep US and coalition partners ahead of the enemy. The determination and funding available to the enemy may be a high hurdle, but not one that can't be overcome with US speed of development. We must ensure that the US is the first to develop this capability using it as both a defensive and offensive cyber warfare/cyber dominance advantage.

RFI Response to National Cyber Leap Year

15 December 2008

WHO WE ARE

TecSec®, Incorporated, founded in 1990, is a privately held company located just outside of Washington, DC. Through a large library of patents and ever-growing intellectual property, TecSec provides (1) Information Assurance products for the network and desktop, (2) Information Management and Dynamic, Assured Information Sharing through cryptographically enforced Role Based Access Control (RBAC) and (3) CKM Enabled® Solutions, for example, for Digital Rights Management (e.g. secure distribution of newly released, first run movies) and for Critical Infrastructure Protection (e.g. SCADA, Utilities).

PROBLEM UNDERSTANDING

The Web has gone through rapid morphing stages. The Web 2.0 is upon us. In the Web 2.0 world, the data stores/repositories are so media rich and so real that the predators behind the veil of the Internet can jump out to take a bite of us. Mashups using the virtual databases from the Internet are now deployed under the Service Oriented Architecture (SOA) offered as web services in the business enterprise. The trend is that mashups have oozed into the financial sectors.

Deleterious and harmful artifacts from the virtual world should not be left unabated; cyber vulnerabilities and threats must be mitigated in the next cyber war. It is necessary that a hierarchy for governance, risk, and compliance (GRC) be provided as a minimum for users of the Internet. Our citizenry must have better protection against the prowlers and predators in conducting his/her personal and business activities. Critically, augmentation and enhancements to include accountability and reporting on the international level must be implemented by the financial institutes to guard against a potential digital harbor: the penultimate war of wealth.

CYBER WARFARE

It is well documented that more cyber attacks have ramped up over the past years. A recent study found that typically, vulnerability exists for up to hundreds of days before it is publicly disclosed. Over a hundred thousands of vulnerabilities are discovered annually – but not reported publicly.

In the Web 2.0 world, a browser with its zeal to ingest rich media has added much functionality and features. Alternatively, it is becoming as the opted-in agent for distributing malware. Often, the sequence of events would entail the use of the browser to infect the host platform. With the convergence of wired and wireless infrastructure, the spreading of the attacks and infestation is further amplified. Additionally, the blurring of the personal and business processing and transactions is more prevalent. It is highly probable an attack such as Distributed Denial of Service (DDOS), carefully and strategically orchestrated can be used to attack the weakest link in the composite Defense in Depth (DID) concept. DDOS attacks can be leveraged against the less sophisticated businesses with minimally deployed and implemented security. This can subsequently impact the businesses' end customers – the US government. This domino effect is indeed a viable scenario.

WAY FORWARD GAME CHANGER

What are we doing about this? The Internet is never turned off—allowing instantaneous access to information anywhere, anytime, anyone, and any device. Our Nation's information infrastructure is so connected to the cloud and open to the world. Do we just disconnect our Internet connection and live under a rock? Obviously, that is not feasible. Our nation has built up this information superhighway to enable the nation to compete economically. The bedrock of a knowledge society is our ability to aggregate data into information which is then turned into knowledge and actions.

We run into the classic n(n-1)/2 problem with communications security (COMSEC). The Public Key infrastructure (PKI) was introduced and is lately choked down due to limitations to scale. The locking down the access point (network access control and network access point) is not the solution either. COMSEC needs to be buttressed with information security (INOFSEC).

It is necessary to shift to the protection of information/content under an "information centric" paradigm. Data protection and security within a distributed computing environment must condition the data at the points of origin; as well as data in transit and in process carried over transport mechanisms including the Internet. In a distributed computing environment, information creators/content generators must exercise protection upon creation. Data access is then bounded to a Unique User Identification (UUID) and is attribute-based where roles and responsibilities are defined in terms of the domains within an enterprise the user resides.

Multi-owner and multi-application card offers distinct "silos" of encryption enforced data space to maintain separation of domains. Different entities, or divisions, can own different silos on the same card. An attribute-based, multi-silo card shall offer high security enablement with large memory for use in an enterprise-wide environment. The card uses the UUID as the foundation building block; it hosts and stores attributes for multiple domains within an enterprise. Depending on the attributes of the person holding the card, the user's role and responsibility stored in the silo are defined and enforced within the specific domain the user is entitled to access. With the built-in hardware firewall within the attribute-based multi-silo card, there is no risk in data breach among silos. This is critically needed for privacy and liability.

Additionally, biometric and password authentications are required for the card to be utilized to effect a robust identification and authentication methodology. Silos can be added and deleted without a kiosk. This multi-silo card must be certified under the stringent testing and certification under the Common Criteria in meeting the requirements such as FIPS 140-2 for cryptography and tested within NIST such as Secure Biometric Match on Card activity to comply with FIPS 201 for Personal Identification Verification (PIV). The combination of the multi-silo card with biometrics and silos provides a high assurance that the authorized person is accessing only that information designated to him or her.

A dynamic key management framework is emplaced within the multi-silo card in providing flexible differential access. This card provides a secure front end to an overarching end-to-end security solution. The dynamic key management framework is extended and applied in an enterprise environment by providing data protection at the object level with encryption exists at multiple access controls (i.e., physical, logical, functional, and content). This object-level data protection with fine grained access controls can be used for sharing data among partners and communities of interest. Automatic "Redaction" of data where data creators can differentially encrypt their data, treating different data types differently (restricting a word, phrase, sentence, paragraph, page, etc). Data access can be granted on a differential basis as well. The same document will appear to different users differently. Unless a user has a full slate of relevant permissions, the data received will be automatically redacted. The result is a scalable dynamic key management that can bind security to information so that security can move and reside with its respective information and the encryption paradigm enforces the assess policies.

In summary, a multi-silo, attribute-based card offers a strong secure front end in enabling an information centric security within an enterprise operating in a distributed environment. The addition of the card/silos and enforcement with encryption should be considered a leap forward technology since this combination offers an advance in technology that can reach a balance between usage and security beyond other possibilities so that resultant security is not considered an impediment.

Research will be needed to apply the combination of the card/silos/biometrics/encryption into a spectrum of network and WEB 2.0 environments to determine how adaptable the combination can be for mitigating risk, costs, and issues such as privacy and liability that are in the forefront of commerce.

Telcordia Technologies, Inc.
One Telcordia Drive
Piscataway, NJ 08854-4182

Secure IP Protocol Stack
Controlling the Communications

Who you are? Scott Alexander - Telcordia Technologies, Senior Scientist

- PhD (1998) in Computer Science from the University of Pennsylvania. Thesis Title: "A Generalized Computing Model of Active Networks"
- MSE (1994) in Computer Science from the University of Pennsylvania
- BA (1986) in Computer Science from Rice University
- Rubinoff Award for innovative applications of computer technology for Ph.D. dissertation
- D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. "A Secure Active Network Architecture: Realization in SwitchWare," IEEE Network Special Issue on Active and Controllable Networks, v. 12, no. 3, pp 37--45.
- D. Scott Alexander, William A. Arbaugh, Angelos D. Keromytis, and Jonathan M. Smith. "Security in Active Networks," Secure Internet Programming: Security Issues for Mobile and Distributed Objects}, Springer-Verlag Lecture Notes in Computer Science State-of-the-Art series, LNCS 1603, pp. 433--451.

Game-changing dimension --- Change the rules by leading to the morphing of the game board.

Concept --- The IP protocol stack was designed for ease of debugging and implementation with little thought given to security. Based on current knowledge, that stack can be replaced with one designed to maximize security. Thus, we change the protocols to include the ability to control the communications that one will respond to and to require security measures such as signatures and encryption on all packets.

Vision --- The vision is to replace the IP protocol stack with a new protocol stack designed with security as a primary consideration. With reinvention of the protocol stack, we have the opportunity to both provide direct security (such as encryption for confidentiality and signatures for integrity) and to simplify monitoring tasks.

Encryption and signatures serve to foil a large class of attacks that rely on either eavesdropping on legitimate traffic or the ability to forge traffic that appears to be from a legitimate source. End-to-end encryption is used to protect the payload of the message. Hop-by-hop encryption is used to protect the control fields in the message. Thus, an intermediate site can be trusted to carry traffic, but not to read the traffic. Wireless networks can be used without concern about eavesdroppers even in public areas.

The signatures used can be based on trusted or untrusted certificates as agreed to by both ends of the transaction. An online retailer may require that I present a certificate that proves that I am entitled to access my account. An anonymous forum on the Internet may be content to allow use of a certificate created for a single flow. In either case, throughout the flow, the signature allows the remote end to be certain that the connection has not been hijacked.

Telcordia

Telcordia Technologies, Inc.
One Telcordia Drive
Piscataway, NJ 08854-4182

We additionally propose to extend our concept of the Dynamic Community of Interest (DCoI) currently being developed under the DARPA IAMANET program. The DCoI provides a security locus for network connections. It requires that an entity that wants to use a service request permission to use that service first. The service provider can set the terms desired for use of the service. Most simply, this allows creation of the equivalent of a whitelist or blacklist for any service. Additionally, though, the service owner can set other terms of service. Rate of requests, services that can be requested, and other policies are easily applied in a uniform way.

Because the system knows the terms of service for flows within a DCoI, monitoring also becomes easier. Only traffic to authorized applications is allowed within a DCoI. Because traffic from unauthorized sources is easily identified (it does not decrypt within the target DCoI), many classes of probing and related attacks cannot be successfully run. Unauthorized traffic from authorized sources can more easily be filtered as the volume of traffic to be scanned is smaller and the types of authorized traffic have been described. This allows tools such as intrusion detection systems to work much more effectively and at much lower cost in terms of CPU cycles.

Returning to our previous examples, an anonymous forum would create a DCoI that allowed web accesses to its site. Since the DCoI is created for human users to access the site, the rate of requests can be limited as part of the terms of service. Any certificate can be accepted to preserve anonymity, but that certificate is used throughout the session. If the user attempts to escalate his privileges by making an ftp connection, the packets forming that attempt will be dropped since they do not constitute a valid access within that DCoI.

In the case of a retailer, the DCoI to access their website is likely to require a certificate from a known source. That source could be a well-known certificate supplier (something like Verisign), a source of payment guarantee (such as American Express), or the retailer itself (issued during a prior session with a DCoI set up purely to allow new users to create accounts). Again only website access would be allowed. However, the rate of requests might be higher to allow for robots to gather information. (Alternatively, a second DCoI could be created for that sort of of transaction if the terms of service are sufficiently different.)

Method --- Initial work in this area has been conducted under the DARPA IAMANET program. That work focuses on the mobile ad hoc network environment and on the military environment. However, we believe that many aspects of our approach could be extended to communications service providers and to commercial enterprise environments.

Dream team --- DHS, NSA, DOD or any government or commercial organization concerned with security threats towards large scale enterprise networks.

The various aspects of the work would require different elements. Initial work is ideally done by a team combining expertise in network security, network protocols, and IP protocols. To socialize the ideas and transition them into the Internet, we would require support from end host software manufacturers as well as network device manufacturers. Gaining that support is likely to require work within standards bodies, particularly the IETF.

Telcordia Technologies, Inc.
One Telcordia Drive
Piscataway, NJ 08854-4182

Optical Encryption for Photonic Layer Security

Who you are? Telcordia: Dr. Shahab Etemad is with the Applied Research at Telcordia Technologies to which he moved at its inception from Bell Laboratories. Dr. Etemad has a B.Sc in Physics from Imperial College, London University, and a Ph.D. in Physics from the University of Pennsylvania. He has more than 35 years of academic and industrial experience in leading and managing research, development, and deployment of novel technologies. In recognition of his research contributions in optics including photon localization, nonlinear optics and all optical switching, and applications of free electron lasers he has been elected a Fellow of the American Physical Society and a Fellow of the Optical Society of America. He is currently leading optical networking research based on controlling the optical phase including phase-locked communication. He is the Principle Investigator of the Optical-CDMA project funded by DARPA.

Game-changing dimension—Change the rules

Concept: Currently confidentiality of the communicated data is guaranteed through the Advanced Encryption Standard (AES) that operates in protocol layers one and above where data exists as ones and zeros in the electronic domain. AES is widely used by the government and financial institutions and it is supported by NIST for certification purposes. Commercially available AES encryptors are limited to 10 Gb/s for SONET/SDH and 1 Gb/s for Ethernet data rates. There is no commercially available AES for 40 Gb/s data rates already deployed by all major carriers. As demand for higher data rates materializes there will not be security support for 100 Gb/s data rates that will be standardized in 2010. What if we carry the encryption in the optical domain at protocol layer zero where higher rates are possible?

Vision: There are different flavors of encryption in the optical domain. However, compatibility with the existing optical networks is paramount in their usefulness. Current optical networks are designed around wavelength division multiplexing (WDM) where each data stream is channeled in a given transparent window. WDM is widely used to increase capacity by major carriers and some enterprises that need security for their data transmission. For example, a major financial institution has WDM windows carrying 40 Gb/s data from Delaware to NYC over public fiber. The Office of Comptroller of Currency (OCC) has mandated that some financial data must be encrypted once in public domain. As a result the 40 Gb/s data rate is first encrypted at lower rates and at high cost and then aggregated to occupy the WDM window. Optical code division multiplexing (OCDM) has been a candidate optical encryption, but until recently it was not compatible with WDM networks and had very low spectral efficiency. Recent research break through by Telcordia funded by DARPA is promising for three reasons: 1) It is compatible with WDM networks. 2) It is scalable to 100 Gb/s and beyond. 3) Through special optical phase scrambling technique it has provable security. The proof on the concept has been demonstrated at 40 Gb/s over 400 km. What needs to be done is construction of a prototype and bringing about acceptance of optical encryption to be supported by NIST for certification and as a methodology acceptable to government agencies such as OCC.

Method: A high-level view of the operation of our proposed OCDM-based security solution is schematically shown in Figure 1. A high data rate 100 Gb/s RZ optical signal can be inverse

multiplexed into a multitude of lower rate tributaries (e.g. 10x10 Gb/s or 4x25Gb/s), each of which is coded by its unique OCDM code, and the combined coded tributaries are injected into a common optical phase scrambler. Conceptually, the coherent summation of these optically encoded tributaries can then be passed through shared phase scrambler before exiting the secure location. The scrambler acts as the key and is a crucial element of the system security due to its large number of possible phase settings. The authorized recipient with the correct key retrieves the ones and zeros of the several decoded signals. The unauthorized eavesdropper does not see ones and zeros to decipher or record the cipher text. Since the scrambler/descrambler setting can be changed at will and the search space for guessing the setting of the key is large, an exhaustive attack is unlikely to be successful. Archival or forensic attack is also difficult since no ones and zeros can be seen in the tapped signal shown in (D) of Figure 1. Furthermore, spoofing of data is made considerably more challenging, since without the key the signal received by an authorized recipient would look like that in panel (D) with no ones and zeros present.

