THE QUILT SECURITY COOKBOOK

"The Quilt community of research and education networks strive to be good network citizens by implementing security best practices, exchanging ideas, and leveraging knowledge and resources."

Preface

In February 2016, The Quilt hosted a Security Round Table to discuss network security for Quilt members and for the institutions those members connect. The priority project resulting from the discussion was to create a cookbook of network security best practices for both the regional / enterprise network as well as additional member institutions that have an interest and may benefit from its content.

The practices and suggestions captured herein are intended only for the purposes of exchanging ideas and best practices for network security among the research and education networking community. It is not intended to be used as the authoritative source on network security nor does it specify standards or policy. Readers of the Cookbook should understand that the intention of the cookbook is to inform and educate on network security practices of contributing research and education networks but not to prescribe any one particular technology, strategy or policy to mitigate security risks. Please note that this Cookbook may contain errors. Users of this Cookbook are solely responsible for ensuring their network security practices are adequate for their needs. Any references to third party products are not an endorsement. The Quilt and its contributors are not responsible for such network security practices or third party products, even if based on this Cookbook. This Cookbook will continue to be developed; i.e. a 'living' document.

TABLE OF CONTENTS

REGIONAL NETWORK CHECKLIST
MANAGEMENT PLANE
Management Plane Network Protections
Management Plane Client Protections
DENIAL OF SERVICE
DDOS MITIGATION STRATEGIES
UTRS
RTBH (remotely triggered black hole) aka BGP Null Route
DDOS VENDORS
DDOS VENDOR TERMINOLOGY
DATA PLANE / CONTROL PLANE SECURITY
BCP38
BCP84
uRPF Traffic Black Hole Techniques
uRPF with a Default Route
uRPF Strict Mode Configuration Examples
Role-based IP Addressing
Segmentation of ResNet
Segmentation of State-Aware Devices
LANGUAGE TO SET EXPECTATION
Sample Language – from MOREnet
PRACTICES IN SAMPLING RATE, TOOLS AND ARCHIVING FOR FLOWS
APPROPRIATE WHITELISTS AND THEIR RATIONALE (CDNS, ETC.)
ENTERPRISE BEST PRACTICES
NETWORK SWITCH AND FIREWALL PLACEMENT
ROLE-BASED SEPARATION OF IP SPACE
ARCHITECTURE
Segmentation of ResNet
PLACEHOLDER FOR SCIENCE DMZ INFORMATION
PLACEHOLDER FOR FACILITATING & MANAGING EDUROAM DEPLOYMENTS
RESOURCES
READILY AVAILABLE TOOLS
KANREN, MORENET, AND UEN SECURITY RT PRESENTATIONS
NETWORK DEVICE HARDENING GUIDES
Cisco IOS Hardening Guide
SELF ASSESSMENT TOOLS
GLOSSARY / TAXONOMY

Regional Network Checklist

> "It is in our missions as R&E networking organizations to educate our members. We are committed as a community of peers to accomplish our mission."

Techniques for protecting the network itself and protecting downstream members/clients, i.e., how you control access to your routers is critically important to the regional network, not end-user.

Management Plane

It is critical to security that R&E networks have access control between enterprise networks and management plane - keeping management plane isolated.

Management Plane Network Protections

I. Scope:
   Applicable to management plane of network; i.e., the systems, devices, and networks that are used for management of the configuration of the WAN. This includes client machines of network engineers, servers/services used to store and propagate WAN configurations, etc.

II. Purpose: The business purpose of these recommendations is to:
   a. Protect the network management plane from attacks/unauthorized access

III. Recommendations:
   a. Management network isolation.
      i. The management network should be isolated/segmented from the data plane, control plane, and any enterprise LAN services (e.g. print services, enterprise desktop, Internet, etc.)
      ii. Access to the management network should be via an encrypted VPN (TLS 1.2+, AES 256+)
      iii. Further segmentation within the management systems as required. For example, development and test systems should be isolated from production systems.
      iv. Access to Bastion/Jump Hosts is limited to IP addresses owned by the REN. If engineer is not in the REN's network space, they need to first connect via VPN.
   b. Two-factor authentication should be required to access the resources – e.g. passphrase + SMS code, passphrase + OTP token (hardware or software).
   c. Example workflow:
      i. Engineer initiates connection to SSL VPN concentrator
      ii. Engineer enters user name, password, realm
      iii. VPN challenges user for 2 factor token
      iv. Engineer provides response to challenge

v. VPN connection to management network is complete
vi. Engineer requires access to a production (non-dev/test) system
vii. Engineer initiates SSH connection to jump host/bastion host
viii. Engineer supplies user name, passphrase, SSH key; verifies server SSH key.
ix. Connection to production system is complete.

Management Plane Client Protections

I. Scope:
Applicable to staff that interact with network management, e.g. network operations, engineering, security, support.

