The Science DMZ

Eli Dart, ESnet

Focused Technical Workshop for Life Sciences

Berkeley, CA

July 17, 2013

Outline

Science DMZ background

Science DMZ Architecture

PerfSONAR

The Data Transfer Node

Science DMZ Security

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

Motivation

Science data increasing both in volume and in value

- Higher instrument performance
- Increased capacity for discovery
- Analyses previously not possible

Lots of promise, but only if scientists can actually work with the data

- Data has to get to analysis resources
- Results have to get to people
- People have to share results

Common pain point – data mobility

- Movement of data between instruments, facilities, analysis systems, and scientists is a gating factor for much of data intensive science
- Data mobility is not the only part of data intensive science – not even the most important part
- However, without data mobility data intensive science is hard

We need to move data – how can we do it consistently well?

Motivation (2)

Networks play a crucial role

- The very structure of modern science assumes science networks exist – high performance, feature rich, global scope
- Networks enable key aspects of data intensive science
  - Data mobility, automated workflows
  - Access to facilities, data, analysis resources

Messing with the network is unpleasant for most scientists

- Not their area of expertise
- Not where the value is (no papers come from messing with the network)
- Data intensive science is about the science, not about the network
- However, it's a critical service – if the network breaks, everything stops

Therefore, infrastructure providers must cooperate to build consistent, reliable, high performance network services for data mobility

Here we describe one blueprint, the Science DMZ model – there are certainly others, but this one seems to work well in a variety of environments

TCP Background

Networks provide connectivity between hosts – how do hosts see the network?

- From an application's perspective, the interface to "the other end" is a socket
- Other similar constructs exist for non-IP protocols
- Communication is between applications – mostly over TCP

TCP – the fragile workhorse

- TCP is (for very good reasons) timid – packet loss is interpreted as congestion
- Packet loss in conjunction with latency is a performance killer
- Like it or not, TCP is used for the vast majority of data transfer applications

TCP Background (2)

It is far easier to architect the network to support TCP than it is to fix TCP

- People have been trying to fix TCP for years – limited success
- Here we are – packet loss is still the number one performance killer in long distance high performance environments

Pragmatically speaking, we must accommodate TCP

- Implications for equipment selection
  - Equipment must be able to accurately account for packets
- Implications for network architecture, deployment models
  - Infrastructure must be designed to allow easy troubleshooting
  - Test and measurement tools are critical – they have to be deployed

A small amount of packet loss makes a huge difference in TCP performance

A Nagios alert based on our regular throughput testing between one site and ESnet core alerted us to poor performance on high latency paths

No errors or drops reported by routers on either side of problem link

- only perfSONAR bwctl tests caught this problem

Using packet filter counters, we saw 0.0046% loss in one direction

- 1 packet in 22000 packets

Performance impact of this: (outbound/inbound)

- To/from test host 1 ms RTT : 7.3 Gbps out / 9.8 Gbps in
- To/from test host 11 ms RTT: 1 Gbps out / 9.5 Gbps in
- To/from test host 51ms RTT: 122 Mbps out / 7 Gbps in
- To/from test host 88 ms RTT: 60 Mbps out / 5 Gbps in
  - More than 80 times slower!

A small amount of packet loss makes a huge difference in TCP performance

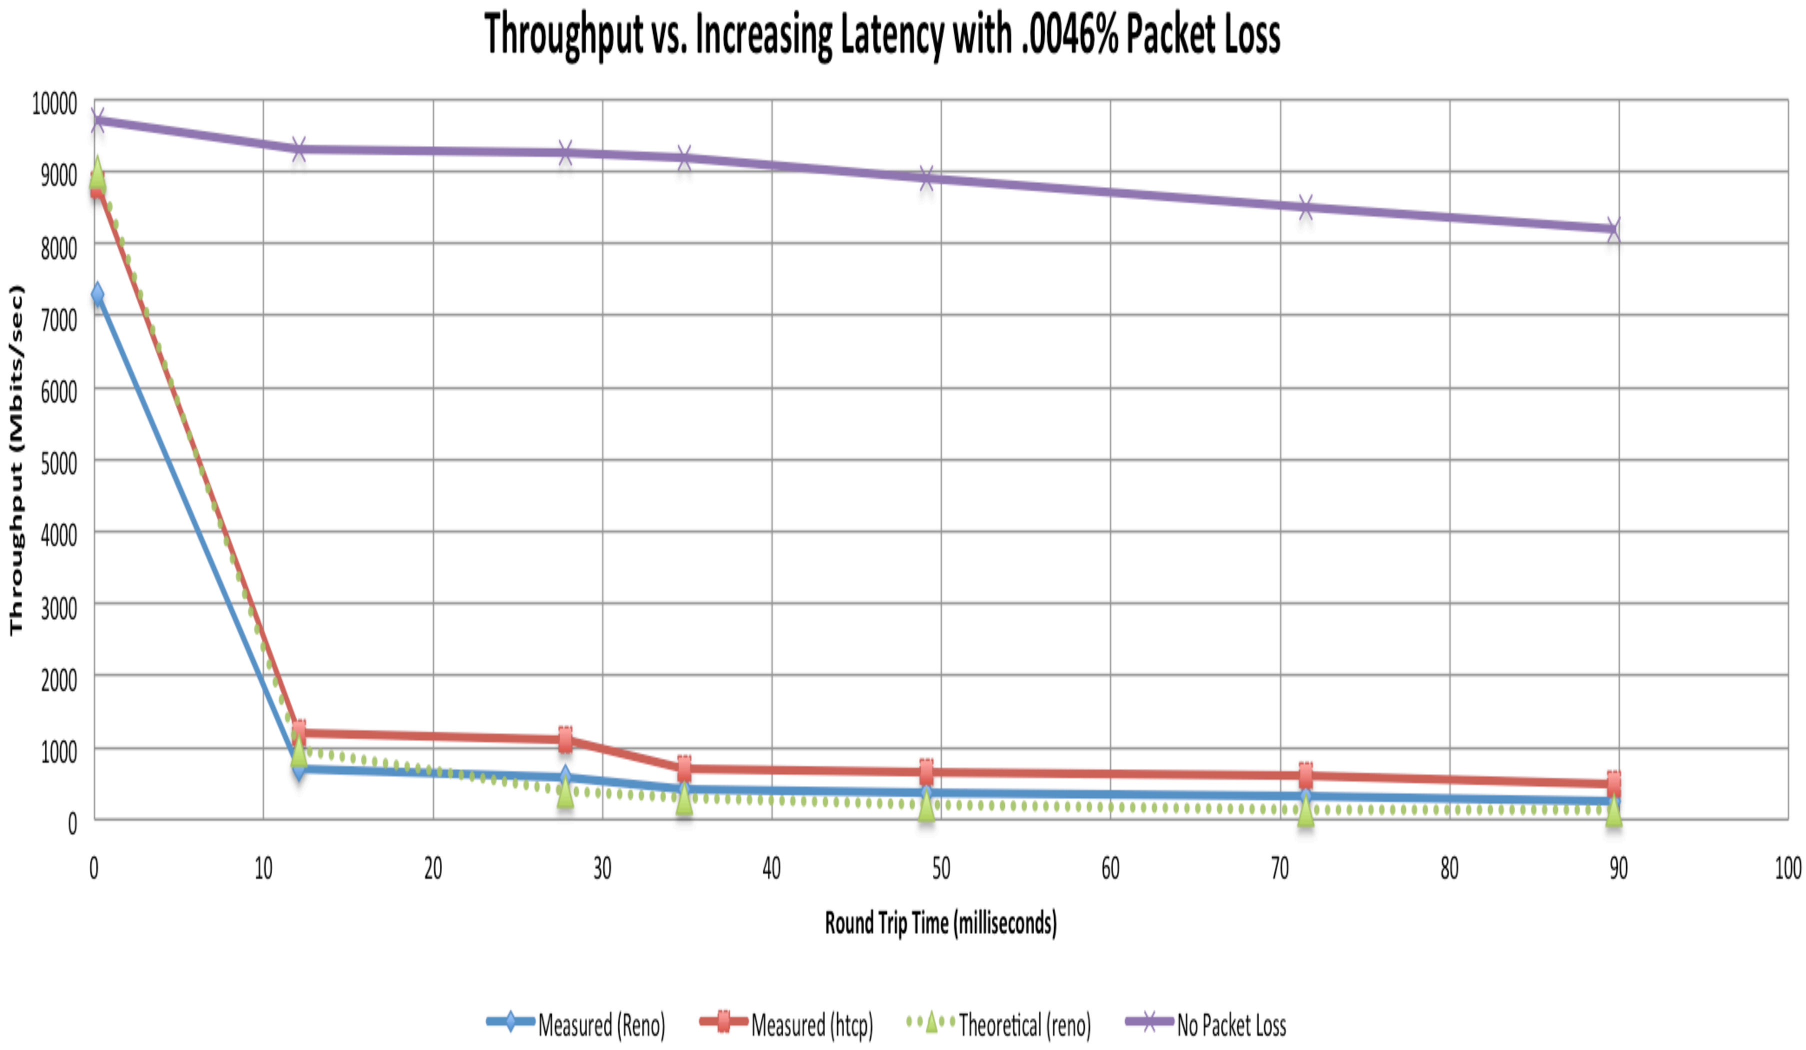

How Do We Accommodate TCP?

High-performance wide area TCP flows must get loss-free service

- Sufficient bandwidth to avoid congestion
- Deep enough buffers in routers and switches to handle bursts
  - Especially true for long-distance flows due to packet behavior
  - No, this isn't buffer bloat

Equally important – the infrastructure must be verifiable so that clean service can be provided

- Stuff breaks
  - Hardware, software, optics, bugs, …
  - How do we deal with it in a production environment?
- Must be able to prove a network device or path is functioning correctly
  - Accurate counters must exist and be accessible
  - Need ability to run tests - perfSONAR
- Small footprint is a huge win – small number of devices so that problem isolation is tractable

Traditional DMZ

DMZ – "Demilitarized Zone"

- Network segment near the site perimeter with different security policy
- Commonly used architectural element for deploying WAN-facing services (e.g. email, DNS, web)

Traffic for WAN-facing services does not traverse the LAN

- WAN flows are isolated from LAN traffic
- Infrastructure for WAN services is specifically configured for WAN

Separation of security policy improves both LAN and WAN

- No conflation of security policy between LAN hosts and WAN services
- DMZ hosts provide specific services
- LAN hosts must traverse the same ACLs as WAN hosts to access DMZ