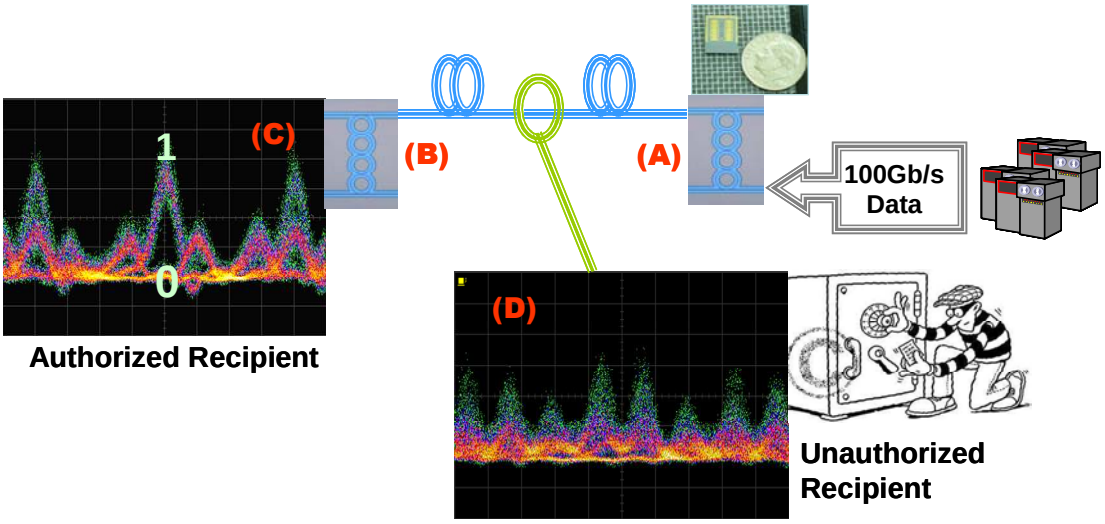

Figure 1: Operation of OCDM-based PLS

Inverse multiplexed tributaries of high data rate digital optical signal are injected into multiple passive phase coder/scrambler (A). Above (A) is the picture of a portion of the passive programmable optical phase coder/scrambler device. At the receiving end, the authorized recipient has the matched decoder/descrambler (B) and can view the RZ signal as shown in (C). The unauthorized recipient sees no ones and zeros in the tapped signal shown in (D) since he does not have the correct setting for his phase coder/descrambler.

Dream team: NIST, NSA, DoD, FCC, OCC, DHS, DISA, TIA.

MORPHINE (MOPHing Identities in NEtworks)

Who you are – Telcordia Technologies is a leading global provider of telecom software and services for IP, wireline, wireless, and cable. Telcordia has more than 2,500 employees worldwide. Telcordia Applied Research holds approximately 560 U.S. and 330 non-U.S. patents in networking hardware and software, computer sciences, cryptography, information privacy, and security. Telcordia is one of the U.S. government's first stops for new ideas on networking, information security and assurance, software architecture, and collaborative knowledge management. Telcordia has successfully worked in a large number of DARPA, ARL, CERDEC, ONR research programs. Telcordia actively participate in international renowned conferences, workshops and standards bodies.

Changing dimension – Morph the game board.

Concept – The main concept is to use Morphing Identities to provide fundamentally improved security to social networks, information sharing and communications. The idea is to use time varying, crypto based identities to identify services, hosts, interfaces, networks and users at the different layers of the protocol stack.

Our approach changes the rules of the game because it changes how to access people (e.g. e-mail addresses social networks), groups (e.g. multicast groups), services (e.g. access to servers, DNS), hosts (e.g. nodes ID) and interfaces (e.g. IP addresses). Current applications and networking protocols work with permanent or semi-permanent identities/addresses instead of dynamic and time varying ones. For example, the web page for NITRD is identified by the following URL http://www.nitrd.gov/. As this identity does not change over time, it can be compromised and http packets destined to this URL could be analyzed, intercepted and/or re-directed to a different sever. Similarly, e-mail addresses are permanent entities that are vulnerable to spam. We can find another example at the network layer. In current IP networks, nodes change IP addresses only when connectivity to the network has been interrupted or when nodes connect/move to a different access network. Therefore, a node/interface can have the same IP addresses for a long period of time. It is therefore easy to perform traffic analysis by spoofing network traffic and correlate data from/to a particular IP address. Our approach introduces a fundamental change on how identifiers and addresses used at the different layers from application to network. The idea is that there is nothing permanent used by applications and protocols that can be used to help an attacker.

Vision – In our approach, we will change radically how people, services and hosts are identified; but, the current communications infrastructure and protocols can remain largely the same. For example, we would give a permanent identifier to our friends for them to contact us and send us an e-mail. Such identifier would be used as a key to obtain the e-mail address to be used by the SMTP protocol. The e-mail address used by the SMTP protocol will then change over time eliminating a major weakness (i.e. permanent e-mail addresses) exploited by those who direct spam to our e-mail accounts.

Our vision is to use morphing identities and addresses for all types of identities and addresses used in the social, information and communication networks. Thus there is no permanent identity

that can be attacked. This will provide means to prevent attacks rather than developing costly piece meal techniques that diminish the effect/impact of the attacks. We envision that the benefits will have some added costs, such as more processing power required at all nodes in the network and increased signaling overhead. The benefits of this technology, however, would also go beyond the guaranteed anonymity, no traceability or traffic analysis, because many existing functions would be greatly simplified, such as less costly firewalls and security programs.

Method - The challenge in having a time varying address/identity is to maintain communication between entities and maintain services (e.g. DNS service) as addresses/identities change. Solutions such as those designed for mobility support, i.e. maintain sessions when nodes change IP addresses due to node mobility, are not suited for this scenario as we would like identities to change frequently to difficult traffic analysis and tracing. Mobility solutions have an extra overhead associated to every IP address change as a signaling message exchange between different entities in the network is required. On the contrary, if nodes knew the identity of the destination at a certain time and the identity change pattern, there would be no need for a signaling exchange between involved parties upon an identity change.

Our proposal is therefore to use crypto based identity generation functions similar to secure ID cards used to access virtual private networks (VPN) to generate random identities and IP addresses. When a node wants to communicate with another node (or server) it obtains the identity at the time of query and address generation function from a trusted server (similar to a DNS server in the current Internet). Authentication and authorization mechanisms are put in place so only trusted nodes are given the address generation functions. Different levels of security may exist in the network, so servers may return functions that will be valid for a short period of time (e.g. functions that generate one valid identity to send a single packet or message) or for a longer period of time (e.g. functions that generate valid identities that last for the duration of a session). Clock synchronization between the requesting nodes and the server providing the functions and the destination node is not needed, as the IP address at time of query is provided. However, the server providing the function and the destination node need to be synchronized. This can be achieved at service activation, network deployment or during maintenance.

Dream team - Our dream team would be formed by experts on the following technical areas: cryptography and security, Internet protocols, mobile ad hoc network protocols (as these networks have dealt with a larger degree of dynamics than Internet), mathematicians and algorithms experts.

From: Tess Rossi
Sent: Monday, April 13, 2009 11:04 PM
To: Leapyear
Subject: Comment only

by John Perry Barlow

Governments of the Industrial World, you weary giants of flesh and steel, I come from Cyberspace, the new home of Mind. On behalf of the future, I ask you of the past to leave us alone. You are not welcome among us. You have no sovereignty where we gather.

We have no elected government, nor are we likely to have one, so I address you with no greater authority than that with which liberty itself always speaks. I declare the global social space we are building to be naturally independent of the tyrannies you seek to impose on us. You have no moral right to rule us nor do you possess any methods of enforcement we have true reason to fear.

Governments derive their just powers from the consent of the governed. You have neither solicited nor received ours. We did not invite you. You do not know us, nor do you know our world. Cyberspace does not lie within your borders. Do not think that you can build it, as though it were a public construction project. You cannot. It is an act of nature and it grows itself through our collective actions.

You have not engaged in our great and gathering conversation, nor did you create the wealth of our marketplaces. You do not know our culture, our ethics, or the unwritten codes that already provide our society more order than could be obtained by any of your impositions.

You claim there are problems among us that you need to solve. You use this claim as an excuse to invade our precincts. Many of these problems don't exist. Where there are real conflicts, where there are wrongs, we will identify them and address them by our means. We are forming our own Social Contract . This governance will arise according to the conditions of our world, not yours. Our world is different.

Cyberspace consists of transactions, relationships, and thought itself, arrayed like a standing wave in the web of our communications. Ours is a world that is both everywhere and nowhere, but it is not where bodies live.

We are creating a world that all may enter without privilege or prejudice accorded by race, economic power, military force, or station of birth.

We are creating a world where anyone, anywhere may express his or her beliefs, no matter how singular, without fear of being coerced into silence or conformity.

Your legal concepts of property, expression, identity, movement, and context do not apply to us. They are all based on matter, and there is no matter here.

Our identities have no bodies, so, unlike you, we cannot obtain order by physical coercion. We believe that from ethics, enlightened self-interest, and the commonweal, our governance will emerge . Our identities may be distributed across many of your jurisdictions. The only law that all our constituent cultures would generally recognize is the Golden Rule. We hope we will be able to build our particular solutions on that basis. But we cannot accept the solutions you are attempting to impose.

In the United States, you have today created a law, the Telecommunications Reform Act, which repudiates your own Constitution and insults the dreams of Jefferson, Washington, Mill, Madison, DeToqueville, and Brandeis. These dreams must now be born anew in us.

You are terrified of your own children, since they are natives in a world where you will always be immigrants.

Because you fear them, you entrust your bureaucracies with the parental responsibilities you are too cowardly to confront yourselves. In our world, all the sentiments and expressions of humanity, from the debasing to the angelic, are parts of a seamless whole, the global conversation of bits. We cannot separate the air that chokes from the air upon which wings beat.

In China, Germany, France, Russia, Singapore, Italy and the United States, you are trying to ward off the virus of liberty by erecting guard posts at the frontiers of Cyberspace. These may keep out the contagion for a small time, but they will not work in a world that will soon be blanketed in bit-bearing media.

Your increasingly obsolete information industries would perpetuate themselves by proposing laws, in America and elsewhere, that claim to own speech itself throughout the world. These laws would declare ideas to be another industrial product, no more noble than pig iron. In our world, whatever the human mind may create can be reproduced and distributed infinitely at no cost. The global conveyance of thought no longer requires your factories to accomplish.

These increasingly hostile and colonial measures place us in the same position as those previous lovers of freedom and self-determination who had to reject the authorities of distant, uninformed powers. We must declare our virtual selves immune to your sovereignty, even as we continue to consent to your rule over our bodies. We will spread ourselves across the Planet so that no one can arrest our thoughts.

We will create a civilization of the Mind in Cyberspace. May it be more humane and fair than the world your governments have made before.

Davos, Switzerland

February 8, 1996

NSF RFI Proposal --National Cyber Leap Year

Who you are – the principal organizer ***The Security Network*** (www.thesecuritynetwork.org) is a San Diego-based nonprofit "Fostering Innovation through Collaboration" in security including extensive collaboration of IA world-wide. In this RFI, we're teaming with: academia (National University-NU), international small business (Nuparadigm (USA), ObjectSecurity (UK) & Spark Integration (Canada)), government (US Navy - SPAWAR 5.1.8 & SSC-PAC 5.5.6) and nonprofits (ISSA & The Security Network) = ***IA/CND/Cyber Collaboration TEAM!***

We believe the IA/Security vision is a *Trifecta*, a major thrust in each of the three dimensions – as there is no one "silver bullet" - rather a symbiosis of progressive, innovative, and integrated designs providing that IA "best value" in affordable, flexible and "good enough" security for all!

P.M. - Michael B. Jones – President, The Security Network - ***P.I***. - Michael H. Davis – US Navy IA Technical Authority

1 - Game-changing dimension -- *Change the rules*.

Concept – Reinvent the approach to internet security by incorporating it into the network and services fabric that remains compatible / transitionable with current applications / architectures.

Vision – Effectively leverage Internet capabilities to provide secure frameworks to conduct operations/business across and inter-connected with the Internet with the flexibility mandated in dynamic information sharing environments, without the current pervasive security risks. Thus move from a passive, forensic-based defense to an active posture using real-time threat updates to: a) dynamically adjust our protection levels, prevent break-ins, rather than just monitor or report them; and b) incorporate virtualization of the "perimeter" and decentralized control.

Include the following features: a) designed in enterprise, top-down trust/security model which is a modified open source approach (i.e., LINUX with more controls) with a lead federal oversight organization (NSA?); b) change security layers alignment within an encapsulated net-centric, web services, SOA environment to form an integrated, cohesive proactive/dynamic defense against all threats; c) dynamic security capability where the network proactively, simultaneously adapts to both emerging threats and changing business needs; and d) re-instantiate / finesse parts of the TCP/IP stack to resolve endemic security deficiencies.

Break away from centralized control, as we can't control everything, everywhere anymore, but rather: a) design distributed, transitive trust methods that *accommodate the unanticipated user in unlikely environments* and b) employ a policy-based, contextual security model with security platforms that share and correlate information rather than point or P2P solutions. With this adaptive nature comes: a) finer-grained contextual access controls, b) secure and scalable automation for self-aware, self-healing networks and services, c) integrated DLP, DRM and ID-aware zoning models (going beyond HSPD-12), and d) secure real-time services - all dynamically adjusting to changing policy, threat conditions or quality of trusted environment.

2 - Game-changing dimension – ***Morph the gameboard*.**

Concept – Reduce the overall complexity of the security environment; thus the IA, CM and governance processes are also minimized, both lowering overall TOC and increasing security.

Vision – Engineer an enterprise end-state use case that embodies both simplicity and capability. Segregate essential, critical "C2" information on a more protected enclave (i.e. SIPRNET-like, with improved security), while all other information is protected within a known, trust level environment. Envision an inexpensive HAIPE / TPM capability on all endpoint devices, and at domain intersections, with tamper resistant keystores and touchless communications; cryptographics with stronger tokens (or equivalent multi-factor authentication) along with policy guards that self-configure depending on current vulnerabilities and the

assurance / trust level requirements of their environment (including threat levels / INFOCONs), while integrating ICS/SCADA safeguards as well. Build a transaction authorization binding environment that allows the use of ZBAC (access control that works cross domain) and distributed policy capabilities - using an external protected device for object security binding.

Instead of a "safe" Internet, focus on safe "communities" using: a) ZBAC authorization and a distributed policy with a mesh of proxy based object filters/guards; b) operate in principal on cryptographically protected messages and not solely on VPN or SSL pipes; and c) facilitate easier enterprise alignment with data, applications and COI agreements. Thus virtualize the security, application synchronization and object delivery responsibilities at the application messaging layer where the interfaces are simpler and objects are easily encapsulated. Develop this mesh by wrapping security bindings and applicable object meta-structure around secure messages (like SOAP/SAML, others) built into the fabric. This encapsulation allows built-in redundancy and delivery guarantee by virtualizing P2P communications by the grid proxies with fine-grained, contextual IA policy management in agile environments (that is, using novel "bolt in" mechanisms like: ZBAC, model-driven security, TSG - an ESB/SOA alternative, etc.); and integrated governance and the IA policy requirements therein. Data/content security protection measures will use simpler more effective encryption (e.g., "lattice" methods), ZBAC access control "E2E", resolve IA metadata issues, etc - all as engrained/designed into the architecture.