II. Purpose: The business purpose of these recommendations is to:
a. Protect the network management plane from attacks/unauthorized access that use the client as a vector
b. Minimize dependencies on non-REN systems for access to the network management plane
c. Maintain the ability for network engineers to access the network management plane quickly and reliably, and when LAN and/or "Internet" are not available or when remote.

III. Recommendations:
a. Restricted access into client systems:
i. Remote access, remote control, and remote execution into client machines is not permitted
ii. Client machines should have all unnecessary ports disabled and not run unnecessary services
iii. Routing through the VPN is not permitted. Built-in video cameras disabled or verified "off" when not actively in use
iv. Screen lock/password timeout should be enabled
v. Auto-login into machines disabled
vi. Guest accounts disabled
b. Reduced system dependencies on outside systems
i. Access to management networks requires the 2-factor access which has reduced network dependencies
ii. Auto-download/install of updates disabled
c. Client Configuration
i. Standard client configuration should meet CIS level 1 or alternative benchmark
ii. Only implement the minimal configuration necessary
iii. Formal change management for configs, including change validation

iv. Patch management (patches applied within 30 days, exceptions noted, reported to central console)
v. Supportability (current, vendor-supported software & firmware in use)
vi. Logging (system and security events should be logged)
vii. Time synch (time should be synchronized to an accurate time source)
viii. Host firewall or HIPS should be enabled
ix. Vulnerability scanning every 30 days and remediation within 30 days
x. Local disks should be encrypted (with a mechanism for central key management)
xi. Local admin accounts should have strong (see NIST 800-63B) passwords, and not used for routine work, and two-factor authentication should be considered
xii. Anti-malware should be enabled and reported to central console
xiii. Annual asset inventory
xiv. Annual software inventory

d. To reduce the risk profile of the client access the following measures should be implemented
    i. Automated scan of local system upon connection to VPN/management network to verify recommendations are met (via Host Checker and/or execution of a script on connect)
    ii. Standardization of secure password management software to reduce dependency on "out of band" or other insecure password maintenance
    iii. Audit trail of outbound connections
    iv. Local scanning of client networks on a daily to weekly basis using vulnerability scanning tool
    v. Updates should be installed and verified if auto-update is turned off
    vi. Central client management tools should treat these clients as "remote" connections rather than "always-local" connections.
    vii. Each individual user covered under these recommendations should execute a document that they understand the risks and agree to follow the policies.

Denial of Service

Distributed denial of service (DDoS) attacks, are malicious attempts to render an online site or application unavailable to users by overwhelming the site with an enormous amount of traffic, causing the site to crash or operate very slowly. DDoS attacks have been increasing in number and size.

Mitigation strategies involve implementing best practices, implementing a commercial solution, or both.

DDoS Mitigation Strategies

UTRS

UTRS is a system that helps mitigate large infrastructure attacks by leveraging an existing network of cooperating BGP speakers such as ISPs, hosting providers and educational institutions that automatically distributes verified BGP-based filter rules from victim to cooperating networks.

UTRS practices: http://www.team-cymru.org/UTRS/

RTBH (remotely triggered black hole) aka BGP Null Route

1. See also OARnet-DDos-RTBH.pdf
    a. BGP Null Route
    b. Client can send a specific host route to ISP tagged NNN:MM (or OOOO:PP for Internet2) via BGP
    c. ISP Backbone router sets up next hop to discard
    d. Once BGP policy configured at ISP, client can send prefixes without involving ISP support/engineering
    e. Sample configuration shown in Table 2

Table 2

```
mark@router> show configuration policy-options policy-statement CLIENT-ASQQQQ-IN
term CANDIDATE-NULL {
        from {
        community ASNNN-NULL;
        route-filter 2.2.176.0/20 prefix-length-range /32-/32;
        route-filter 200.200.232.0/23 prefix-length-range /32-/32;
        }
        then {
                next-hop discard;
        accept;
        }
}
```

1. Enhancements
    a. ISP could potentially distribute the route internally so the attack is blocked closer to the source(s)
    b. ISP could also re-tag and advertised to upstream providers

2. Considerations
    a. The Victim IP is no longer able to communicate with the Internet.
    b. If that was the intent of the attack, then the attack will succeed
    c. However, it may be necessary in order to keep the rest of your network up
3. Related ideas –Feel free to add content.
    a. 3rd party BGP black hole
    b. Host routes
    c. Scrubber infrastructure
    d. Flowspec

DDoS Vendors

Several commercial vendors provide solutions to help mitigate DDoS attacks. Table 1 provides a list of vendors, the types of service they provide, and a link to their websites.

