EPOC
Engagement and Performance Operations Center

Science DMZ Overview: Practical Designs and Use Cases

Jason Zurawski
zurawski@es.net
ESnet / Lawrence Berkeley National Laboratory

Empowering Secure Data-Driven Research
January 14, 2026

TACC TEXAS ADVANCED COMPUTING CENTER
ESnet ENERGY SCIENCES NETWORK

Outline

- ***Introduction***
- Solution Space
- Conclusions / QA

EPOC
Engagement and Performance Operations Center

Motivation/Background

- Networks are an essential part of data-intensive science
  - Connect data sources to data analysis
  - Connect collaborators to each other
- Performance is critical, but **often** overlooked
  - Exponential data growth
  - Constant human factors
  - Data movement and data analysis must keep up
- Effective use of wide area (long-haul) networks by scientists has historically been difficult

EPOC
Engagement and Performance Operations Center

Network as ~~Infrastructure~~ *Instrument*

Connectivity is the first step – ***usability*** must follow

EPOC
Engagement and Performance Operations Center

Outline

- Introduction
- ***Solution Space***
  1. ***Understanding the Solution Space (Users, Use Cases, Long Term Impacts)***
  2. Preliminaries (e.g. Network Protocols 101)
  3. Architecture & Design
  4. Data Mobility
- Conclusions / QA

EPOC
Engagement and Performance Operations Center

Common Theme / New Mindset

- We aren't building a "Network Architecture", we want a "Data Architecture"
  - A lot of the items that will be thrown at you transcend the traditional network space.
- To get there:
  - Understand the data pipeline for your target user/use case – cradle to retirement home
  - This implies all the things:
    - Creation
    - Usage
    - Transfer/Share
    - Curation

EPOC
Engagement and Performance Operations Center

Common Theme / New Mindset

- What you build must be
  - ***Usable*** – if this becomes a ‘walled garden’, what’s the point? Make it such that people can be easily onboarded and integrated.
  - ***Defensible*** – it is not, nor should it be, the wild west. Control the users and use cases, but don’t impact the usage.
  - ***Scalable*** – as demand grows. Think cornfields and baseball diamonds.
  - ***an institutional capability / source of pride*** – this is something that will draw more users / research dollars if created/marketed/operated correctly. Treat it as such.

EPOC
Engagement and Performance Operations Center

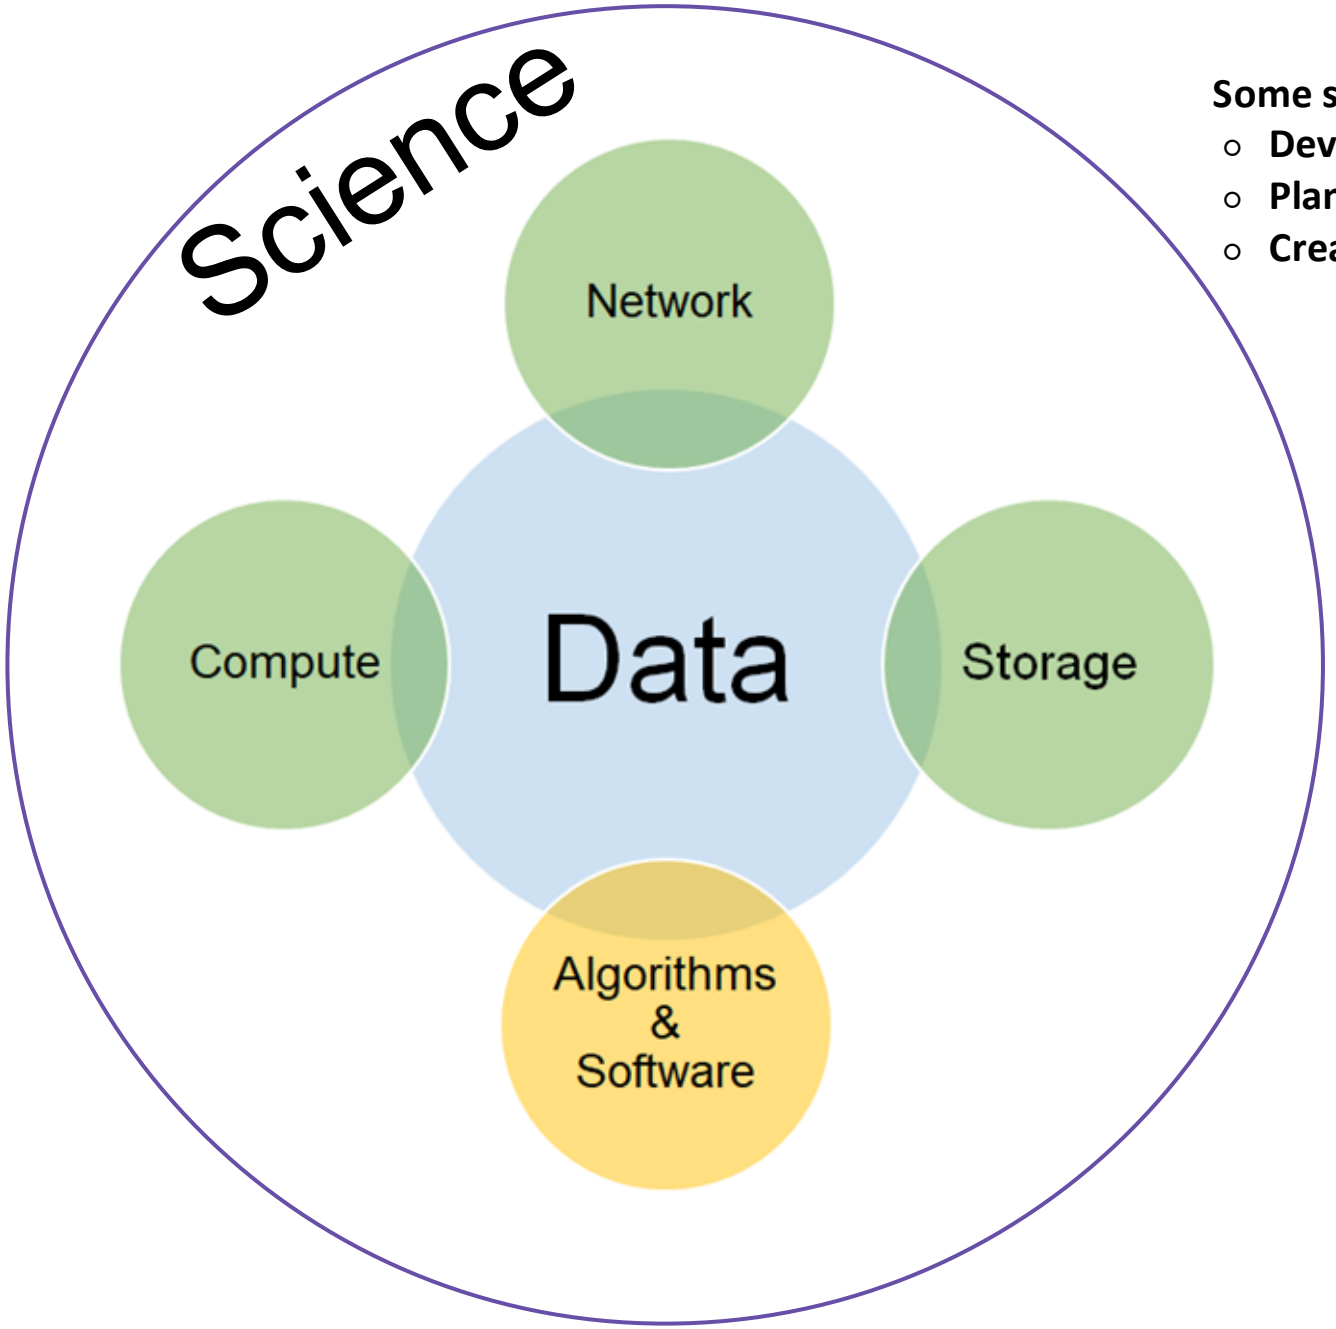

Some specific issues for networks are

- **Development of services**
- **Planning capacity growth**
- **Creation of collaborations**

EPOC
Engagement and Performance Operations Center

Outline

- Introduction
- ***Solution Space***
  1. Understanding the Solution Space (Users, Use Cases, Long Term Impacts)
  2. ***Preliminaries (e.g. Network Protocols 101)***
  3. Architecture & Design
  4. Data Mobility
- Conclusions / QA

EPOC
Engagement and Performance Operations Center

Data Movement / TCP 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, FTP)
  - Wait-for-trouble-ticket operational models for network performance

EPOC
Engagement and Performance Operations Center

TCP – Ubiquitous and Fragile

- 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
  - Communication is between applications – mostly over TCP
  - ***Congestion*** dictates performance – back off when danger is sensed to preserve/protect resources
- 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 (more than 95% of ESnet traffic is TCP)