3 - Game-changing dimension – ***Raise the stakes*****.**

Concept – Do this inversely by lowering costs, incentivizing positive security results, making the "stakes" hugely asymmetrical – thus inexpensive for users and exorbitant for attackers.

Vision – Fully mass-produced, common, ubiquitous IA - aka "COTS IA" - that is "good enough" so the security is pervasive, integrated, dynamic and relatively cheap ($10-50 / device). Build this level of IA in through secure products / services with "pedigrees" (PPL building blocks with known: metrics, residual risks, and C&A V&V), minimized supply chain security issues (using tamper-proof device ID - like a "Secure MAC" - and penetrating T&E methods). Where both privacy and security are built in qualities, minimizing the security complexity, so then an advanced security monitor can automatically isolate threats and disable their malfeasance vectors – shifting to an end-point controls focus with higher assurance identities. Effective C&A is then institutionalized by common testing, verification, certification, accreditation of these pedigreed capabilities (IA building blocks) in emerging agile, composed environments with fine-grained IA/security T&E /V&V methods/processes. A rigorous architectural approach, using object (or field) based security, will lead to common, provable, known assurance levels.

METHOD – We held a "CyberSecurity Collaboration Summit" to address IA issues, vision, leap ahead technologies and including first responder needs. We collaborated with local ISSA, NDIA, AFCEA chapters and other technical professional & government organizations – including the State policy board (working HIPPA/HIE), California DHS S&T, and local university IA/Security leads – to enhance our participative approach – linking government, industry and academia. We distilled the issues and options into three areas, followed by additional brainstorming, collaborations and discussions to produce this trifecta.

DREAM TEAM – additional IA players: (1) R&D/S&T entities/labs (NIST/NSA, DISA, OSTP / NITRD, DARPA/IARPA, ONR/NRL) & universities (NU, UCSD, SDSU, NPS); (2) savvy / proactive systems integrators (Cubic, LMCO, Raytheon, etc); (3) threat specialists; and (4) the team for this input – small, innovative, collaborating entities who get the bigger, long-term IA picture and mandate to reduce complexity (costs) and update security methods to align IA with future business needs! *All while embracing CNCI principles to get us a secure D.I.M.E.*

Who we are--www.trustdigital.com, a private company that develops Enterprise Mobility Management software to secure and manage wireless devices. Our customers include commercial enterprises, US Gov and US DOD. Our platform secures and manages mobile devices across multiple operating systems, including iPhone, Windows Mobile, Palm OS, and Symbian. We enable our customers to get the expected business value from the devices, to manage them efficiently, and to maintain the security and integrity of the device, its' data and the enterprise network and applications to which it connects.

Game-changing dimension-- Mobility is the next computing frontier and smartphones, such as the iPhone, are rapidly becoming the new PC. Their capabilities enable companies and agencies to move mission productivity beyond their IT networks to the point of action. Smartphone technology converges voice and data capabilities on an easy-to-use handheld device, giving organizations a flexible and robust platform for mobile applications. These capabilities can also combine Internet, Web 2.0 technologies and traditional voice services into mash-up applications. This new breed of applications combines multiple sources of new and existing information that are easily tailored to the diverse application requirements of an organization. The powerful combination of device convenience, laptop-like capabilities and high-value applications virtualizes the user's desktop for maximum productivity in and out of the office.

Concept--There are many tensions inherent within the cellular infrastructure. How do you align an individual's privacy rights with the carrier's need to authenticate users to bill for services? Another is the networks' need to know who you are calling or where you are in order to complete calls? How do you separate the individual's personal use of the device, from their employers' need for ownership and control over its information on the device, not to mention the credentials authorizing access to that information? Who manages the device and the applications on it? How do you separate the carrier's need to manage the device for network connectivity from the user's need to personalize the device and from the enterprise's need to maintain the integrity of the device and applications attached? How do we give law enforcement the access it is allowed to have? How can intelligence agencies take advantage of weaknesses in the network, and how do we decide when to implement countermeasures?

These tensions derive from two areas: joint ownership and information sharing necessary for system operation. Who owns the device, the carrier, the individual, or his employer? Who owns the radio on the device, the user or the network? Who owns the data on the device, the employer or the individual? How do you share calling information with the network in order to complete calls without enabling traffic

analysis? How do you share the identity and location information that is necessary to route calls without leaving footprints of where you were? ***We suggest changing the board so joint ownership is eliminated, and changing the rules so the consequences of sharing are mitigated and well defined.***

Vision--Imagine a world where users need carry only one device: that is their wallet, music player, personal email, and their corporate laptop, but gives the enterprise ownership and control over its data and the individual confidence that his information is private. Imagine a world where instead of the carrier controlling your smartphone, the carrier's access to your smartphone was no more intrusive than the cable company's router in your home. Imagine a world where you could control how the information you share with the network was used--calls with records on an unlimited plan? Imagine a world where we understood what information could be available to law enforcement, while policy makers and stakeholders could debate what should be available? Imagine a world where government users could roam on foreign networks without worrying about traffic analysis and location tracking? Imagine a world where intelligence could operate by taking advantage of both network internals and endpoint vulnerabilities.

Some of this is possible today. What is generally possible is hardening that focuses on strengthening individual elements of the system: the smartphone, the cell tower, the cell company switch. What is difficult today is integrating system-wide technical understanding and policy issues.

Method--We are deeply engaged in the cellular ecosystem which includes carriers, OS and handset manufacturers. These stakeholders cooperate, but do not have an end-to-end view of the world. Enterprises view smartphones as small laptops. Carriers are concerned about protecting the integrity of their network, but focus slightly on the security of the endpoint. Handset manufacturers are device but not service focused. Government stakeholders including NSA, DISA, and DOD work together to identify vulnerabilities and define best practice, but this is typically separated from policy discussions and law enforcement and intelligence.

Dream team-- A group of experts and stakeholders who together can understand the end-to-end system and policy objectives: carriers, network equipment, handset, smartphone security and device management experts, as well as stakeholders from enterprises, privacy, law enforcement, intelligence, and DOD. They should baseline the current infrastructure, identify the policy objectives of a future network, and map a course that incrementally evolves the current network to that future network.

Who we are –

http://www.cse.buffalo.edu/caeiae/ – We are the Center of Excellence in Information Systems Assurance Research and Education (CEISARE) at University at Buffalo. Our center is multidisciplinary across four academic disciplines – Computer Science and Engineering, Management Information Sciences, Mathematics and Law school. Our group consisting of more than 10 faculty members with a number of graduate students doing Ph.D. in IA is engaged in several externally sponsored research projects.

Game-changing dimension –

Raise the stake. Impose an automated gracefully operating penalty system to counter the problem of data breaches.

Concept –

Today's business model encourages data sharing but, unfortunately, this also contributes to security threats in more than one ways. According to a recent article in the Wall Street Journal http://online.wsj.com/article_email/SB122093405633914081-lMyQjAxMDI4MjEwMTkxMzE0Wj.html, September 9, 2008), companies in the U.S. have reached a disgraceful milestone when the number of data breaches at their companies attained an embarrassingly high level. The study shows that there were 449 publicly disclosed data breaches as of August 2008 which had surpassed last year's total of 446. According to a security expert, this epidemic is a serious one and will continue because, today, there is no mechanism to punish companies for this kind of security lapses. The best action by far the companies take is the disclosure of data breaches if required by their state law and beyond that, they have no incentives to investigate or get to the root of the problem. Ultimately, the client becomes a helpless victim with no power to punish the company that exposed his records.

Having recognized the harmful impact of this data breach problem on citizens, we propose a game-changer idea where the companies that are accountable for data breaches, and yet do not take any significant action will be levied a penalty in a way that will hinder them from doing their very business.

Vision –

The vision is to advocate a graceful degradation of the rendered application-specific quality of service (QoS) that the company perceives in the face of a conspicuous lack of cooperation. We propose to develop a QoS-Throttling system that could be viewed as a contractual requirement by the customer of businesses; it may be viewed as a mechanism to correct complacency by corporate members "in-situ." Any complacency by businesses (and their employees) in applying appropriate security measures towards data protection would lead to a lowering of QoS, which in turn, would directly affect productivity (and hence, would affect the "sacred" bottom-line). Thus conformance to security measures will not be limited to merely a moral code but enforced with a monetary means. The automated mechanism will thus raise the stakes and make the businesses take responsibility for data breaches.

We also envision that one might seek to circumvent the game's rules by turning off the game. The QoS Throttling mechanism would have be protected from tampering. Our previous work in user-level intrusion detection would directly apply as a means to counter attempts to disable or evade the throttling of QoS.

Method –

This idea is an extension of our research on human centered security. In the cyber security domain, the human becomes the weakest link because normal users, unaware of the implications of their actions, often attempt to bypass or relax the security mechanisms in place, seeking instead increased performance or ease of use. This shortcoming adds a level of uncertainty unacceptable in highly critical information systems. Merely educating the user to adopt safe security practices is limited in its effectiveness; there is a need to implement a technically sound measure to address the weak human factor across a broad spectrum of systems. We have developed a game theoretic model to elicit user cooperation with the security mechanisms in a system. We argue for a change in the design methodology, where users are persuaded to cooperate with the security mechanisms after suitable feedback. Users are offered incentives in the form of increased QoS in terms of application and system level performance increase. User's motives and their actions are modeled in a game theoretic framework using the class of generalized pursuit-evasion differential games.

This idea was developed as part of a project entitled "Inferring the Loss of Service Quality in a Disadvantaged Network – A Game Theoretic Perspective" funded by AFRL. This work was well received within the research community as is evident from a large number of publications in conferences and journals. The Ph.D. student who worked on this research was hired by Microsoft in 2008. There is considerable interest within AFRL to pursue this research by enhancing the scope. The flip-side of the problem of offering incentive for good behavior is the problem of levying penalty for lack of cooperation and non-compliance. A combination of penalties for non-cooperation and incentives for good behavior will be more productive and we believe that incorporating our ideas into the business workflow will hold promise towards gracefully improving the problem of data breaches.

Dream team –

Participants are from universities (Shambhu Upadhyaya, CSE and H.R. Rao, MIS), government labs (Kevin Kiwat, ARFL, Rome, NY) and companies (S. Vidyaraman, Microsoft). This dream team has worked together for the last 8-10 years in various capacities. Upadhyaya and Rao are associated with the CAE/IAE at University at Buffalo. Vidyaraman was a former Ph.D. student of Upadhyaya and had worked on the proposed concept as part of his dissertation work and he is currently employed by Microsoft. Kwiat is a principal engineer at AFRL Rome and funded several research projects at University at Buffalo and worked closely with the research team and the students.

RFI being submitted by Shambhu Upadhyaya, H.R. Rao

Precision Surveillance with Smart Cameras

University of Illinois at Urbana Champaign

Proposed Research Team

Narendra Ahuja

Professor, Computer Vision and Robotics

Masooda Bashir

Research Scientist, Information Trust Institute

Jay Kesan

Professor, Director, Program in Intellectual Property & Technology Law

Himanshu Khurana

Principal Research Scientist, Information Trust Institute

Rajiv Shah

Post Doctoral Associate

Who you are: We are researchers at the University of Illinois with expertise in computer vision, systems, hardware/software design, law and privacy. We have extensive technology transfer experience.

Game-changing Dimension: Morph the Gameboard

Concept

Surveillance systems are the core of protection and security of critical infrastructures and citizens around the world. Yet, they provide weak precision in identifying adversarial and potentially harmful actions. If we can enhance surveillance systems with "smart cameras" and associated intelligence then we will be able to significantly enhance precision. This would involve refining the hardware, software, and overall goals of current camera systems in a way that incorporates security and privacy into the design.

Vision

Over the last five years, security specialists have forecasted an explosion of smart cameras. Smart cameras rely on video analytics and sensors to provide the next generation of video surveillance. Some of these technologies include license plate recognition, facial recognition, and object detection. Some examples of these technologies include cameras that detect one-way motion at airport security corridors or the massive virtual fence panned for creation along the border with Mexico by using cameras and sensors.

The reality of smart cameras has been underwhelming. The received wisdom is that these cameras have trouble with real world environments. When the cameras move outside of the laboratory, they have problems. One key problem is that they generate too many false positives when placed in real world environments. The excessive alerts dramatically reduce the value of these cameras.

Our research team aims to address a number of core problems and issues with smart cameras. Our ultimate goal is enhancing smart cameras while incorporating security and privacy into the design. The end result of such an effort would be a morphing of the gameboard in that fundamentally better defensive capabilities in surveillance would be realized. The following six fundamental enhancements are envisioned:

1. Improving hardware (eyes)

Improving the camera hardware so it can acquire more information or retain its information content under a broader range of adverse conditions. Acquiring richer data at the front end would increase the effectiveness of the entire surveillance system.

2. Defining Events (world)

Interfacing the hardware and software of camera systems with real world environments requires language and ontologies to define events of interest. Research here aims to redefine and represent events to improve breadth of coverage and performance.

3. Improving software (brain)

The software used by the camera system can be further improved. Software is used to identify events from the raw camera data. These algorithms serve as the brain and can

offer many different functions. They range from object detection and optical character recognition to facial recognition. Software is a key feature that is easy to manipulate and offers great enhancements.

4. Integrating Privacy

Privacy in smart camera systems is at best an afterthought. Most systems are not designed with privacy features that ensure data can be limited to certain users or contain auditing functions. We believe that by rethinking camera systems and incorporating privacy from the initial design will lead toa fundamentally new achievement.

5. Considering Security Issues

Security in smart camera systems is typically limited to physical security concerns. There is little research on how smart camera systems may be "tricked". Our research will explore the potential weaknesses of smart camera systems using the principles of computer security analysis.

6. Modifying the Environment

The current approach towards smart cameras has been far too focused on modeling human characteristics and much less so on the unique technological capabilities of computers. We propose to focus on areas where smart cameras can excel and then reshape the environment to meet the needs of the cameras. For example, computer vision has performed well in the task of license plate recognition. This is a powerful tool and is currently being widely deployed. However, for this tool to be most effective, license plates also need to be redesigned so that they are easier for computers to read. This example shows how modifying the external environment can strengthen the use of cameras. This research proposes to assess the technical strengths of cameras and then consider redesigning the environment to match these strengths. In doing so, we will be careful to ensure that privacy and security are incorporated into the design.

Methods

Authors of this white paper met several times to discuss and revise this response.

Dream Team

Researchers at the University of Illinois with expertise in vision, computer systems, law, privacy and security along with industry partners would work together to realized the outlined vision. To this end, we will work with many small as well as large companies, too numerous to list here. Some of the candidates are those companies that we have already interacted with. As some examples of such interactions and companies, we have transferred some technologies that we have developed to a number of companies including Northrop-Grumman, A&T Systems, Westinghouse, Honeywell, Eastman Kodak, Lockheed, SAIC and HRL. Some of our cameras are under commercialization by a start-up company in Champaign. We have also worked with government agencies in formulating and implementing surveillance systems. These and other camera, analytics and software companies could be our partners in the proposed work.

Comprehensive National Cybersecurity Initiative (CNCI)
Request for Input (RFI) – National Cyber Leap Year

Unisys contact: Glenn Becker

Who we are – Unisys Corporation is an international systems integrator with 30,000 employees. Unisys provides design, development, and support services to both government and private sector clients. Unisys has many clients in government and banking for whom cybersecurity is critical.