The Data Transfer Trifecta: The "Science DMZ" Model

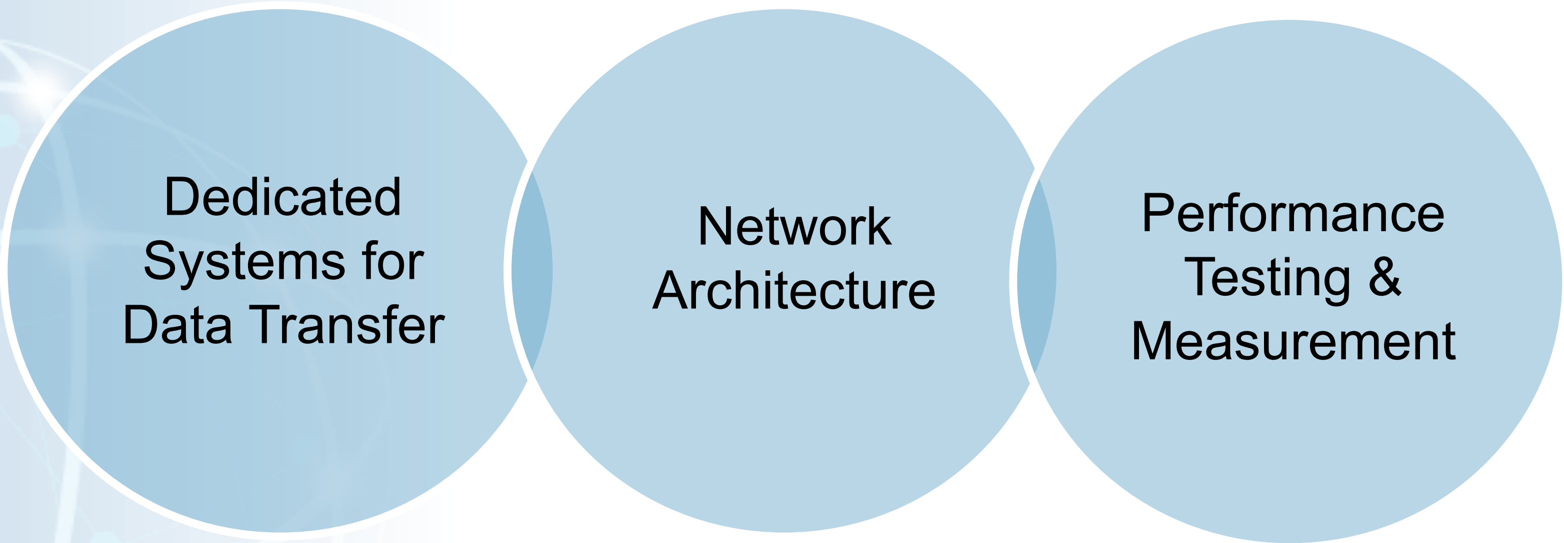

Dedicated Systems for Data Transfer

Network Architecture

Performance Testing & Measurement

Data Transfer Node

- High performance
- Configured for data transfer
- Proper tools

Science DMZ

- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network
- Additional info: http://fasterdata.es.net/

perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

Science DMZ Takes Many Forms

There are a lot of ways to combine these things – it all depends on what you need to do

- Small installation for a project or two
- Facility inside a larger institution
- Institutional capability serving multiple departments/divisions
- Science capability that consumes a majority of the infrastructure

Some of these are straightforward, others are less obvious

Key point of concentration: High-latency path for TCP

The Data Transfer Trifecta: The “Science DMZ” Model

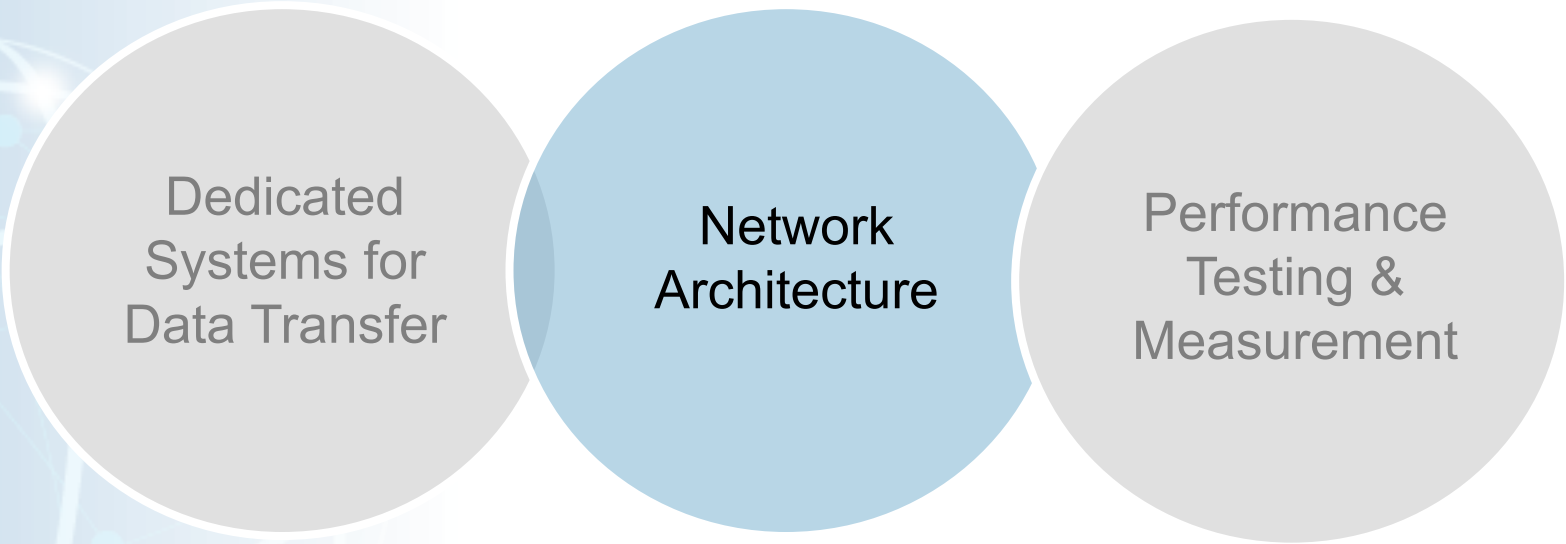

Dedicated Systems for Data Transfer

Network Architecture

Performance Testing & Measurement

Data Transfer Node

- High performance
- Configured for data transfer
- Proper tools

Science DMZ

- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network
- Additional info: http://fasterdata.es.net/

perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

Ad Hoc DTN Deployment

This is often what gets tried first

Data transfer node deployed where the owner has space

- This is often the easiest thing to do at the time
- Straightforward to turn on, hard to achieve performance

If present, perfSONAR is at the border

- This is a good start
- Need a second one next to the DTN

Entire LAN path has to be sized for data flows

Entire LAN path is part of any troubleshooting exercise

Ad Hoc DTN Deployment

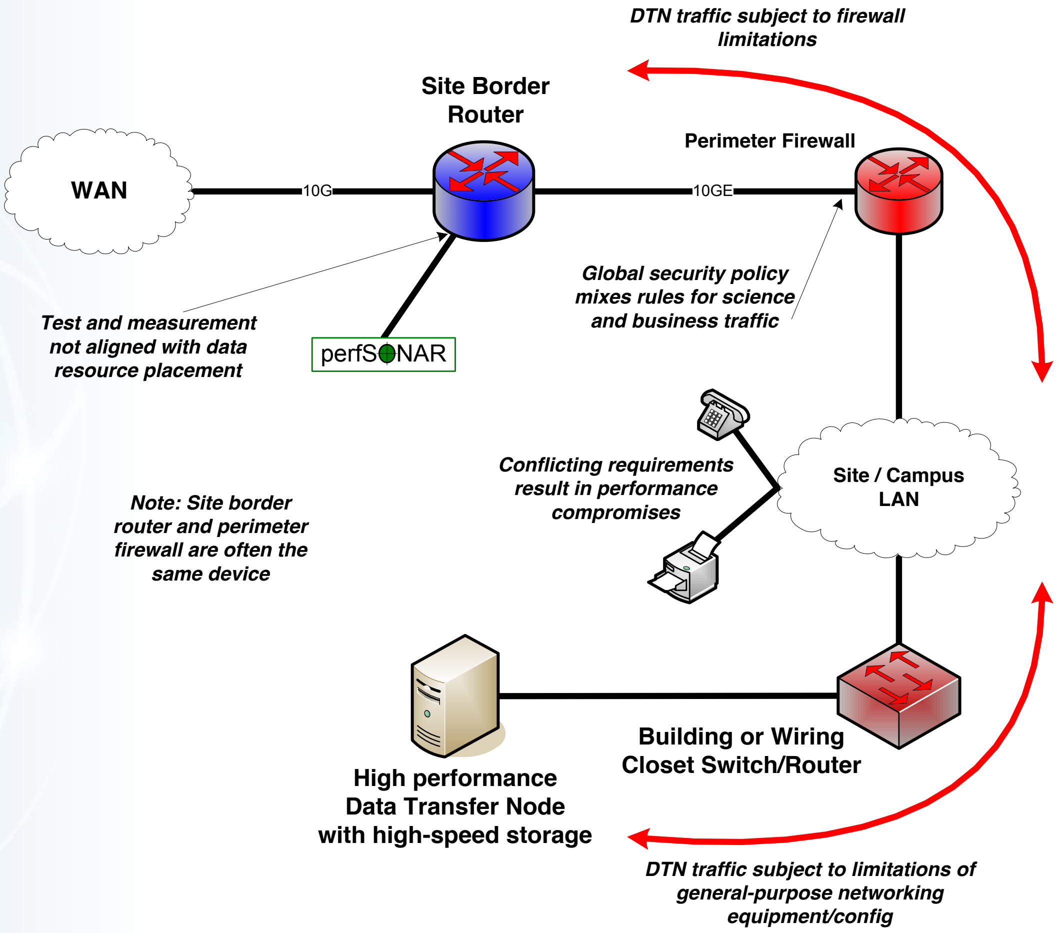

DTN traffic subject to firewall limitations

Site Border Router

Perimeter Firewall

Test and measurement not aligned with data resource placement

Global security policy mixes rules for science and business traffic

Conflicting requirements result in performance compromises

Note: Site border router and perimeter firewall are often the same device

High performance Data Transfer Node with high-speed storage

Building or Wiring Closet Switch/Router

DTN traffic subject to limitations of general-purpose networking equipment/config

Small-scale or Prototype Deployment

Add-on to existing network infrastructure

- All that is required is a port on the border router
- Small footprint, pre-production commitment

Easy to experiment with components and technologies

- DTN prototyping
- perfSONAR testing

Limited scope makes security policy exceptions easy

- Only allow traffic from partners
- Add-on to production infrastructure – lower risk

Prototype Science DMZ

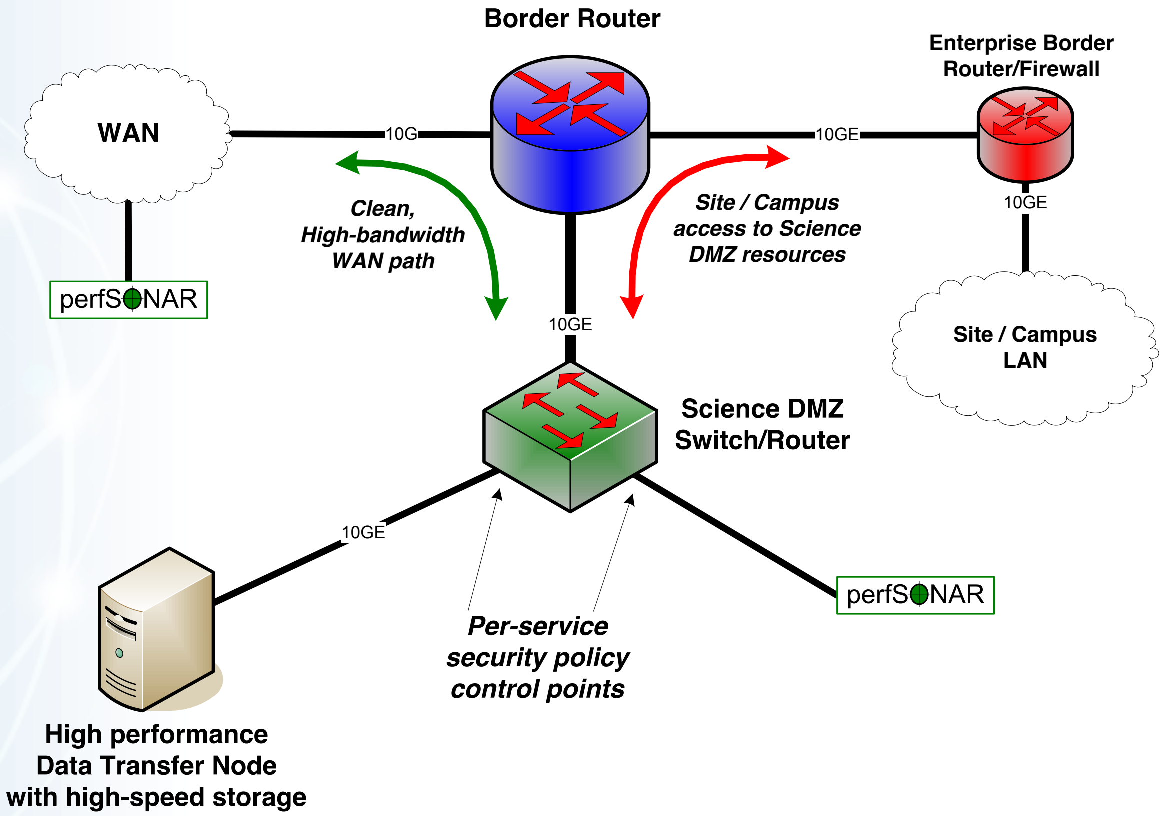
Border Router
Enterprise Border Router/Firewall
WAN
10G
10GE
10GE
perfSONAR
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
10GE
Site / Campus LAN
Science DMZ Switch/Router
10GE
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage

Prototype Science DMZ Data Path

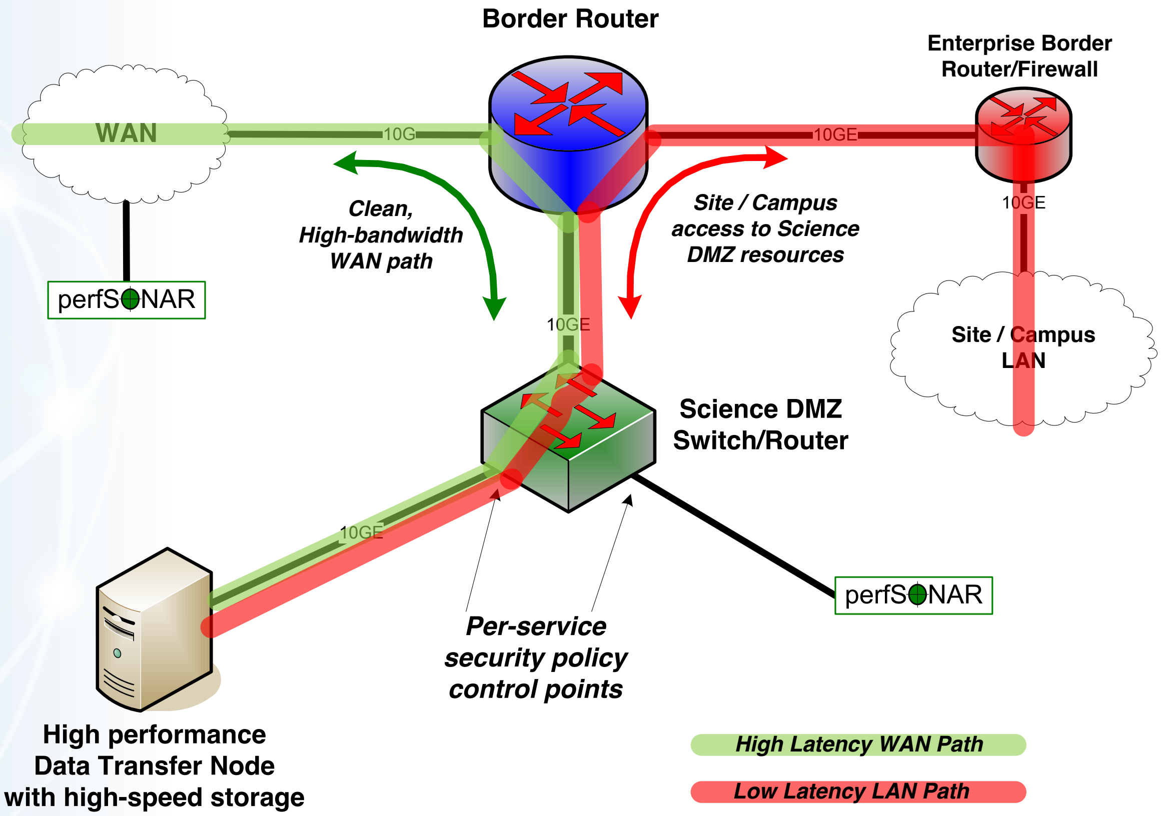
Border Router
Enterprise Border Router/Firewall
WAN
10G
10GE
10GE
perfSONAR
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
10GE
Site / Campus LAN
Science DMZ Switch/Router
10GE
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage
High Latency WAN Path
Low Latency LAN Path

Prototype With Virtual Circuits

Small virtual circuit prototype can be done in a small Science DMZ

- Perfect example is a Software Defined Networking (SDN) testbed
- Virtual circuit connection may or may not traverse border router

As with any Science DMZ deployment, this can be expanded as need grows

In this particular diagram, Science DMZ hosts can use either the routed or the circuit connection

Virtual Circuit Prototype Deployment

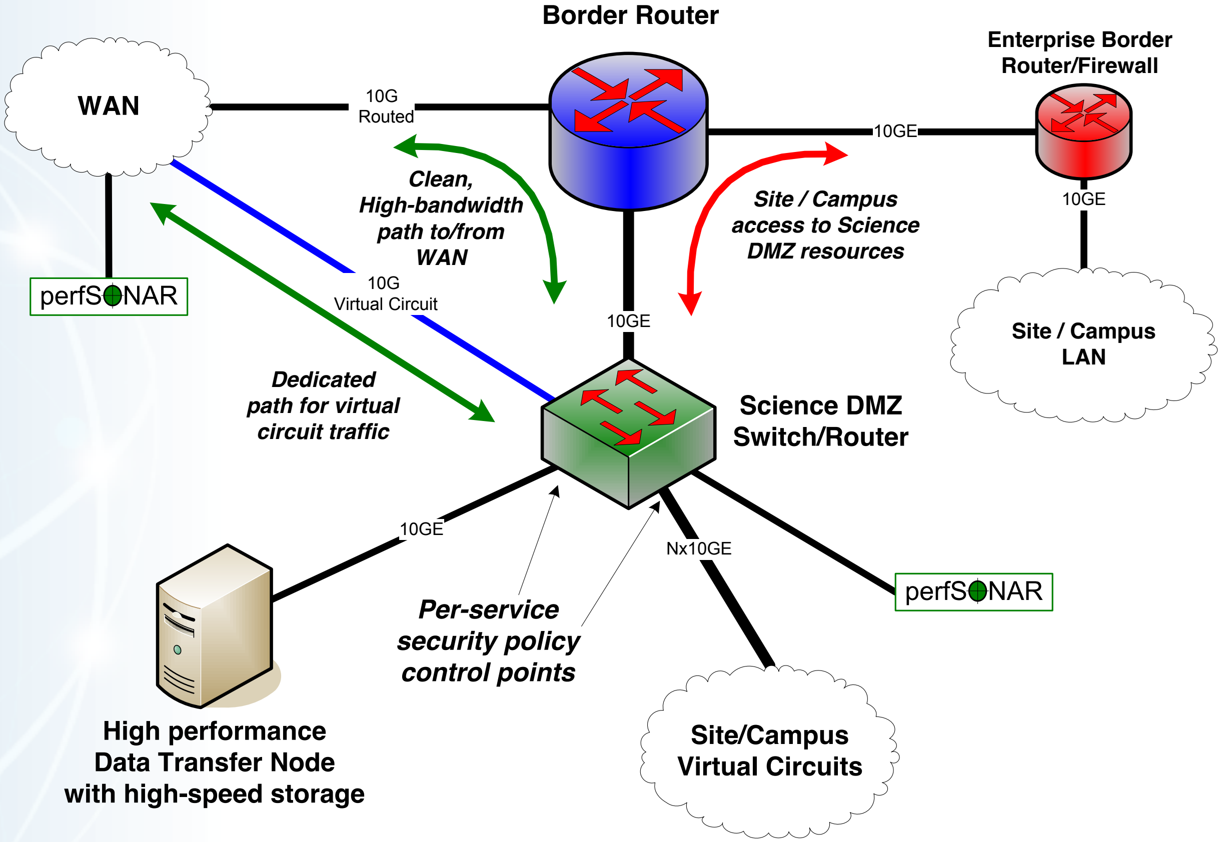
Border Router
Enterprise Border Router/Firewall
WAN
10G Routed
10GE
Clean, High-bandwidth path to/from WAN
Site / Campus access to Science DMZ resources
10GE
perfSONAR
10G Virtual Circuit
10GE
Site / Campus LAN
Dedicated path for virtual circuit traffic
Science DMZ Switch/Router
10GE
Nx10GE
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage
Site/Campus Virtual Circuits

Virtual Circuit Prototype Data Path

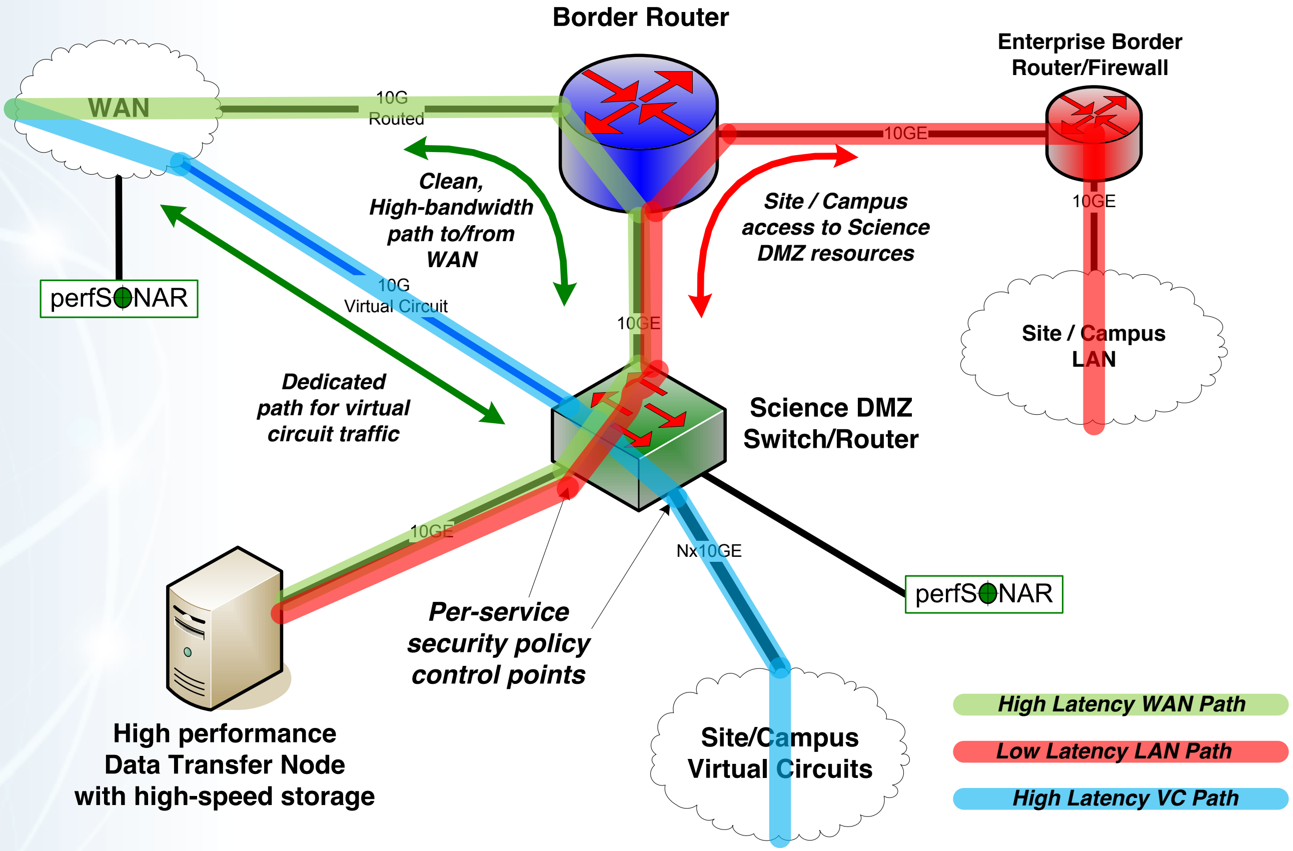
Border Router
Enterprise Border Router/Firewall
WAN
10G Routed
10GE
Clean, High-bandwidth path to/from WAN
Site / Campus access to Science DMZ resources
10GE
perfSONAR
10G Virtual Circuit
10GE
Site / Campus LAN
Dedicated path for virtual circuit traffic
Science DMZ Switch/Router
10GE
Nx10GE
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage
Site/Campus Virtual Circuits
High Latency WAN Path
Low Latency LAN Path
High Latency VC Path

Support For Multiple Projects

Science DMZ architecture allows multiple projects to put DTNs in place

- Modular architecture
- Centralized location for data servers

This may or may not work well depending on institutional politics

- Issues such as physical security can make this a non-starter
- On the other hand, some shops already have service models in place

On balance, this can provide a cost savings – it depends

- Central support for data servers vs. carrying data flows
- How far do the data flows have to go?