EPOC
Engagement and Performance Operations Center

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

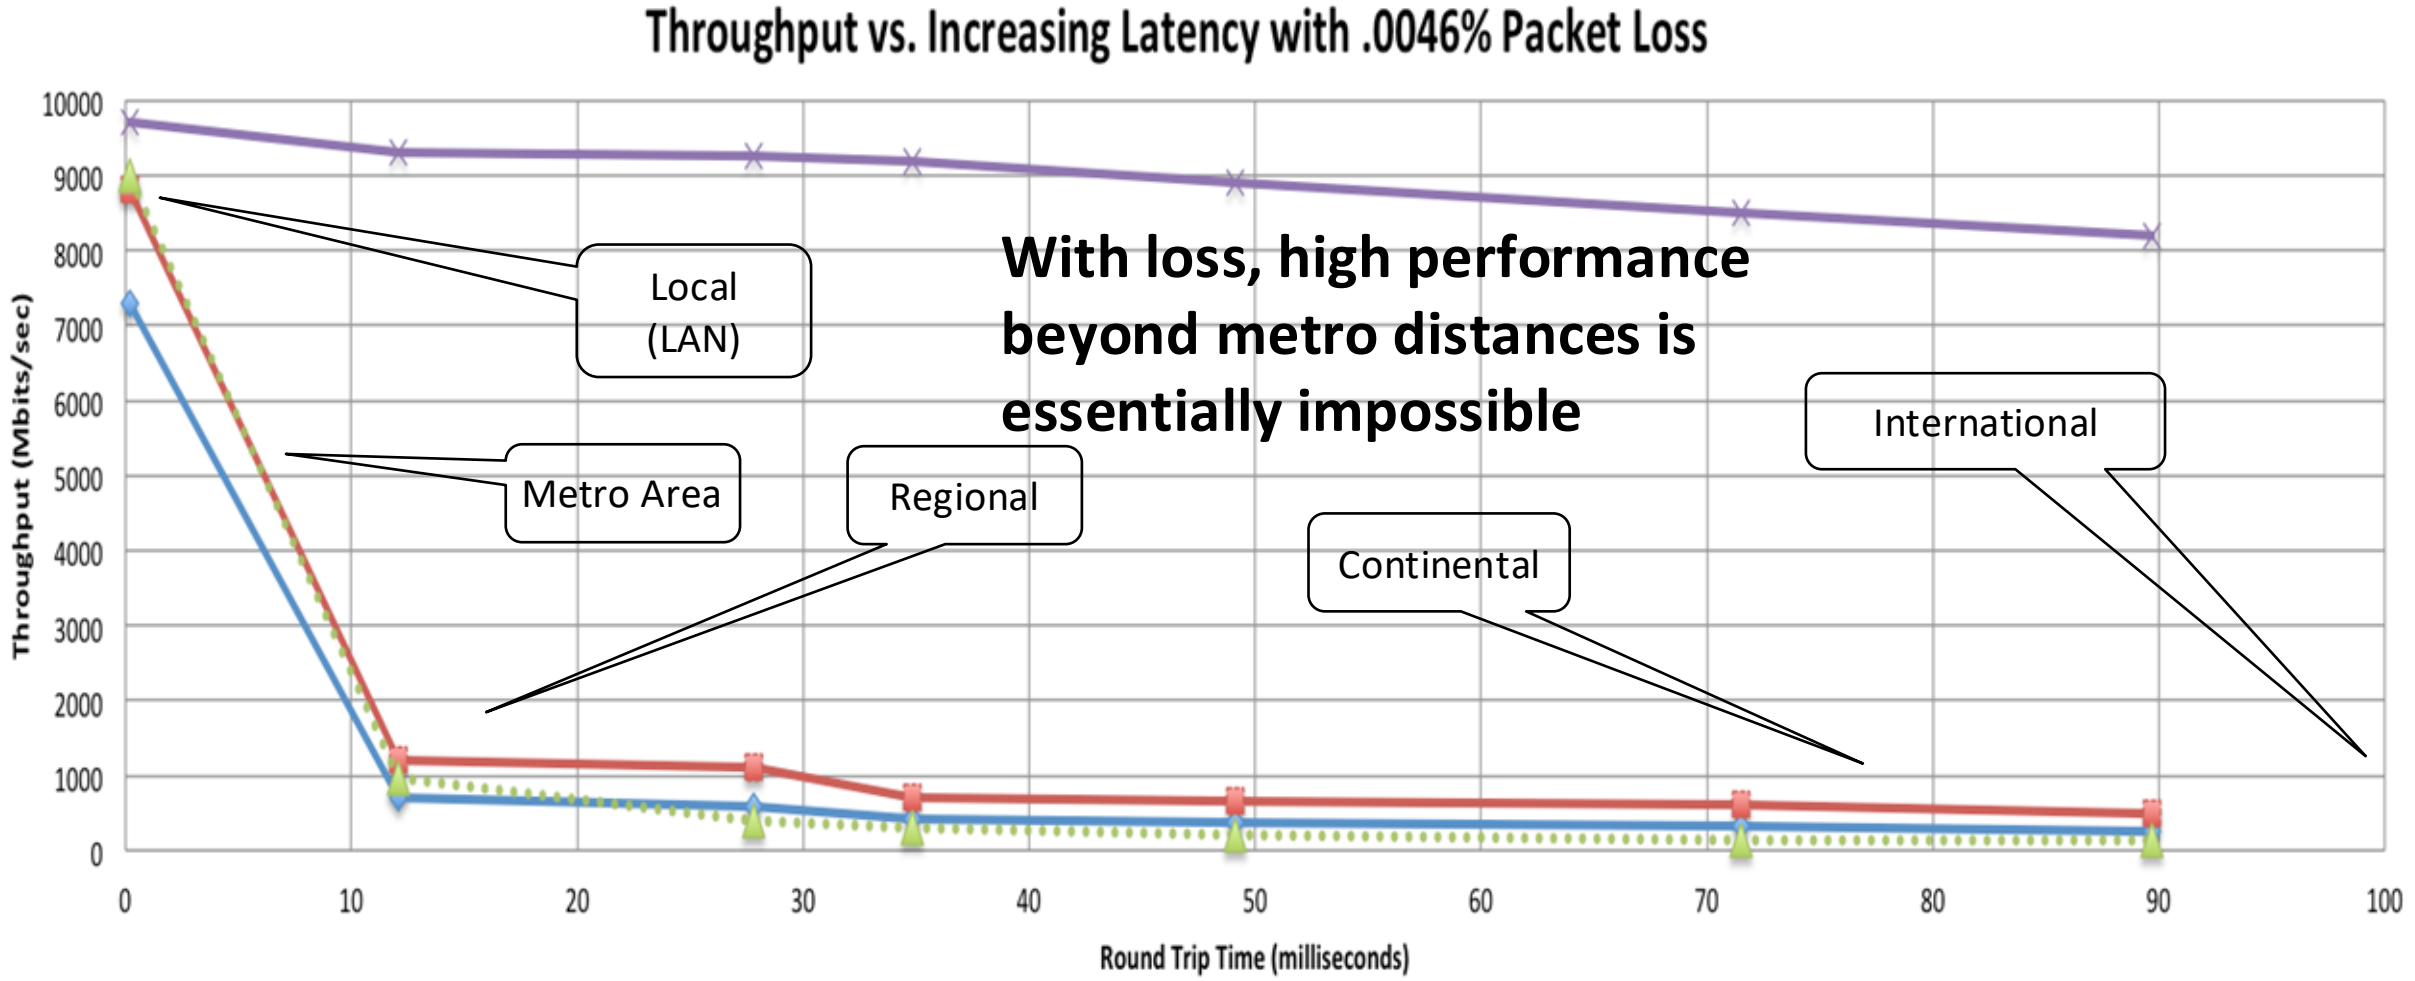

Data Movement / TCP Background

- 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

EPOC
Engagement and Performance Operations Center

The

Consi

- "Fri
  -
  -
  -
  -
- Dec
  -
  -
- Per
  -
- Eng

Outline

- Introduction
- ***Solution Space***
  1. Understanding the Solution Space (Users, Use Cases, Long Term Impacts)
  2. Preliminaries (e.g. Network Protocols 101)
  3. ***Architecture & Design***
  4. Data Mobility
- Conclusions / QA

EPOC
Engagement and Performance Operations Center

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: eliminate sources of packet loss / packet friction

EPOC
Engagement and Performance Operations Center

Legacy Method: 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 lucky, 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
- This usually fails to provide the necessary performance.

