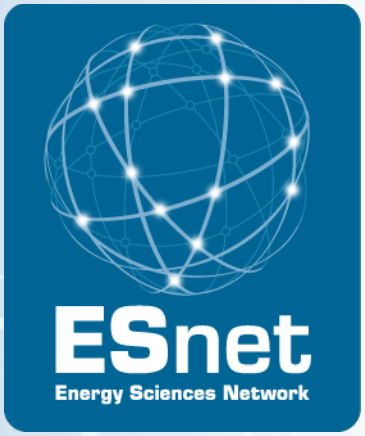

Science DMZ Security

Eli Dart, Network Engineer

ESnet Network Engineering Group

Joint Techs, Winter 2013

Honolulu, HI

January 15, 2013

Outline

- Quick background
- Firewall issues
- Non-firewall security options
- Touch on organizational structures

Science DMZ Background

The data mobility performance requirements for data intensive science are beyond what can typically be achieved using traditional methods

- Default host configurations (TCP, filesystems, NICs)
- Converged network architectures designed for commodity traffic
- Conventional security tools and policies
- Legacy data transfer tools (e.g. SCP)
- Wait-for-trouble-ticket operational models for network performance

The Science DMZ model describes a performance-based approach

- Dedicated infrastructure for wide-area data transfer
  - Well-configured data transfer hosts with modern tools
  - Capable network devices
  - High-performance data path which does not traverse commodity LAN
- Proactive operational models that enable performance
  - Well-deployed test and measurement tools (perfSONAR)
  - Periodic testing to locate issues instead of waiting for users to complain
- Security posture well-matched to high-performance science applications