Multiple Projects

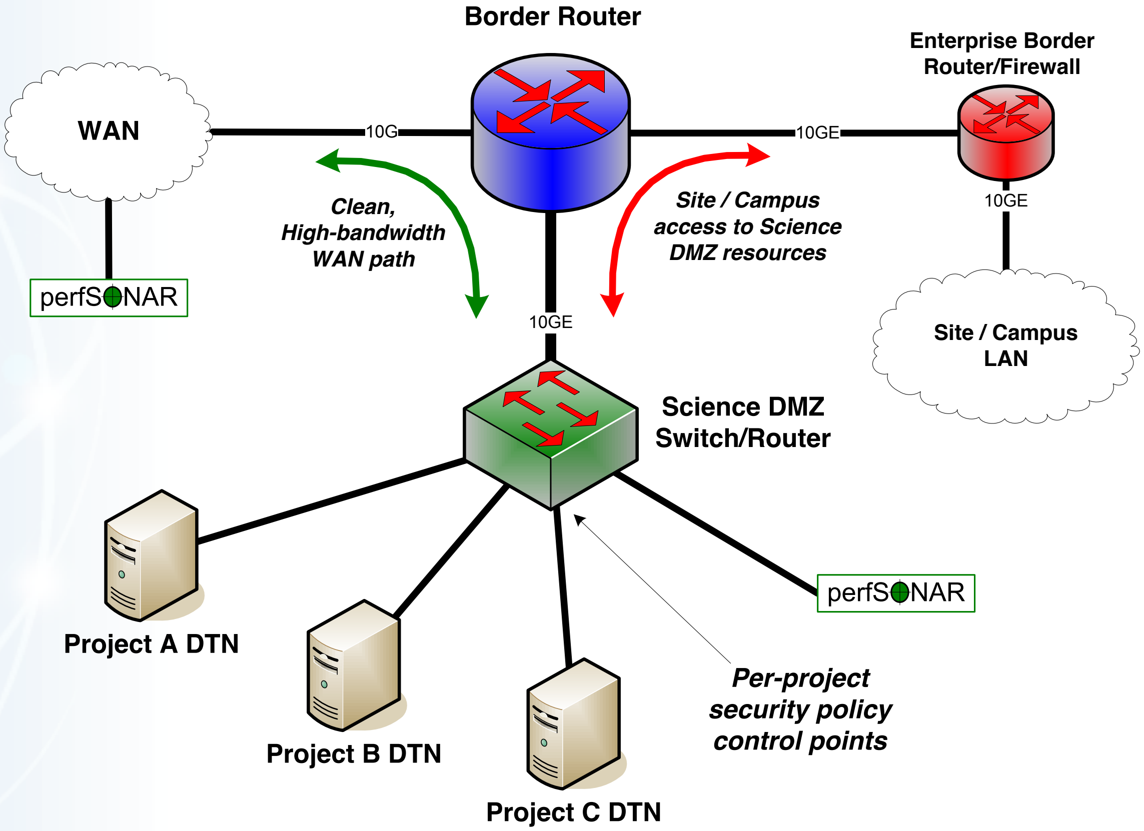
Border Router
Enterprise Border Router/Firewall
WAN
10G
10GE
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
10GE
perfSONAR
10GE
Site / Campus LAN
Science DMZ Switch/Router
Project A DTN
Project B DTN
Project C DTN
perfSONAR
Per-project security policy control points

Supercomputer Center Deployment

High-performance networking is assumed in this environment

- Data flows between systems, between systems and storage, wide area, etc.
- Global filesystem often ties resources together
  - Portions of this may not run over Ethernet (e.g. IB)
  - Implications for Data Transfer Nodes

“Science DMZ” may not look like a discrete entity here

- By the time you get through interconnecting all the resources, you end up with most of the network in the Science DMZ
- This is as it should be – the point is appropriate deployment of tools, configuration, policy control, etc.

Office networks can look like an afterthought, but they aren’t

- Deployed with appropriate security controls
- Office infrastructure need not be sized for science traffic

Supercomputer Center

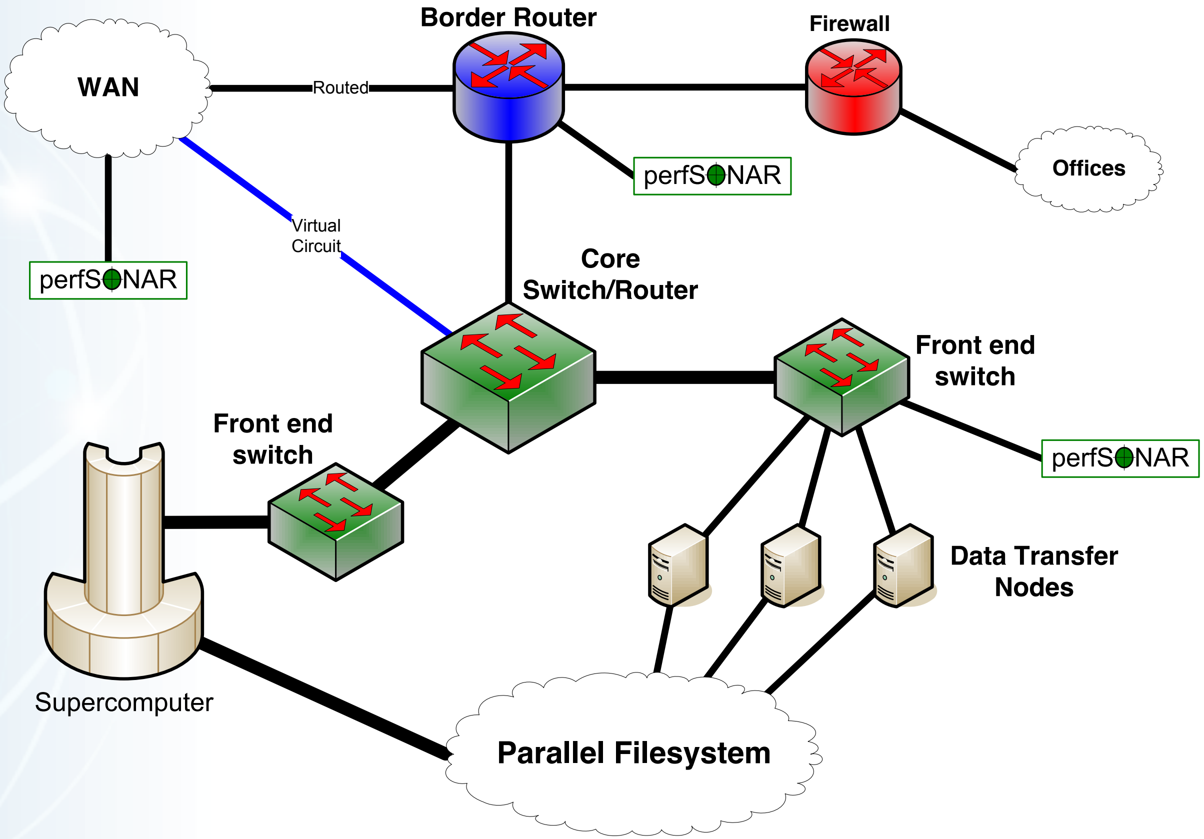

Supercomputer Center Data Path

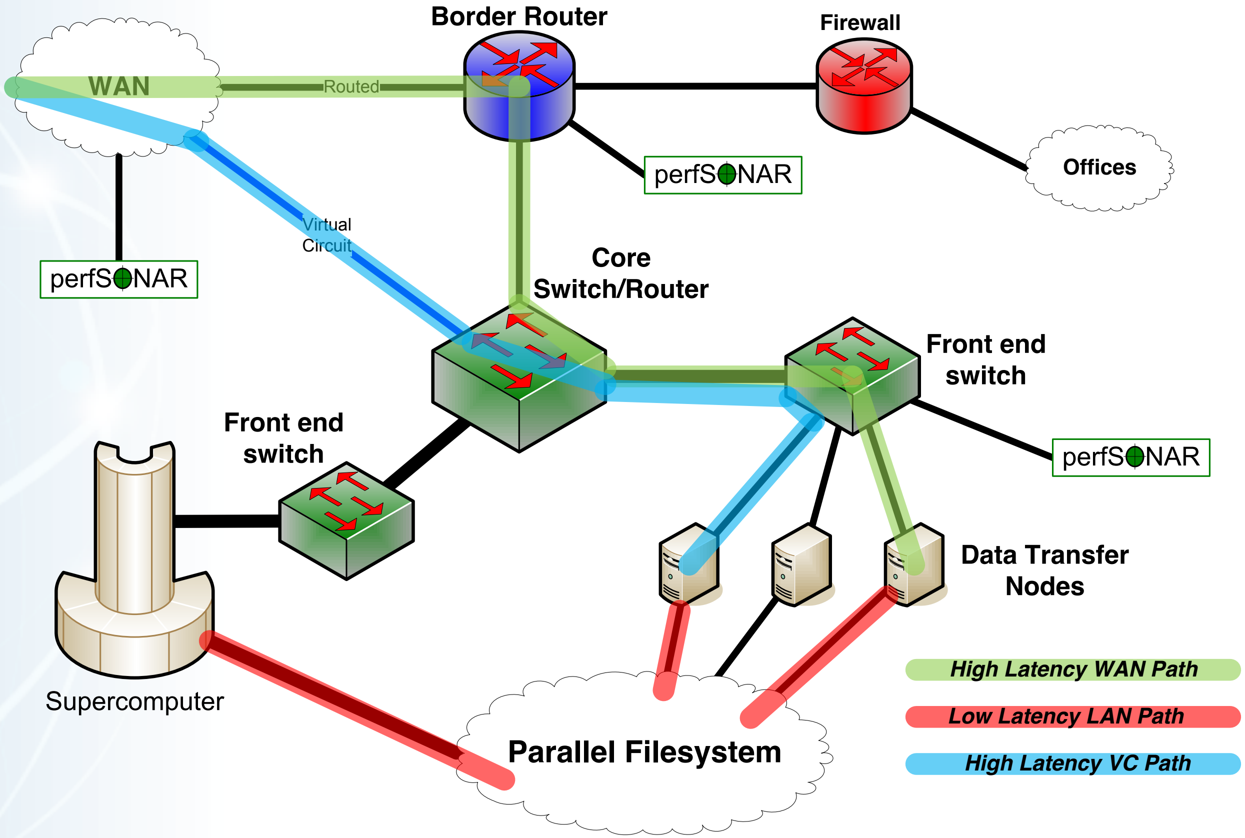

Major Data Site Deployment

In some cases, large scale data service is the major driver

- Huge volumes of data – ingest, export
- Individual DTNs don't exist here – data transfer clusters