Table 1

Vendor	Offering types	Link
Arbor Networks	Appliance, Cloud scrubbing, Hybrid, Service provider	https://www.arbornetworks.com/ddos-protection-products
Neustar	Cloud scrubbing, Hybrid (using Arbor appliances)	https://www.neustar.biz/services/ddos-protection
Radware	Appliance, Cloud scrubbing, Hybrid, Service Provider	http://www.radware.com/
Verisign	Cloud scrubbing, Hybrid (using 3rd party devices)	https://www.verisign.com/en_US/security-services/ddos-protection/index.xhtml
A10 Networks	Appliance, Cloud scrubbing, Hybrid	https://www.a10networks.com/products/thunder-series/ddos-detection-protection-mitigation
F5 Networks		https://f5.com/
Akamai (Prolexic)		https://www.akamai.com/us/en/solutions/products/cloud-security/prolexic-routed.jsp
Cloudflare		https://www.cloudflare.com/
Google Project Shield		https://projectshield.withgoogle.com/public/

DDoS Vendor Terminology

”Appliance”: CPE-based device which detects and/or mitigates DDoS activity at network edge. Form factor may be physical or virtual.”

”Cloud scrubbing”: Network infrastructure (hosted by provider) to which DDoS traffic can be routed, malicious traffic removed, and then delivered back to client “clean” (typically through a nailed-up GRE tunnel).

”Hybrid”: solution consisting of both appliance and cloud scrubbing. Appliance typically includes capability to automatically determine when to “swing” the dirty traffic to the Cloud scrubbing center based on on heuristics.

”Service provider”: Solution that is intended to be deployed inside an ISP’s own network

Data Plane / Control Plane Security

BCP38

Generally aimed at enterprise and campus networks, BCP38 (RFC2827) describes methods and strategies for filtering traffic entering a routed interface based on the IP packet’s source IP address. The goal being to only allow traffic sourced from the assigned IP subnet and discarding traffic with improper source addresses. Both Unicast Reverse Path Forwarding (uRPF) and ingress ACLs (or firewall filter) may be used to achieve this goal. Within an enterprise network uRPF is generally seen as simpler to maintain and operates at linerate on most platforms. If uRPF is performed on all edge-facing routed interfaces (physical or SVI) a simple, final check may be applied to all traffic leaving the campus network as an egress ACL.

See also BCP38 Tool funded by DHS/MIT: https://www.caida.org/projects/spoofer/

BCP84

BCP84 (RFC3704) aims to refine BCP38 and further define possible solutions for multi-homed networks such as service providers where routing is often asymmetric due to remote networks’ routing policy and path selection. Complex enterprise and campus networks will likely find many useful suggestions in BCP84.

uRPF Traffic Black Hole Techniques

uRPF comes in two distinct flavors; Strict Mode and Loose Mode, the two are described below.

Strict Mode Checks:

1. Does the local router have a valid route (directly connected or static route to a destination on the receive interface) to the source IP of the packet?
2. If yes, does the route use the same interface on which the packet was received?

If the answer to both questions is “Yes” then the packet is accepted and forwarded to the next hop. If either check fails, the packet is dropped; logging or flow sampling may also be configured for dropped packets.

In the diagram below, the router is configured with uRPF strict mode on all of its user-facing ports. Two packets are received with with the same destination (8.8.8.8), one from ResNet and one from the Servers network. The ResNet packet passes both checks (active route on the same interface) and is permitted. The Servers packet passes one check (active route) but fails the second check (on the same interface) and is thus discarded.