Science DMZ – Simple Abstract Cartoon

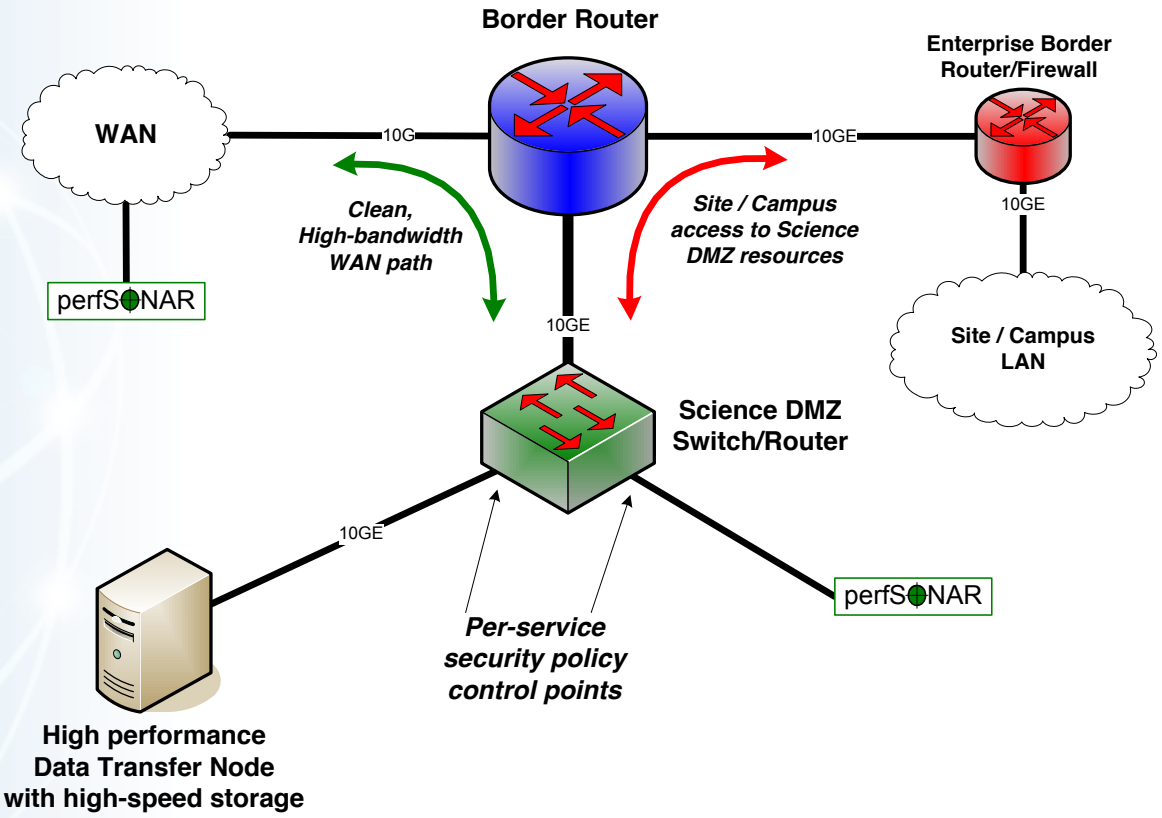

Science DMZ With Virtual Circuits/Openflow

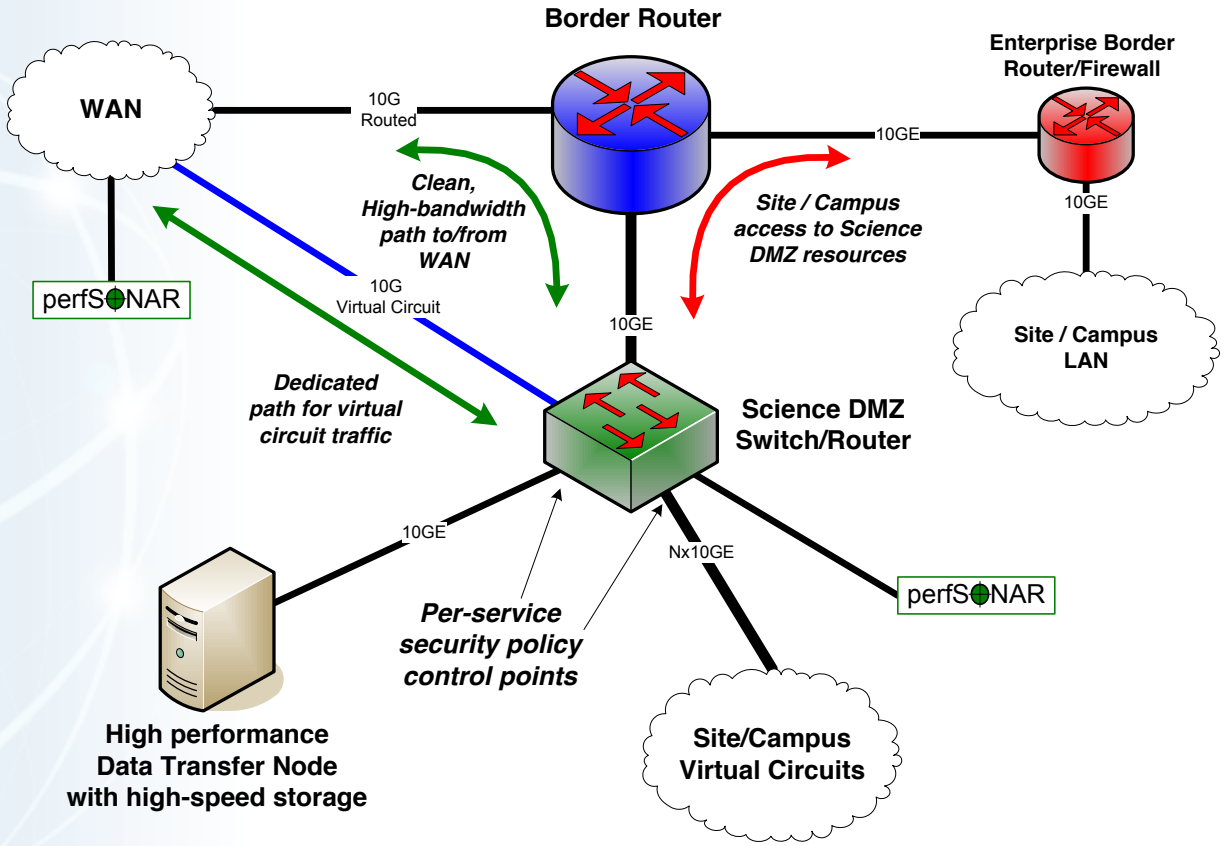

Science DMZ Supporting Multiple Projects

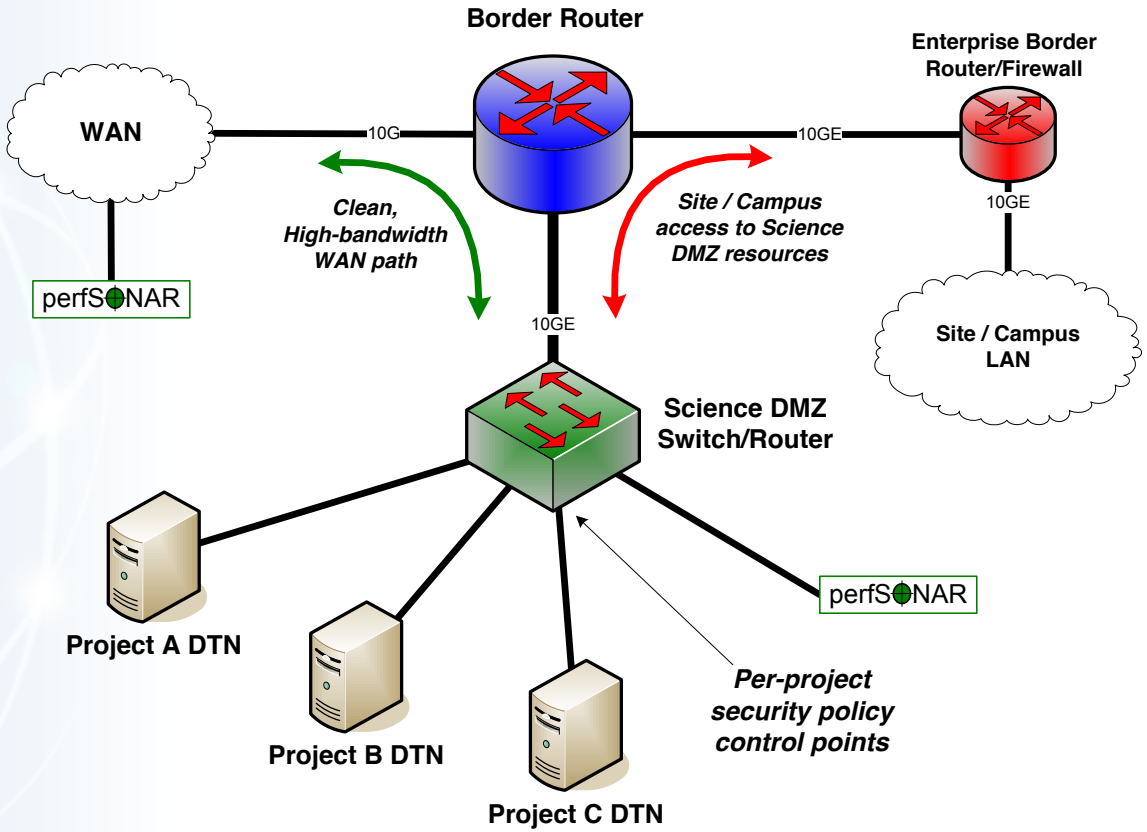

Science DMZ Security Model

Goal – disentangle security policy and enforcement for science flows from security for business systems

Rationale

- Science flows are relatively simple from a security perspective
- Narrow application set on Science DMZ
  - Data transfer, data streaming packages
  - No printers, document readers, web browsers, building control systems, staff desktops, etc.
- Security controls that are typically implemented to protect business resources often cause performance problems

Performance Is A Core Requirement

Core information security principles

- Confidentiality, Integrity, Availability (CIA)
- These apply to systems as well as to information, and have far-reaching effects
  - Credentials for privileged access must typically be kept confidential
  - Systems that are faulty or unreliable are not useful scientific tools
  - Data access is sometimes restricted, e.g. embargo before publication
  - Some data (e.g. medical data) has stringent requirements

In data-intensive science, performance is an additional core mission requirement

- CIA principles are important, but ***if the performance isn't there the science mission fails***
- This isn't about "how much" security you have, but how the security is implemented
- We need to be able to appropriately secure systems in a way that does not compromise performance or hinder the adoption of advanced services

Placement Outside the Firewall

The Science DMZ resources are placed outside the enterprise firewall for performance reasons

- The meaning of this is specific – ***Science DMZ traffic does not traverse the firewall data plane***
- This has nothing to do with whether packet filtering is part of the security enforcement toolkit

Lots of heartburn over this, especially from the perspective of a conventional firewall manager

- Lots of organizational policy directives mandating firewalls
- Firewalls are designed to protect converged enterprise networks
- Why would you put critical assets outside the firewall???

The answer is that firewalls are typically a poor fit for high-performance science applications

Let's Talk About Firewalls

A firewall's job is to enhance security by blocking activity that might compromise security

- This means that a firewall's job is to prevent things from happening
- Traditional firewall policy doctrine dictates a default-deny policy
  - Find out what business you need to do
  - Block everything else

Firewalls are typically designed for commodity or enterprise environments

- This makes sense from the firewall designer's perspective – lots of IT spending in commodity environments
- Firewall design choices are well-matched to commodity traffic profile
  - High device count, high user count, high concurrent flow count
  - Low per-flow bandwidth
  - Highly capable inspection and analysis of business applications

Thought Experiment

- We're going to do a thought experiment
- Consider a network between three buildings – A, B, and C
- This is supposedly a 10Gbps network end to end (look at the links on the buildings)
- Building A houses the border router – not much goes on there except the external connectivity
- Lots of work happens in building B – so much so that the processing is done with multiple processors to spread the load in an affordable way, and aggregate the results after
- Building C is where we branch out to other buildings
- Every link between buildings is 10Gbps – this is a 10Gbps network, right???

Notional 10G Network Between Buildings

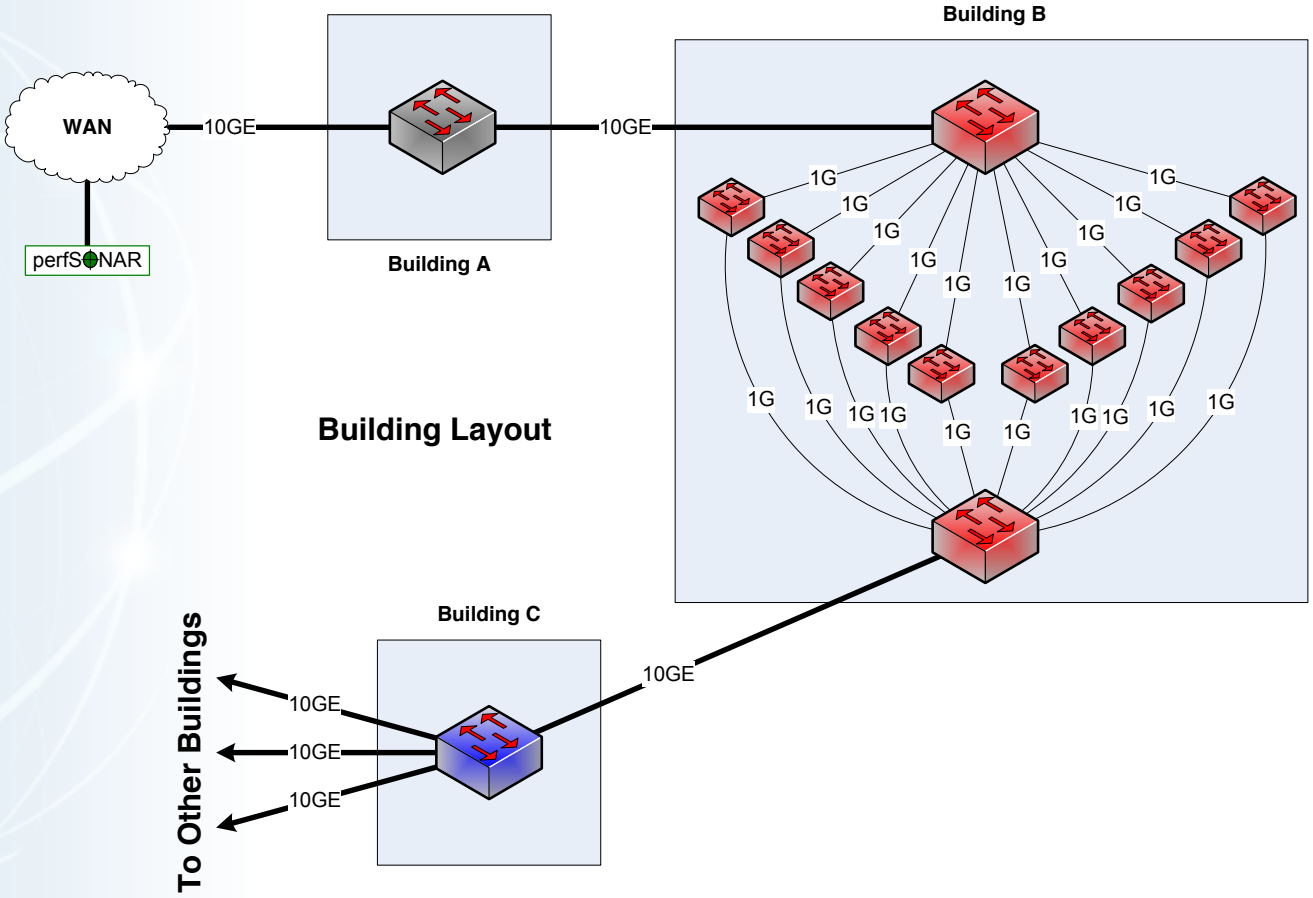

Clearly Not A 10Gbps Network

If you look at the inside of Building B, it is obvious from a network engineering perspective that this is not a 10Gbps network

- Clearly the maximum per-flow data rate is 1Gbps, not 10Gbps
- However, if you convert the buildings into network elements while keeping their internals intact, you get routers and firewalls
- What firewall did the organization buy? What's inside it?
- Those little 1G "switches" are firewall processors

This parallel firewall architecture has been in use for years

- Slower processors are cheaper
- Typically fine for a commodity traffic load
- Therefore, this design is cost competitive and common

Notional 10G Network Between Devices

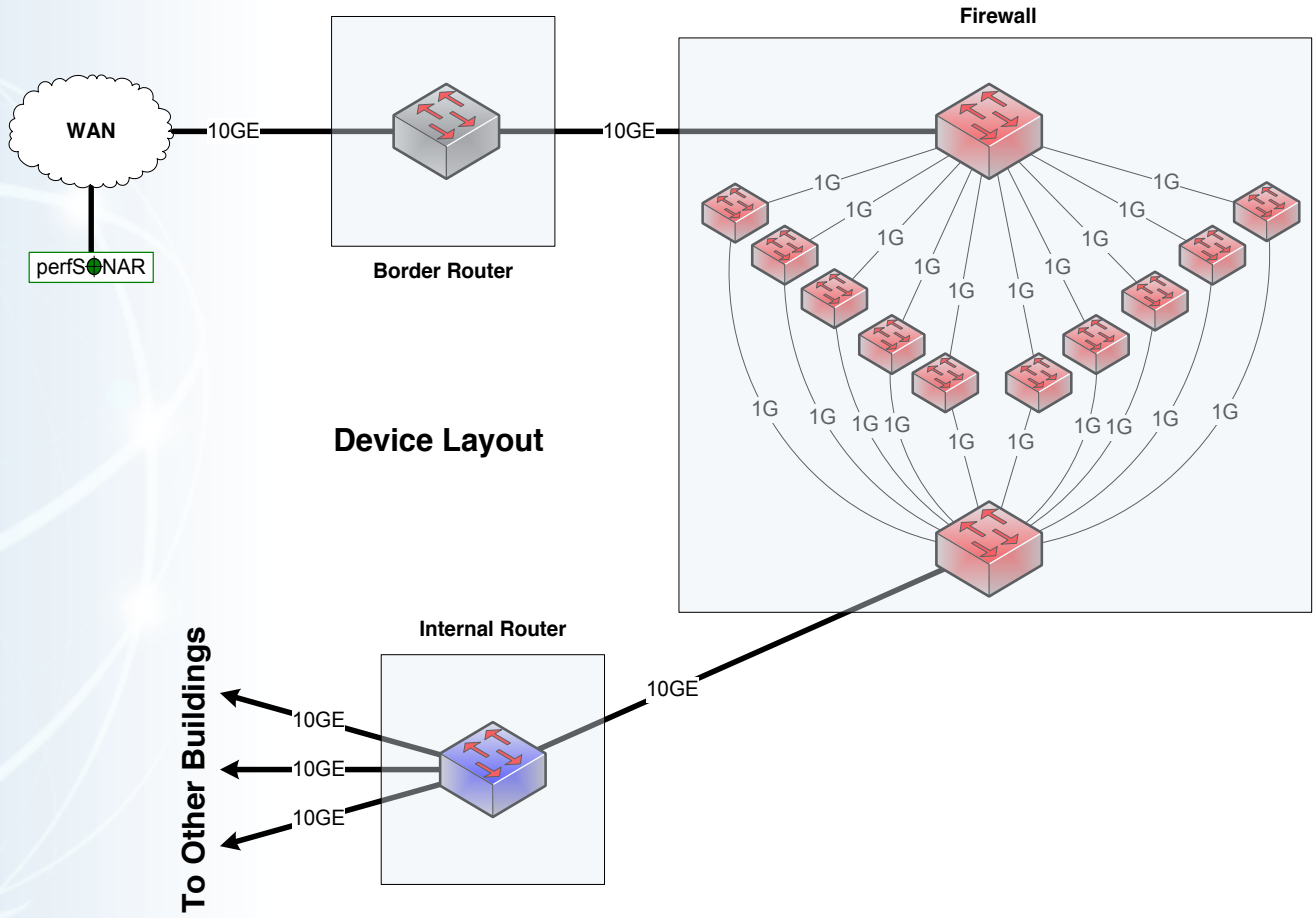

Notional Network Logical Diagram

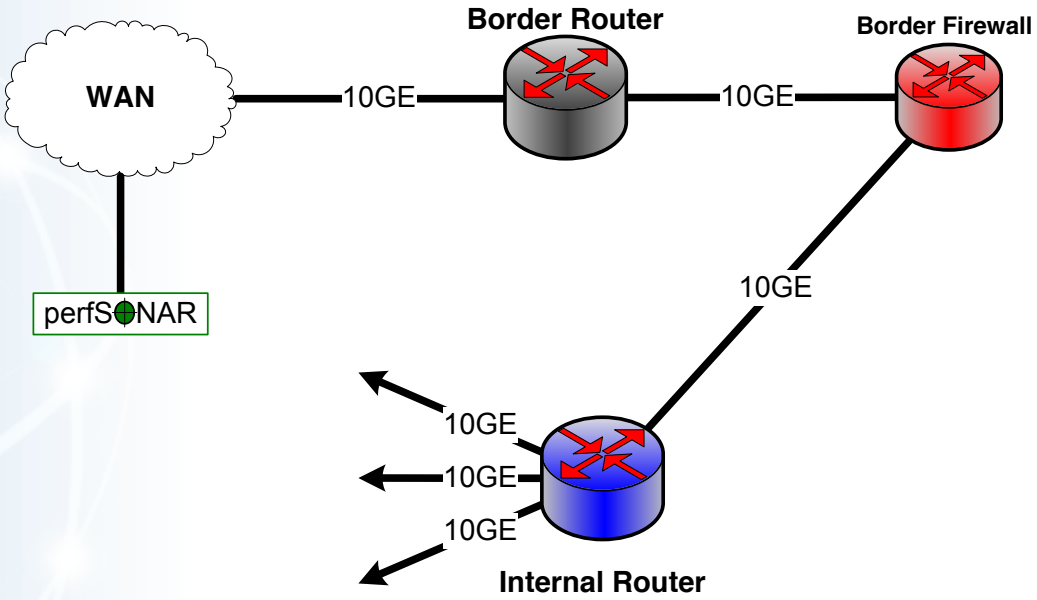

What's Inside Your Firewall?

"But wait – we don't do this anymore!"

- It is true that vendors are working toward line-rate 10G firewalls, and some may even have them now
- 10GE has been deployed in science environments for over 10 years
- Firewall internals have only recently started to catch up with the 10G world
- 100GE is being deployed now, 40Gbps host interfaces are available now
- Firewalls are behind again

In general, IT shops want to get 5+ years out of a firewall purchase

- This often means that the firewall is years behind the technology curve
- Whatever you deploy now, that's the hardware feature set you get
- When a new science project tries to deploy data-intensive resources, they get whatever feature set was purchased several years ago

The Firewall State Table

Many firewalls use a state table to improve performance

- State table lookup is fast
- No need to process entire ruleset for every packet
- Also allows session tracking (e.g. TCP sequence numbers)

State table built dynamically

- Incoming packets are matched against the state table
- If no state table entry, go to the ruleset
- If permitted by ruleset, create state table entry
- Remove state table entry after observing connection teardown

Semantically similar to punt-and-switch model of traffic forwarding used on many older routers

State Table Issues

If the state table is not pruned, it will overflow

- Not all connections close cleanly
  - I shut my laptop and go to a meeting
  - Software crashes happen
- Some attacks try to fill state tables

Solution: put a timer on state table entries

- When a packet matches the state table entry, update the timer
- If the timer expires, delete the state table entry

What if I just pause for a few minutes?