EPOC
Engagement and Performance Operations Center

Ad Hoc DTN Deployment

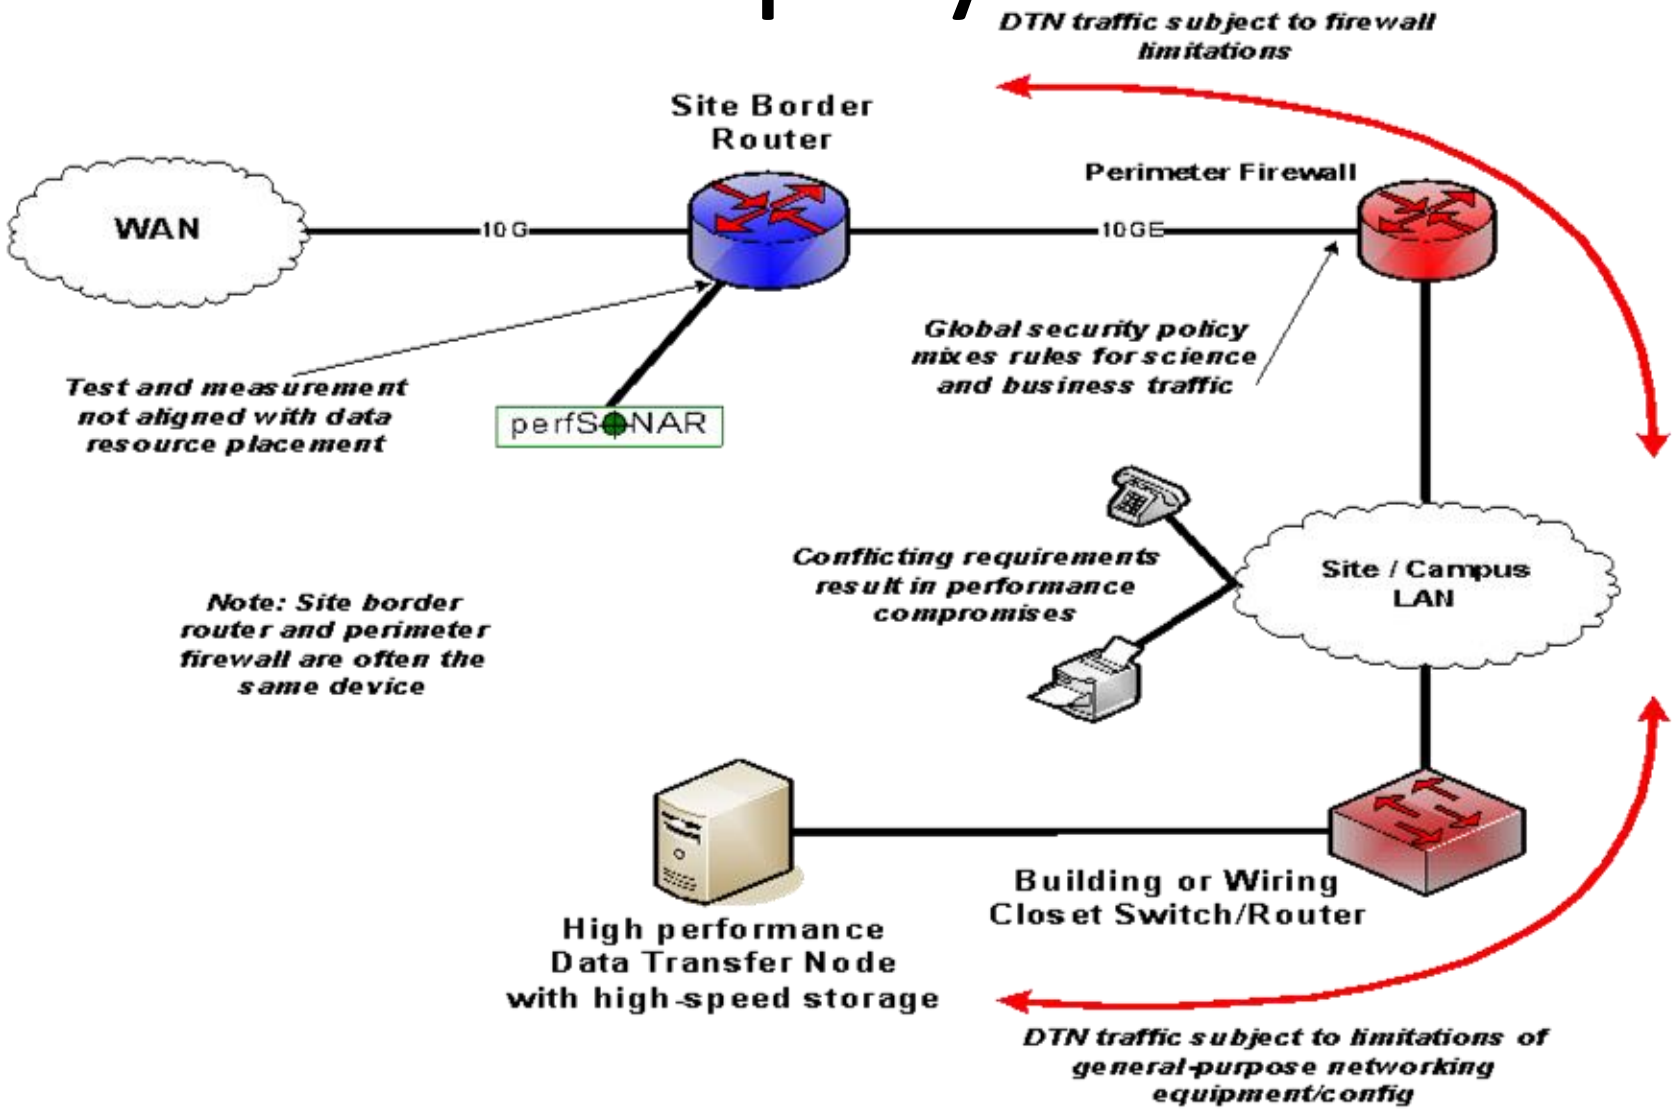

EPOC
Engagement and Performance Operations Center

A better approach: simple Science DMZ

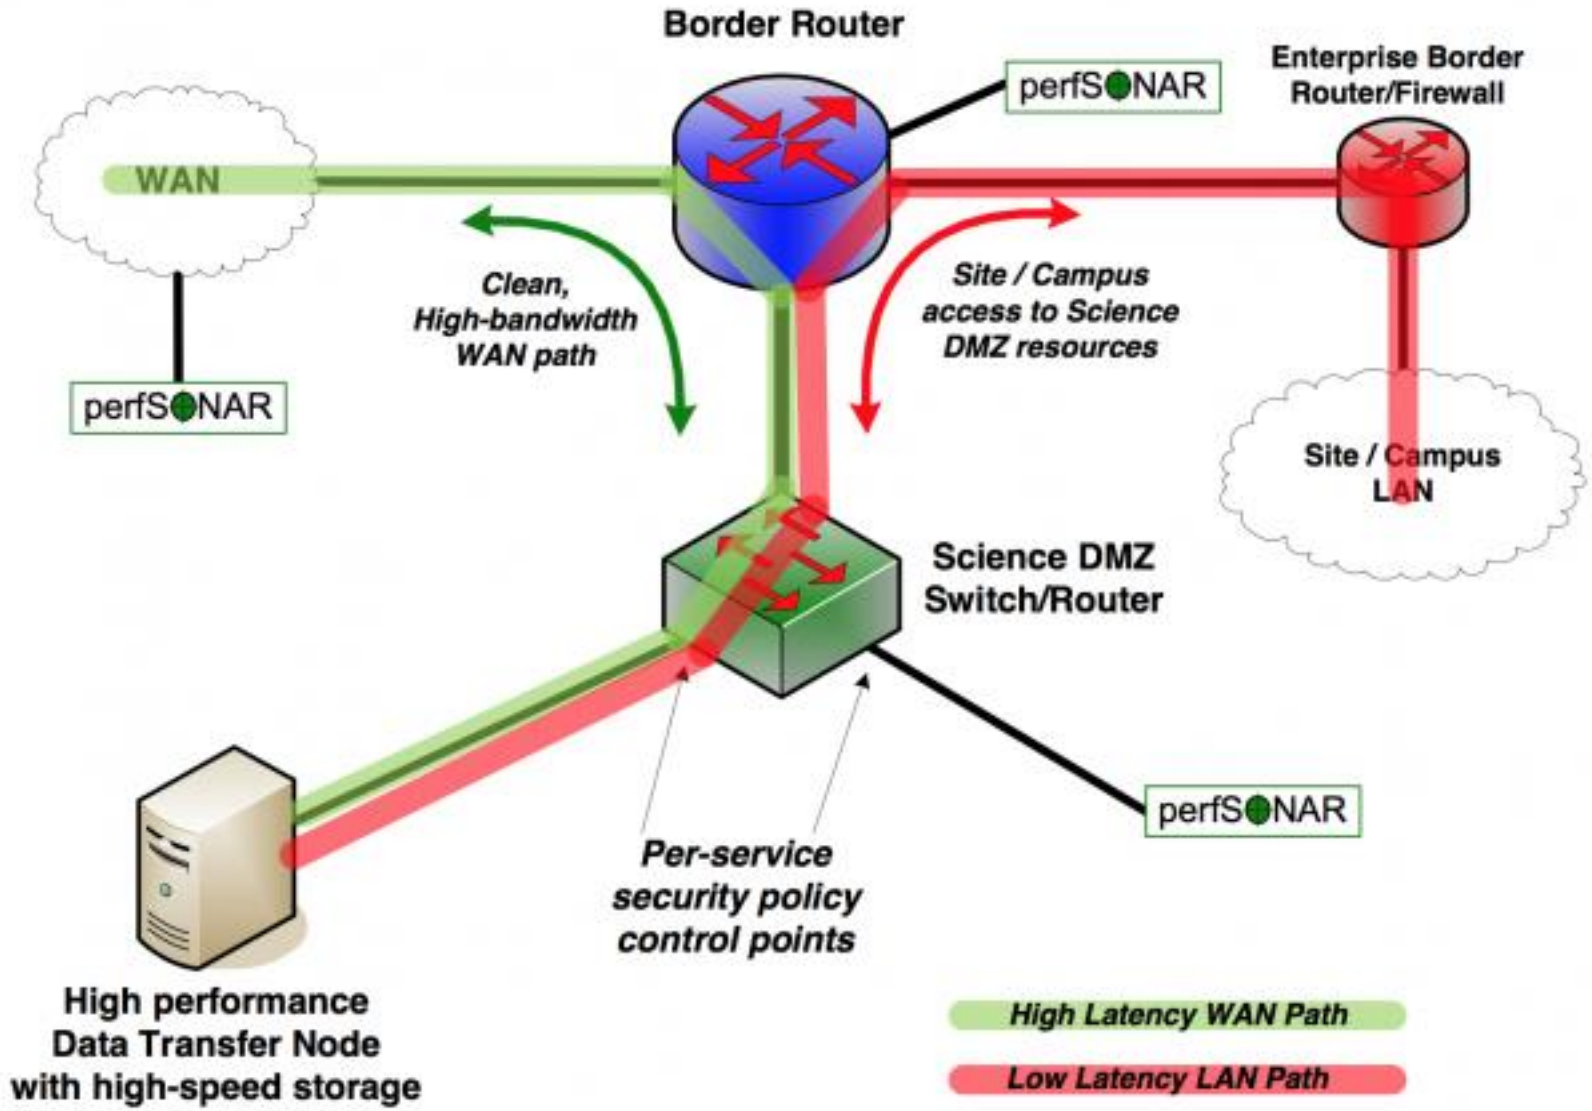

EPOC
Engagement and Performance Operations Center

Distributed Science DMZ – Dark Fiber

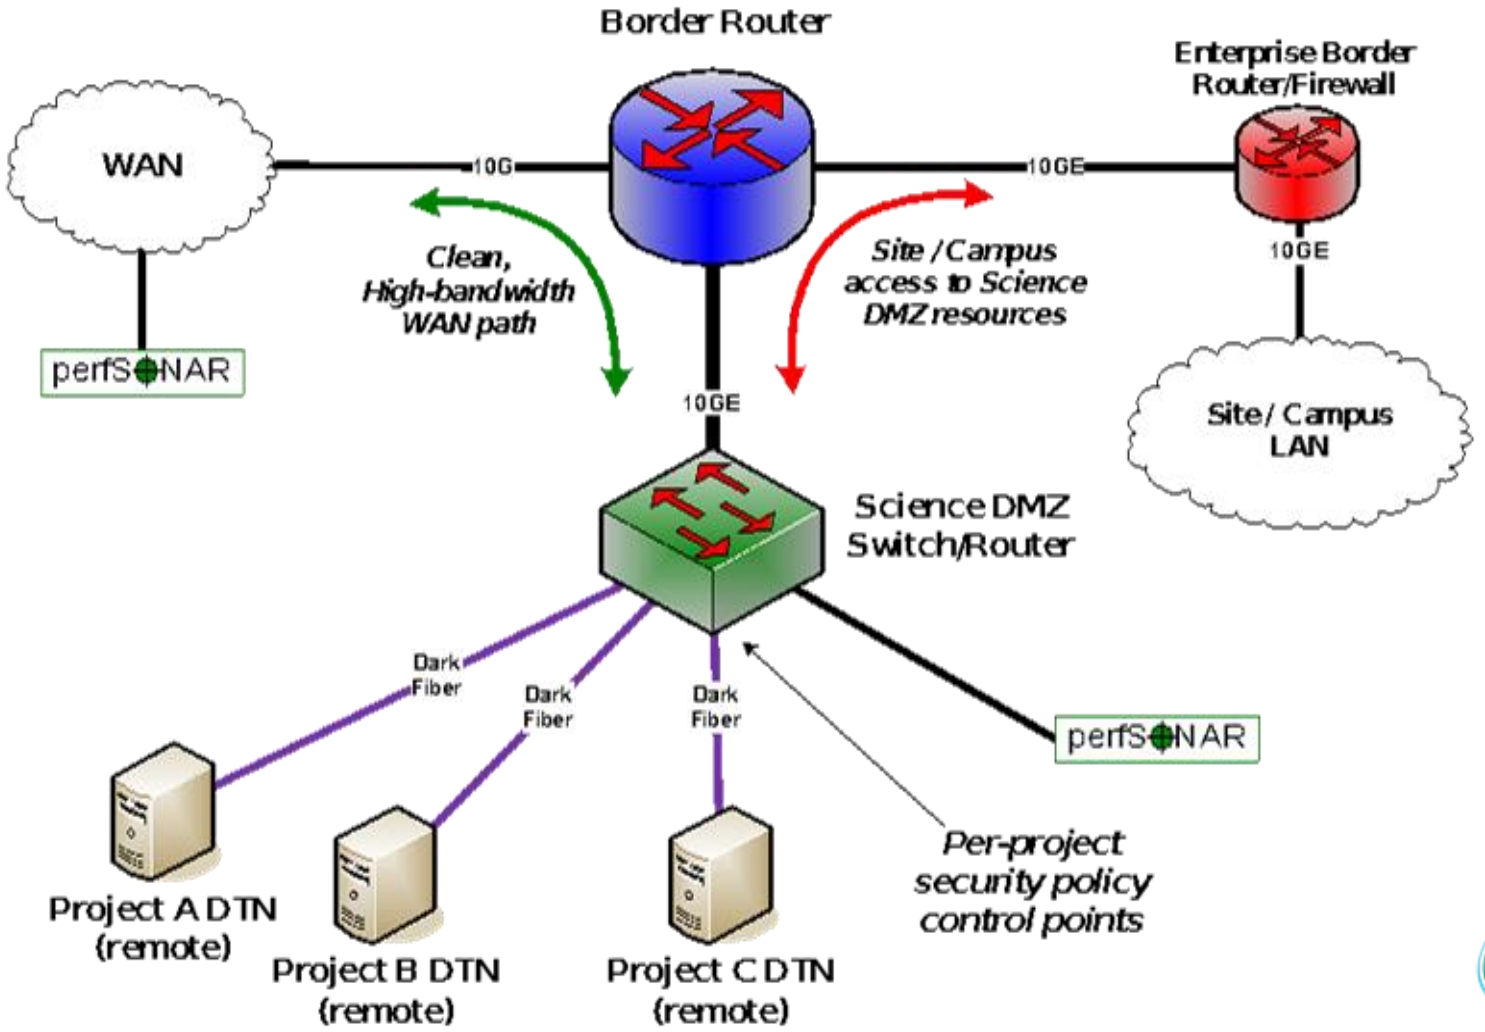

EPOC
Engagement and Performance Operations Center

Multiple Science DMZs – Dark Fiber to Dedicated Switches

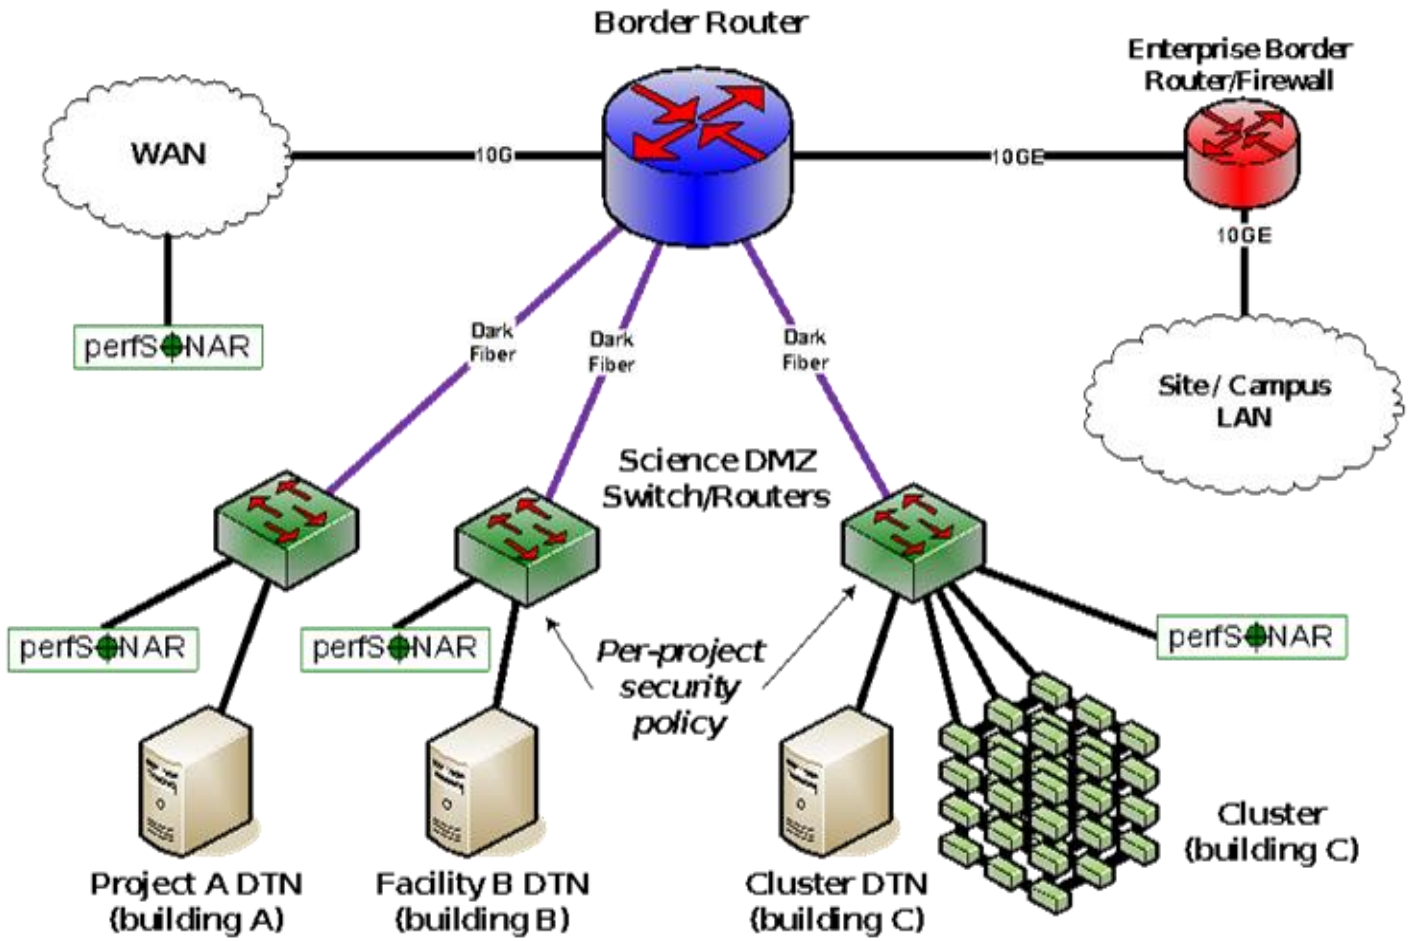

EPOC
Engagement and Performance Operations Center

Science DMZ Model in HPC Facility

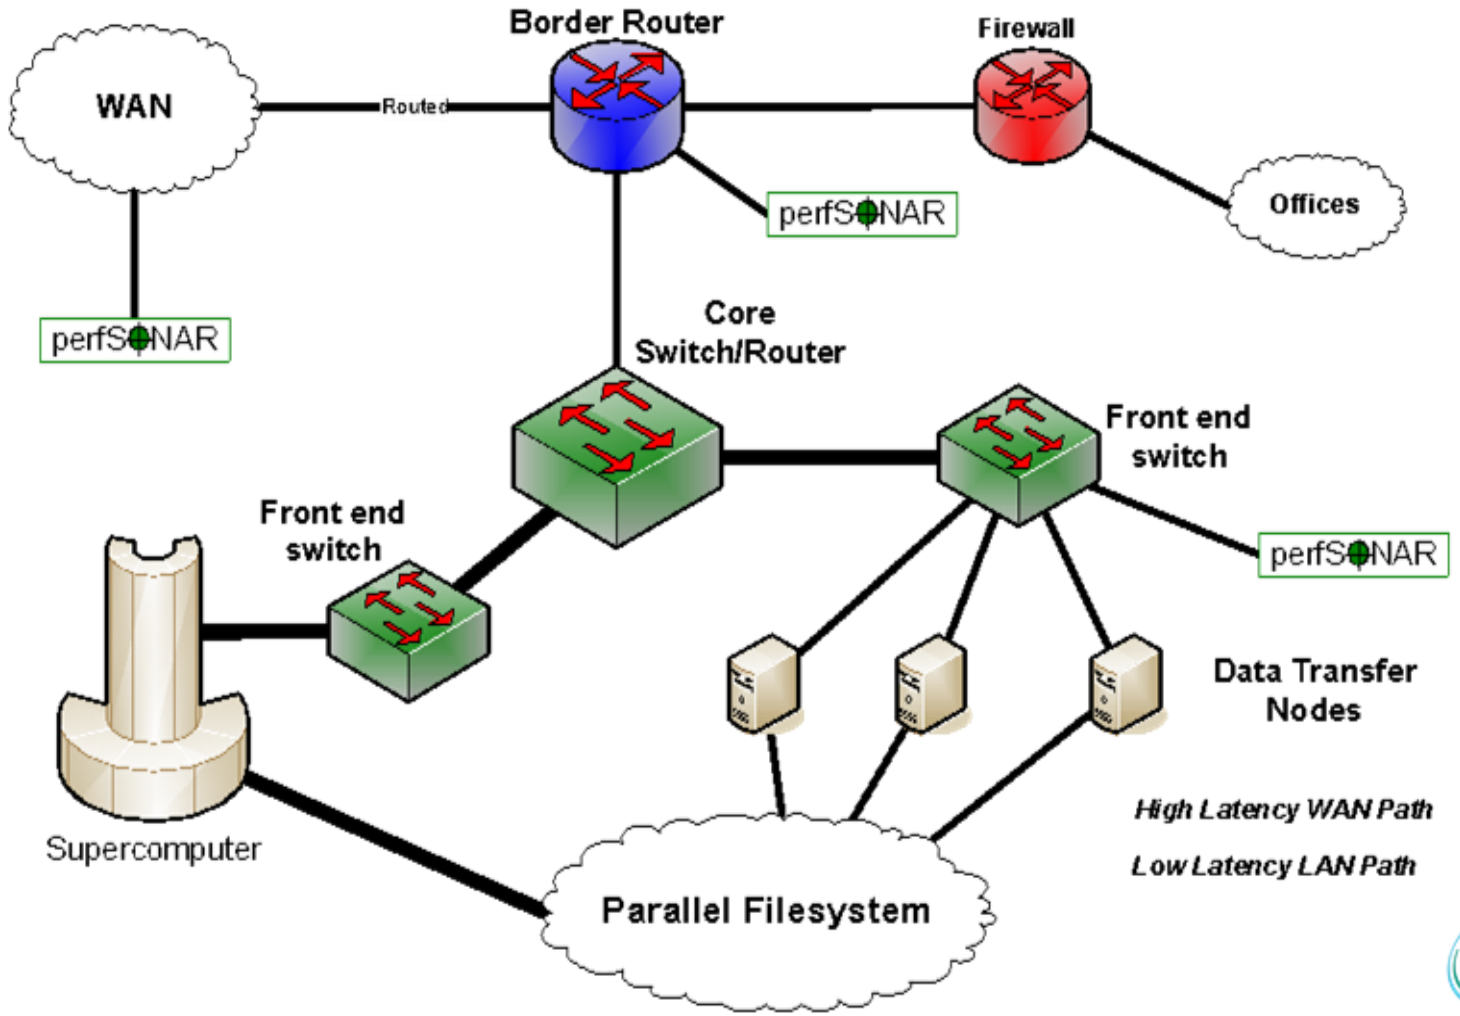

EPOC
Engagement and Performance Operations Center

Equipment – Routers and Switches

- Requirements for Science DMZ gear are different than the enterprise
  - No need to go for the kitchen sink list of services
  - A Science DMZ box only needs to do a few things, but do them well
  - Support for the latest LAN integration magic with your Windows Active Directory environment is probably not super-important