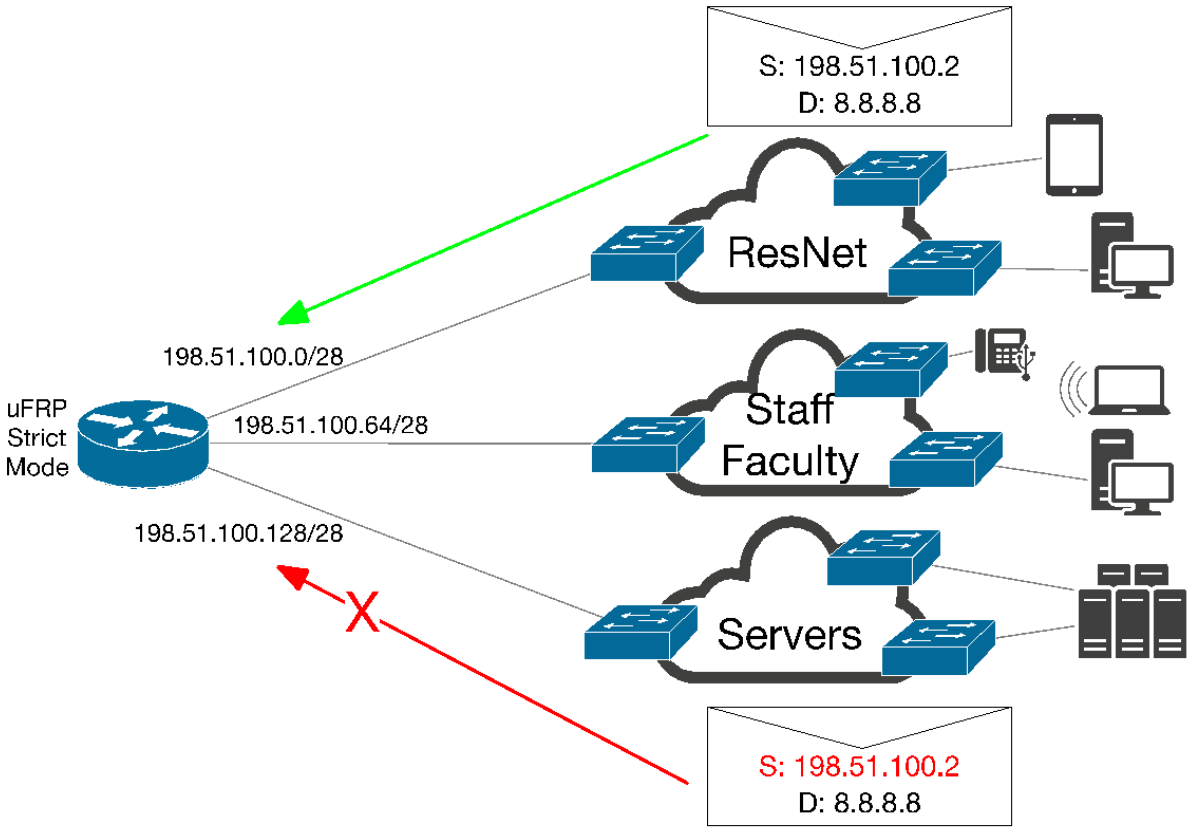

uRPF Loose Mode relies on a single check:

1. Does the local router have a valid route to the source IP of the packet?

Loose Mode allows for greater flexibility in multi-homed networks or networks with multiple paths to the same destination via Equal Cost Multipath (ECMP) or if routing is asymmetric. Note that some platforms require extra knobs to check “feasible paths”, check your platform’s documentation before proceeding.

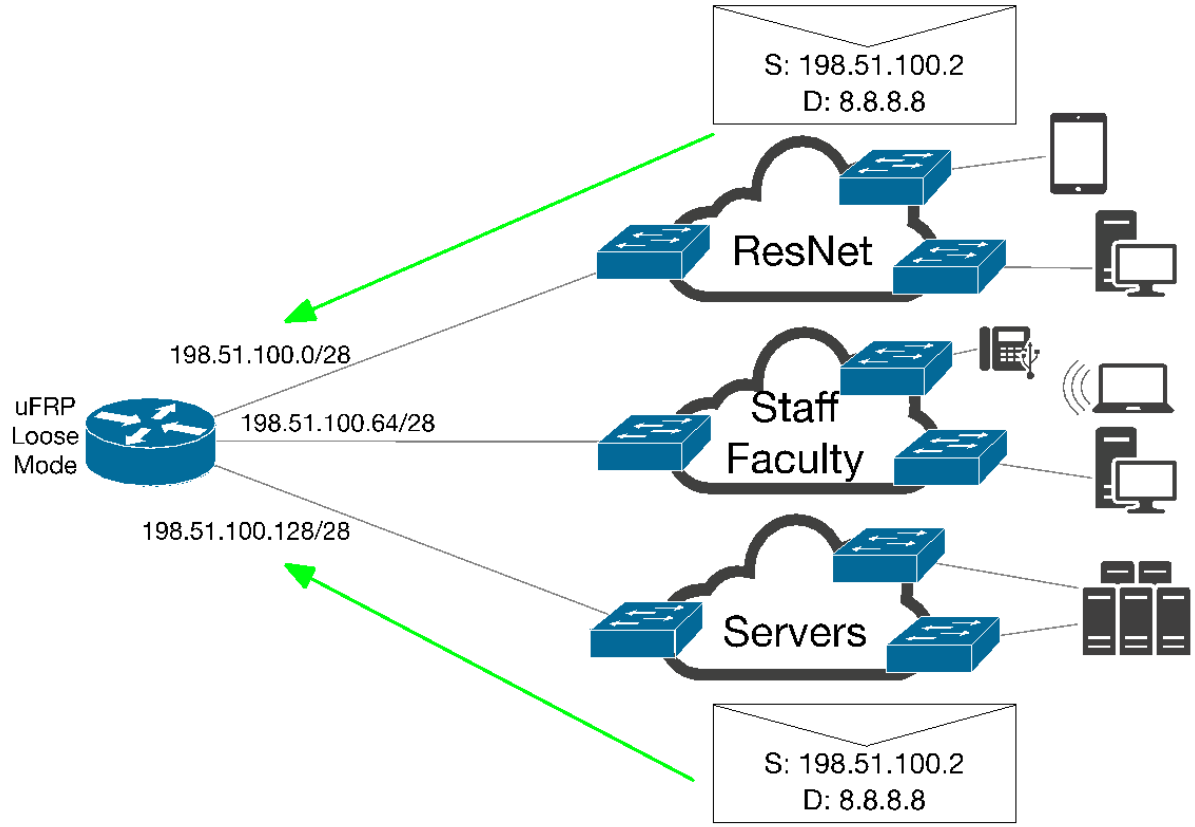

In the example above the router is configured with uRPF in Loose Mode. Two packets are received, one from ResNet and another from the Servers network. Because uRPF Loose Mode only performs a simple check both packets are accepted and forwarded to their next hop. In most user-facing, edge configurations this is not desirable.

