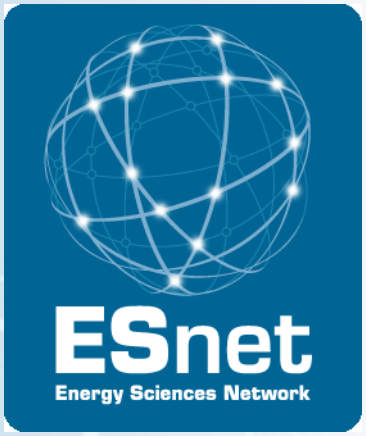

The Science DMZ

Eli Dart, Network Engineer

Joe Metzger, Network Engineer

ESnet Engineering Group

LHCOPN / LHCONE meeting

Internet2, Washington DC

June 13 2011

Overview

Science Needs

- Data Deluge, new science processes
- Barriers to scientific productivity, and their consequences

Science DMZ architecture

- Lower cost and increased productivity for data-intensive science
- Science DMZ ideas already proven in production

Security Concerns Addressed

Closing Points

Science Needs – the "Data Deluge"

"Data Deluge" is a nice cliché

- Dramatic and catchy
- It's been used many times before
- That's because it's a good way of describing the issue succinctly

Data volume is causing many disciplines to re-think their strategies, collaboration structures, etc. – Examples:

- Genomics
- Materials science (e.g. light source users)
- Medicine

Most disciplines do not have deep internal networking expertise – they need help to use networks well

New Science Processes

Data Deluge (data volume) is not the whole story

New Science processes rely on networking for success

- Remote instrument control
- Real-time analysis for experiment health-check, collaboration, etc.
- Increasing scope of collaboration
  - Fewer instruments, more collaborators → increased need for collaboration tools, data sharing, etc.
  - Networks are the fabric the holds modern science together

Baseline performance requirements for wide-area networking are increasing dramatically – scientific productivity in many disciplines is gated on this

Benefit of networks is often unrealized

The network serves many scientists poorly or not at all

- Without in-house (or in-collaboration) support, scientists are typically left to figure networking out for themselves
- The typical kit is composed of hosts with default configurations, and the SSH toolset
- If problems are encountered, the scientist is expected to find the right networking organization, explain the problem in networking terms, and help troubleshoot the problem
  - This model has demonstrably failed
  - Typical problem resolution time is on the scale of weeks
- Most scientists that have tried to use the network for data transfer or visualization have failed and have stopped trying – they “know” it can’t be done (this also means there are no trouble reports)

Current Problems – Technical Causes

Network device and host issues → poor performance

- Many networks still have packet loss (incorrect config, cheap hardware, poor design, etc)
- System defaults are wrong – hosts still need to be tuned

Wrong tool for the job → poor performance

- Use of SSH-based tools is common
- SSH has built-in, protocol-level performance limitations (10x to 50x slower than GridFTP on systems with high-performance storage)

Security blocks scientists at every turn

- Tools are blocked at the network layer or disallowed by policy
- Firewalls cause poor performance
- High-performance tools are often incompatible with system access technologies (e.g. SSH)

Consequences if these problems persist

Science will proceed with or without networks

Networks will decrease in relevance for most scientific disciplines and the institutions that support the science unless networks can be made more useful

- Lower return on investment in scientific infrastructure
- Longer time to discovery, loss of institutional leadership in key fields
- Reduction of technological and scientific leadership for USA

If the high-performance networks built for science are not useful – if the scientists can't use them effectively – then we have built a facility that is of no value to science

- Productivity/collaboration benefits from networking not realized
- Network-enabled modes of discovery unavailable
- For many disciplines, THIS IS THE CURRENT STATE

All Is Not Lost

Bad outcomes are not certain

There are successes – they just need to be generalized

Leadership by the networking community is required

- The old model of providing a toolkit and expecting scientists to learn networking has demonstrably failed
- The networking community must provide useful services and useful documentation for those services

Remember – *Most users are not networking experts, and it is unreasonable to expect them to become experts*

How to enable scientific use of networks?

It must be easy for scientists to use the network!

In many cases, the data-intensive part of a scientific experiment can run on one machine

Build a simple enclave for data-intensive services

- Near site perimeter
- Separate security policy
- No need to burden converged, multi-service network infrastructure with high-data-rate WAN flow requirements
  - Switches and routers with deep packet buffers are expensive
  - Debugging performance problems in converged networks is labor-intensive

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 Science DMZ

Science DMZ – a well-configured location for high-performance WAN-facing science services

- Located at or near site perimeter on dedicated infrastructure
- Dedicated, high-performance data movers
- Highly capable network devices (wire-speed, deep queues)
- Virtual circuit infrastructure
- perfSONAR

DYNES project is a specific example of this general architectural theme

- Dedicated infrastructure for data movers
- Virtual circuit termination
- Many high-bandwidth science sites have moved to this architecture already as a matter of necessity – DYNES is expanding on this because it works

Science DMZ – Conceptual Diagram

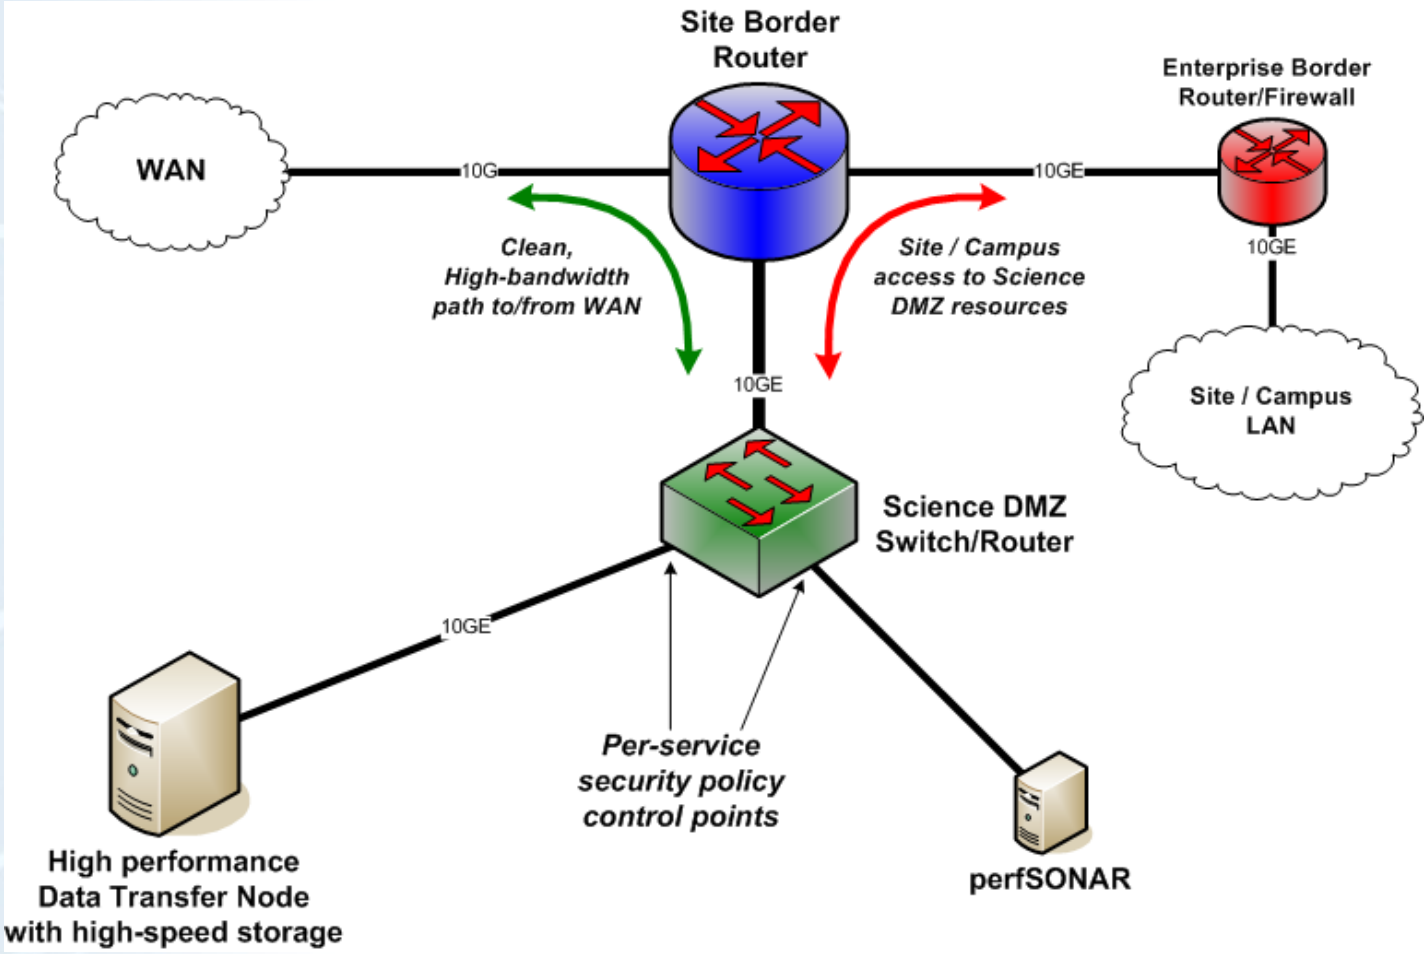

Science DMZ - Advanced

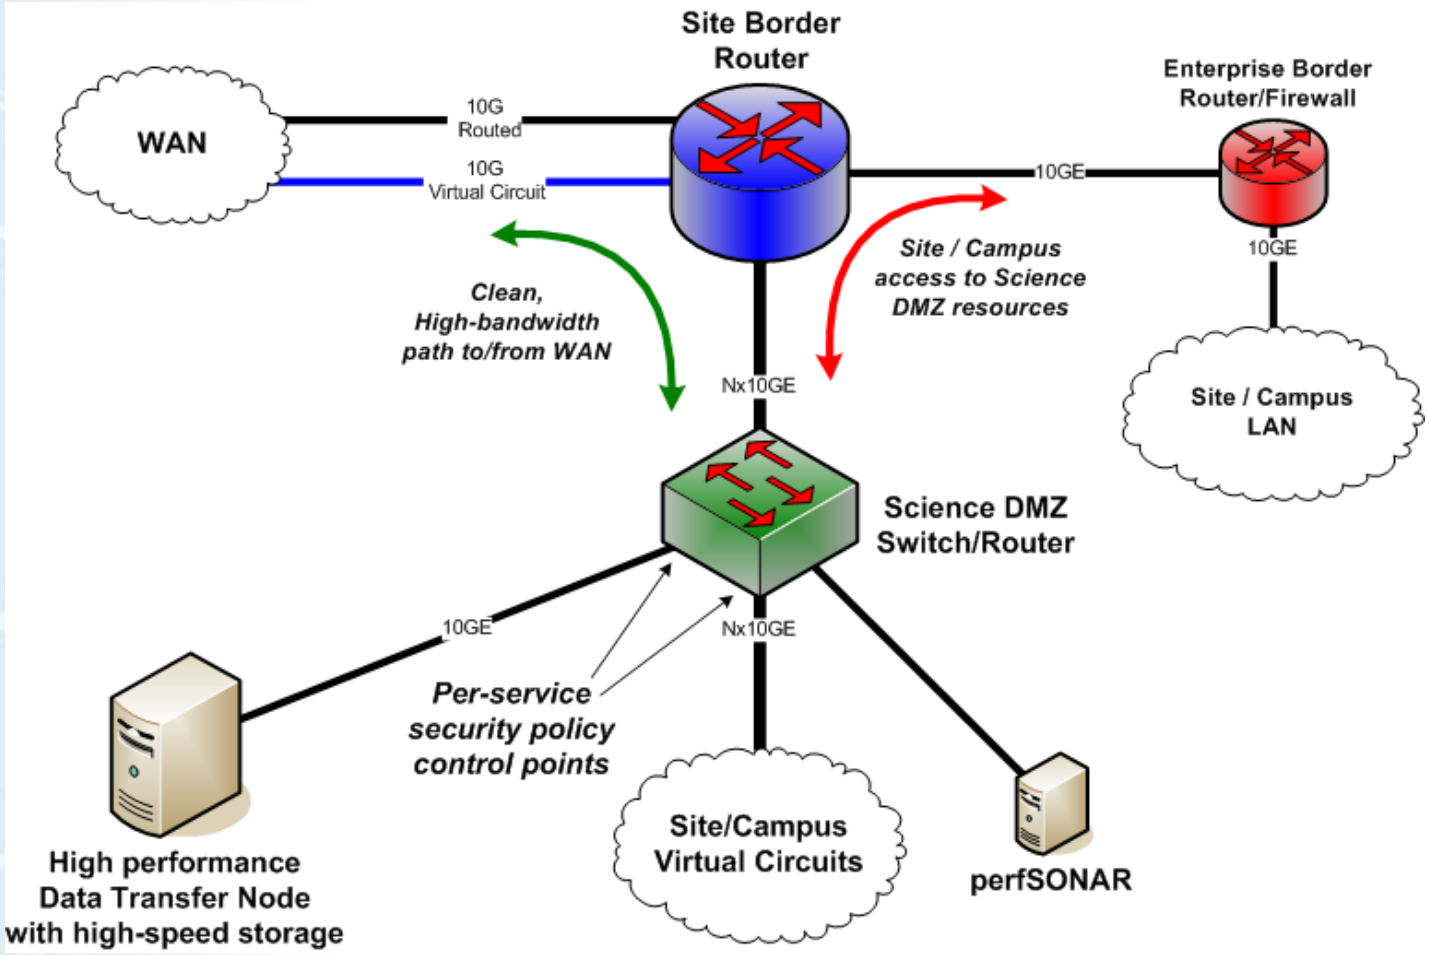

Science DMZ – Separate Circuit Connection

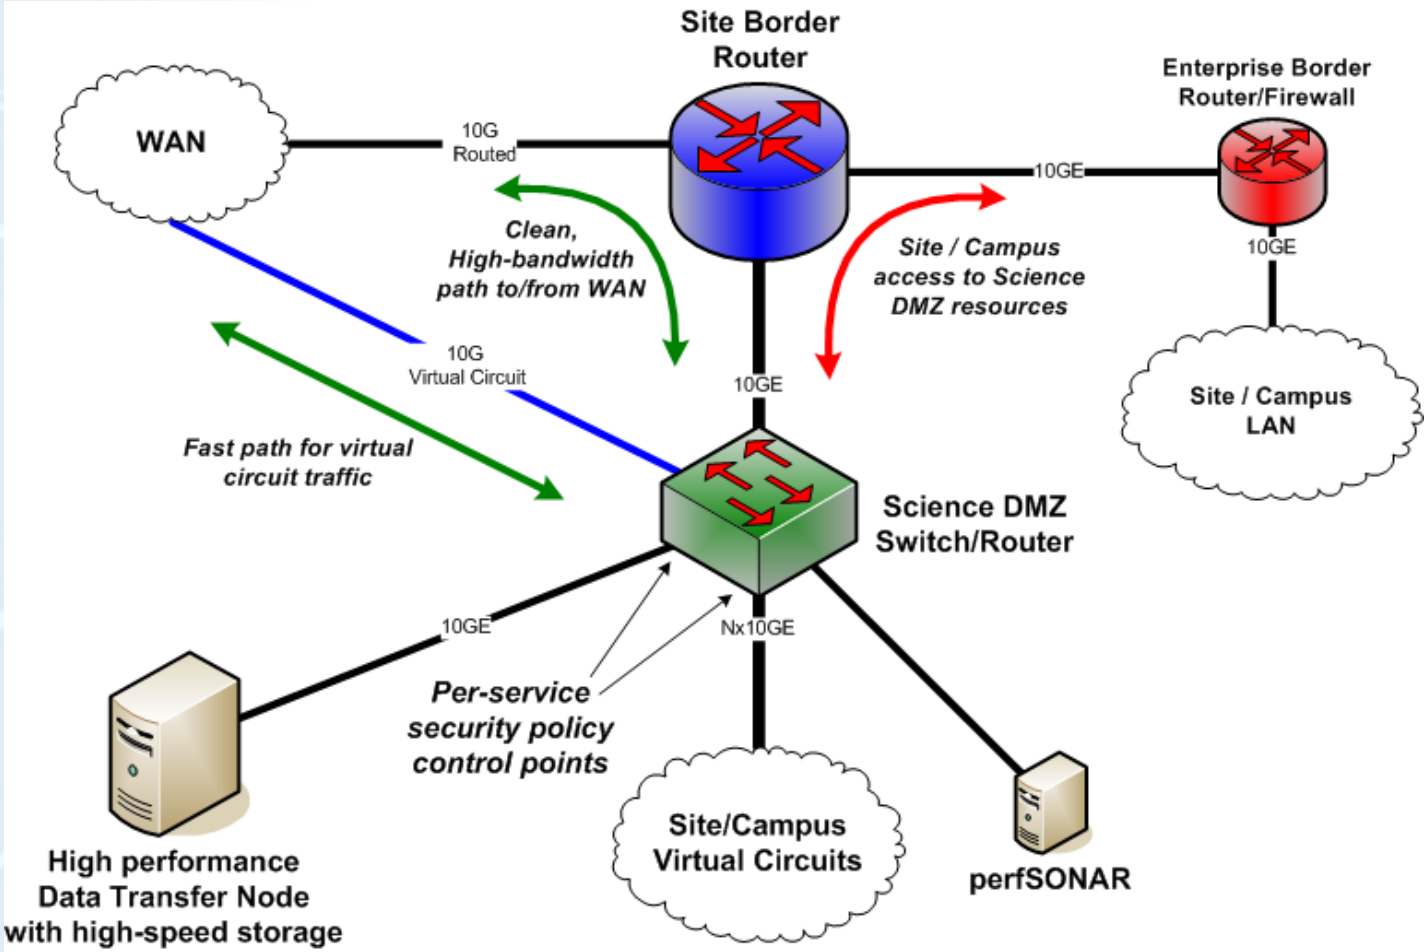

Science DMZ Features and Components

Direct connection to site perimeter

- LAN devices eliminated from the path
  - LAN infrastructure need not be sized for science flows (reduced hardware costs)
  - LAN infrastructure need not be configured to support science flows (e.g. deep output queues can conflict with VOIP requirements)
  - LAN infrastructure need not implement features necessary for virtual circuits (reduced costs, reduced complexity)
- Security policy for science data movers is not conflated with policy for business systems, wireless devices, printers, VOIP, etc

Science DMZ Features and Components

Dedicated infrastructure for science applications

- LAN devices are not part of the troubleshooting mix
- Dedicated devices are easier to configure properly and maintain
- Data Transfer Node for high-performance data movement

Test and measurement deployed on same network infrastructure as science resources

Note also that networks without a monolithic firewall need not add a monolithic firewall just because it's in the picture – it is still very beneficial to move large-scale data transfers close to the site perimeter

Note well – this is *not* a research project! It is based on hard-won experience

The Data Transfer Node (DTN)

Dedicated, high-performance host for long-distance data transfer

High performance disk, for example:

- High-speed local RAID
- Fibrechannel attachment to SAN, if available
- Lustre or GPFS filesystem mount (e.g. when deployed at supercomputer center)

High-speed network connection (1G or 10G)

- Connected to Science DMZ
- Separate security policy from business traffic

Multiple sites and facilities are deploying DTNs (Supercomputer centers, labs, experiments, etc)

Significant performance gains from DTN deployment

Science DMZ Security

Goal – disentangle security policy and enforcement for science flows from that of 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
- Sizing security infrastructure on business networks for large science flows is expensive

Science DMZ Security Implementation

Security policy for Science DMZ can typically be implemented with router filters/ACLs

- No need for an expensive firewall (and LAN firewall doesn't have to have WAN capabilities such as large buffers)
- Modern routers can filter packets at wire speed
- Routers do not rewrite packet headers and track state like firewalls do → entire classes of bugs/issues eliminated

Security policy for Science DMZ does not affect business systems (e.g. inbound ports for support of parallel data transfers)

- Desktops, printers, etc. not on Science DMZ
- Science DMZ systems are not subject to business security policies or policy enforcement devices
- This is sane – appropriate policies and controls are applied in both cases

Science DMZ Benefits

LAN infrastructure need not carry wide area science traffic

- Science traffic has different characteristics than business traffic
- Deep output queues, dedicated interfaces, etc. are expensive
- Accurate counters, per-filter counters, etc. are expensive

LAN transfers are much easier than WAN transfers

- Internal transfer of data from the site/campus LAN to/from the local Science DMZ will be much easier to debug, and is much more tolerant of the loss typically found in LANs
- LAN losses do not affect WAN transfers

Science DMZ Benefits

Separation of security policy

- Science DMZ security policy need not protect desktops, printers, etc
- Router filters do not have the problems firewalls have