Game-changing dimension – Change the rules.

Concept – Establish "secure web zones" using Virtual Private Network (VPN) technologies where users are willing to go through some registration process and give up some anonymity in return for a secure on-line environment. All of the technologies (VPNs, biometrics, smartcards) required to implement this concept already exist and have been used to build user communities as large as tens of millions. The challenges and costs of scaling these solutions to the web with billions of users will be justified by the rapidly growing costs of on-line identity theft and fraud.

Vision – The vision is to establish secure web zones where users can access email, messaging services, and a wide variety of on-line information and shopping sites in a secure environment. People will pay more to live or travel to locations where they can live and shop in a more secure environment. This proposal would extend this concept to on-line communities.

Users would be required to register and have biometric samples collected (probably either iris or fingerprints) in order to join the community. The registration information would be encoded onto a smartcard which would be required to access the community. The biometric registration would allow bad-actors to be identified and potentially have their access restricted. The biometric would also help prevent identity theft and users from establishing multiple identities. Large scale deployments of biometric identification technology, such as ration cards in Andhra Pradesh, India (iris recognition) and driver's licenses in the state of Illinois (face matching), have dramatically reduced fraud.

Similar to shopping malls and libraries, these secure web zones would have to work with web information portals and vendors to establish virtual store fronts, making their services available to the members of the secure community. These secure web zones may evolve geographically, by interest groups, or at various levels of security.

Method – This concept evolved from experience deploying identification systems in Andhra Pradesh, the state of Illinois, and other similar cases. Although these deployments were not on-line, they serve as examples of many of the same problems; including how to identify valid members of a group, how to prevent identity theft and fraud, and how to implement non-repudiation (verify the identity of a person).

In the Andhra Pradesh case, the state government needed a program to control and manage the distribution of nearly 80 million state-issued food ration cards. These ration cards provide citizens with necessities – including electricity, petrol, and food – and the program, historically, has been laden with fraud. The Government of Andhra Pradesh wanted a solution to eliminate fraudulent cards and theft of goods and services, and to reduce costs and ensure its citizens are receiving the entitlements they are qualified to receive. In addition to providing access to goods and services, the ration card is also a pseudo national ID card, helping citizens get passports, admission into college, and other privileges.

After significant testing, the government selected iris recognition technology as the best solution for its current and future needs. They found enrollment to be easy and very fast, and the technology to be highly accurate in a one-to-many search mode. To date, there has been no push-back from the Andhra Pradesh citizens regarding the use and implementation of the iris recognition-based solution. In fact, indications are that the citizens are willing to participate, since they cannot receive their benefits if they do not. As of October 2006, over 20 million ration cards were distributed in a 16-month timeframe, with fraudulent cards being eliminated in tandem.

The state turned over the running of 600+ enrollment sites to private entities who charge a small fee for enrollment into the ration card program. The Andhra Pradesh government provides limited support to these enrollment stations, as they are privately managed and run. The manager/owner of the enrollment station keeps the profits and shares a portion of the ration card fee with the government. As of October 2006, enrollments are 85% complete.

Initial calculations in some Andhra Pradesh districts indicate the government has already benefited from substantial savings by deploying the technology, in terms of reduced fraud and subsidies, which extends beyond the primary ration card application into district stores, youth hostels, and low-income housing.

Moving this concept into the web environment could be done by deploying pilot secure web zones in places where strong biometric identification systems are already available. This would simplify the registration process. These pilot implementations would be used to refine the secure web zone concept in an operational environment.

Dream team – The team should include a couple of internet service providers (ISPs), several biometrics vendors (e.g., LG, L-1 Identity Solutions, Cogent), several web-based vendors and information portals (e.g., Wikipedia, Google, Amazon), and some virtual community building expertise.

Who We Are

Unspam Technologies, Inc.
P.O. Box 57265
Murray, UT 84157-0265
(888) 4-UNSPAM
www.unspam.com

Contact:
Matthew Prince
CEO, Unspam Technologies, Inc.

Unspam Technologies, Inc. is a Utah-based company dedicated to tracking and stopping online malicious behavior. As part of these efforts, we created Project Honey Pot (*www.projecthoneypot.org*). The Project is made up of more than 50,000 volunteers in more than 120 countries worldwide who have installed software on their web servers to track suspicious online activity. Every day the Project tracks over 1 million IP addresses engaged in email address harvesting, spamming, phishing, fraudulent comment posting, cross site scripting, or other attacks on web servers.

Game Changing Dimension

Morph the gameboard.

Concept

Most companies, organizations, and governments place their web servers in a DMZ, outside the protective cover of a network firewall. This makes a web server especially vulnerable to attack.

Our concept is to empower web servers to change their behavior depending on whether a particular visitor is a known attacker. If web servers can access data on what visitors are likely to engage in malicious activity then they can alter the information they return and the services they support in order to minimize the threat. Successful protection of web servers has a ripple effect that dramatically reduces downstream treats including spam, phishing, and virus propagation.

Web servers empowered in this way can also create an opportunity to help eliminate a vector for further attacks. Many attacks online today are carried out through proxy machines that have been turned into so-called "zombies" through computer viruses. The legitimate users of these machines are often unaware that their computers have been compromised. Alerting these users to their infections a critical first step to reducing zombies.

Our concept allows web servers to not only protect themselves but also help educate these legitimate users of compromised machines. Web servers that can access data on known zombie machines can change the web pages they return to include a warning that the user's computer appears to be compromised along with instructions on how to eliminate the infection. This provides a unique opportunity to contact the legitimate users of compromised machines and help them clean up their infection.

Vision

Our vision is to widely deploy software systems to protect web servers using the data gathered by Project Honey Pot and other similar sources. We propose building systems that sit in front of the web server, like a firewall, and recognize known malicious users. If a

known malicious computer attempts to access a web page served by the web server, the system can restrict the data that would typically be returned as well as the services the web server supports. For example, if a visitor to a protected website was a known cross-site scripting attacker, the system could restrict the ability of the visitor to issue POST commands.

Legitimate users of compromised machines could be prompted through additional information included in a frame at the top of the web page they requested. The frame would include information on why the user's computer is suspected of being compromised, steps to clean up the infection, and a mechanism to temporarily access the site unrestricted if a CAPTCHA-like challenge is completed.

We propose that the system be rolled out initially on public-facing government websites. If effective, we would encourage other high-traffic websites to protect themselves through a similar mechanism as well as to adopt the same trusted standard of notifying users that their machines appeared to be infected.

Method

To date we have gathered data on online attackers through Project Honey Pot. The Project has been online since 2004. During that time we have received more than a half a billion spam and phishing messages, tracked more than 40 million machines engaged in malicious behavior, and helped protect tens of thousands of websites through early versions of the software envisioned above. The results of our studies suggest that this is a promising approach worth pursuing more broadly.

Protecting web servers from known-malicious users leads to several unexpected and dramatic decreases in other types of attacks. For example, spammers typically build their lists of email addresses through "harvesting." This process entails a software program run by the spammer visiting web pages and retrieving email addresses. Data we have collected indicate that each email address a single harvester retrieves from a web page will, on average over the next 36 months, receive more than 2,000 spam and phishing messages. Web servers that can recognize a visitor as a known harvester and remove all email addresses listed on pages displayed to that visitor dramatically decrease the ultimate volume of spam and phishing messages the otherwise-listed email addresses would have received.

Dream Team

Department of Homeland Security; Federal Trade Commission; Microsoft; Apache Software Foundation; Yahoo, Google, MySpace, Facebook, eBay, Wikipedia, Amazon.com, and other high-traffic websites; Symantec, McAfee, TrendMicro, and other anti-virus firms.

VEGA: Finding root causes of root kits

S. Clark, M. A. Blaze and J. M. Smith
CIS Department, University of Pennsylvania

Introduction

The current model for ideal computer system security is one of sets of properties that, if and only if *proven correct,* will result in a secure system. Much effort has been devoted since the 1960s to trying (and unfortunately, failing) to develop workable *Trusted Computing Systems* that are *provably* secure. The problem with this seductive model is that these security properties and proofs are based on axiomatic assumptions that are independent from their environment. Alas, computers and networks are not only intrinsically dependent on their environment, the environment itself is always evolving, which can results in unexpected interactions violating assumptions. The strategic goal of an attacker is to discover these interactions, exploit them, and gain control of a system.

Viewed correctly, the environment for computers systems is an *ecosystem*, analogous to a biological ecosystem, wherein both predators and prey must continuously adapt to survive. In this peculiar ecosystem, attackers are free to evolve, adapting their attacks to fit individual situations, but defenders are static, relying on one of two methods for protecting themselves. One is the 'Center for Disease Control' model. A user pro-actively inoculates a host against all known viruses, then waits helplessly to see if it becomes infected. A second method is the 'Castle and Moat' model, where institutions deploy computer systems behind firewalls ("moats"), trying to limit contact with incoming and (less commonly) outgoing traffic ("castle gates") in and (more rarely) out. Intrusion detection systems send alarms when a breach in security is detected. After the security appliances are deployed, institutions sit and wait helplessly to be attacked. Neither defensive method provides a measure for likelihood of attack, neither method can be used to adapt to new threats, neither method can provide any way to make security proactive, to take security on the offensive. There is no science and no models for the "cyber" environment (the Network/Computer Security "ecosystem") and therefore, no ability to intelligently design adaptive, predictive or offensive capabilities.

Prior Explorations – Root Causes

Work has been done to identify and model attacker methodology, and some work to identify the time-frame of particular vulnerabilities, and even some work to try and understand the economics driving the attacker "black market". The focus of the available research is much too narrow and is designed to incrementally enhance one of the two defensive methods mentioned above. Some seed efforts to look to the biological and physical world for examples of successful adaptive and evolving security models have been pursued but have not yet proven fruitful, and lack an ecosystem worldview.

Our Approach – Vulnerability Exploit Genesis Analysis (VEGA)

VEGA will build a taxonomy of computer security used to develop a model of the Network/Computer Security ecosystem using a 3-fold approach.

First, model the attacker: Collect and analyze data regarding attacker methodology, economics and motivations. Our model will answer the questions, "Who is doing the attacking?" "What tools are they using?", "What approaches are they using?", "What

previously ignored areas are now being attacked?" (*e.g.*, supply-chain, hardware vulnerabilities, fuzzing), "How are attacks evolving over time".
Two, model the vulnerability and exploit life cycle: Collect and analyze data about zero-day (unknown) attacks and the 'Honeymoon Period' (*i.e.*, on average, how much time is there between the release of software and the discovery of the first exploitable vulnerability), collect and analyze data about the exploits (*i.e.*, how much time is there from first exploit to automated script, expanding on work by Dr. William Arbaugh. "Which systems/applications are the most vulnerable to particular attacks", Analyze the effectiveness of automated patching.
Three, with the data from the above models, use adaptive learning techniques to develop an evolutionary environmental model that predicts attacker behavior so that it can be countered.

Expected results of VEGA
We expect the results of our research to provide users with the tools they need to perform *Risk Analysis Based Deployment* (RABD). RABD is a model for tolerating insecure computer systems as fielded. Using our model a user will be able to access where in the vulnerability life cycle she stands, identify predators and potential future vulnerabilities, and adapt her defensive strategies to make best use of available resources. RABD will help users to take their security "on the offensive", or at least “knowledgeably defensive”.

VEGA Timeline
VEGA will require several years of effort to develop and test the model against real software systems, test its predictive power, generate data from the predictions, and build infrastructure that will quantitatively predict when and where to deploy defensive effort. The current model is flying in the dark without radar. Subtleties include research to see what Honeymoon periods exist and how they can be used by defenders, perhaps, for example, by generating new “versions” of software by automated methods on a quicker timecycle than attackers can learn it to build exploits.

We should use Metasploit and the CVE (vulnerabilities database) to correlate the times between 0-day and readily available scripts, we need some experiments to determine the efficacy of automated patching. We need some experiments that test fuzzing. Our estimate is that this would require about three years, starting now.

VEGA Missing Pieces and “Dream Team”

Hacker/Black Marketplace data - Hugh Thompson, Chief Strategist- People Security or Rob Davies, Cymru
Economics – Dr. Andre Odlyzko, U. Minnesota
Risk Engineering – Dr. Peter Neumann, SRI
AI - Professor Sal Stolfo, Columbia University
Statistics: Simon Byers, AT&T Labs, Murray Hill
Hardware/Supply-Side/Embedded Systems: Karsten Nohl, University of Virginia

Who we are: Dr. Mohamed Eltoweissy (Director, Center for Cyber Assurance and Trust (CyCare)) and Dr. Saifur Rahman (Joseph Loring Professor of electrical and computer engineering and director, Advanced Research Institute), Virginia Tech.

Morph the Gameboard: *Shuffle and evolve the software system implementation decks.*

Concept: *Confuse and Enhance* – ***Confuse*** the attacker by non-determinism through shuffling of software system component implementations; and ***Enhance*** the software system by survival of the fittest (evolution) through trust-based selection in an online software component marketplace.

Preamble: We perceive a software system comprised of a set of interconnected components that act and interact together to achieve a desired functionality. Each component adheres to certain specifications defining its explicit behavior. At the same time, each component exhibits an inherent implicit behavior due to the potential individuality and variability across different execution environment and different implementations of the same specifications.

Accordingly, we can classify the types of vulnerabilities of a software system into:

- *Specification vulnerabilities*: Those exist due to flaws in the system specifications. We are not concerned with this type of vulnerabilities as they are likely to be detected early on by the careful analysis of the specifications documents, which is a common and necessary process in systems engineering; and
- *Implementation vulnerabilities*: Those result from implementation flaws as a side effect of implementing the specifications of the system or any other component of the execution environment. This class of vulnerabilities is much harder to detect, since auditing the implementation requires much more effort than auditing the specifications and is usually infeasible, especially for commercial off-the-shelf components. Obviously, this type of vulnerabilities is implementation dependent and almost unavoidable.

Different implementations of the same specifications preserve the explicit behavior of the system, but exhibit variations in the implicit behavior, and hence have different sets of implementation vulnerabilities. Attacking a system through implementation vulnerability involves two steps:

- *Step 1*: discovering the vulnerability by observing and analyzing the implicit behavior of the system when given a specific input; and
- *Step 2*: Exploiting the vulnerability by providing the input learned from *Step 1* to instances of the system having the same implementation.

It can be seen that both steps desire the system to possess a deterministic implicit behavior.

Vision: *By introducing non-determinism of implicit behavior to the system, the effort of discovering and exploiting implementation vulnerabilities becomes fruitless. Further, continual enhancement in components is achieved by fostering an electronic marketplace for component implementations.*

Method: We coin the term *implementation redundancy* to refer to a technique of using multiple different implementations (provided by different individuals or vendors) for the same specifications. The system maintains a set of alternative implementations for each component fulfilling its specifications. We call the set of alternative implementations for a component; the component's *working set*. Throughout the lifetime of the system, implementation *shuffling* takes place, whereby the on-duty implementation (the one currently in use) of each component is continually replaced by a different one from the component's working set according to a certain policy. Explicit connectors[1] govern the interaction between system components, allowing the shuffling to take place transparently without disturbing the system operation. Consequently, we are leveraging the inherent differences between the implicit behaviors of individual implementations, effectively injecting non-determinism.