Single-pipe deployments don't work

- Everything is parallel
  - Networks (Nx10G LAGs, soon to be Nx100G)
  - Hosts – data transfer clusters, no individual DTNs
  - WAN connections – multiple entry, redundant equipment
- Choke points (e.g. firewalls) cause problems

Data Site – Architecture

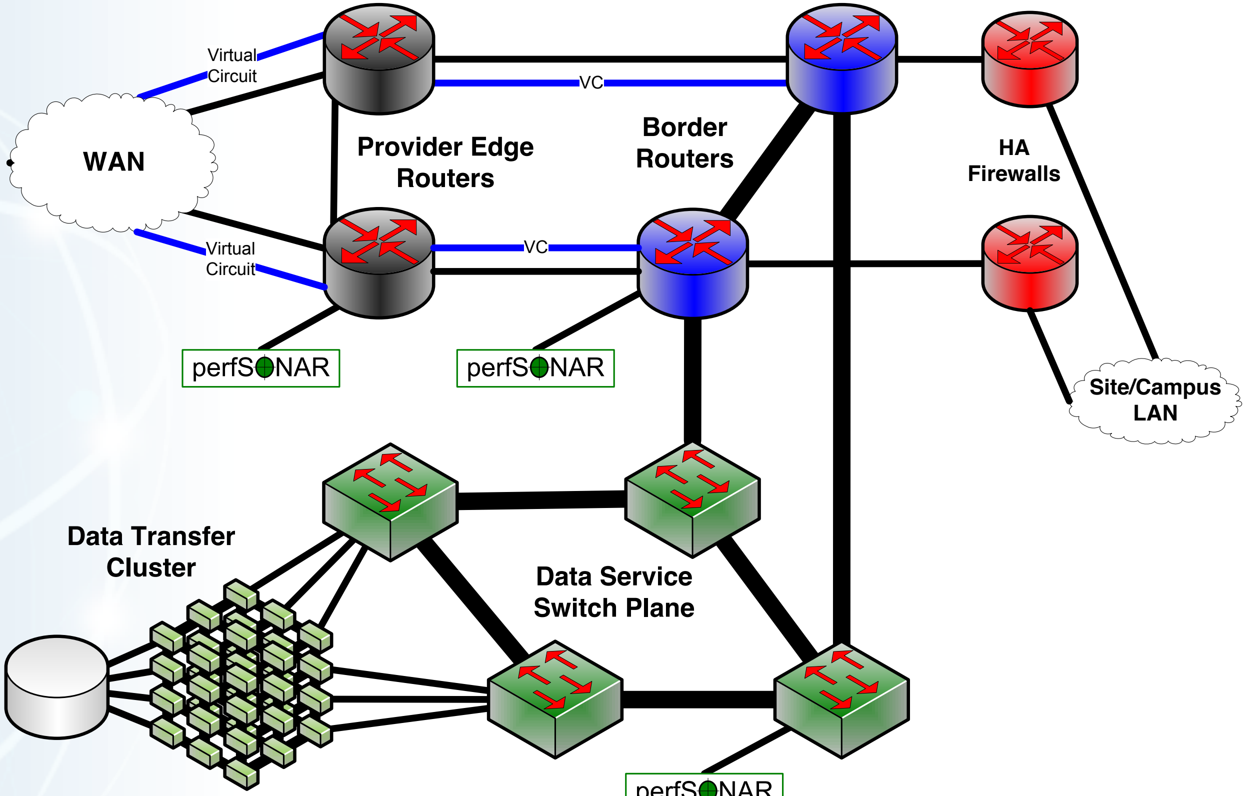

Data Site – Data Path

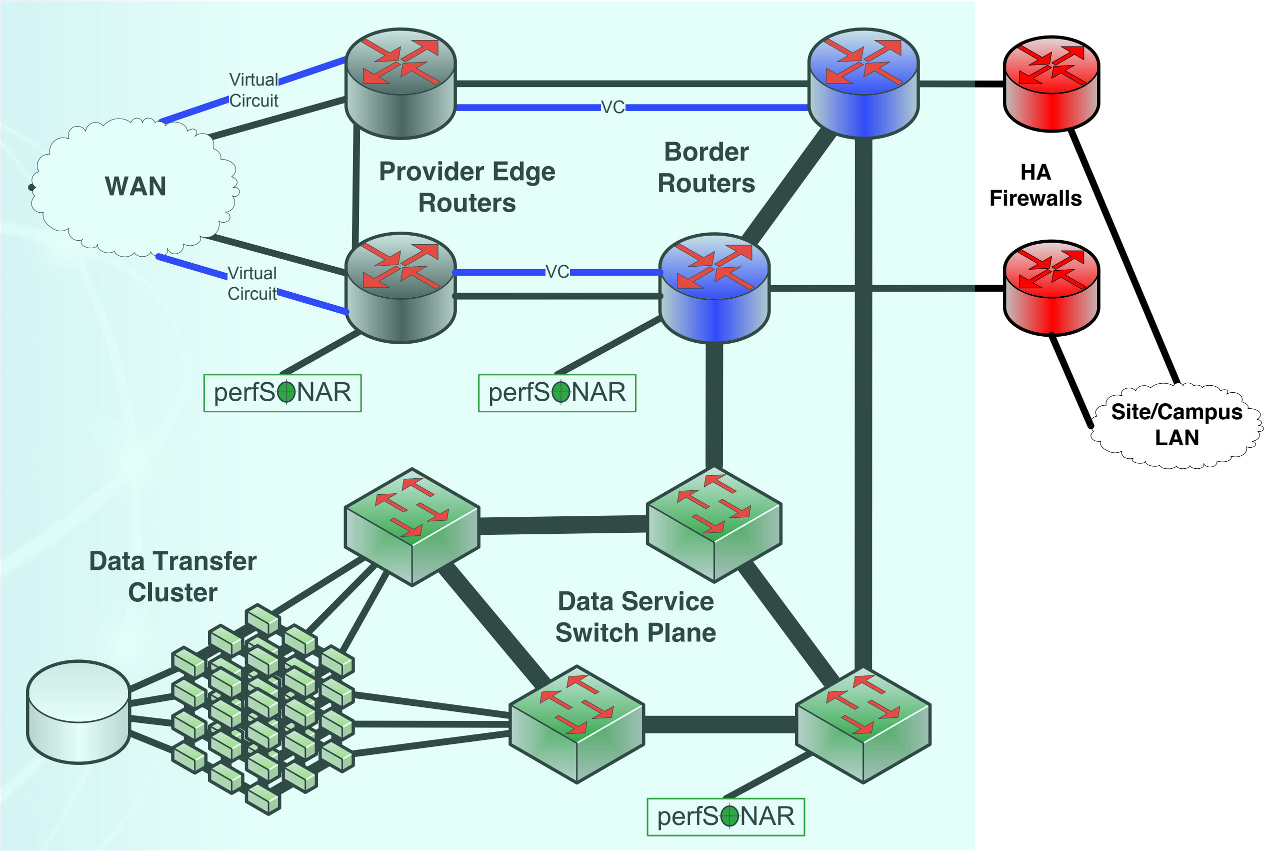

Distributed Science DMZ

Fiber-rich environment enables distributed Science DMZ

- No need to accommodate all equipment in one location
- Allows the deployment of institutional science service

WAN services arrive at the site in the normal way

Dark fiber distributes connectivity to Science DMZ services throughout the site

- Departments with their own networking groups can manage their own local Science DMZ infrastructure
- Facilities or buildings can be served without building up the business network to support those flows

Security is made more complex

- Remote infrastructure must be monitored
- Several technical remedies exist (arpwatch, no DHCP, separate address space, etc)
- Solutions depend on relationships with security groups

Distributed Science DMZ – Dark Fiber

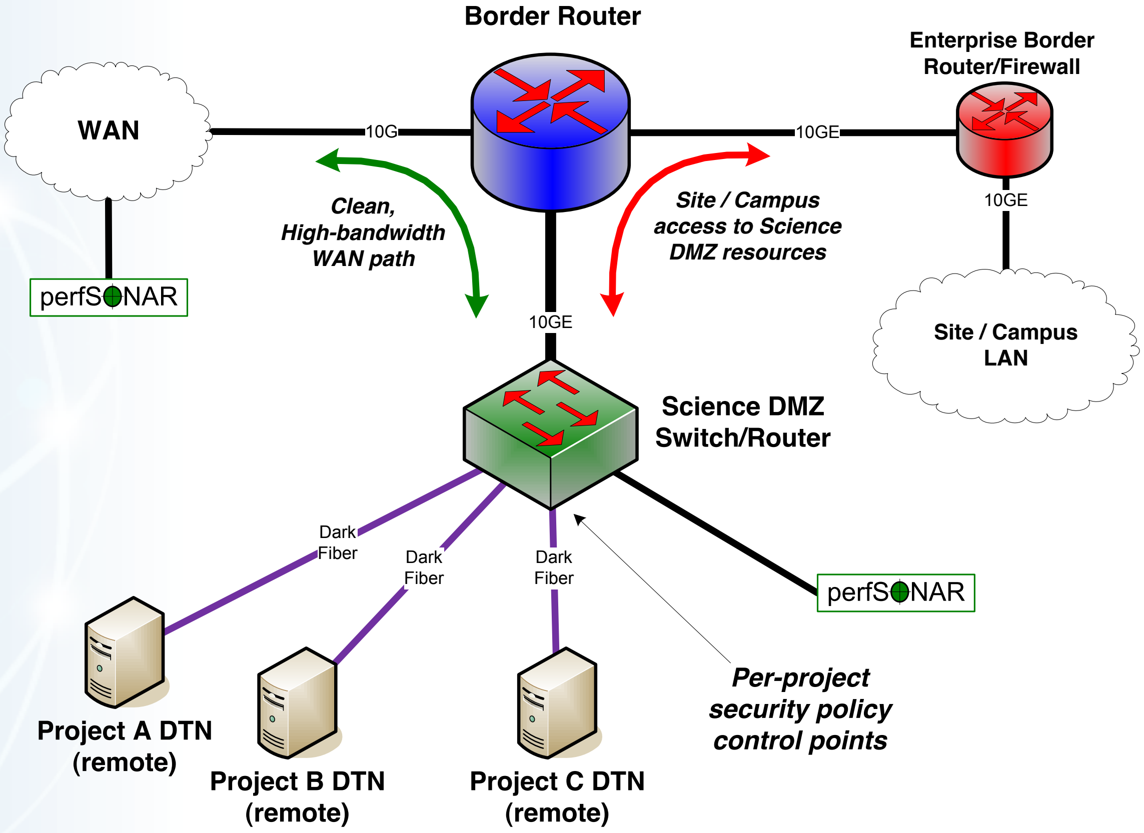

Border Router
Enterprise Border Router/Firewall
WAN
10G
10GE
10GE
perfSONAR
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
10GE
Site / Campus LAN
Science DMZ Switch/Router
Dark Fiber
Dark Fiber
Dark Fiber
perfSONAR
Per-project security policy control points
Project A DTN (remote)
Project B DTN (remote)
Project C DTN (remote)

Multiple Science DMZs – Dark Fiber

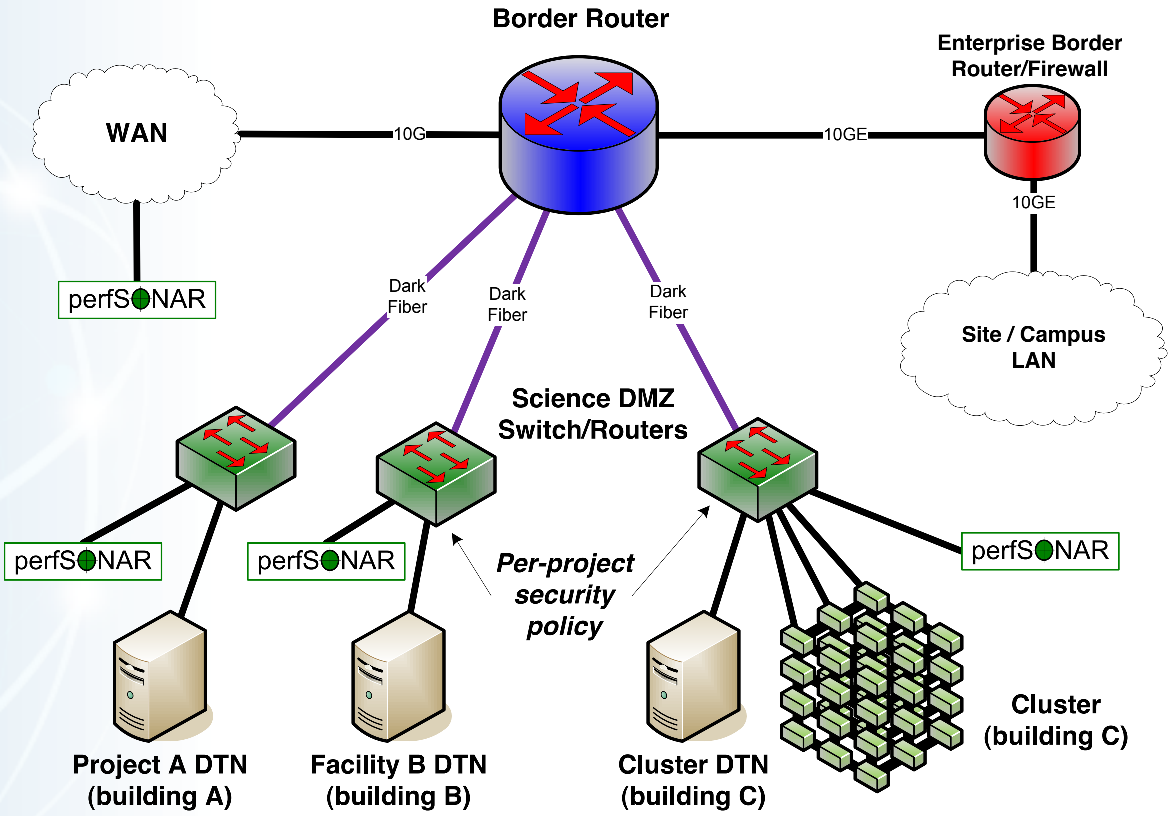

Common Threads

Two common threads exist in all these examples

Accommodation of TCP

- Wide area portion of data transfers traverses purpose-built path
- High performance devices that don’t drop packets

Ability to test and verify

- When problems arise (and they always will), they can be solved if the infrastructure is built correctly
- Small device count makes it easier to find issues
- Multiple test and measurement hosts provide multiple views of the data path
  - perfSONAR nodes at the site and in the WAN
  - perfSONAR nodes at the remote site

The Data Transfer Trifecta: The "Science DMZ" Model

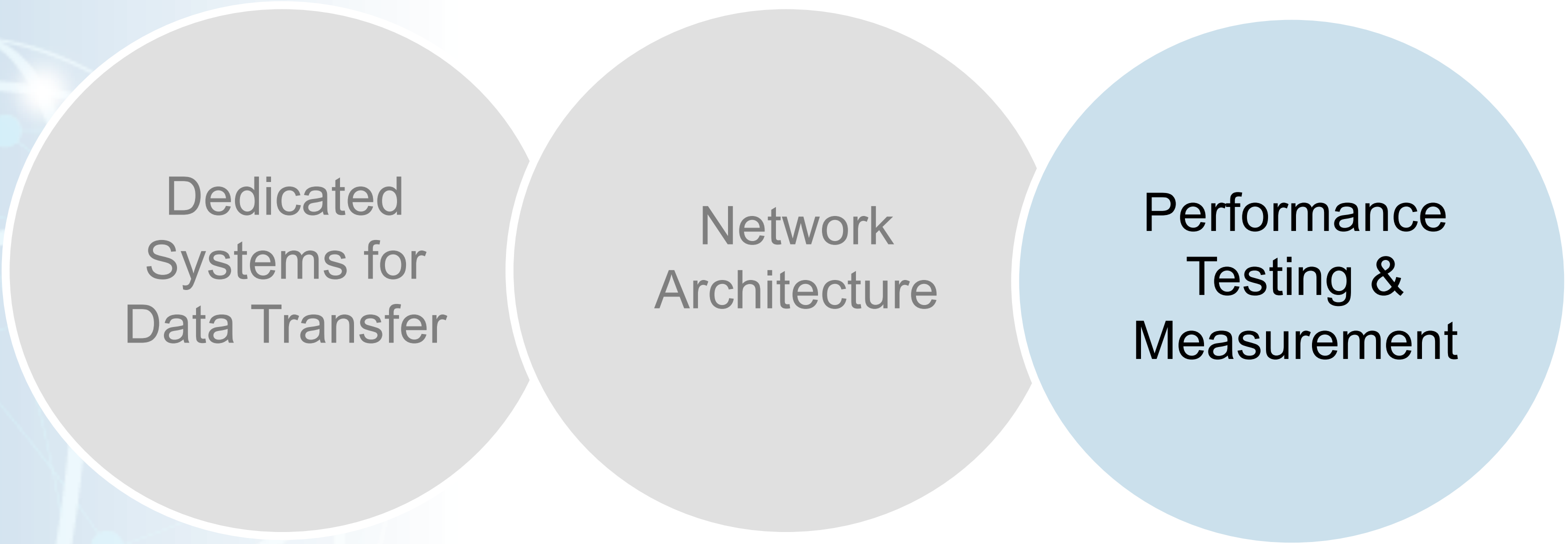

Dedicated Systems for Data Transfer

Network Architecture

Performance Testing & Measurement

Data Transfer Node

- High performance
- Configured for data transfer
- Proper tools

Science DMZ

- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network
- Additional info: http://fasterdata.es.net/

perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

Test and Measurement – Keeping the Network Clean

The wide area network, the Science DMZ, and all its systems can be functioning perfectly

Eventually something is going to break

- Networks and systems are built with many, many components
- Sometimes things just break – this is why we buy support contracts

Other problems arise as well – bugs, mistakes, whatever

We must be able to find and fix problems when they occur

Why is this so important? Because we use TCP!

Where are common problems?