Dedicated hosts

- Data Transfer Node is not some desktop, interactive node, cluster head-end, etc. with a bunch of other stuff running on it
- Fewer hosts to tune for WAN transfers
- Configuration is stable, CPU and I/O resources dedicated to moving data

Science DMZ scales with science need

- As science data intensity increases, Science DMZ can be upgraded without burdening entire LAN
- High-bandwidth instruments, connections to facilities, etc. can be added
- We know these needs are coming soon – it is time to prepare

Closing points

The data intensity of science is rapidly increasing – this is true of many disciplines (humanities to high energy physics)

Increased network utility is critical, both because of data volume and because of the scientific efficiencies made possible by new network services

Today, it is common for scientists to be unable to use networks to their full potential

The Science DMZ is an element of network architecture that allows science data flows and science data services to use the network to its full potential

The Science DMZ idea is cost-effective and proven

Network Performance Knowledge Base

http://fasterdata.es.net/

Host tuning information:

- http://fasterdata.es.net/fasterdata/host-tuning/

Data transfer tools (including SCP/SFTP issues):

- http://fasterdata.es.net/fasterdata/data-transfer-tools/

Data Transfer Node, including sample hardware config:

- http://fasterdata.es.net/fasterdata/data-transfer-node/

Questions?

Thanks!

Extra Slides

Other thoughts/options

- Network and IT security already has its advocates
  - security officers and engineers
  - firewalls and VPNs
- What about the *usability* of the network?
  - Joe St. Sauver proposes a Network Usability Officer
    - Advocate for the network's users, who ensures that the network is meeting users' needs.
    - In the science context, this means identifying the needs of science—both data-driven and interactive—and making the network meet those needs…
    - …*and* ensuring that scientists understand that the network *can be made to meet their needs.*
  - Marc Wallman: Firewalls less useful in a collaborative arena.

Example – Light Source Users

Current model – scientists take data at light source, data sets are relatively small (scale of ~1TB) – use USB hard drives today

Changes – data rates

- Next generation of instruments produce 10x to 100x data
- Instruments with data rates of 250MB/sec being deployed
- Physical transport on portable media will shortly become unworkable for many collaborations that have relied on this method for many years

Changes – science process

- Experiment automation bringing added efficiency, but requires experiment health checks
- Near-real-time data analysis needed
- Beamline computational resources inadequate – need supercomputer access
- Reliable, high speed data transfer to supercomputer centers necessary

What is Important?

Ease of use by non-expert users

- Scientists are scientists – most are not network experts
- From the perspective of users, “the network” is the aggregate of applications, computers, storage systems, campus networks, regional networks, long-haul networks that participate in WAN data movement and similar tasks.
- In order for “the network” to be a useful tool, this complex assembly of components must somehow present a sane interface to non-experts

Zero packet loss between data movers

- High performance is impossible in the presence of packet loss
- This is different from commodity networking

Test and measurement for troubleshooting and repair

- Even though it works today, it will break tomorrow or the next day
- This is infrastructure – provisions must be made for maintenance and repair

Accomplishing the Important Things

Ease of use – reduce complexity

- Reduce number of devices in the path (fewer things to troubleshoot, configure, etc – more on this in a minute)
- Dedicated infrastructure for data movement

Zero packet loss

- Again, reduce device count
- Use appropriate network devices (e.g. with deep output queues) – this means eliminating LAN devices from WAN data path

Test and measurement

- Need well-defined location for well-configured test and measurement gear (e.g. perfSONAR)
- Locate test and measurement devices near data movers

Security – decouple science and business security policy and control points

Modern Networks are Complex

Security issues (firewalls, network access control, HR data, PII, etc.)

VOIP

- New phone systems are typically VOIP
- VOIP deployments often require network rebuilds

Wireless networking

- Access/security
- Service quality (smartphones → device count doubles!)

Data Center issues (space, power, cooling, etc)

Who has recently launched major initiatives in these areas?

Data Intensive Science

Many science disciplines rely on data analysis

- Materials science
- Biomedicine
- Genomics
- High Energy Physics
- Climate

Data transfer and data sharing are critical to scientific collaborations – in fact, scientific productivity is often determined by the ability to transfer/stream/share data

Who has launched a major effort to increase the productive use of the network by scientists?