We define the term *behavioral state* as the unique configuration of the on-duty implementations of system components resulting in a unique overall implicit behavior for the system. The higher the number of different behavioral states in the system, the higher is the potential for non-determinism. Assuming a component-oriented design, we show that it is possible to achieve a large behavioral state space in a cost-efficient manner by employing the above shuffling technique.

---

[1] Schreiner, D. and Göschka, K. M. 2007. Explicit Connectors in Component Based Software Engineering for Distributed Embedded Systems. In *Proceedings of the 33rd Conference on Current Trends in theory and Practice of Computer Science.*

Given that all different implementations of the same component are functionally equivalent as dictated by the specifications, by random independent shuffling of component implementations, we achieve a higher number of behavioral states while maintaining the same overall explicit behavior of the system. The number of behavioral states, N_{bs}, is equal to $\prod_{c=1}^{N_c} \|W_c\|$, where W_c is the working set of component c, and N_c is the total number of components in the system.

Normally, a single operation may not involve all system components. Usually, only a subset of components along a certain path is involved in the execution of any specific operation. From the perspective of executing a specific operation, this path is called the operation's *execution path*, and the components residing on it are called the operation's *active set.* We define the *perceived behavioral state* for a specific system operation as the unique configuration of the on-duty implementations of system components belonging to the operation's *active set*. Accordingly, the perceived behavioral state of an operation is only affected if the shuffling step taken by the system involves at least one of the components in the operation's active set.

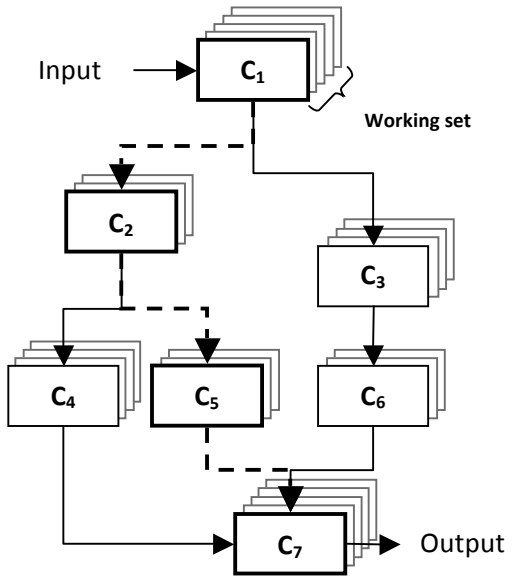

In the system shown in the figure, we assume an operation whose execution involves the active set {C_1, C_2, C_5, C_7} (the ones residing on the path indicated by dashed lines). In a **traditional system**, the implementations used for the components in the active set never change. Thus performing the same operation will always induce the system to show the same implicit behavior. However, with **implementation shuffling,** the implementations selected by the system for the active set are -- with a high probability -- different each time the same operation is performed. For instance, if at time t_1, the system had the on-duty implementations {I_1, I_2, I_5, I_7} for the shown active set, then at time t_2 the on-duty implementations for the active set might be {I'_1, I'_2, I'_5, I'_7}, where the two sets are equal only with a very low probability. Switching between implementation sets doesn't affect the explicit behavior of the system (i.e. the system has the same *functionally*), nevertheless, inherent differences across the various implementations of the same component will result in a different overall implicit behavior for the system.

The need for multiple implementations leads to the necessity of dealing with multiple software vendors. We envision an electronic *marketplace* where the demand for component implementations is met by offers from vendors. Such marketplace is driven by the following stakeholders (players):

- *Customers:* Those are the system owners or operators that have a need for certain specifications to be designed and implemented in order to enable the system to perform a certain function.
- *Specification Design Vendors:* Those put the specifications for functions demanded by the customers.
- *Implementation Vendors*: Those supply customers with component implementations adhering to the required specifications.
- *Trust Management Authority (TMA)*: that evaluates component implementations and provides unbiased valuations and feedback about each implementation available in the marketplace to guide customers to choose from best available candidates according to their own preferences (security, performance, power efficiency, etc).

The systems in production continually update their working sets of components from the marketplace according to the customers' policies in order to enhance their favorable attributes and progressively evolve to better serve their operators' needs.

Dream team: Cyber Assurance and Trust, Software Engineering, Networking, Distributed Operating Systems, Modeling, Formal Methods, and Economics.

PC- Integrated Highly Secure Net-Terminal

Who we are: Professor Saifur Rahman (Joseph Loring Professor of electrical and computer engineering and director, Advanced Research Institute, Virginia Tech), Professor Bernhard Hämmerli (President, Acris GmbH / University of Applied Sciences Lucerne, Switzerland), and Dr. Mohamed Eltoweissy, (director CyCare Center, Virginia Tech).

The Challenge

Today's PCs are connected to the internet in two ways: Direct via IP and over security software like SSL, TLS etc. This dual connectivity enables attackers to intercept the application or the operating system at places where the data are not yet encrypted. Therefore, a lot of money is stolen in E-Banking applications and secure applications are more fiction than reality today.

Today, PC applications and operating systems have a very large and diverse array of functions. This makes it increasingly prohibitive for software systems to remain error-free and without an attack surface regardless of enhancements in both the software engineering side and the software implementation side. Due to these conditions the net-connected PC will remain vulnerable and prone to attack through the insecure Internet connection.

Vision

Raise the Bar: **PC- Integrated Highly Secure Net-Terminal (PC-IHS-NT)**

The PC will be enhanced with an additional display and a highly secure module such that proven highest security standards will be met, with theoretically no attack surface. When this PC- Integrated Highly Secure Net-Terminal is used, a fail-safe security can be provided to enable risk-free E-banking, electronic voting, document signing, etc.

In addition to its technical contribution, the PC-IHS-NT will also have an enormous economic impact. The PC will qualify for a completely new set of applications with a huge potential. Banks, trusted partners and even governments can use the same unique electronic identity of citizens. With PC-IHS-NT customizing processes, provisioning and distribution will be done only once for a person and not for every transaction he/she engages in.

Technically, the encryption keys may be certificate based such as what Austria has given its citizens, in a separate highly secure compartment on the PC. This compartment has only one trivial interface to the PC (e.g. USB), runs under its own and dedicated operating system, using exclusively a strong encrypted channel to the other secure endpoint at business partners premises. The functionality has to be defined in detail, but it should be very simple:

- Displaying text (one line, maximum a couple lines)
- Sending verification codes to enable already defined transactions
- has a Yes and No button to prove or reject verification codes

The simplicity of the device enables a complete security evaluation. Today such devices already exist, but exclusively as a security solution for one specific business. The big value of the PC-IHS-NT solution will be the pervasive usage in multiple domains.

Mission

Existing security devices will be analyzed according to SWOT methodology. A stakeholder group will be set up to verify commonly accepted functionality and the design of the PC-IHS-NT. Field observations will be used for connection testing. A pilot series will be used for field testing and for the necessary interaction suites on the remote server side.

Following is a metric to judge the success of this project. If five years down the road, more than 70% of the PCs have PC-IHS-NT integrated, then the project will be considered to be a success.

Concept

Initial situation as described above:

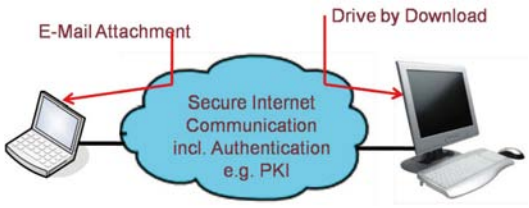

Fig. 1: Secure and insecure channels

Figure 1 shows that in addition to a protected channel over the internet, there also exist insecure channels which may provide an attack path to insecure points in the OS / application software. The most challenging attacks are Trojan horses in E-mail attachments and drive by downloads.

The PC-IHS-NT concepts overview:

Fig. 2: Security Server and PC-IHS-NT

The green elements are PC-IHS-NT, representing a minimal software system set exclusively for confirmation/ verification of secure transactions of other systems. The attack surface is nearly zero, because of the simplicity of the device. Exclusively secure sessions are possible from and out of PC-IHS-NT. The Internet is used only for transportation of security containers, which will never be packed or unpacked in public Internet connected systems. The server will make transactions (e.g. banking, voting in an election) via https in a regular way. The PC-IHS-NT serves exclusively for reconfirmation purposes of displayed massages, e.g. the amount of a banking transaction or the name of the person voted for.

Why is PC-IHS-NT necessary?

Today's security devices are ready to provide such functionality in a B2C situation. The cost of replicating this functionality for each business is enormous and is one of the biggest obstacles for secure computing. Since most of IT research today focuses on providing solutions and rationalization, this proposed idea can make a huge difference and move the security level to a higher plane.

Dream Team

Cybersecurity, embedded systems, data storage and retrieval, and secure communication systems.

Who you are: Bruce Sturk, Director of Federal Facilities Support, City of Hampton, City Managers Office, 22 Lincoln Street, Hampton, Virginia, 23669-3522 (BIO included)

Game-changing dimension: Create a world-class "Municipal Cyber Lab" supporting the Comprehensive National Cyber security Initiative. This lab would simulate a municipality under a cyber attack. Currently municipalities have very robust teams and infrastructure to support emergencies of the physical type, i.e. floods, hurricanes, fires, etc. However, cyber attacks are usually addressed as administrative IT staffs function without the same level of attention, resources and processes that are associated physical threats. Risks to the community for cyber attacks are not identified nor incorporated in municipal emergency operations plans. This program focuses on changing the game by incorporating cyber attacks into the mainstream of emergency operations at the local level and creating a "virtual municipality" of randomly generated internet protocol addresses. The concept would allow capabilities, processes and procedures to be developed.

Concept: Create a node or laboratory that would allow local governments and first responders to plug into state and federal entities and participate in simulated cyber attacks. The lab would be both physical and virtual in structure allowing many types of simulation and also be a place to share ideas amongst stakeholders across the nation. Participants would brainstorm and identify cyber attack scenarios. Scenarios would then be chosen to run in the physical and virtual environments in order to assess the impact of the attack on the community from a physical and virtual prospective. Running the cyber scenarios would help identify the processes, procedures, capabilities and gaps in protection.

Vision: Cyber security must be addressed and understood at local governmental levels in order to provide all citizens within a community the means to continue normal life function during potential cyber attacks. High level strategies in Federal government often fail to incorporate what the consequences of cyber attack would do to a local community. How does the local community government prepare and respond to cyber attacks? What resources, procedures and polices will be implemented or are in place to support local communities under cyber attack? Cyberspace has changed the fundamental assumptions of everyday life all individuals. With the advent of networks and interactions between machines and humans, software is now required to address a wealth of new aspects which potentially impact the ways municipalities react in cyber related incidents to serve their citizens. Procedures, policies and capabilities which enable localities to operate in the cyber domain must be considered simultaneously with state and federal government agencies responsible for cyber security. These aspects include: real-time communications, security, privacy and trust as well as economic, social and societal implications. Affordability is a crosscutting concern, a major constraint in future cyber domain capabilities, designs, procedures and policies. These all require new collaborative insights and approaches. In a world of spontaneously evolving cyber related activities which ultimate impact all individuals, communities must be prepared to implement procedures to ensure basic services are provided during potential cyber related attacks. Of particular interest would be a process where analyzing approaches that unify seemingly disparate levels of local, state and federal government entities responsible for implementing cyber security could take place. The value of a collaborative (municipal cyber lab) approach is its anticipated benefit and hopefully significant improvement in producing procedures, processes and capabilities which enable better a understanding of potential cyber security architectures, for trust and affordability. "When it comes to cyber security, government and the private sector need to recognize that an individual vulnerability is a common weakness." Melissa Hathaway, DNI Cyber Security
Sample scenarios are seen below:

- Cyber Denial of Service (DOS) attacks during or immediately following a physical natural disaster. How to identify the attack in a crisis? How to respond and track this during an event with human resources under stress.

- Simultaneous cyber attacks on multiple critical organizations in a community i.e. local government's public safety, local schools, traffic lights and management, water, power, telecom.
- Attacks aimed at providing misinformation to the public from several sources. Governments are relying more and more on electronic methods and the public media to communicate with the public to communicate critical information to the public especially during times of crisis. Cyber attacks coupled with the lack of authentication of critical messages could lead to putting the public at risk. For instance e-mail, message broadcasts, website information to the press, major employers, community leaders, etc. that appears to come from City officials but is actually from another source. For instance the message could instruct citizens to evacuate a community at the wrong time and into a dangerous situation.

Method: Virginia's Operational Integration Cyber Center of Excellence (VOICCE)

An **environment** that:

- Simulates a city or community (municipality)
- Tests cyber related events or impacts
- Evaluates and assess current cyber processes and procedures
- Supports National level organizations cyber capabilities development

An **operational facility** that:

- Hosts systems
- Hosts technical expertise
- Hosts the physical and virtual facilities
- Creates scenarios

A **training facility** to facilitate:

- Experimentation
- Structured evaluations
- User workshops

The facilities would look at cyber attacks both with traditional systems and also through new types of technologies being deployed to the public like Web 2.0. New technologies and applications like social networking, blogs, virtual communities, Wikis, RSS, podcasting, microblogging, video sharing, photo sharing, wireless networks, new handheld devices and others are growing and being used extensively by the citizenry. The hosted facilities would incorporate these types of new technologies into the facility and scenarios.

Dream Team: (Specific offices and individuals will be identified)

- Department of Homeland Security
- DoD
- Federal installations located in or near participating local and state governments
- Industry located in or near participating local and state government
- Telecommunications carriers and ISPs serving participating local and state government
- Academia
- Individual State Government Participants
- Individual Local Government Participants
- State Organizations like Multi-State Information Sharing and Analysis Center, National Association of State CIOs (NASCIO), National Association of State Technology Directors (NASTD)
- Local Government Organization like MuniGov 2.0, Mix, etc.
- Public Safety Organizations

Who you are -- Voltage Security (www.voltage.com) -- we are an encryption technology company that focuses on creating innovative applications of encryption that help our customers protect their sensitive data.

Game-changing dimension -- Change the rules.

Concept – Make encryption ubiquitous to protect sensitive data from cybercriminals.

Vision -- Encryption is an easy way to protect sensitive information, but encryption is often too difficult to use, particularly in legacy IT environments. Encrypting data typically changes both the format and the size of data. In legacy environments, this means that there's a very good chance that some part of an older system will be unable to handle encrypted data, which makes encryption difficult to fit into the "data-centric" architectures that are being proposed today.

It's almost like a return to the Y2K problem, in which many applications were hardwired to only handle a two-digit number. In legacy computing environments it's not uncommon to have systems that are hardwired to handle a nine-digit quantity because it's a Social Security number. In cases like these, encryption that changes the format of data or increases its size can cause trouble, and traditional ways to encrypt data typically do both of these.

An encryption technology that preserves the format of encrypted data will make it possible to protect sensitive data in any IT environment – even those using older legacy systems. There are ways to accomplish this that have rigorous mathematical proofs of their security, so that there's no concern about a lack of security in the technology.

We have developed the basic tools that are needed to make this a reality, but there not yet integrated with common applications. Integrating this technology with popular applications will be a significant first step towards keeping sensitive data safe and reducing the amount of identity theft that many people suffer from today.

Method – We have created our first set tools that implement format-preserving encryption after consulting with the CSOs and CIOs responsible for large IT operations, both in the public and private sector. Their overwhelming support for the idea of format-preserving encryption led us to invest in the development of the first implementations of the technology.

Dream team – Department of Commerce/National Institute for Standards and Technology, Department of Homeland Security.

Who you are -- Voltage Security (www.voltage.com) -- we are an encryption technology company that focuses on creating innovative applications of encryption that help our customers protect their sensitive data.

Game-changing dimension -- Change the rules.

Concept – Make Government to citizen secure communications a reality by using identity-based encryption.

Vision – There are many cases where it would be useful for Government agencies to send messages that contain sensitive information to private citizens, but to do this, the sensitive information needs to be encrypted if it is sent over public networks. This has traditionally been difficult because the technologies that Government agencies use to communicate securely both internally and with each other are often not well suited for use by the general population. Although Government employees may have a PIV card that lets them digitally sign and encrypt e-mail messages, the typical citizen does not have the necessary cryptographic keys that are needed. This makes secure communications outside the Government community difficult.

The military health care system provides an example of why the current technology does not solve the complete problem that the Government faces. It is easy to use the keys from Government employees' PIV cards to encrypt information to them, but military dependents and retirees do not have PIV cards that can be used in this way. This means that it is currently infeasible to electronically deliver many health care documents to dependents and retirees because it is infeasible to encrypt the documents that contain the sensitive information. Identity-based encryption provides an alternative to traditional encryption technologies that makes such communications feasible.

Although the problem of creating a practical and secure identity-based encryption scheme was first posed in 1984 by Adi Shamir, it was not solved until 2001, when Professors Dan Boneh and Matt Franklin invented what is now known as the Boneh-Franklin identity-based encryption scheme. Their work was sponsored by DARPA and the NSF.

Identity-based encryption uses an identity for an encryption key. This makes distributing keys unnecessary. If an e-mail address is used for an identity, for example, then it is possible to use the recipient's email address as their public key and to use that key to encrypt sensitive information to them. So even though the typical private citizen does not have a PIV card, they do have an identity of some sort, and this identity can be used to encrypt information to them.

By working with Federal agencies that need to communicate sensitive personal information to private citizens, this project will let the Government leverage their DARPA and NSF research into a useful way to provide secure communications that will let the Government realize significant cost savings and efficiency gains.

Method – The first projects using identity-based encryption used the technology in Government applications to communicate securely to first responders to a disaster. Based on the positive feedback from these Government exercises, Voltage commercialized their identity-based encryption technology and began selling it to the private sector. Subsequent discussions with Government representatives at the Identity-Based Encryption Workshop that was recently held by NIST indicated that there is significant interest in the technology from Government agencies.

Dream team – Department of Commerce/National Institute for Standards and Technology, Department of Homeland Security.

Who you are. Michael Walfish, J.D. Zamfirescu, Hari Balakrishnan, David Karger, and Scott Shenker. We are computer scientists and one entrepreneur. Walfish is an assistant professor at UT Austin. Zamfirescu is a principal at AppJet, Inc. Balakrishnan and Karger are professors at MIT. Shenker is a professor at UC Berkeley and head of the networking group at the International Computer Science Institute (ICSI).

Game-changing dimension. Raise the stakes.

Concept. We control spam with email quotas, using "Distributed Quota Enforcement" (DQE). Today's gameboard looks like this: spammers and legitimate people send emails, and the recipient's spam filter makes a guess about whether each email is "good" or "bad". When the spam filter miscategorizes email from legitimate clients, which we argue is inevitable sometimes, the cost is extremely high: a missed opportunity, a missed apology, a misunderstanding, etc. Worse, the *fact* of filtering—the *shape* of the email gameboard—means that email is no longer a reliable communication medium: a sender cannot be certain that a receiver saw a given email. Thus, filtering exacts a steep price.

Under DQE, depicted in Figure 1, the gameboard looks different: each sender gets a quota of stamps and attaches a stamp to each email. Receivers' computers do not make judgments about the content of a message; instead, they communicate with a well-known quota enforcer to verify that the stamp on the email is fresh and to cancel the stamp to prevent reuse. The receiving host delivers only messages with fresh stamps to the human user; messages with used stamps are assumed to be spam. The intent is to set quotas such that, unlike today, no one can send more than a tiny fraction of all email. Because spammers need huge volumes to be profitable, such quota levels would probably drive them out of business. However, even if they remain solvent, the system does its job: the fraction of our inboxes that is spam will be negligible, and filtering is absent.

Email postage has been proposed before, so why DQE? DQE makes email postage practical, via two innovations. First, DQE separates the allocation step (which can happen at infrequent time scales, like once yearly per user) and the enforcement step (which needs to happen "online", as an email is sent). The result of this separation is that a range of allocation policies become possible, some of which are "analog" and require great care (e.g., the allocator checks users' drivers licenses, or verifies that they have paid a certain amount, or checks that the user has paid an ISP, etc.) Second, prior to DQE, there was no way to prevent double-spending of email postage at the volume of the world's email, which is roughly 200 billion email messages per day, or several million checks per second. DQE's enforcer, in contrast, can scale to this workload with just a few thousand machines (which isn't that many, given the sizes of data centers). Moreover, DQE's enforcer tolerates faults in its constituent hosts, is untrusted by its clients, does not require its constituent hosts to trust each other, resists external attack, and avoids heavyweight cryptography. As a result, DQE makes email postage technically viable.

Vision. With DQE deployed, spam filters (which we argued above are harmful) are no longer needed. Spam becomes a tiny fraction of all email in our inboxes. The needed components would be deployed at mail servers, so users would never need to "see it", though technical users could interact with quotas and stamps, at their option. We think that it is possible because we have built a prototype and because we have whittled the technical components to an unadorned, simple-to-

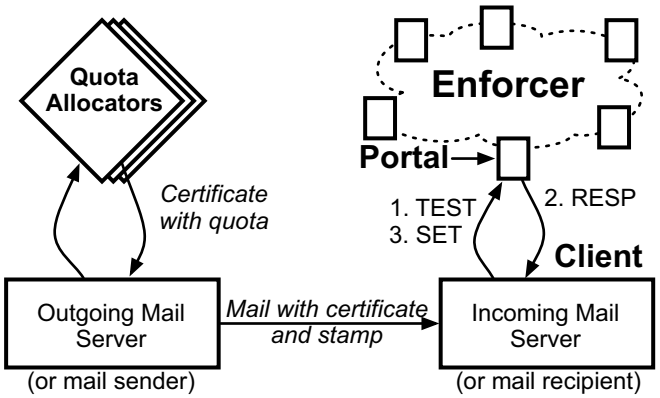

Figure 1: DQE architecture.

deploy infrastructure. The part that needs to come into being is the non-technical part: one or more quota allocators that would accumulate the required trust. There are multiple options here, including a consortium of email providers agreeing on an allocator, a competitive market in which allocators could compete based on their reputations and prior performance, etc.

Method. After working out the design, we prototyped it, analyzed it, and experimentally evaluated our prototype. We published the results in NSDI, a top-tier conference in networked systems [2]. In addition, a significant fraction of Walfish's dissertation is on this topic [1]. Thus, our work has undergone rigorous peer review (NSDI reviewers and Walfish's dissertation committee). As a result, we believe that the idea is sound. The main dependencies are as follows: (a) DQE requires quota allocators, described above; and (b) for DQE to be maximally useful, it needs widespread adoption. Addressing these dependencies will require political will.

Dream team. Senior and operational staff from (a) the big email providers (Yahoo, Hotmail, etc.); (b) some medium-sized email providers; (c) large companies (over 5000 people); and (d) ISPs.

References

[1] M. Walfish. *Defending Networked Resources Against Floods of Unwelcome Requests*. PhD thesis, Massachusetts Institute of Technology, Nov. 2007. `http://nms.csail.mit.edu/~mwalfish/diss.html`.

[2] M. Walfish, J.D. Zamfirescu, H. Balakrishnan, D. Karger, and S. Shenker. Distributed quota enforcement for spam control. In *NSDI*, May 2006. `http://nms.csail.mit.edu/papers/index.php?detail=145`.

Who you are. Michael Walfish, Mythili Vutukuru, Hari Balakrishnan, David Karger, and Scott Shenker. We are computer scientists. Walfish is an assistant professor at UT Austin. Vutukuru is a PhD student at MIT. Balakrishnan and Karger are professors at MIT. Shenker is a professor at UC Berkeley and head of the networking group at the International Computer Science Institute (ICSI).

Game-changing dimension. Raise the stakes.

Concept. To defend against *application-level* denial-of-service, in which attackers cripple a server by sending legitimate-*looking* requests to overwhelm the server's computational resources (e.g., CPU, disk), we use "Fighting Fire with Fire", or "Defense by Offense". Rather than trying to reduce the sending rate of attackers, a server so victimized (a) encourages all clients to send it *more* traffic and (b) randomly selects a fraction of the incoming requests for service. With this defense, an attacker's standard for success moves from "sending more traffic than good guys" to "having more aggregate bandwidth than good guys", which is far harder to meet.

Why does this defense work? Attackers are already using most of their bandwidth (that is what it means to be an attacker) so cannot react to the encouragement. Good clients, however, have spare bandwidth and react to the encouragement with drastically higher volumes of traffic. As a result, the traffic into the server inflates, but the good clients are much better represented in the traffic mix than before and thus capture a much larger fraction of the server's resources than before.

When is this defense called for? When the denial-of-service gameboard is "stacked" against the server in the sense that an attacked computer, when presented with a request, cannot tell whether the request is from a "good" or "bad" client. Reasons include well-formed requests from bad clients and the fact that differentiating good and bad *computers* may not be feasible because a bad computer may adopt multiple identities and appear to be several hundred good computers. In those cases, a very small number of bad clients can claim a very large fraction of the server's resources. "Defense by Offense" changes this mix without requiring the server to differentiate good and bad.

Vision. With "Defense by Offense" in place, sites could avoid massive over-provisioning (which is costly) or detect-and-block defenses (which are error-prone). More generally, denial-of-service would be less effective and thus it would stop. To deploy the defense, all that is required is for a server owner to place a middlebox in front of the server. The solution works with today's Web browsers, unmodified. We think that the defense is viable because we have built prototypes for all of the parts, and there is no fundamental technology obstacle. For the defense to "become real", a production version must be developed, and sites must deploy it.

Method. After working out the design, we prototyped it, analyzed it, and experimentally evaluated our prototype. We published the results first in the HotNets workshop [2] and then in SIGCOMM, the top conference in computer networking [3]. In addition, a significant fraction of Walfish's dissertation is on this topic [1]. Thus, our work has undergone rigorous peer review (HotNets reviewers, SIGCOMM reviewers, and Walfish's dissertation committee). As a result, we believe that the idea is sound. It is most useful under three assumptions: adequate good client bandwidth, adequate server bandwidth, and bad clients are using most of their bandwidth. There are ways to make these assumptions hold, but even if they do not, the defense is still useful.

Dream team. Senior and operational staff from (a) ISPs; (b) DoS-prone organizations; and (c) big Web sites (Google, Amazon, etc). Also, staff from the NSA (what kind of attacks are launched on their computers?)

References

[1] M. Walfish. *Defending Networked Resources Against Floods of Unwelcome Requests*. PhD thesis, Massachusetts Institute of Technology, Nov. 2007. `http://nms.csail.mit.edu/~mwalfish/diss.html`.

[2] M. Walfish, H. Balakrishnan, D. Karger, and S. Shenker. DoS: Fighting fire with fire. In *HotNets*, Nov. 2005. `http://nms.csail.mit.edu/papers/index.php?detail=138`.

[3] M. Walfish, M. Vutukuru, H. Balakrishnan, D. Karger, and S. Shenker. DDoS defense by offense. In *SIGCOMM*, Sept. 2006. `http://nms.csail.mit.edu/papers/index.php?detail=148`.

Who you are. Michael Walfish, David Mazières, and Jad Naous. We are computer scientists. Walfish is an assistant professor at UT Austin. Mazières is an associate professor at Stanford. Naous is a PhD student at Stanford.

Game-changing dimension. Change the rules.

Concept. Today, the network security gameboard is highly favorable to attackers. First, attackers are free to evolve new exploits, but their victims—being constrained by legacy protocol fields, installed infrastructure, and the difficulty of worldwide coordination—have no such freedom to evolve defenses. Second, the landscape contains a mismatch: from the attacker's perspective the network is highly connected (as just one example, incorrect routing advertisements in Pakistan can knock a Web site in New York offline), but from the defender's perspective it is highly fragmented (as just one example, defending against a denial-of-service attack often requires help from a disinterested third-party ISP). Third, while attackers enjoy many different lines of attack, there are only so many defense mechanisms that one can deploy in the fabric of the network. The literature is full of attack-specific defense mechanisms, but deploying them all would lead to partial redundancy and performance costs. Realistically, the network must make do with one or a small number of protection mechanisms that can address a much larger number of attacks.

Thus, we need a fixed set of defense mechanisms that can address both today's and tomorrow's threats. We also need a network infrastructure that supports these mechanisms. Such an infrastructure has several requirements. First, policies must be expressible in *servers* (which can be updated), not core forwarding infrastructure (which cannot). Second, policies must be expressible in globally meaningful terms. For example, if a hostile country attacks a US Dept. of Defense network through US ISPs, the DoD's policy must be expressed in such a way that US ISPs stop the attack at the borders, long before the attacking traffic reaches the DoD network. Third, the rules must change from "any sender can send traffic to any destination, and destinations can try to block traffic based on the apparent sender and various ad hoc heuristics" to "*traffic only flows if all entities (hosts and providers) along the path approve of the entire path*". We call this last requirement consent-to-connectivity (C2C); it unifies many aspects of network security with a single mechanism.

Vision. We believe that the principle of C2C described above should be incorporated into a new network layer that we call ICING (Incorporating Consent in the Internet's Next Generation). Realizing ICING would allows us both to defend against a range of attacks in the present *and* to evolve solutions for future threats without needing further modifications to forwarding infrastructure. In the present, ICING would rule out attacks such as denial-of-service, route hijacking, source IP forging, silent dropping of packets by misbehaving intermediate providers, and more. More generally, the requirement that traffic have permission from *all* entities creates a high burden for those who would send unwanted traffic.

But more importantly, ICING would adapt. The reasons are that (a) C2C incorporates the opinions of all of the stakeholders; and (b) those stakeholders' policies are expressed in a modular way, outside the core of the network. For example, if a new class of traffic or particular senders become problematic, receivers or providers withhold their consent, thereby preventing the traffic from flowing through upstream providers (where today such filtering would require out-of-band communication between network administrators).

We think that ICING is feasible because we are currently prototyping it, and there appear to be no fundamental technology obstacles. The heart of our prototype is a forwarder that carries ICING traffic. Our implementation of the forwarder consists of a software control plane, a software slow path, and a hardware fast path. The focus of our implementation effort is making the needed checks operate at line speed.