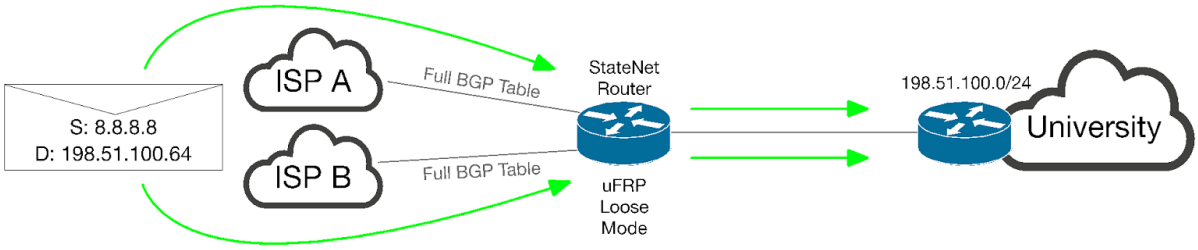

In this example a statenet (RON) router is configured for uRPF Loose Mode and is multi-homed with "ISP A" and "ISP B", each sending the statenet router a full BGP table. A packet is received with source IP 8.8.8.8 from each ISP. Assuming Google is announcing the supernet containing 8.8.8.8 correctly the uRPF Loose Mode check will pass and both packets will be forwarded to their destination.

uRPF with a Default Route

Note that uRPF's two modes were designed with the following assumptions:

1. The network operator will use a default route to a host on the same interface configured for uRPF Strict Mode.
2. The network will not configure uRPF Loose Mode if a default route (static or dynamic) is configured or learned on the interface.

The presence of a default route will cause uRPF Loose Mode checks to always pass on the interface which it is configured. Most routing platforms have options to either include or exclude default routes in uRPF checks. Each platform is configured differently so be sure to check your vendor's documentation.

uRPF Strict Mode Configuration Examples

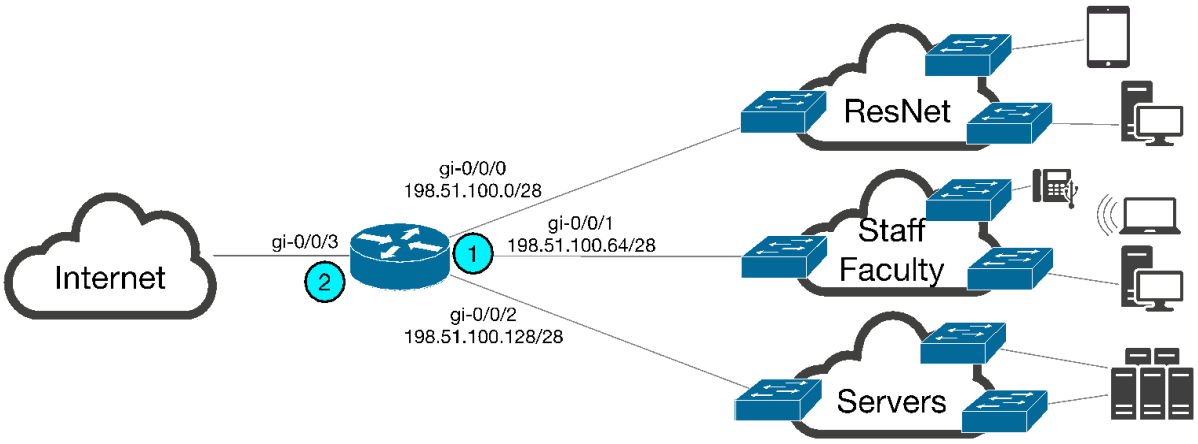

Configuration of uRPF depends on the routing platform in use; consult documentation regarding possible extra memory use or other implications before proceeding.

Example Cisco Configuration for user-facing ports ("1"):

```
interface gigabitEthernet 0/0/0
 ip address 198.51.100.1 255.255.255.240
 ip verify unicast source reachable-via rx
interface gigabitEthernet 0/0/1
 ip address 198.51.100.65 255.255.255.240
 ip verify unicast source reachable-via rx
interface gigabitEthernet 0/0/2
 ip address 198.51.100.129 255.255.255.240
 ip verify unicast source reachable-via rx
```

Example Cisco Configuration for Internet-facing port ("2"):

```
interface gigabitEthernet 0/0/3
 ip address 1.1.1.1 255.255.255.252
 ip access-group bcp38-out out
!
```

```
ip access-list extended bcp38-out
 10 permit ip 192.51.100.0 0.0.0.255 any
 20 deny ip any any
```