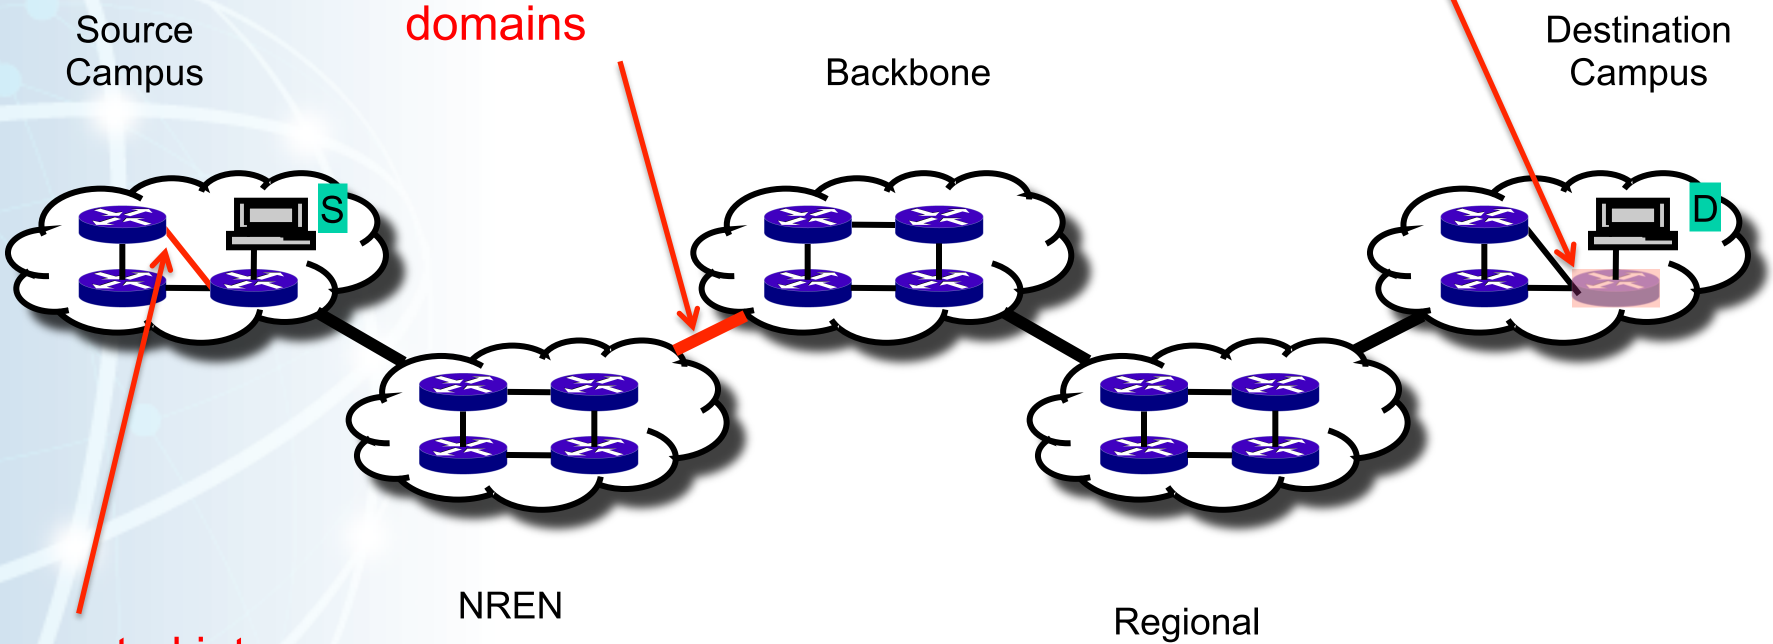

Congested or faulty links between domains

Latency dependant problems inside domains with small RTT

Congested intra-campus links

Local testing will not find all problems

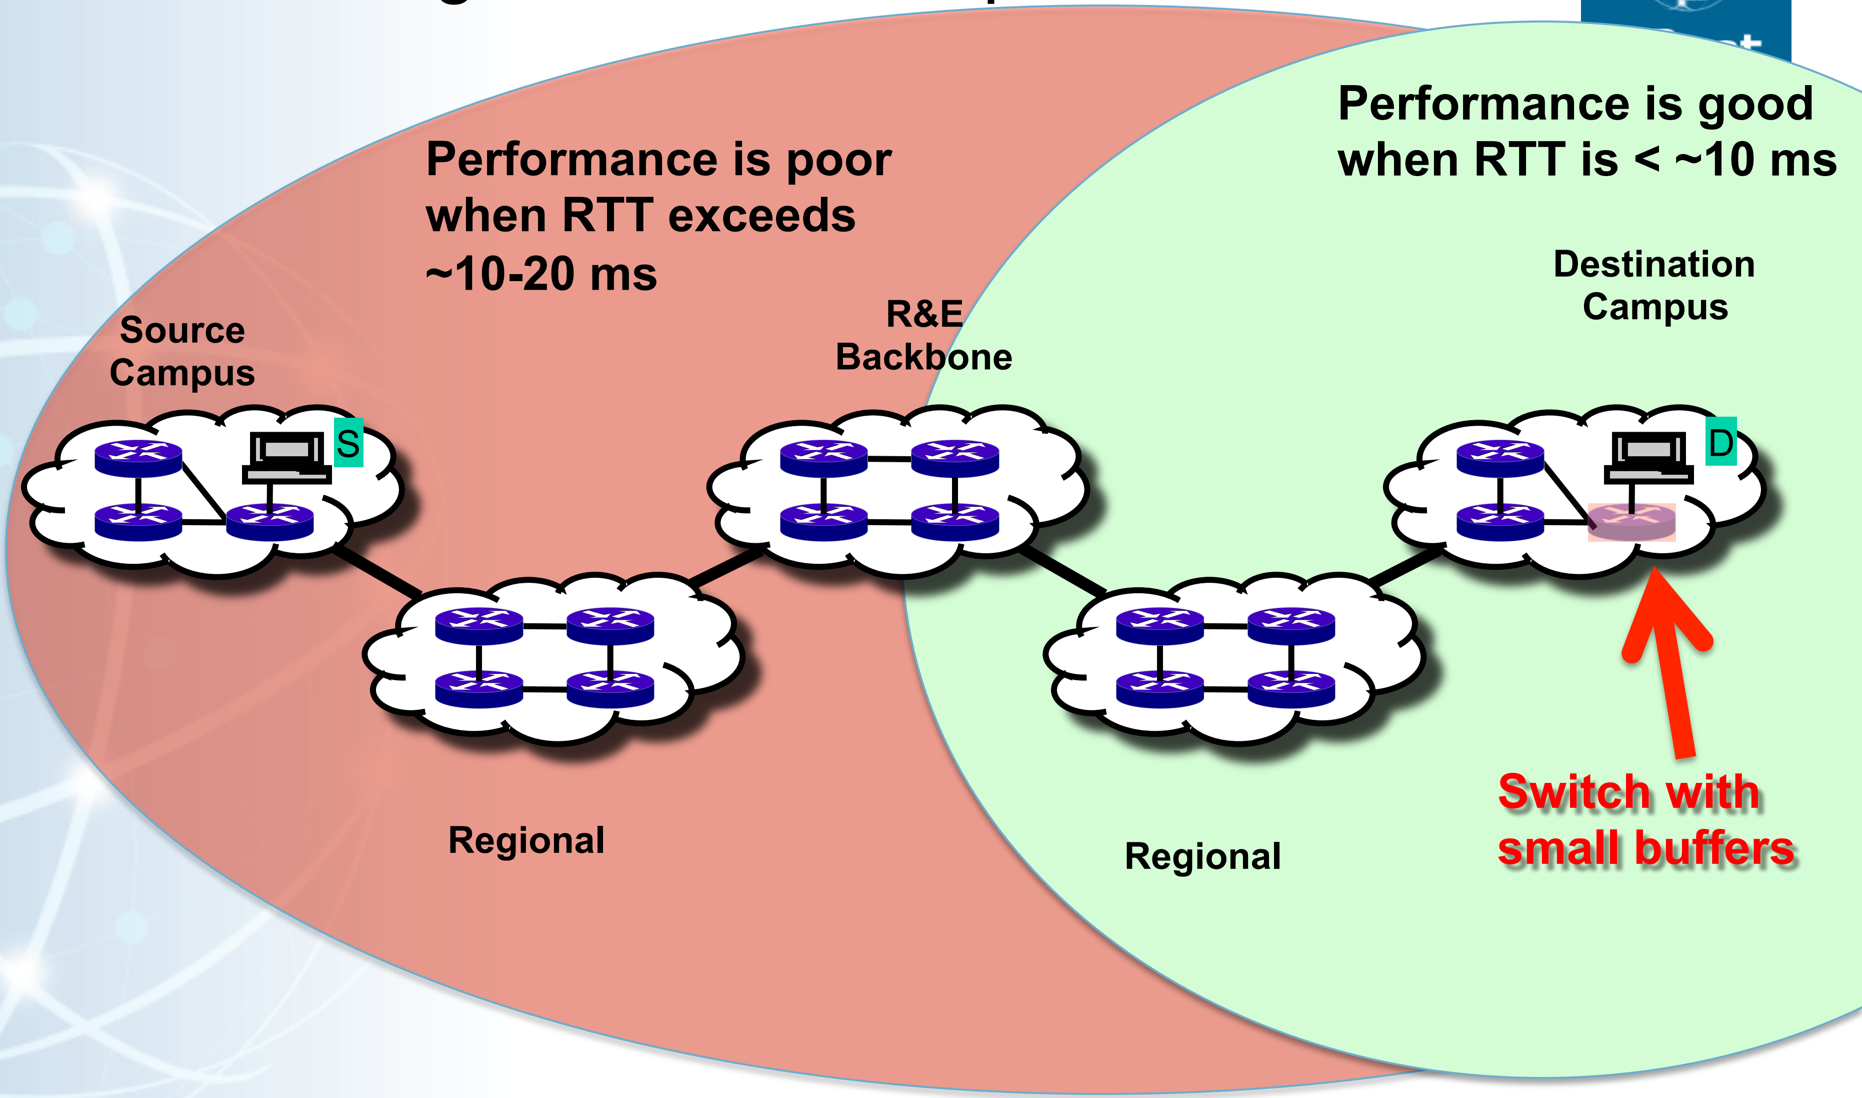
Performance is poor when RTT exceeds ~10-20 ms
Performance is good when RTT is < ~10 ms
Source Campus
S
R&E Backbone
Destination Campus
D
Regional
Regional
Switch with small buffers

Soft Network Failures

Soft failures are where basic connectivity functions, but high performance is not possible.

TCP was intentionally designed to hide all transmission errors from the user:

- “As long as the TCPs continue to function properly and the internet system does not become completely partitioned, no transmission errors will affect the users.” (From IEN 129, RFC 716)

Some soft failures only affect high bandwidth long RTT flows.

Hard failures are easy to detect & fix

- soft failures can lie hidden for years!

One network problem can often mask others

Common Soft Failures

Random Packet Loss

- Bad/dirty fibers or connectors
- Light levels too low or too high
- Duplex mismatch

Small Queue Tail Drop

- Switches not able to handle the long packet trains prevalent in long RTT sessions with cross traffic present

Un-intentional Rate Limiting

- Processor-based switching on routers due to faults, acl's, or mis-configuration
- Security Devices
  - E.g.: 10X improvement by turning off Cisco Reflexive ACL

Sample Results: Finding/Fixing soft failures

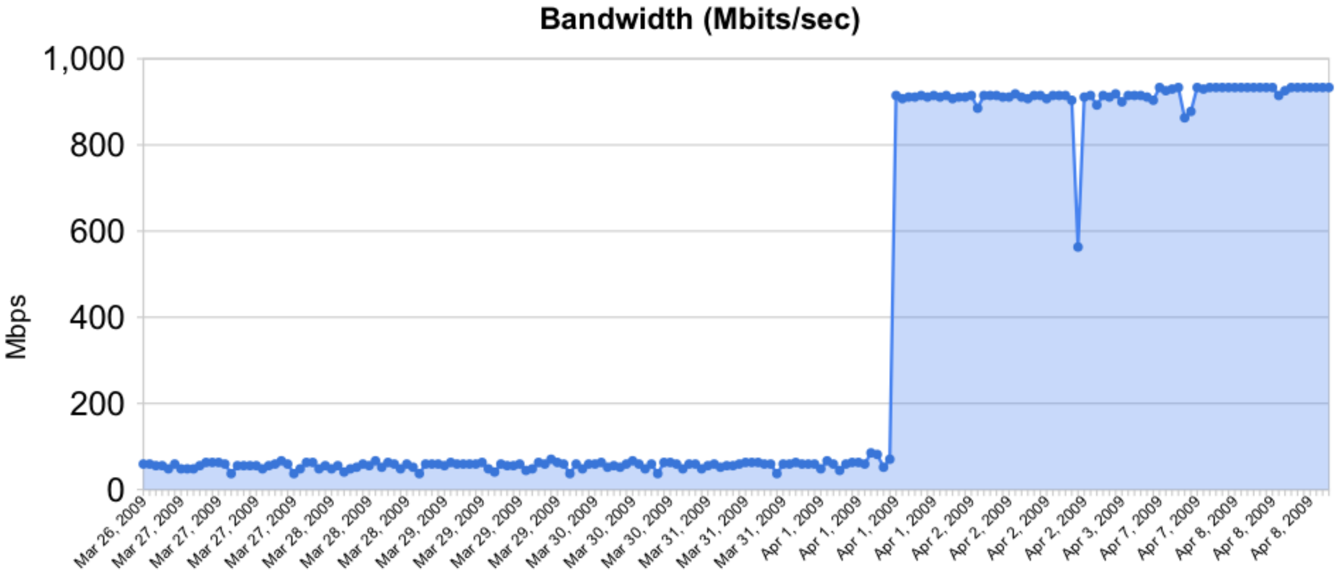

Rebooted router that was process switching after route table overload

Gradual failure of optical line card

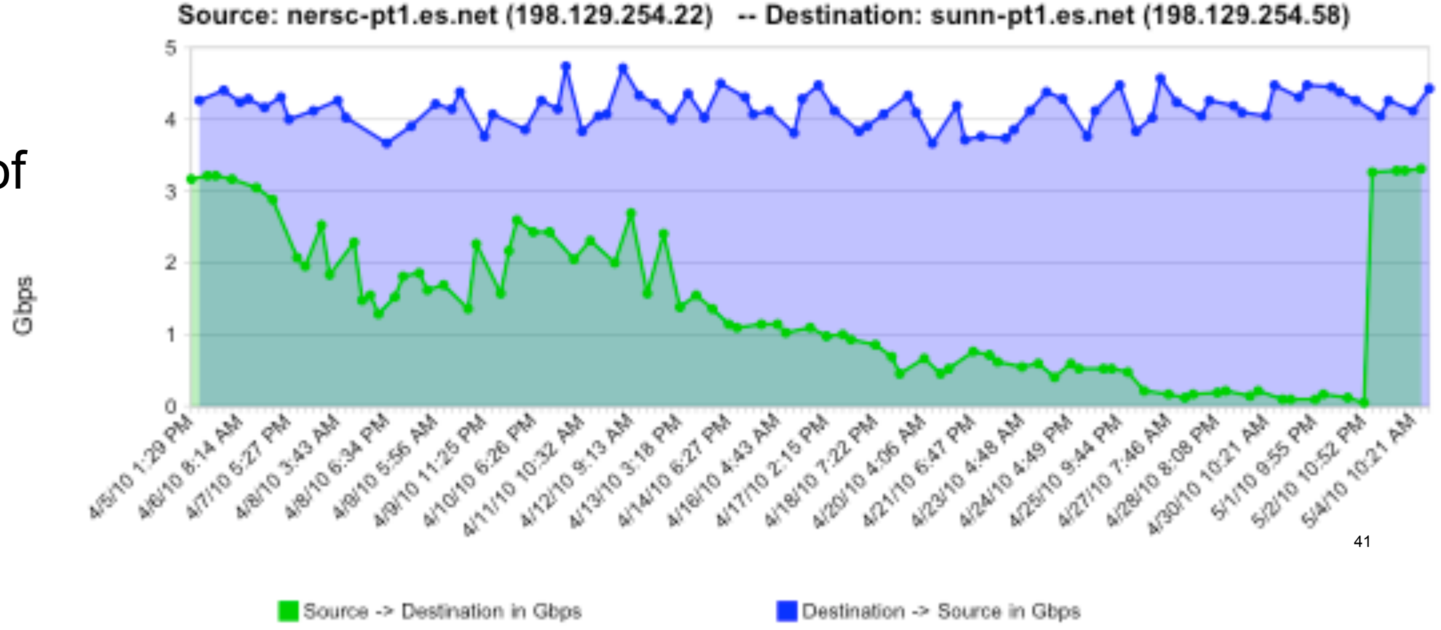

What is perfSONAR?

The perfSONAR Consortium is a joint collaboration between science networks

- ESnet
- GÉANT
- Internet2
- RNP

Decisions regarding protocol development, software branding, and interoperability are handled at this organization level

There are multiple independent efforts to develop software frameworks that are perfSONAR compatible.