For ICING to "become real", providers and organizations must adopt it. In the short term, an organization can install an ICING forwarder as a gateway at the edge of its network, while not touching the rest of its network. Our long term vision is that a significant fraction of network traffic will be ICING packets.

Method. We started by observing that many different research and industry groups have proposed many different security policies, but that the various mechanisms are incompatible. We asked whether there was a way to unify them and at the same time to enable *additional* notions of security. Our hypothesis is that ICING and C2C are a way to do so and that it is practical to uphold C2C. To validate this hypothesis, we are developing the prototype ICING forwarder described above. The main dependency or assumption is that organizations (ISPs, companies, universities, etc.) want the stronger security that ICING provides, relative to the status quo, and that they are willing to install ICING forwarders, as described above.

Dream team. Senior and operational staff from (a) ISPs; (b) large organizations (universities, corporations, etc.); and (c) network equipment providers (Cisco, Juniper, etc.)

Contact: Don Ritzman
200 Bachelder Ranch Rd
Santa Cruz, CA 95065
(831)459-8199
(888)459-8199

Who Are We? – www.whitenoisesystems.com – We are a small business with 20 members dedicated to providing complete data protection solutions that absolutely assure data cannot be lost, stolen or compromised and that your data is there when you need it.

Our data security professionals have dedicated over 30 years of their professional lives addressing the data protection problem by developing comprehensive solutions to the issues. Our experience is broad, deep and detailed in all aspects of data security.

Game-changing Dimension – Morph the game board

Concept – Cyber-thieves are well funded, highly resourceful and tremendously successful at pulling off "big hits" – stealing mass amounts of confidential data in quick-in quick-out events. According to a recent study, the Average Tangible Cost (ATC) per corporate breach is $7.5M – a 40% increase over the previous two year averages.

In today's volatile and dangerous marketplace, it's not enough to rely on perimeter security mechanisms. Many security experts are saying that perimeter security has outlived its effectiveness and some even go so far as to suggest it is "Dead". The new business mantras are: SaaS, Cloud Computing and unified data protection.

SecureNOW!™ is a family of SaaS data protection products and services that brings the positive aspects of SaaS and Cloud Computing to users by adding "Opaqueness" to user data which protects that data at the point of creation and maintains total data protection throughout the entire data lifecycle.

SecureNOW!™ supports individual users, multi-national corporations and government entities with the same core technologies and provides users with Data Confidentiality, Data Integrity, Continuous Data Availability, Business Continuity, Disaster Recovery, and a new Total Cost Ownership (TCO) for crypto-key management.

SecureNOW!™ is comprised of four tightly integrated components. Two of the company's proprietary FIPS 140-2 certified technologies – Data Shred and Stitch Technology (DSST)™ and Constructive Key Management (CKM)® – complement a fully integrated user interface and the company's Opaque Data Cloud™.

With our Data Shred and Stitch Technology (DSST)™ and Constructive Key Management (CKM)®, SecureNOW!™ turns clear data files into eight non-deterministic shreds, then wraps each shred with (CKM)®, creating Opaque Data Objects that are able to protect themselves and still be readily available to authorized users when and where they need to have access. The shreds are then showered onto eight geographically dispersed data centers in the company's Opaque Data Cloud™ where they remain as Opaque Data Objects. When the owner needs to use the data, only four of the original eight shreds are required for SecureNOW!™ to successfully re-stitch the data into its original format.

Users access the WhiteNoise Opaque Data Cloud™ infrastructure to use a secure environment for sharing information with communities of interest (COI's) confident knowing that SecureNOW!™ supports cross enterprise data protection (CEDP) by providing a mechanism that enforces data protection policies on data that traverses the corporate firewall, preventing unauthorized access or theft.

One of the most critical issues to making COI's and CEDP's practical, is an efficient way to manage the crypto-keys. SecureNOW!™ uses our proprietary (CKM)® technology to provide an efficient way to manage crypto-keys supporting a new total cost of ownership (TCO) for crypto-key management and making SecureNOW!™ the most comprehensive data protection solution available in the market today.

Deploying SecureNOW!™ data protection solutions, eliminates expensive (and embarrassing) data breaches and losses, provides business continuity and disaster recovery, while supporting government regulations such as Sarbanes and HIPPA and helps to ensure customer confidence, at a fraction of the cost/user/year of a data breach.

Vision – The vision is to provide the market a highly secure data repository and virtual data exchange environment that is uniquely theft-proof. All data in the environment is opaque – available ONLY to the data owner and those individuals to which the data owner chooses to make the data available.

Our approach was to build this highly secure solution from the perspective of the data itself and not the infrastructure. Since the data is the jewel, we determined that the ONLY best way to protect the jewel is to create an environment where the Jewel is able to protect itself throughout its entire life-cycle no matter where it resides on or off the infrastructure and no matter who gains access to the infrastructure on which it resides.

Users can purchase a SaaS subscription and download the SecureNOW!™ client application onto their desktop or laptop from the company's website. For large-base users, SecureNOW!™ can be licensed and fully managed by the organizations' IT department. Once the SecureNOW!™ application is loaded, the user can begin securing their data at the point of creation: they do so in exactly the same manner as they save files from their applications today – "save", "save as", or drag and drop within Windows Explorer.

With SecureNOW!™, users are no longer hampered by cyber-thieves – users' data is ALWAYS secure and available ONLY to the owner. As more and more users move to SecureNOW!™ solutions for their data protection, the data they own and share with their (COI's) becomes more and more resilient to cyber-theft, ultimately eliminating the "big-hits" of stolen data. This creates a uniquely safe environment for users to do their business and personal activities in the cyber-world.

The complete SecureNOW!™ data protection solution is currently available on the Windows XP platform. Work has begun to add Windows Vista, Mac OS-X and Linux and eventually to provide the SecureNOW!™ client for mobile devices as well.

Method – WhiteNoise spent two years in development including six months in Alpha and twelve months in Beta testing of SecureNOW!™. Fifty tier-1 prospective customers from the company's target markets participated during the Beta test phase, which finished in October, 2008. We took the feedback from the Beta customers and designed in many of the suggested changes to create the best possible data protection solution for the industry.

Dream Team – Decision-maker(s) from: The Transglobal Secure Collaboration Program www.tscp.org, Dept. of Homeland Security, Iron Mountain, EDS/HP, MXI, VirtualBox, McAfee, and The Department of Defense.

Government Deployed Identity Management System

One person–one citizen–one identity management key

This is possible because of the unique characteristics of new generation, deterministic key streams.

Identity management key

- > 400 trillion bytes long key streams
- > 240,000 bit strength

The life of a Canadian

A Canadian is born and the government issues an electronic identity management key that is first associated with their birth certification which is represented by a specific range on a key.

The newborn is then issued a health and insurance card which is represented by a different, specific range on the same key.

Throughout their lives different government services are represented by different, unique ranges on the same key for passports, drivers' licenses, tax numbers etc.

The keys are far, far more unique than an individual's DNA. The uniqueness of the keys allow for different, dynamic, distributed topologies.

These topologies overcome the traditional stoppers that have historically been associated with distributed key networks.

DIVA identity management system

The highly secure Identity Management system called Dynamic Identity Verification and Authentication (DIVA™) utilizes unique features of Whitenoise™ and provide an integrated security system that people will use because it doesn't slow them down.

The Whitenoise Identity Management key provides continuous, state based identity verification and authentication of a user throughout the session and not just at login. Dynamic Identity Verification and Authentication [DIVA] provides inherent intrusion detection because the offsets must remain in sync, and automatic denial of network access to hackers and spoofing. This is a technological capability not seen to date.

This identity management system provides multi-layered user access security. Users are issued a unique Identity Management access key that employs user ID and password protection as well as the requirement for the physical presence of a key or device. Any other layers of authentication can be used in conjunction. The key allows access to the system both locally and remotely. Keys can be issued to partner companies as well.

Secured information is exchanged with predefined rules. Information can be sent to one or more recipients, or predefined groups, in one operation.

DIVA Identity Management is simple to implement and is designed to fit within existing security schemes. A very important feature of the system design is the ability for the system administrator to deactivate lost or stolen keys immediately.

Traditional problems solved

Key management of these systems explodes into an exponential headache.

(Historically the number of keys to manage is the square of the number of secure endpoints on a network.) DIVA Identity Management has a one-to-one relationship between the number of keys and endpoints on a secure network.

Key storage – long keys are a better source of identification and security but storing large keys is a nightmare.

Whitenoise creates keys that will generate unique key streams on the order of 10^{60} bytes in length. However, only the internal key structure and the offset are required to recreate any key segment. This is a small amount of data. For example, 158 bytes of this information will generate a random key stream over 1 billion bytes long. You can learn about multiplicity in conference presentations on our technology page.

Key distribution is a major problem for distributed key systems.

This is not true any longer – Whitenoise topologies allow distributed keys to in turn securely generate and distribute more encrypted keys. It allows the easy creation of secure tiered networks.

With all the traditional problems solved, DIVA Identity Management provides a secure digital network architecture that is far easier and less expensive to use than asymmetric key systems and there is NO reliance on Trusted Third Parties (outside of the government/law enforcement) for your security.

Distributed symmetric systems have always been the prevalent architecture and are the approach that has the least impact on user behaviour and is the architecture that consumers worldwide are familiar with. This is evidenced by all your important documents that you carry daily: drivers' licenses, credit cards, employee ID cards and passports are all examples of distributed keys that you rely on daily.

The flexibility of the DIVA Identity Management architecture allows the systems to be used with existing public key systems to add continuous authentication, 100% accurate inherent intrusion detection, and automatic denial of network access to criminals. Add the DIVA to your security protocols without replacing existing systems and without the need for additional hardware. All you require is an Internet connection.

Government sets the standard and expectations

The government sets the standard for which key segments are used for which identification or service. It is a token, key segment – it is a ***verifiable subset*** of a user's identity that the government/law enforcement can verify.

These kinds of keys are deployed in **Dynamic Distributed Key** tiered architectures. These systems are distinguished by the ability of distributed keys to dynamically create and distribute more keys securely and electronically.

So we can easily issue such keys.

What do secure networks require?

Only three things:

1. All components of the network are identified by a unique key
2. All persons on the network are identified by a unique key
3. All usage is logged

Why is such a key secure?

To break keys when they are used for **encryption** there are three pieces:

1. Plain text
2. Cipher text
3. Encryption key

One needs sufficient information from two of the three components in order to break the key.

Identity management keys **<u>are not used</u>** for encryption. There is only one piece – the key itself.

There are only three ways to break a key in this context:

1. One must discover internal linearity characteristics or a mathematical relationship. These keys are structural in nature so mathematical techniques do not work.

2. One must capture at least 50% of the key stream.

   This leads to the question of "How can a criminal capture that much key volume from aggregating the tokens (key segments) that represent different services?

   - The smallest identification key is 10^{60} bytes in length.
   - To break the key this way, a criminal needs to capture 10^{30} of key stream or 10,000,000,000,000,000,000,000,000,000,000 bytes of key stream from various sources.
   - If the average token for a service is 1,000,000 bytes then the hacker would need to capture those tokens from a minimum of 10,000,000,000,000,000,000,000,000 different services. This is not feasible.

3. One can do brute force attacks.

   Brute force to guess the key is the only alternative left if one had the computing power. That is the process of testing every possible key. To apply brute force to a single key is not feasible.

   > *"Exhaustive key search is not a threat.*
   >
   > *Whitenoise uses keys with at least 1600 bits of randomness. ... Even if we hypothesized the existence of some magic computer that could test a trillion trillion key trials per second (very unlikely!), and even if we could place a trillion trillion such computers somewhere throughout the universe (even more unlikely!), and even if we were willing to wait a trillion trillion years (not a chance!), then the probability that we would discover the correct key would be negligible (about $1/2^{1340}$, which is unimaginably small).*
   >
   > *In this report, I tried every attack I could think of. All of them failed. This provides evidence for the hypothesis that Whitenoise is cryptographically secure."*
   >
   > **-Professor David Wagner, University of California, Berkeley, October 2003**

So the keys are secure.

What is the problem?

It is scary that after three years that governments cannot even agree on terminology. Instead of trying to define the terms, we need to define the outcomes we want to avoid.

- disparate systems, conflicting standards, decades of unorganized implementation, non-interoperability
- theft, crime, etc

What is the conceptual solution?

Government is the biggest fish and sets the rules (democratically)

Government IS THE TRUSTED THIRD PARTY IN ALL CASES

First - encapsulate the problem

Encapsulate the problem like surrounding an oil slick so it doesn't spread. We are at a critical period of time in regards to security of critical infrastructures etc.

It is a top down solution.

Top tier

Government issues keys for every telecommunications provider and networks (link keys) and they reside on the government authentication server.

Government issues keys for every citizen. That is about 35 million keys easily stored on a government authentication server.

Next tier

The carrier issues keys for every business providing a service. This is a token or subset of their carrier identifying link key, which itself is a subset of a master key which the government issues and regulates.

Final tier

Citizens/clients have their Government issued electronic identity management key. As they use this key for any possible service, only the token for that particular service is accessed. These can be identified and used for registration or subscription to electronic

services as belonging to a specific citizen since **the Government is the ONLY PARTY THAT HAS THE COMPLETE KEY OR KEY STRUCTURE.**

Law enforcement is the tier linking the government and the public to resolve disputes and address crime

LAW ENFORCEMENT CAN ACCESS BOTH THE TOKEN IN QUESTION THROUGH SERVICE PROVIDERS AND COMPARE AGAINST GOVERNMENT REPOSITORY in court ordered scenarios.

This tiered approach encapsulating the entire ID Management problem and all networks will enable organized, secure co-existence among dysfunction family members/networks. Over time, of its own volition and evolution, redundant aspects of networks will be removed and disparate networks types will become harmonized as they want the ability to safely communicate with more networks and individuals under the umbrella.

Implementing the system

Government requires that all telecommunications and network providers use Dynamic Identity Verification and Authentication on all login and transaction procedures. It can be a redirect to a government run authentication server for this service. DIVA can be integrated into existing systems of any kind; it can be used in parallel to any kind of network systems; it can be used in lieu of any other kind of network system.

In encapsulating the entire insecure network issues you are mandating that electronic citizens are adding one additional layer in electronic authentication. At the carrier this is the one time addition of three database fields to their client records: unique identifier of person, unique identifier of a device, and current offset. On device firmware/software it is the addition of a small identity application. At today's network speeds the extra step does not impact network performance. Government is simply saying that login protocols will use this system in parallel with any other existing process.

No entity is being asked to change anything in their existing architectures other than this step that ensures continuous authentication, 100% accurate intrusion detection, and automatic denial of network access to criminal behavior.

The government has just become the biggest encapsulating Russian doll and networks and communications are secured and allow accurate identification of everyone and everything on networks. They have simply added one encapsulating protocol at the TOP of the network food chain.

This paradigm allows government to easily and inexpensively address the identity management issues that are part of its legitimate mandate.

The system is simple for everyone involved. The addendum below shows the two ways to configure such a system and the level of “intrusion” or “effort” required.

The system is simple for citizens because they only have one key, just like they only have one identity, and this simplifies or eliminates all problems associated with compliance and password and key fatigue. This system when applied to keymail would eliminate spam as a nice bonus.