  - A clean architecture is important
    - How fast can a single flow go?
    - Are there any components that go slower than interface wire speed?
- There is a temptation to go cheap
  - It only needs to do a few things, right?
  - ”You get what you pay for”
- There is also a temptation to put it ‘everywhere’ - remember we want to optimize a single path, not all the paths
  - Helps keep $$$ in check, also helps keep security a primary concern

EPOC
Engagement and Performance Operations Center

Common Circumstance: Multiple Ingress Data Flows, Common Egress

Hosts will typically send packets at the speed of their interface (1G, 10G, etc.)

- Instantaneous rate, not average rate
- If TCP has window available and data to send, host sends until there is either no data or no window

Hosts moving big data (e.g. DTNs) can send large bursts of back-to-back packets

- This is true even if the average rate as measured over seconds is slower (e.g. 4Gbps)
- On microsecond time scales, there is often congestion
- Router or switch must queue packets or drop them

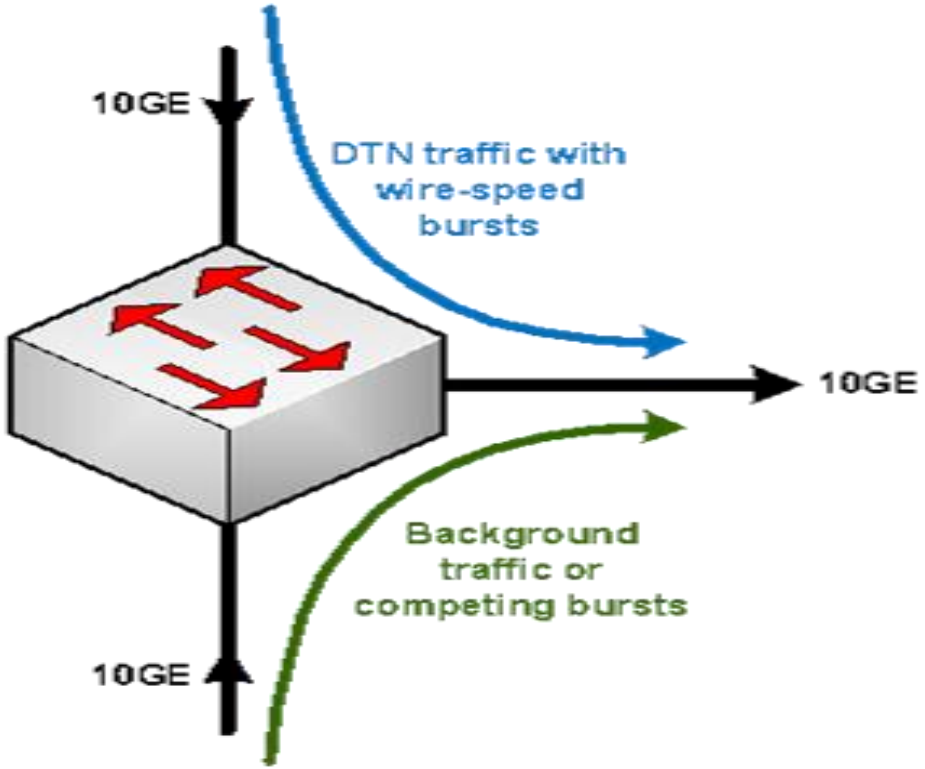

EPOC
Engagement and Performance Operations Center

Some Stuff We Think Is Important

- Deep interface queues (e.g. ***buffer***)
  - Output queue or VOQ – doesn't matter
  - What TCP sees is what matters – fan-in is *not* your friend
  - No, this isn't buffer bloat
- Good counters
  - We like the ability to reliably count *every* packet associated with a particular flow, address pair, etc
    - Very helpful for debugging packet loss
    - Must not affect performance (just count it, don't punt it)
    - sflow support if possible
  - If the box is going to drop a packet, it should increment a counter somewhere indicating that it dropped the packet
    - Magic vendor permissions and hidden commands should not be necessary
    - Some boxes just lie – run away!
- Single-flow performance should be wire-speed

EPOC
Engagement and Performance Operations Center

All About That Buffer (No Cut Through)