Example Juniper Configuration for user-facing ports ("1"):

```
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter urpf-fallback;
        address 198.51.100.1/28;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family inet {
        rpf-check fail-filter urpf-fallback;
        address 198.51.100.65/28;
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family inet {
        rpf-check fail-filter urpf-fallback;
        address 198.51.100.129/28;
      }
    }
  }
}
firewall {
  family inet {
    filter urpf-fallback {
      interface-specific;
      term allow-dhcp {
        from {
          source-address {
            0.0.0.0/32;
          }
        }
        then {
          count dhcp-traffic;
          accept;
        }
```

```
            }
            term discard-other {
              then {
                log;
                discard;
              }
            }
          }
      }
    }
```

Example Juniper Configuration for Internet-facing port ("2"):

```
    interfaces {
      ge-0/0/3 {
        unit 0 {
          family inet {
            filter {
              output internet-out;
            }
            address 1.1.1.1/30;
          }
        }
      }
    }
    firewall {
      family inet {
        filter internet-out {
          term allow-legit-sources {
            from {
              source-address {
                198.51.100.0/24;
              }
            }
            then accept;
          }
          term discard-and-log-others {
            then {
              count spoofed-source;
              log;
              discard;
            }
          }
        }
      }
    }
```

Example Brocade Configuration for user-facing ports ("1"):

```
interface ethernet 0/0/1
 ip address 198.51.100.1/28
 rpf-mode strict log
interface ethernet 0/0/2
 ip address 198.51.100.65/28
 rpf-mode strict log
interface ethernet 0/0/3
 ip address 198.51.100.129/28
 rpf-mode strict log
```

Example Brocade Configuration for Internet-facing port ("2"):

```
interface ethernet 0/0/4
 ip address 1.1.1.1/30
 ip access-group bcp38-out out
!
ip access-list extended
 permit ip 198.51.100.0/24
 deny ip any any
```

Note: Some platforms process ACL log statements in software which could increase strain on management and control plane resources. If possible consider sampling via NetFlow, IPFIX, sFlow, or other hardware-accelerated technologies in place of local logging. Be sure to consult documentation for the specific hardware platform and software release before deploying uRPF.

Role-based IP Addressing

During a Distributed Denial of Service (DDoS) attack aimed at your network or connected user networks it is often desirable to install a Real-Time Blackhole Route (RTBHR); however, many networks NAT/PAT overload vast user counts to a small handful of public IPs. When one of these IPs is the target of a DDoS attack any RTBHR will take all the users behind the IP offline. Most university and education networks have access to large IP allocations and may find it advisable to segment users for NAT/PAT purposes to help mitigate impact of DDoS events.

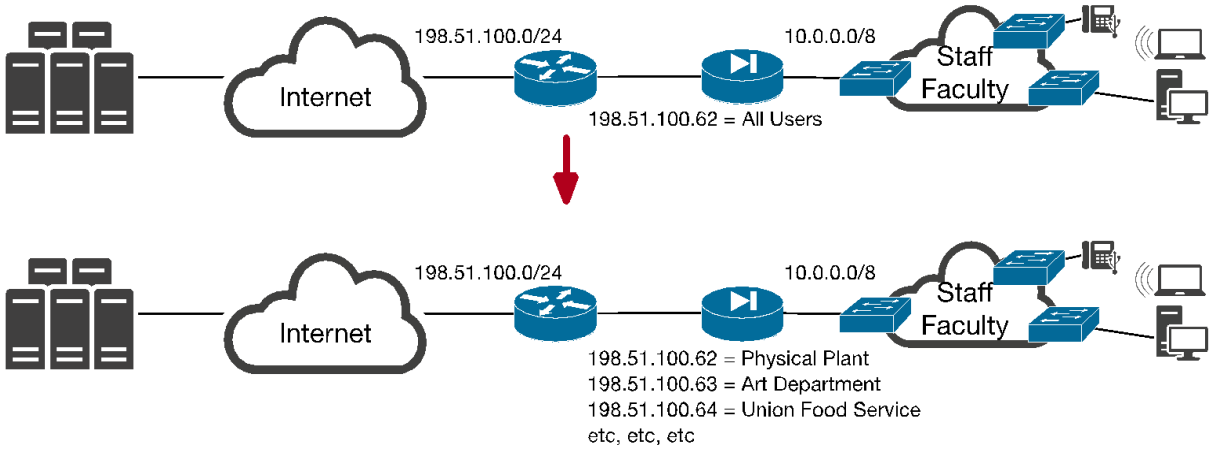

In this example the university is assigned the public IP block 198.51.100.0/24. The university employs 10.0.0.0/8 from RFC1918 space for staff and faculty hosts. Originally all users are NAT/PAT behind a single IP (198.51.100.52). Role-based IP addressing involves segmenting users' public NAT/PAT translation for greater resiliency. In this case each department is assigned its own public NAT/PAT address.

While the example shows single public IP addresses, use of a pool (or range) will provide greater resiliency. If a pool is used the attack target IP can be blackholed, removed from the NAT/PAT range, and most users will not even notice the change.