- This turns out to be a problem – state table timers are typically in the 5-15 minute range, while host keepalive timers are 2 hours
- If a connection pauses (e.g. to wait for a large transfer), the firewall will delete the state table entry from under it – connection hangs
- I have seen this in production environments

Firewall Capabilities and Science Traffic

Firewalls have a lot of sophistication in an enterprise setting

- Application layer protocol analysis (HTTP, POP, MSRPC, etc.)
- Built-in VPN servers
- User awareness

Data-intensive science flows don't match this profile

- Common case – data on filesystem A needs to be on filesystem Z
  - Data transfer tool verifies credentials over an encrypted channel
  - Then open a socket or set of sockets, and send data until done (1TB, 10TB, 100TB, …)
- One workflow can use 10% to 50% or more of a 10G network link

Do we have to use a firewall?

Firewalls As Access Lists

When you ask a firewall administrator to allow data transfers through the firewall, what do they ask for?

- IP address of your host
- IP address of the remote host
- Port range
- ***That looks like an ACL to me!***

No special config for advanced protocol analysis – just address/port

Router ACLs are better than firewalls at address/port filtering

- ACL capabilities are typically built into the router
- Router ACLs typically do not drop traffic permitted by policy

Security Without Firewalls

Data intensive science traffic interacts poorly with firewalls

Does this mean we ignore security? ***NO!***

- We **must** protect our systems
- We just need to find a way to do security that does not prevent us from getting the science done

Key point – security policies and mechanisms that protect the Science DMZ should be implemented so that they do not compromise performance

If Not Firewalls, Then What?

- Remember – the goal is to protect systems in a way that allows the science mission to succeed
- I like something I heard at NERSC – paraphrasing: "Security controls should enhance the utility of science infrastructure."
- There are multiple ways to solve this – some are technical, and some are organizational/sociological
- I'm not going to lie to you – this is harder than just putting up a firewall and closing your eyes

Other Technical Capabilities

Intrusion Detection Systems (IDS)

- One example is Bro – http://bro-ids.org/
- Bro is high-performance and battle-tested
  - Bro protects several high-performance national assets
  - Bro can be scaled with clustering: http://www.bro-ids.org/documentation/cluster.html
- Other IDS solutions are available also

Netflow and IPFIX can provide intelligence, but not filtering

Openflow and SDN

- Using Openflow to control access to a network-based service seems pretty obvious
- There is clearly a hole in the ecosystem with the label “Openflow Firewall” – I really hope someone is working on this (it appears so)
- This could significantly reduce the attack surface for any authenticated network service
- This would only work if the Openflow device had a robust data plane

Other Technical Capabilities (2)

Aggressive access lists

- More useful with project-specific DTNs
- If the purpose of the DTN is to exchange data with a small set of remote collaborators, the ACL is pretty easy to write
- Large-scale data distribution servers are hard to handle this way (but then, the firewall ruleset for such a service would be pretty open too)

Limitation of the application set

- One of the reasons to limit the application set in the Science DMZ is to make it easier to protect
- Keep desktop applications off the DTN (and watch for them anyway using logging, netflow, etc – take violations seriously)
- This requires collaboration between people – networking, security, systems, and scientists

Collaboration Within The Organization

All stakeholders should collaborate on Science DMZ design, policy, and enforcement

The security people have to be on board

- Remember: security people already have political cover – it's called the firewall
- If a host gets compromised, the security officer can say they did their due diligence because there was a firewall in place
- If the deployment of a Science DMZ is going to jeopardize the job of the security officer, expect pushback

The Science DMZ is a strategic asset, and should be understood by the strategic thinkers in the organization

- Changes in security models
- Changes in operational models
- Enhanced ability to compete for funding
- Increased institutional capability – greater science output

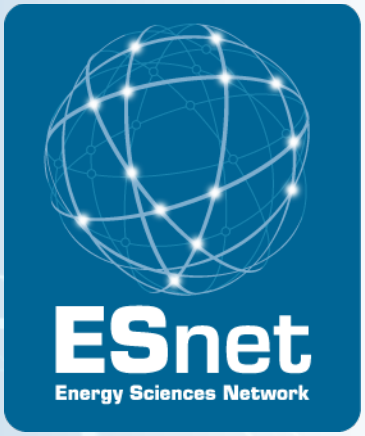

Questions?

Thanks!

Eli Dart - dart@es.net

http://www.es.net/

http://fasterdata.es.net/