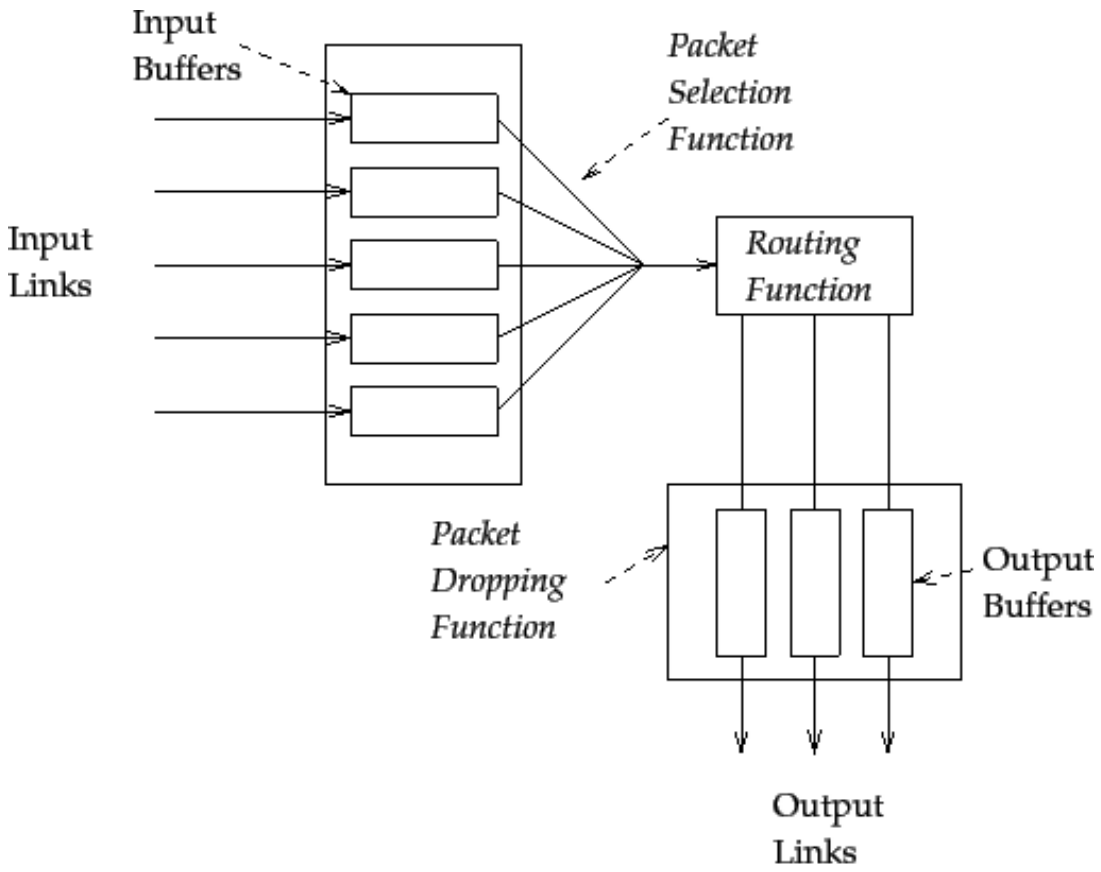

Figure 1: Basic Router Architecture

EPOC
Engagement and Performance Operations Center

All About That Buffer (No Cut Through)

- Data arrives from multiple sources

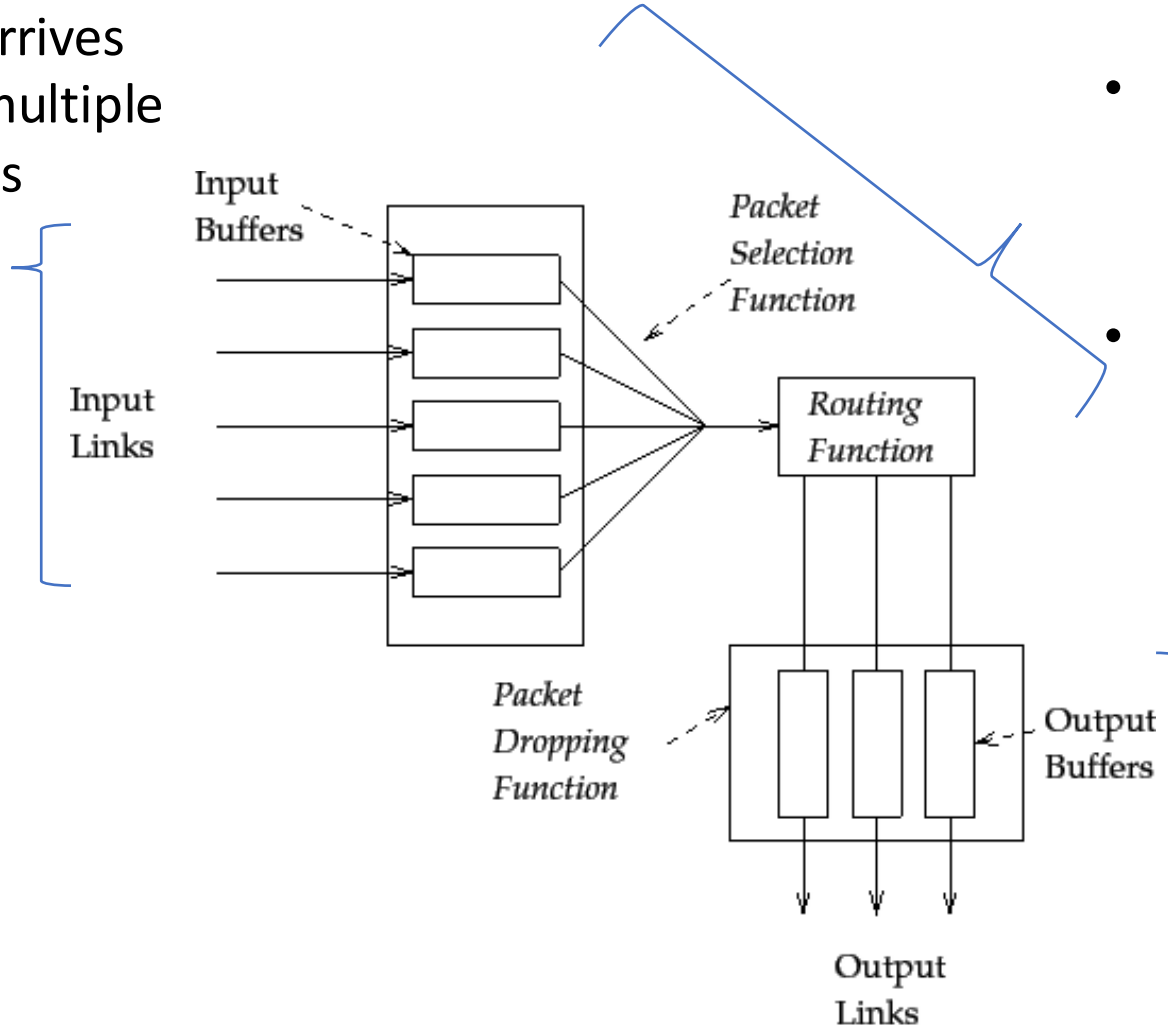

Figure 1: Basic Router Architecture

- Buffers have a finite amount of memory
  - Some have this per interface
  - Others may have access to a shared memory region with other interfaces
- The processing engine will:
  - Extract each packet/frame from the queues
  - Pull off header information to see where the destination should be
  - Move the packet/frame to the correct output queue
- Additional delay is possible as the queues physically write the packet to the transport medium (e.g. optical interface, copper interface)

EPOC Engagement and Performance Operations Center

All About That Buffer (No Cut Through)

- ***The Bandwidth Delay Product***
  - The amount of “in flight” data for a TCP connection (BDP = bandwidth * round trip time)
- Example: 10Gb/s cross country, ~100ms
  - 10,000,000,000 b/s * .1 s = 1,000,000,000 bits
  - 1,000,000,000 / 8 = 125,000,000 bytes
  - 125,000,000 bytes / (1024*1024) ~ ***125MB***
- Ignore the math aspect: its making sure there is memory to catch and send packets
  - *At **ALL** hops*
  - As the speed increases, there are more packets.
  - If there is not memory, we drop them, and that makes TCP react, and the user sad.

EPOC
Engagement and Performance Operations Center

All About That Buffer (No Cut Through)

- Buffering isn't as important on the LAN (this is why you are normally pressured to buy 'cut through' devices)
  - Change the math to make the Latency 1ms and the expectation 10Gbps = ***1.25MB***
  - 'Cut through' and low latency switches are designed for the data center, and can handle typical data center loads that don't require buffering (e.g. same to same speeds, destinations within the broadcast domain)

- Buffering ****MATTERS**** for WAN Transfers
  - Placing something with inadequate buffering in the path reduces the buffer for the entire path. E.g. if you have an expectation of 10Gbps over 100ms – don't place a 12MB buffer anywhere in there – your reality is now ~10x less than it was before (e.g. 10Gbps @ 10ms, or 1Gbps @ 100ms)

All About That Buffer (No Cut Through)

- What does this “look” like to a data transfer? Consider the test of iperf below
  - See TCP ‘ramp up’ and slowly increase the window
  - When something in the path has no more space for packets – a drop occurs. TCP will eventually react to the lost packet, and ‘back off’
  - In the example, this first occurs when we reach a buffer of around 6-8MB. Then after backoff the window is halved a couple of times
  - This happens again later – at a slightly higher buffer limit. This could be because there was cross traffic the first time, etc.