Segmentation of ResNet

ResNet is a group which should be targeted for segmentation since so called "booter" services are designed for gamers to DDoS one another; causing a timeout which "boots" the target from the game. Due to ease of use and low cost such services have become popular in recent years. Since ResNet users are so often the target of DDoS attacks, segmenting them early may provide the greatest immediate benefit to the campus network.

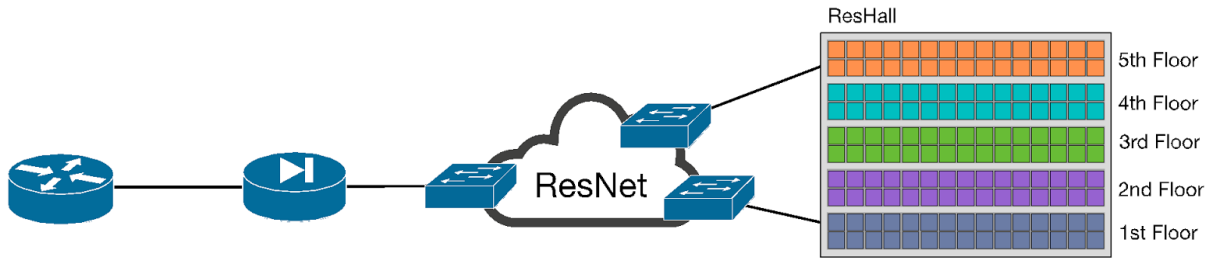

In this example a single residence hall with five floors is segmented into five different public IPs. If a user living on 1st Floor is attacked by a DDoS event the only impacted users are on the same floor in the same ResHall. If multiple IPs are assigned to each group in a pool the attack target can be removed from the pool and 1st Floor users impacted will be back online. Once the attack subsides the RTBHR can be removed and the NAT/PAT pool returned to its original configuration.

The natural progression of this practice is called microsegmentation.

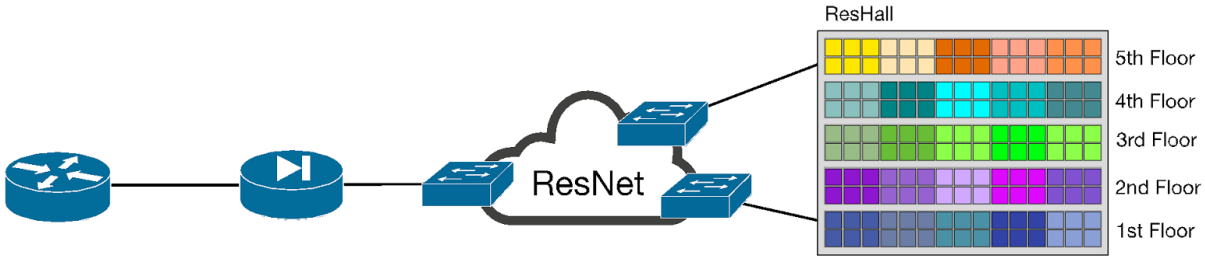

In this example each floor of each residence hall is further segmented into five groups; making a total of 25 segments. This example uses physical rooms for segmentation; however, a better method is segmentation by type of host using DHCP fingerprinting. Implementation of those technologies is beyond the scope of this document and highly network-specific.

Segmentation of State-Aware Devices

Any devices in the data path which maintain a state table (e.g.: firewalls) are susceptible to resource exhaustion during DoS attacks. Network engineers should consider physical segmentation of these devices to provide greater segmentation and resiliency of services.

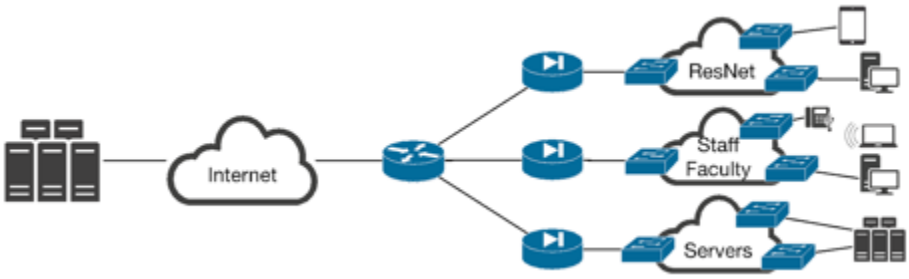

In this example each primary campus group sits behind a dedicated firewall. This allows Servers and Staff/Faculty to remain online when a ResNet gamer provokes a DDoS. If the firewall platform supports hard segmentation of resources a single physical device may be used; however, be careful to verify the platform can truly physically divide the hardware in this manner before proceeding.

Language to Set Expectation

The goal of R&E networking organizations is to mitigate risk (best effort) but not eliminate risk. The language will be slightly different for each organization, but all R&E networks should have the narrative to protect itself whether or not they provide additional security services. As organizations move into providing more security services, the language may need to evolve.