- perfSONAR-MDM
- perfSONAR-PS

Each project works on an individual development roadmap and works with the consortium to further protocol development and ensure compatibility

Importance of Regular Testing

We can't wait for users to report problems and then fix them (soft failures can go unreported for years!)

Things just break sometimes

- Failing optics
- Somebody messed around in a patch panel and kinked a fiber
- Hardware goes bad

Problems that get fixed have a way of coming back

- System defaults come back after hardware/software upgrades
- New employees may not know why the previous employee set things up a certain way and back out fixes

Important to continually collect, archive, and alert on active throughput test results

The Data Transfer Trifecta: The “Science DMZ” Model

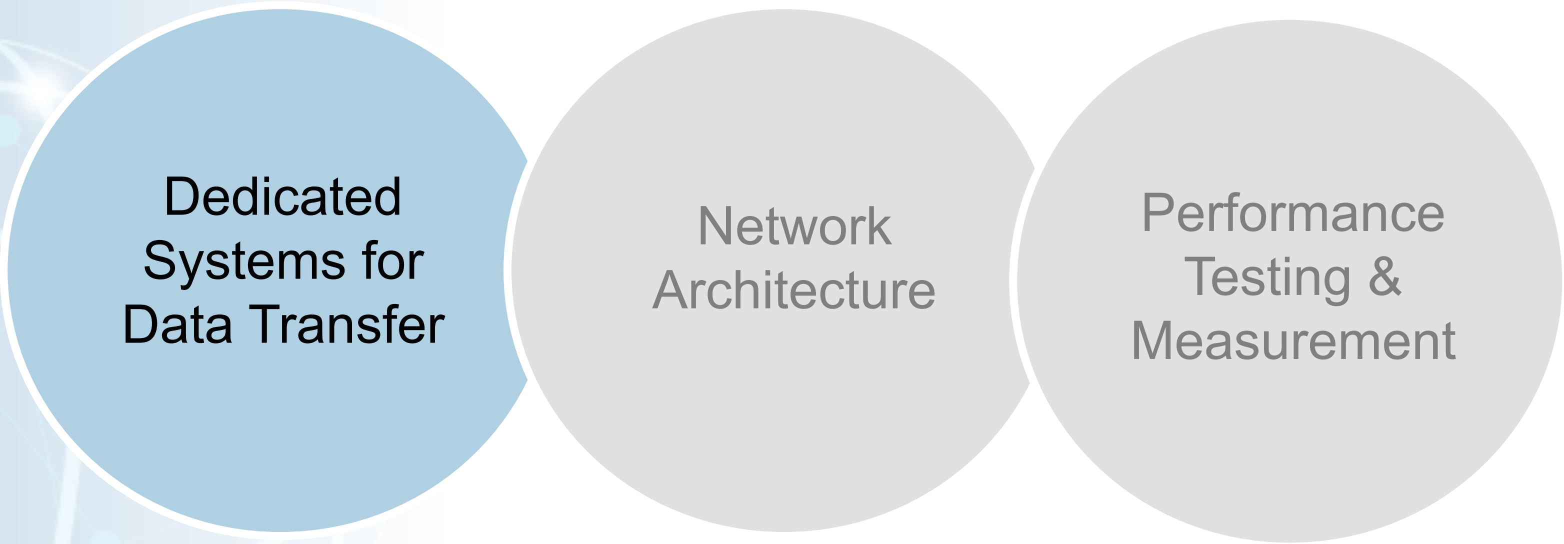

Dedicated Systems for Data Transfer

Network Architecture

Performance Testing & Measurement

Data Transfer Node

- High performance
- Configured for data transfer
- Proper tools

Science DMZ

- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network
- Additional info: http://fasterdata.es.net/

perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

The Data Transfer Node

A DTN server is made of several subsystems. Each needs to perform optimally for the DTN workflow:

Storage: capacity, performance, reliability, physical footprint

Networking: protocol support, optimization, reliability

Motherboard: I/O paths, PCIe subsystem, IPMI

Chassis: adequate power supply, extra cooling

Note: the workflow we are optimizing for here is sequential reads/ write of large files, and a moderate number of high bandwidth flows.

We assume this host is dedicated to data transfer, and not doing data analysis/manipulation

Tuning

Defaults are usually not appropriate for performance.

What needs to be tuned:

- BIOS
- Firmware
- Device Drivers
- Networking
- File System
- Application

- See: http://fasterdata.es.net/science-dmz/DTN/tuning/

Tuning Example: Network Tuning

```
# add to /etc/sysctl.conf
net.core.rmem_max = 33554432
net.core.wmem_max = 33554432
net.ipv4.tcp_rmem = 4096 87380 33554432
net.ipv4.tcp_wmem = 4096 65536 33554432
net.core.netdev_max_backlog = 250000
```

Add to /etc/rc.local

```
# increase txqueuelen
/sbin/ifconfig eth2 txqueuelen 10000
/sbin/ifconfig eth3 txqueuelen 10000
```

```
# make sure cubic and htcp are loaded
/sbin/modprobe tcp_htcp
/sbin/modprobe tcp_cubic
# set default to CC alg to htcp
net.ipv4.tcp_congestion_control=htcp
```

```
# with the Myricom 10G NIC increasing
    interrupt coalencing helps a lot:
/usr/sbin/ethtool -C ethN rx-usecs 75
```

And use Jumbo Frames!

Data Transfer Tools For DTNs

Parallelism is key

- It is much easier to achieve a given performance level with four parallel connections than one connection
- Several tools offer parallel transfers

Latency interaction is critical

- Wide area data transfers have much higher latency than LAN transfers
- Many tools and protocols assume a LAN
- Examples: SCP/SFTP, HPSS mover protocol

Tools: Sample Data Transfer Results

Using the right tool is very important

Sample Results: Berkeley, CA to Argonne, IL (near Chicago). RTT = 53 ms, network capacity = 10Gbps.

Tool	**Throughput**
scp:	140 Mbps
HPN patched scp:	1.2 Gbps
ftp	1.4 Gbps
GridFTP, 4 streams	5.4 Gbps
GridFTP, 8 streams	6.6 Gbps

Note that to get more than 1 Gbps (125 MB/s) disk to disk requires RAID.

Why Not Use SCP or SFTP?

Pros:

- Most scientific systems are accessed via OpenSSH
- SCP/SFTP are therefore installed by default
- Modern CPUs encrypt and decrypt well enough for small to medium scale transfers
- Credentials for system access and credentials for data transfer are the same

Cons:

- The protocol used by SCP/SFTP has a fundamental flaw that limits WAN performance
- CPU speed doesn’t matter – latency matters
- Fixed-size buffers reduce performance as latency increases
- It doesn’t matter how easy it is to use SCP and SFTP – they simply do not perform

Verdict: Do Not Use Without Performance Patches

A Fix For scp/sftp

- PSC has a patch set that fixes problems with SSH
  - http://www.psc.edu/networking/projects/hpn-ssh/
- Significant performance increase
- Advantage – this helps rsync too

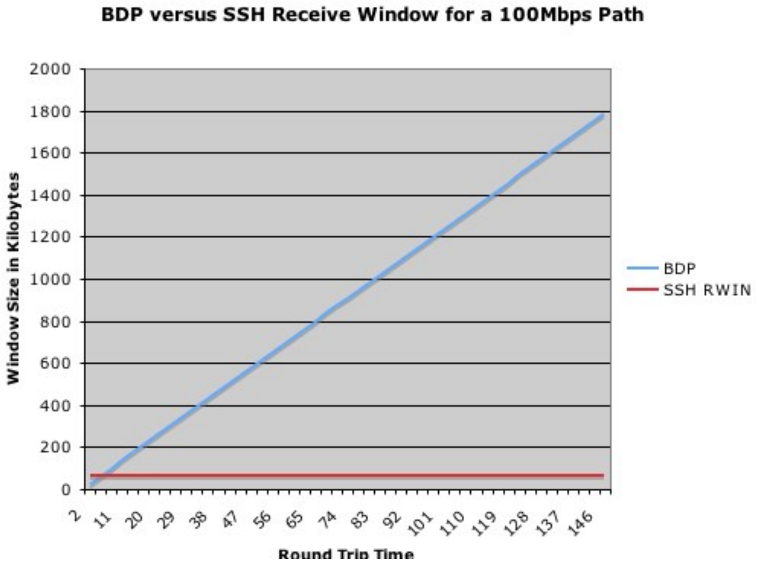

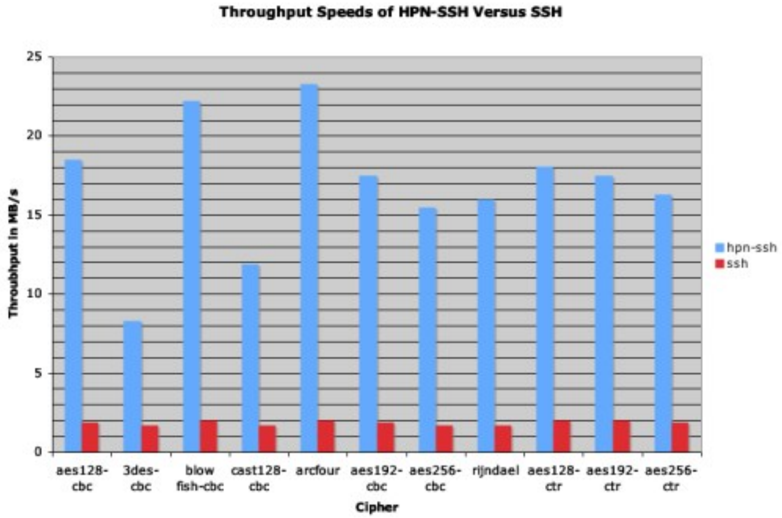

GridFTP

GridFTP from ANL has features needed to fill the network pipe

- Buffer Tuning
- Parallel Streams

Supports multiple authentication options

- Anonymous
- ssh
- X509

Ability to define a range of data ports

- helpful to get through firewalls

Globus Online / GridFTP and the Science DMZ

ESnet recommends Globus Online / GridFTP for data transfers to/from the Science DMZ

Key features needed by a Science DMZ

- High Performance: parallel streams, small file optimization
- Reliability: auto-restart, user-level checksum
- Multiple security models: ssh key, X509, Open ID, Shibboleth, etc.
- Firewall/NAT traversal support
- Easy to install and configure

Globus Online has all these features

Commercial Data Transfer Tools

There are several commercial UDP-based tools

- Aspera: http://www.asperasoft.com/
- Data Expedition: http://www.dataexpedition.com/
- TIXstream: http://www.tixeltec.com/tixstream_en.html

These should all do better than TCP on a congested, high-latency path

- advantage of these tools less clear on an uncongested path

They all have different, fairly complicated pricing models

http://fastdata.es.net

ESnet maintains a knowledge base of tips and tricks for obtaining maximum WAN throughput

Lots of useful stuff there, including:

- Network/TCP tuning information (in cut and paste-friendly form)
- Data Transfer Node (DTN) tuning information
- DTN reference designs
- Science DMZ information
- perfSONAR information

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

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

Science DMZ Benefits

Better access to remote facilities by local users

Local facilities provide better service to remote users

Ability to support science that might otherwise be impossible

Metcalf's Law – value increases as the square of connected devices

- Communication between institutions with functional Science DMZs is greatly facilitated
- Increased ability to collaborate in a data-intensive world

Cost/Effort benefits also

- Shorter time to fix performance problems – less staff effort
- Appropriate implementation of security policy – lower risk
- No need to drag high-speed flows across business network → lower IT infrastructure costs

Thanks!

Questions?

Eli Dart – dart@es.net

http://www.es.net/

http://fasterdata.es.net/