ADDENDUM –

There are two ways to configure a Dynamic Identity Management system. This represents the entire technological implementation requirements.

TRUSTED THIRD PARTY SERVICE CONFIGURATION

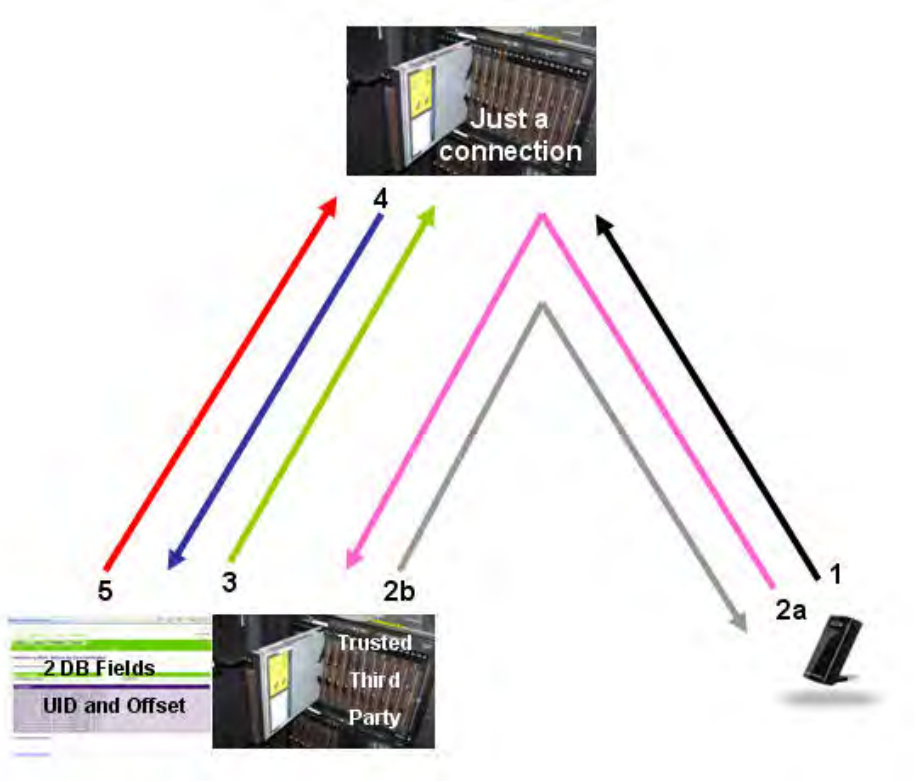

1 - The Aircard or cell phone must connect to the network.

2a – Card sends authentication token through Telco to server

2b – Server sends pass/fail to the Aircard

3 – TTP and TELCO share last login time stamp (prevents bypassing DIVA)

4 – TTP receives timestamp info and makes comparison

5 – On fail, TPP request Telco to terminate connection/deactivate.

With each “Pass” both the endpoint and the server automatically update their current dynamic offset, INDEPENDENTLY BY INCREMENTING THE OFFSET BY THE LENGTH OF THE TOKEN, so that an authentication token is NEVER re-used.

This Third Party authentication configuration can plug the security hole if the carrier shares with the Third Party provider information about when the last time a card using our service was connected. This information is the unique identifier and last login timestamp, neither which is a security risk for the TELCO. This would stop the ability to bypass step 2 and 3 in this authentication process protocol.

IDEAL CONFIGURATION

1. The Aircard begins connection routine to the carrier including an authentication token.

2. The carrier server verifies the authentication token and gives a pass/fail.

    - The carrier needs to add only two fields to the data base managing logins for specific accounts: unique key identifier and last current offset. 64 bit offset for current offset.

- The carrier would add a DIVA authenticating code segment to the firmware or service application manager to thus embed it into the network protocol.

- The carrier would embed the DIVA service into the ATT/network access system on their server.

Note: With each “Pass” both the endpoint and the server automatically update their current dynamic offset, INDEPENDENTLY BY INCREMENTING THE OFFSET BY THE LENGTH OF THE TOKEN, so that an authentication token is NEVER re-used.

All of this can be done with software and electronically with existing systems and so it is the lowest cost approach.

National Cyber Leap Year
RFI

T.R.U.S.T.

Total Reliability Utilizing Standardization & Test

ZTI

Zero to Infinity

Technical POC:
Derek Pearson
705 El Centauro
El Paso, TX 79922

Who you are – ZTI- we are a small business that conceived the concept of T.R.U.S.T. (Total Reliability Utilizing Standardization and Test). This concept has support from both government and commercial entities currently feeling the pain of the status quo.

Game-changing dimension – Morph the game-board.

Concept – Semiconductors are the building blocks of all electronic systems. The rapid migration of semiconductor manufacturing plants to locations outside the United States has resulted in untrustworthy semiconductors and therefore systems becoming a concern. What if we morph the game-board and make the current manufacturers out-going qualification code **portable** so that it can be used as in-coming qualification code at the end-user's facility.

Vision –"Do for the ATE world what the PC did for the Computer world"
ATE (Automated Test Equipment) is a multi-billion dollar market upon which the performance of all semiconductors and therefore manufacturers are dependent. It is on these machines that all semiconductors are tested due to liability issues prior to leaving the factory floor and a paper data-sheet is derived. However, because of the highly proprietary nature of each test platform, portability of the qualification code is currently futile.
An open architecture test system is based on widely-used and commonly-accepted interface specifications, the PC being a perfect example. By establishing an Open-Automated Test Equipment platform with the assistance of the Semiconductor manufacturers, the qualification/test code will become portable and back-ward compatible much like software can be ported from one PC to another. The code will become an electronic version of today's datasheet and allow for continual improvement. By establishing T.R.U.S.T, a reliable source of supply of microelectronics can be ensured and competition will quickly reduce costs and change the game. This idea benefits everyone from DoD having qualified parts, to the contractor developing the system and finally to the ultimate customer, the serviceman/woman in the field who's placing their trust on the system to work.

Method –Talked with end-users of Semiconductor Automated Test Equipment.

Dream team– AT&L- Someone familiar with the '05 DSB report on "High Performance Microchip Supply" **; DMEA** (Defense MicroElectronics Activity) – Fred Fraser www.dmea.osd.mil**; AFRL** – (Air-Force Research Lab) – David Alexander**; ARMY** - RTASSC(Radiation Tolerance Supply and Support Center)-Randy Brady**; SEMI**-(Semiconductor Equipment and Materials International)-Karl Stuber www.semi.org**; IDEA**- (Independent Distributors of Electronics Association)- Debra Eggeman www.idofea.org**; GSA** (Global Semiconductor Alliance) - Lisa Tafoya www.gsaglobal.org; **Honeywell**- Trusted Foundry- Joseph A. Mielke**; Texas Instruments**- Hi-Rel Defense and Aerospace Semiconductor Group- Mont Taylor**; Freescale**- Foundry services- Jeff Todd**; IEEE**- Burnie West**; IC Test Houses**- Amkor, Infiniti Solutions & Test Spectrum

National Cyber Leap Year
RFI-2

T.R.U.S.T.

Total Reliability Utilizing Standardization & Test

ZTI

Zero to Infinity

Technical POC:
Derek Pearson
705 El Centauro
El Paso, TX 79922
(915) 867-9202

Who you are – ZTI- we are a small business that conceived the concept of T.R.U.S.T. (Total Reliability Utilizing Standardization and Test). This concept has support from both government and commercial entities currently feeling the pain of the status quo.

Game-changing dimension – Morph the game-board and Change the rules.

Concept – Semiconductors are the building blocks of all electronic systems. The rapid migration of semiconductor manufacturing plants to locations outside the United States has resulted in untrustworthy semiconductors and therefore systems becoming a concern. http://www.businessweek.com/magazine/content/08_41/b4103034193886.htm

What if we *morph the game-board* and make the current manufacturers out-going qualification code portable, a **digital datasheet**, so that it can be used as in-coming qualification code at the end-user's facility. Simultaneously, *what if we changed the rules* of the game? The rule would state that by a certain date (ex. Oct. 2010) all electronics purchased by the U.S. Government or Contractor will have an accompanying **digital datasheet**.

Vision –"Do for the ATE world what the PC did for the Computer world"

ATE (Automated Test Equipment) is a multi-billion dollar market upon which the performance of all semiconductors and therefore manufacturers are dependent. It is on these machines that all semiconductors are tested due to liability issues prior to leaving the factory floor and a paper data-sheet is derived. However, because of the highly proprietary nature of each test platform, portability of the qualification code is currently futile.

An open architecture test system is based on widely-used and commonly-accepted interface specifications, the PC being a perfect example. By establishing an Open-Automated Test Equipment platform with the assistance of the Semiconductor manufacturers, the qualification/test code will become portable and back-ward compatible much like software can be ported from one PC to another. The code will become an electronic version of today's datasheet and allow for continual improvement. By establishing T.R.U.S.T, a reliable source of supply of microelectronics can be ensured and competition will quickly reduce costs and change the game. The rule change would quickly establish the standardization of portable qualification code, a **digital datasheet**.

This idea benefits everyone from DoD having qualified parts, to the contractor developing the system and finally to the ultimate customer, the serviceman/woman in the field who's placing their T.R.U.S.T. on the system to work.

Method –Worked with end-users of Semiconductor Automated Test Equipment that are currently feeling the pain of the status quo.

Dream team– AT&L- Familiar with the '05 DSB report on "High Performance Microchip Supply" **; NIST**

Support for the concept of T.R.U.S.T.-

LMI (Logistics Management Institute)- Bill Crowder www.lmi.org A non-profit Logistics Company that recognizes this changes the entire Supply Chain.
DMEA (Defense MicroElectronics Activity) – Fred Fraser www.dmea.osd.mil**;**
AFRL – (Air-Force Research Lab) – David Alexander**;**
ARMY -RTASSC(Radiation Tolerance Supply and Support Center)-Randy Brady**;**
SEMI- (Semiconductor Equipment and Materials International)-Karl Stuber www.semi.org**;**
IDEA- (Independent Distributors of Electronics Association)- Debra Eggeman www.idofea.org**;**
GSA (Global Semiconductor Alliance) - Lisa Tafoya www.gsaglobal.org;
Honeywell- Trusted Foundry- Joseph A. Mielke**;**
Texas Instruments- Hi-Rel Defense and Aerospace Semiconductor Group- Mont Taylor**;**
Freescale- Foundry services- Jeff Todd**;**
ATE GURU-Dr. Burnie West **;**
IC Test Houses- Seyed Paransun(Sr. Vice President of Major Test House) , Infiniti Solutions & Test Spectrum

National Cyber Leap Year
RFI-3

T.R.U.S.T.

Total Reliability Utilizing Standardization & Test

ZTI

Zero to Infinity

Technical POCs:

Derek Pearson

Dr. Burnell G. West, LFIEEE

April 15th, 2009

Who you are – ZTI- we are a small business that conceived the concept of T.R.U.S.T. (Total Reliability Utilizing Standardization and Test). T.R.U.S.T is a open-architecture standard for Automated Test Equipment which allows for portable qualification code. T.R.U.S.T. has broad support from both government and commercial entities currently feeling the pain of the proprietary status quo.

Game-changing dimension – Morph the game-board and Change the rules.

Concept – Semiconductors are the building blocks of all electronic systems. The rapid migration of semiconductor manufacturing plants to locations outside the United States has resulted in untrustworthy semiconductors and therefore systems becoming a concern.
http://www.businessweek.com/magazine/content/08_41/b4103034193886.htm

What if we *morph the game-board* and make the current manufacturers out-going qualification code portable, a **"digital datasheet"**, so that it can be used as in-coming qualification code at the end-user's facility, anywhere in the world. This digital datasheet can be used to easily detect counterfeit microelectronics from entering the supply chain. This can be performed to track new IC's as well as re-qualify recycled microelectronics (e-waste) and/or obsolete microelectronics which currently are the majority of counterfeit parts.

Simultaneously, *what if we changed the rules* of the game? The mandate would state that by a certain date (ex. Oct. 2011) all electronics purchased on behalf of the U.S. Government must have an accompanying digital datasheet.

Vision –"Do for the ATE world what the PC did for the Computer world"
ATE (Automated Test Equipment) is a multi-billion dollar market upon which the performance of all semiconductors and therefore manufacturers are dependent. It is on these machines that all semiconductors are tested prior to leaving the manufacturer's facility and a paper data-sheet is derived. However, because of the highly proprietary nature of each test platform, portability of the qualification code is currently futile.

An open architecture test system is based on widely-used and commonly-accepted interface specifications, the PC being a perfect example. By establishing T.R.U.S.T, an Open-Automated Test Equipment platform specification, the qualification/test code becomes portable and back-ward compatible. Again think of the PC, where yesterday's software can be run on today's machine and easily ported from one computer to another. The qualification code will become an electronic version of today's outdated paper datasheet and allow for continual improvements.

By establishing T.R.U.S.T, a reliable source of supply of microelectronics can be ensured and competition will quickly reduce costs and change the game. Portable qualification code on an Open Architecture ATE Platform would significantly reduce U.S. Government's vulnerability to counterfeit micro-circuits and at the same time dramatically reduce the expense, complexity and cycle time of all environmental testing (ionizing radiation, temperature, shake/rattle & roll, etc., etc.)
Conservative Savings: $35 billion/yr [1][2]

Method –ZTI has worked with end-users of Semiconductor Automated Test Equipment that are currently feeling the pain of the status quo.

Dream team–
FBI
***NIST*-**Product Authentication Information Management [PAIM] -David Brown (Intel)
Obeng, Yaw S
Simmon, Eric s
NIST*-**Familiar with AT&L '05 DSB Task Force on: ***HIGH PERFORMANCE MICROCHIP SUPPLY
John S. Suehle, Ph.D.

Support for the concept of T.R.U.S.T.-

ATE GURU's- Mark Roos
CDS Inc- ATE Hardware & Software Specialist- Bill Dunlap
Intel- Tracking-David Brown & ATE standards-Don Edenfeld
Counterfeit IC Facilities- Integra Technologies- Joe Holt;
DMEA (Defense MicroElectronics Activity) – Fred Fraser www.dmea.osd.mil**;**
ARMY -RTASSC(Radiation Tolerance Supply and Support Center)-Randy Brady**; SEMI**-(Semiconductor Equipment and Materials International)-Karl Stuber www.semi.org**;**
IDEA- (Independent Distributors of Electronics Association)- Debra Eggeman www.idofea.org**;**
GSA (Global Semiconductor Alliance) - Lisa Tafoya www.gsaglobal.org;
Honeywell- Trusted Foundry- Joseph A. Mielke**;**
Texas Instruments- Hi-Rel Defense and Aerospace Semiconductor Group- Mont Taylor**;**
Freescale- Foundry services- Jeff Todd**;**
IC Test Houses- Seyed Paransun(Sr. Vice President of Major Test House) , Infiniti Solutions & Test Spectrum

[1]- $25 billion/yr from [Test & Measurement] **Synthetic Instruments Tackle Military Testing**
John Stratton | August 2005
[2]- $10 billion/yr from Lockheed Martin DMSMS Oct. 29th, 2007