Sample Language – from MOREnet

1. MOREnet does not guarantee the security, integrity, confidentiality and/or availability of the Services. MOREnet will make a commercially reasonable effort to protect its Services against attacks or access from malicious or unauthorized users.
2. MOREnet may disable user IDs or remove, without notice, any data associated with a computer intrusion or other threatening activities to protect the Services or other institutions using the Services from damage. As applicable, MOREnet will advise the Member Security Contact as soon as possible should these actions occur.

Practices in Sampling Rate, Tools and Archiving for Flows

Please feel free to add content.

Appropriate whitelists and their rationale (CDNs, etc.)

Please feel free to add content

Enterprise Best Practices

"We all need to be responsible stewards of network resources."

Network Switch and Firewall Placement

Role-based Separation of IP Space

Architecture

Segmentation of ResNet

ResNet is a group which should be targeted for segmentation since so called "booter" services are designed for gamers to DDoS one another; causing a timeout which "boots" the target from the game. Due to ease of use and low cost such services have become popular in recent years. Since ResNet users are so often the target of DDoS attacks, segmenting them early may provide the greatest immediate benefit to the campus network.

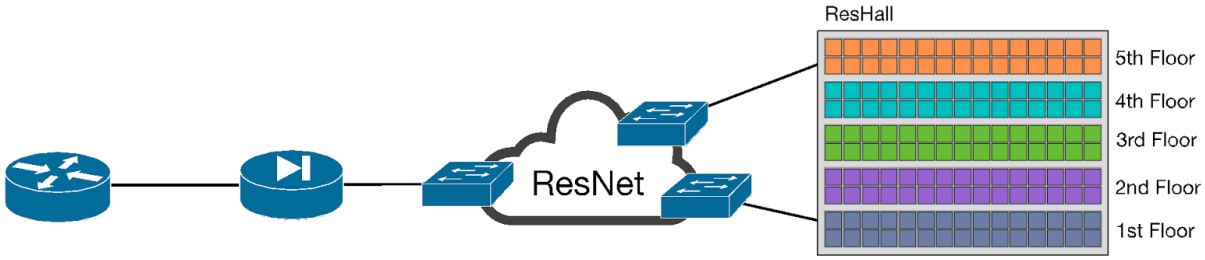

In this example a single residence hall with five floors is segmented into five different public IPs. If a user living on 1st Floor is attacked by a DDoS event the only impacted users are on the same floor in the same ResHall. If multiple IPs are assigned to each group in a pool the attack target can be removed from the pool and 1st Floor users impacted will be back online. Once the attack subsides the RTBHR can be removed and the NAT/PAT pool returned to its original configuration.

The natural progression of this practice is called microsegmentation.

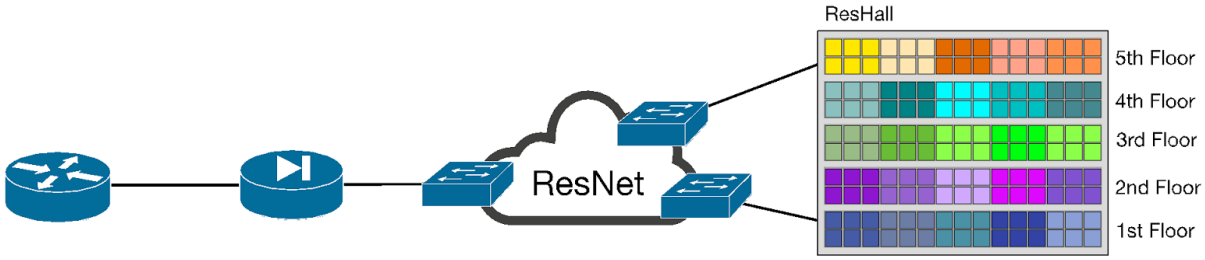

In this example each floor of each residence hall is further segmented into five groups; making a total of 25 segments. This example uses physical rooms for segmentation; however, a better method is segmentation by type of host using DHCP fingerprinting. Implementation of those technologies is beyond the scope of this document and highly network-specific.

Placeholder for Science DMZ information

Placeholder for Facilitating & Managing Eduroam Deployments

Resources

Readily Available Tools

KanREN, MOREnet, and UEN Security RT Presentations

Network Device Hardening Guides

Cisco IOS Hardening Guide

Self assessment tools

1. CIS (Center for Internet Security) has device benchmarks for free download (https://benchmarks.cisecurity.org/downloads/multiform/) including benchmarks for Checkpoint, Cisco, freeradius, ISC BIND, Juniper, and generic routers/wireless devices.

2. CIS/SANS also maintains a standard self-assessment tool for enterprise security, the CIS Cyber Security Controls (formerly known as the SANS Top 20 or CAG): https://www.sans.org/critical-security-controls, available for download at https://www.cisecurity.org/critical-controls/

Glossary / Taxonomy

"NITRD is posting this document with permission from The Quilt because it may have information that is of interest to our users. NITRD is not responsible for material contained in this document. Any information, opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NITRD Program. Posting this document on the NITRD website does not constitute endorsement of any product, service, organization, company, information provider, or content."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov