The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) Meeting Minutes
November 15, 2017, 1:30-3:30 MT
Colorado Convention Center, Rm. 711

Richard Carlson	DOE/SC	Joyce Lee	NCO
Phil Demar	FNAL	Miron Livny	UW-Madison
Devarshi Ghoshal	LBNL	David Martin	ANL
Burt Holzman	FNAL	Linden Mercer	NRL
Erik Huizer	GÉANT	Rajiv Ramnath	NSF
Daniel S. Katz	UIUC	Scott Sellaus	NSF
Ken Klingenstein	I2	Denis Serenyi	Google
Caas de Laat	UvA	Alan Sill	TTU
Eric Lancon	BNL	Derek Simmel	PSC
Brian Lin	FNAL	Todd Tennenbaum	UW-Madison
		Steve Tuecke	UChicago

Proceedings
This MAGIC meeting was chaired by Richard Carlson (DOE/SC) and Rajiv Ramnath (NSF).

SuperComputing 2017 Presentations
Cluster-Level Storage @ Google, Denis Serenyi (Google)
Identity Management, Von Welch (IU/CACR)
SciTokens, Todd Tannenbaum (UW-Madison)
Globus, Steve Tuecke (Argonne)
eduGAIN Roadmap, Erik Huizer (GÉANT)
InCommon and Global Federation Issues, Ken Klingenstein (Internet2)

Cluster-Level Storage @ Google: Denis Serenyi (condensed from key note presentation at INDIS)
Goals: To have a very large number of diverse workloads in a shared cluster environment and to deploy and manage these systems at the lowest possible cost
Typical Google cluster

- Tens of thousands of machines, Petabytes stored on hard drives
- Fast networking: machines with local SSD have a couple TB of storage
- Hundreds to thousands of applications doing serving
- Workloads performing serving:
- Key characteristic: low latency requirements
    - Have peaks and troughs. Users who deploy these services want to provision at peaks in order to have low latency, even at the peaks
    - Latency is not critical for batch style workloads or analysis workloads

Total Cost of Operations (TCO)

- Storage TCO (not disk TCO): costs of data durability, data availability and serving costs
- Goal: minimize storage TCO for all devices deploying in these systems and still meet users' requirements; Wish to share systems durably and have available, serviceable data
    - Buy components with low TCO (e.g. hard drive: major hardware component)
    - Not minimize TCO if spindles on the disks are not at full capacity

Flash

- Used heavily in data centers to store hottest data
    - Remaining data can then be served by disk; buy disks needed to host data (and hope projections about spindle and bite requirements are accurate)
- Workflow characteristics of hosted data
    - If just wish to store cold data, then forced to waste spindles
    - If hot data, then TCO challenges; cannot take advantage of TCO-lowering because drive manufacturers are improving bite densities.
- Disk I/O- becoming more expensive, so applications that are not designed with flash (i.e., cannot localize their hot data) will become pricier over time because dependent on disk I/O.
    - Trend: I/O cost trends are forcing applications and storage systems to evolve.

Discussion

Instead of segmenting clusters and having uniform hardware, a wide variety of old and new commodity hardware are in a single cluster. Software is needed to ensure physical devices are full and busy. Balancing data to avoid bottlenecks also facilitates the usefulness of older hardware. In this context, Denis referred to his INDIS presentation, which described Google's transition from the Google Files System to Colossus, the next-generation cluster-level file system which rebalances cold data and ensures the even distribution of data across spindles.

Identity Management: Von Welch (IU/CACR) and Jim Basney (UIUC/NCSA, contributor)

Attribute release still a problem, mainly due to unclear decision processes at the universities

Research and Scholarship category is still promising

- R&S growth plateau–growing only as fast as the growth of participation in InCommon (15%). Thus, need to make R&S memberships the default for IDPs (Opt-out model)
- NSF Campus Cyberinfrastructure (CC*), PROGRAM SOLICITATION NSF 18-508. Shows more weight behind R&S and should get attention of universities.
- I2 Global Summit (May 2017): Ian Glazier, Senior Director for Identity at Salesforce, keynote speech shows growing significance of IdM in higher education

Globus project

- January 2018 - ending support for open source libraries supporting X.509 proxy certificates). Community is picking up support, but this is another "nail in the coffin" of X.509 (client-side).
- Clarification: client-side distinguished from server-side and not judging X.509 on technical grounds. In 10 years, the growth and adoption of web authentication will push it out. It has successfully built trust in the scientific community.

ORCiD – recently emerged as an identify provider; no longer just an identifier

- OpenId based – appears to be a good fit for solving scientific research and ID access management issues
- Prediction – ORCiD is probably a big deal because: 1) avoids some traps InCommon has fallen into (releases identifier sufficient for most research service providers); 2) institution agnostic; 3) started as a persistent, unique identifiers of researchers - a better fit for university projects

Not covered, but important (later presentations will address some issues)

- Sirtfi– incident response for Fed IdM
- SciTokens – capability model
- Consent – another path to attribute release
- CTSC IAM work: http://trustedci.org/iam/

Updates on 2015, 2016 Opinions

- IdM/IAM will continue to mature
- Ecosystem will be a collection of solutions, instead of a single solution
- Soft-token 2FA is on the path to being a default
- Changes:
    - R&S should be the common default for SAML federations, so they can useful for the research space
    - Web protocols (e.g., OpenIds, OAuth) gaining speed

Real competitions in multi-player ecosystem

- Most scientific collaborations are using the Google applications suite and identities therein; collaborations are part of the research platform space

Globus: Steve Tuecke (UChicago/Argonne)

Why use Globus Auth?

- Served as underpinning for identity delegation and basic authentication, but not for restricted delegation due to usability issues.

Security PaaS Challenges to building science infrastructure:

1) Providing basic login to applications - Identity/authentication problem
2) Protecting REST API communications- cyberinfrastructure challenge :
    a. Easy for Google because of its REST APIs, OAuth2 service, authorization, service with token management are in 1 universe. If 1 token stolen, thieves will not have the "keys to the castle." X.509 proxies in that type of universe would experience delegation problems if the middle link is compromised.

Globus Auth:

- Trying to take next big leap embracing all new security standards (open identity connect etc.)
- In production – new 3rd party services (Jet Stream, NCAR, etc) are using it. This enables interesting scenarios (Python SDK).

Modern research data portal – use Globus Auth, transfer APIs with Science DMZs, open source code, etc. Ultimately, having a social login would be enable logging in via Globus anywhere

Research & Scholarship- Will be moving away from R&S

Use cases driving Globus Auth's current direction:

- Increasing platform use in research automation (e.g., ALS/Berkeley)
- Tying together data transfer, metadata, automated analysis, custom REST APIs and orchestrating/automating process. Supportable due to added security management capabilities.

Native applications (e.g., Command line interface) are well supported.

Securing REST APIs

Globus Auth is unique in its ability to secure Rest APIs and manage tokens (e.g., SSH with Globus Auth)
Moving beyond X.509 to Globus Auth/OAuth2 (and talking to all endpoints)

- SSH server is Globus Auth/OAuth2 resource

- OAuth2 support added to GridFTP. Can be used for any protocol

eduGAIN Roadmap: Erik Huizer (GÉANT)

eduGAIN is a meta system for identity federations:

- Provides meta data distribution service that gathers needed information from ID federations
- provides policy framework and standards to build trust among its members
- Republishes metadata for world-wide use

Fast growth

- 55 production federations
- 48 members
- 6 voting-only members (not yet share attributes)
- 8 known federations (intending to become members)

InCommon

- More InCommon entities are joining eduGAIN (19% have joined)
- Comprise 20% of eduGAIN worldwide

Users

- NIH- collaborating in Uganda & Mali to support federated access (MORE)
- Hong Kong Access Federation recently joined
- CSTNet and CERNET backed federations have recently applied
- RENU Identity Federation (Uganda)

New opportunities

- eduTEAMS – build and deliver identity as a service for entire research coalition so have same level access to eduGAIN and can work together
- Academia- provide privacy sensitive service to commercial service providers. In academia, could submit query to confirm if student with no sensitive information given

Roadmap

- Continual improvements
- Work on resolving GDPR issues (strict EU privacy regulations)
- Opportunities to do more in Q2, 2018 on.

SciTokens: Todd Tannenbaum (UW-Madison)

SciTokens is a federated authorization ecosystem for distributed scientific computing (NSF-funded project, started July 2017). SciTokens does not manage identity or provide authorization service but aims to scale existing solutions out of distributed grid infrastructure.

Goals:

- Introduce a capabilities-based authorization infrastructure for distributed scientific computing
- Provide reference platform with this infrastructure (combining a token library with CILogon, HTCondor, CVMFS, and Xrootd)
- deploy to help science stakeholders (LIGO and LSST) better achieve their scientific aims

In a common grid computing scenario, how does the storage service validate data access request?

- Common grid solution: identity and impersonation via X.509 certificates.
- Not ideal for a few reasons: Not least privilege (what if identity is stolen?). Compare with having to present a token instead of identification.

Capabilities-based authorization infrastructure

- Trend: Focus on capabilities instead of Identity
- Authorization service creates token that describes the capabilities or authorizations of the token bearer who presents it to the resource. Implemented primarily through OAuth2.
- SciTokens Project Reference Platform:
    - Working to integrate an OAuth2 client into HTCondor submit host and make OAuth2 enhancements into CILogin.
    - HTCondor being enhanced to manage token lifetime (refresh as needed) and deliver to job.
    - Data services enhanced to allow use of rights/authorizations by inspecting tokens instead of grid proxy certificates.
    - Endgoal- first time user of HTCondor would navigate to web interface to set up desired permissions
        - User will authorize tokens to be given to HTCondor submission service. On subsequent job submission, HTCondor would create the access token for user. User does not have to do anything.
        - Early milestone: send out job [e.g., on open science grid] and have job output data to storage service (Box.com).

Leverage relevant RFC standards- adopt existing standards

- Workflows for acquiring/using tokens: OAuth 2.0. (how the various parties should interact
- Access Tokens: JWT bearer tokens.
    - The contents (claims) of JWT tokens are not standardized, so we will provide a SciTokens Claims specification and reference library implementation so tokens issued by an organization are understood by a wide variety of resources.
    - SciTokens Library also supports distributed verification and privacy preservation.

Working with:

- Stakeholders: LIGO, LSST,
- Technologies: HTCondor, CILogin, CVMFS, XRootD, FTS
- Discussions and Interest from Open Science Grid (OSG), LHC WLCG (Worldwide LHC Computing Grid), CMS, CERN ITTecLHC, CMS

Discussion:

- Capability model challenge: If the first time user botches the job, omits some capabilities in the token, so job fails. Then next time, user puts in all possible capabilities. Need to explore.
- Job Submission and access tokens:
    - Currently, specify the tokens needed upon job submission. HTCondor will only send the tokens that the job claims it needs when HTCondor sends the job to the grid.
    - Identifying which access tokens are needed. Job submission can help as well as scheduler who is involved in management of tokens. If job is to run 5 days and send output to Box.com, can delay sending the token to box. So can send different tokens at different lifetimes of the job.
    - The job submission point knows what needs to be moved, where and what to do if move fails; it can control token, lifetime and the file. Condor performs many file movements that are described in submission (move to inbox, outbox).
- Granularity of claim:
    - Much will depend on storage service. Each organization can manage its own name space in the claims of the token.

- Exploring limiting granularity in terms of time: only send token that has write access when scheduler knows the job is ready to write.

- Finer grain signaling mechanisms: SciTokens fits an interesting niche regarding portals and data minimization. A portal will ask for all the attributes that any application behind the portal may need. Important capability: Ability to signal at finer grain that what going to access behind the portal needs these particular attributes.

InCommon and Global Federation Issues: Ken Klingenstein (Internet2)

InCommon

- Growth and use is continuing
  - Largely trailing edges coming in (as opposed to large providers (Amazon, Microsoft) already in – trailing edge have different capabilities than leading edge.
- Importance of trailing edge schools
  - Important to fulfillment of major project like LIGO, so must address this challenging space (trailing edge)
- Strengthened InCommon infrastructure: much robustness is behind the scenes since moved all signing mechanisms into the cloud. Integrity of metadata being managed much more
- Moving into self-service federation management
  - E.g.: When security contact changes, screen with multi-factor authentication so institution can change parameters within federation (important for scaling)
  - Agency use:
    - CC* solicitation – asks institutions to discuss participation in R&S and in baseline expectations - first attempt to raise bar of security across the board (how manage metadata, manage site, patch site- nitty gritty of security being addressed).
    - Some of NIH has embraced COmanage and federation
- Created mechanisms that are crossing to eduGAIN space – (cultural issues
  - MFA –how to do on soft phones?
  - cultural issues arising about internationalization
  - Baseline expectations – raise bar, but may limit organizations into InCommon

  Other efforts addressing through IDP of Last Resort, Steward program, but need business models
- Attribute release and R&S continue to be vexing problems – hope consent will be game changer
- Dynamic metadata slowly rolling out. Moving to these Dynamic metadata services.

International dimensions

- Potential federation shopping by SP, IdP
  - Triggers need federation practice statements, need international normalization of behavior as move to inter-federation (patching software, federated operators, sign-in keys, etc.)
- SirTfi for security and incident handling for federated identity –useful tool
  - Serious requirements (e.g., notify community within 72 hours of an incident)
- Snctfi Activity for collaboration-scale authorization
  - Covers policy space around delegation, auditing, responsibilities, separation of duties, etc. Trying to patch everything you would want in healthy collaboration
  - Not cover technology like SciTokens
- Bridging cultures and regulatory regimes

- International cultural issues (e.g., InCommon has faculty, unlike UK)
- Compensating controls
- GDPR

New identifiers

- Coming out to address concerns Steve raised with the reassignment of identifiers
- Good, but will take a long time to make it into the business of an institution

Finland

- Adding OIDC capabilities, so push out JWT instead of SAML, but the rest of infrastructure will look the same (management, etc.)

Snctifi policies

- normalization of policies behind "the curtain" of a large collaboration that wishes to present as unified infrastructure
  - Requires coordination of policies regarding authorization, etc. GEANT is preparing draft, moving into more discussion. May be normalized by end of 2017 or early 2018.
- Research organizations well represented?
  - Authentication and Authorization in the Research Community speaks for research infrastructures in Europe (https://wiki.geant.org/display/AARC/Snctfi)

Federated OIDC- complex, multilateral trust issue

- OIDC (and OAuth) were developed as bilateral protocols for mobile and light trust apps
- Now showing value in multilateral situations- trusting multiple parties in multiple interactions
  - Sweden is working on adding metadata statements to OIDC. Why others should trust our attributes, etc., must be transported from SAML to a JAVA environment (intrinsic to multilateral environment, independent of protocol).
- Client registration becomes the key step: Pilot (February/March 2018)

General Data Protection Regulation (GDPR)

- Created by EU to manage data protection uniformly across the EU that binds every member EU nation (operational May 25, 2018).
- Global impacts (e.g., applies to interactions with EU resident, faculty on sabbatical in EU,
- Covers many issues: tracking, attribute release, right to be forgotten to data breaches, etc.
- Draconian: enforceable penalties of up to 4% of global revenue
- See graphic to describe GDPR component (all welcome to use) - Issues with carrying over preferences, portability, etc.

Added Details - GDPR

- PII - Almost everything is PII (from IP address to persistent identifiers)
  - Some identifiers are not (e.g. ePTID)
- Created Sensitive PII - Requires more special handling (e.g., LGBT (sensitive) vs. U.S. Citizen)
  - How to display and interact regarding Sensitive PII? In consent work, for example, we have adopted a passwords approach.
- Consent – call may be recorded. May not hear or may be given a choice.
- Data breach notifications
- Organizations may need a data protection officer

Next meeting

December 6, 2017 at the National Coordination Office, 490 L'Enfant Plaza, Suite 8001.