[ID]	Interval		Transfer	Bandwidth	Retr	Cwnd
[14]	0.00-1.00	sec	524 KBytes	4.29 Mbits/sec	0	157 KBytes
[14]	1.00-2.00	sec	3.31 MBytes	27.8 Mbits/sec	0	979 KBytes
[14]	2.00-3.00	sec	17.7 MBytes	148 Mbits/sec	0	***5.36 MBytes***
[14]	***3.00-4.00***	***sec***	***18.8 MBytes***	***157 Mbits/sec***	***214***	***1.77 MBytes***
[14]	4.00-5.00	sec	11.2 MBytes	94.4 Mbits/sec	0	***1.88 MBytes***
[14]	5.00-6.00	sec	10.0 MBytes	83.9 Mbits/sec	0	***2.39 MBytes***
[14]	6.00-7.00	sec	16.2 MBytes	136 Mbits/sec	0	3.63 MBytes
[14]	7.00-8.00	sec	23.8 MBytes	199 Mbits/sec	0	5.50 MBytes
[14]	8.00-9.00	sec	38.8 MBytes	325 Mbits/sec	0	8.23 MBytes
[14]	9.00-10.00	sec	57.5 MBytes	482 Mbits/sec	0	11.8 MBytes
[14]	10.00-11.00	sec	81.2 MBytes	682 Mbits/sec	0	***16.2 MBytes***
[14]	***11.00-12.00***	***sec***	***50.0 MBytes***	***419 Mbits/sec***	***35***	***3.93 MBytes***
[14]	12.00-13.00	sec	15.0 MBytes	126 Mbits/sec	0	***2.20 MBytes***
[14]	13.00-14.00	sec	11.2 MBytes	94.4 Mbits/sec	0	***2.53 MBytes***
[14]	14.00-15.00	sec	13.8 MBytes	115 Mbits/sec	1	1.50 MBytes
[14]	15.00-16.00	sec	6.25 MBytes	52.4 Mbits/sec	5	813 KBytes
[14]	16.00-17.00	sec	5.00 MBytes	41.9 Mbits/sec	0	909 KBytes
[14]	17.00-18.00	sec	5.00 MBytes	41.9 Mbits/sec	0	1.37 MBytes
[14]	18.00-19.00	sec	10.0 MBytes	83.9 Mbits/sec	0	2.43 MBytes
[14]	19.00-20.00	sec	17.5 MBytes	147 Mbits/sec	0	4.22 MBytes

EPOC
Engagement and Performance Operations Center

Decoding Specifications

- Consider this architecture
  - 48 Ports
    - 12 ASICS
    - 4 Ports per ASIC
  - ***72MB*** total
    - ***6MB per ASIC***
    - If all ports are in use – expect that each port has access to ***1.5MB***. If only one is in use, it can use 6MB
  - Additional memory is often available in a ‘burst buffer’ in the fabric

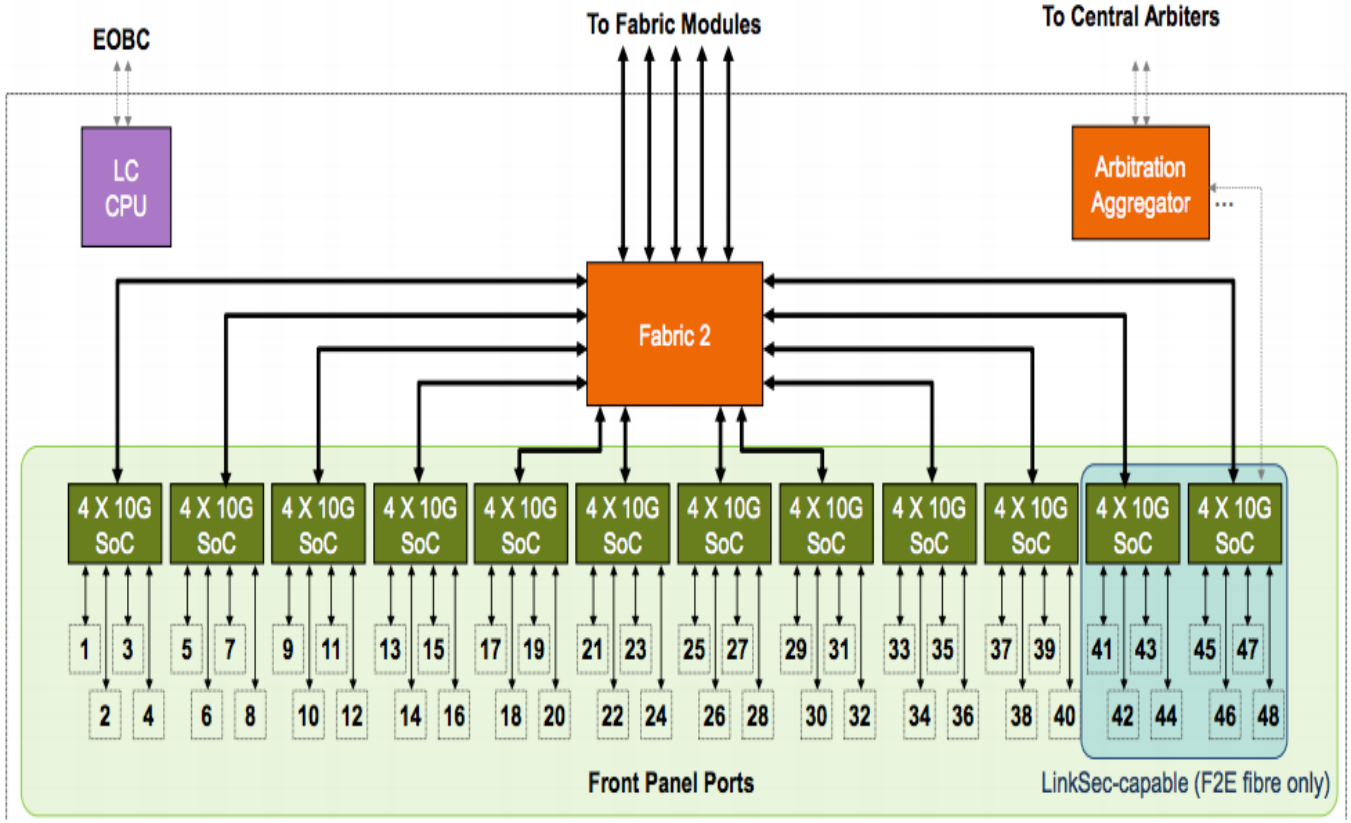

ASIC = application-specific integrated circuit, think 'small routing engine'

EPOC
Engagement and Performance Operations Center

Decoding Specifications

- A useful tool: https://network.switch.ch/pub/tools/tcp-throughput/
  - ***Note: this helps you understand buffer behavior for a single stream, in theory a switch/router should be able to support *many* streams***
- What does 6MB get you?
  - 1Gbps @ <= 48ms (e.g. ½ needed for coast-to-coast)
  - 10Gbps @ <= 4.8ms (e.g. metro area)
- What does 1.5MB get you?
  - 1Gbps @ <= 12ms (e.g. regional area)
  - 10Gbps @ <= 1.2ms (e.g. data center [or more accurately, rack or row)]
- In either case – remember this assumes you are the only thing using that memory … congestion is a more likely reality

EPOC
Engagement and Performance Operations Center

TCP's Congestion Control

50ms simulated RTT
Congestion w/ 2Gbps UDP traffic
HTCP / Linux 2.6.32

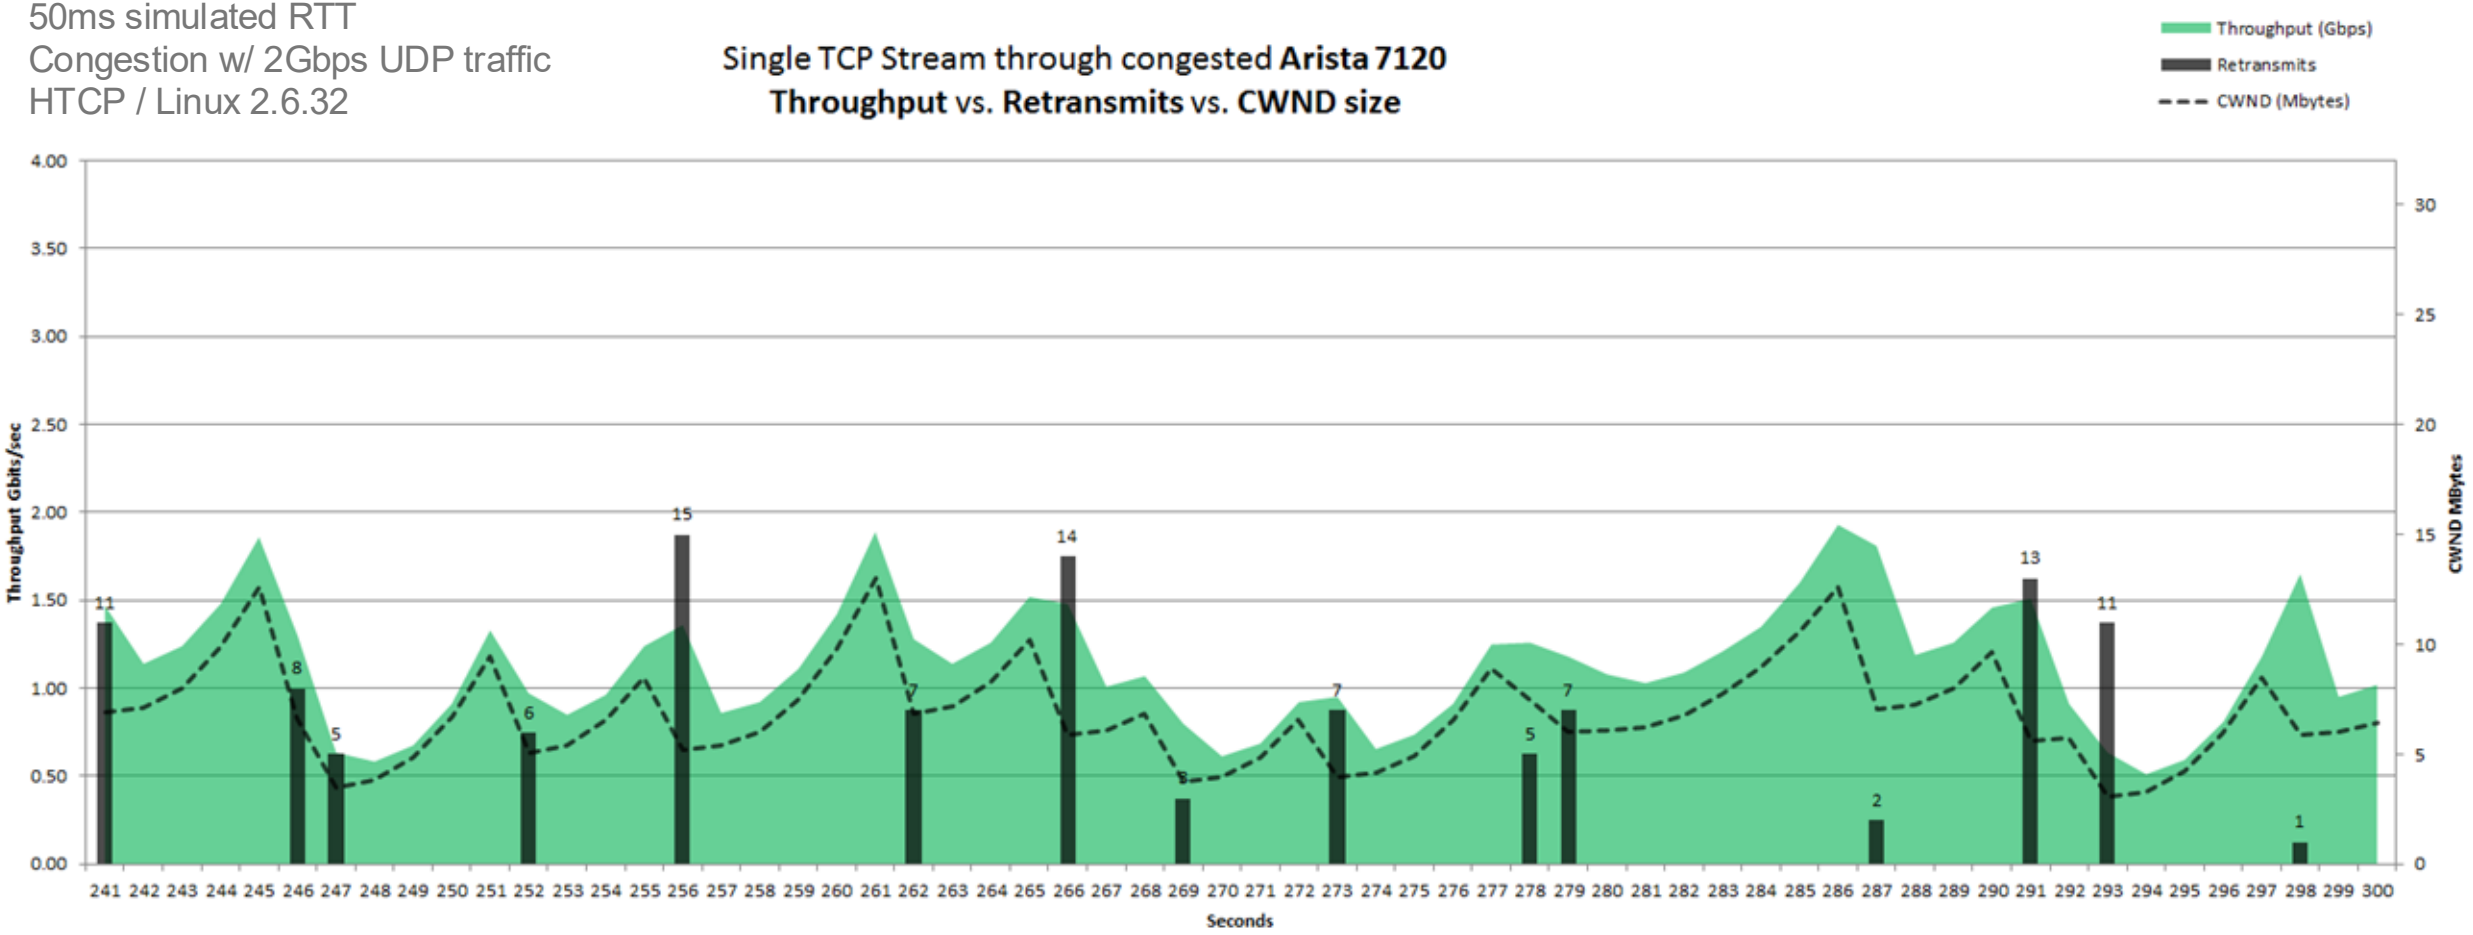

EPOC
Engagement and Performance Operations Center

Slide from Michael Smitasin, LBLnet

Outline

- Introduction
- ***Solution Space***
  1. Understanding the Solution Space (Users, Use Cases, Long Term Impacts)
  2. Preliminaries (e.g. Network Protocols 101)
  3. Architecture & Design
  4. ***Data Mobility***
- Conclusions / QA

EPOC
Engagement and Performance Operations Center

Next Steps – Building On The Science DMZ

- Enhanced cyberinfrastructure substrate exists and it works
  - Wide area networks (ESnet, Internet2, Regionals)
  - Science DMZs connected to those networks
  - DTNs in the Science DMZs
- What does the scientist see?
  - Scientist sees a science application
    - Data transfer
    - Data portal
    - Data analysis
  - Science applications are the user interface to networks and DMZs
- Large-scale data-intensive science requires that we build larger structures on top of those components

EPOC
Engagement and Performance Operations Center

Performance At Different Data Scales

Data set size	1 Minute	5 Minutes	20 Minutes	1 Hour
10PB	1,333.33 Tbps	266.67 Tbps	66.67 Tbps	22.22 Tbps
1PB	133.33 Tbps	26.67 Tbps	6.67 Tbps	2.22 Tbps
100TB	13.33 Tbps	2.67 Tbps		
10TB > 100Gbps	1.33 Tbps	266.67 Gbps		
1TB	133.33 Gbps	26.67 Gbps	6.67 Gbps	2.22 Gbps
100GB 100Gbps	13.33 Gbps	2.67 Gbps	666.67 Mbps	222.22 Mbps
10GB < 10Gbps	1.33 Gbps	266.67 Mbps	66.67 Mbps	22.22 Mbps
1GB	133.33 Mbps	26.67 Mbps	6.67 Mbps	2.22 Mbps
100MB < 100Mbps	13.33 Mbps	2.67 Mbps	0.67 Mbps	0.22 Mbps
Time to transfer				

10G DTN

10G DTN min

- This table available at: http://fasterdata.es.net/fasterdata-home/requirements-and-expectations/

EPOC Engagement and Performance Operations Center

Solution Space – Data Mobility

- DTN History & Purpose:
  - Original concept came from initial Science DMZ Design (~2012)
  - Basic idea:
    - Host(s) dedicated to the task of data movement (and only data movement)
    - Limited application set (data movement tools), and users (rarely shell access)
    - Specific security policy enforced on the switch/router ACLs
      - Ports for data movement tools, most in a ‘closed wait’ state
      - Nothing to impact the data channel
    - Typically 2 footed:
      - Limited reach into local network (e.g. ‘control channel’: shared filesystem, instruments)
      - WAN piece that the data tools use (e.g. ‘data channel’)
  - Position this, and the pS node, in the DMZ enclave near the border

DTN Architecture Considerations

DTNs can be all ‘internal’, e.g. not connected to external storage, or ‘pass through’ where they have access to external storage

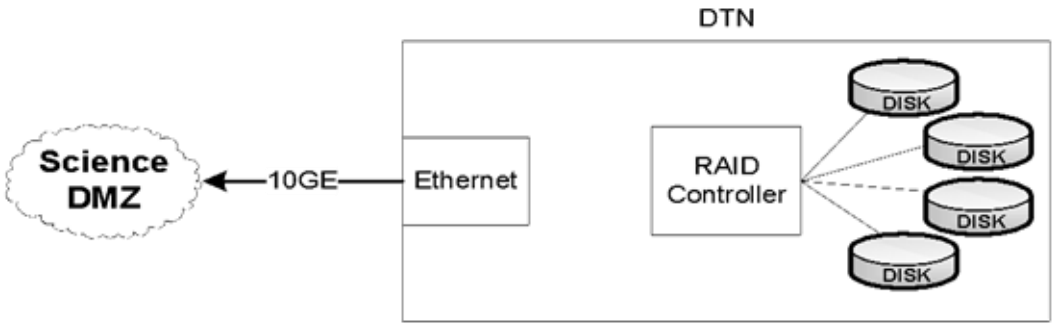

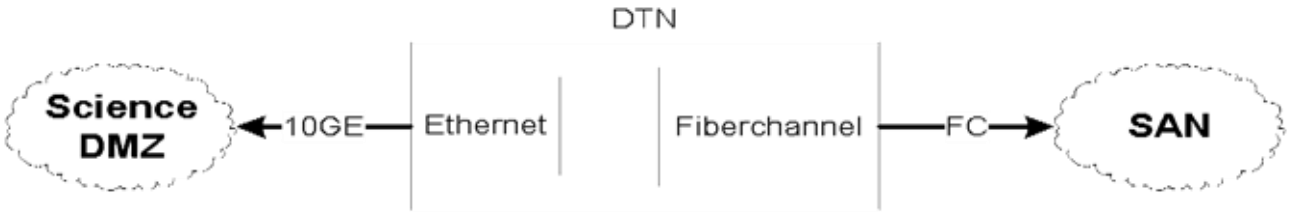

Discuss options with your scientific users – figure out which workflow will work best!

EPOC
Engagement and Performance Operations Center

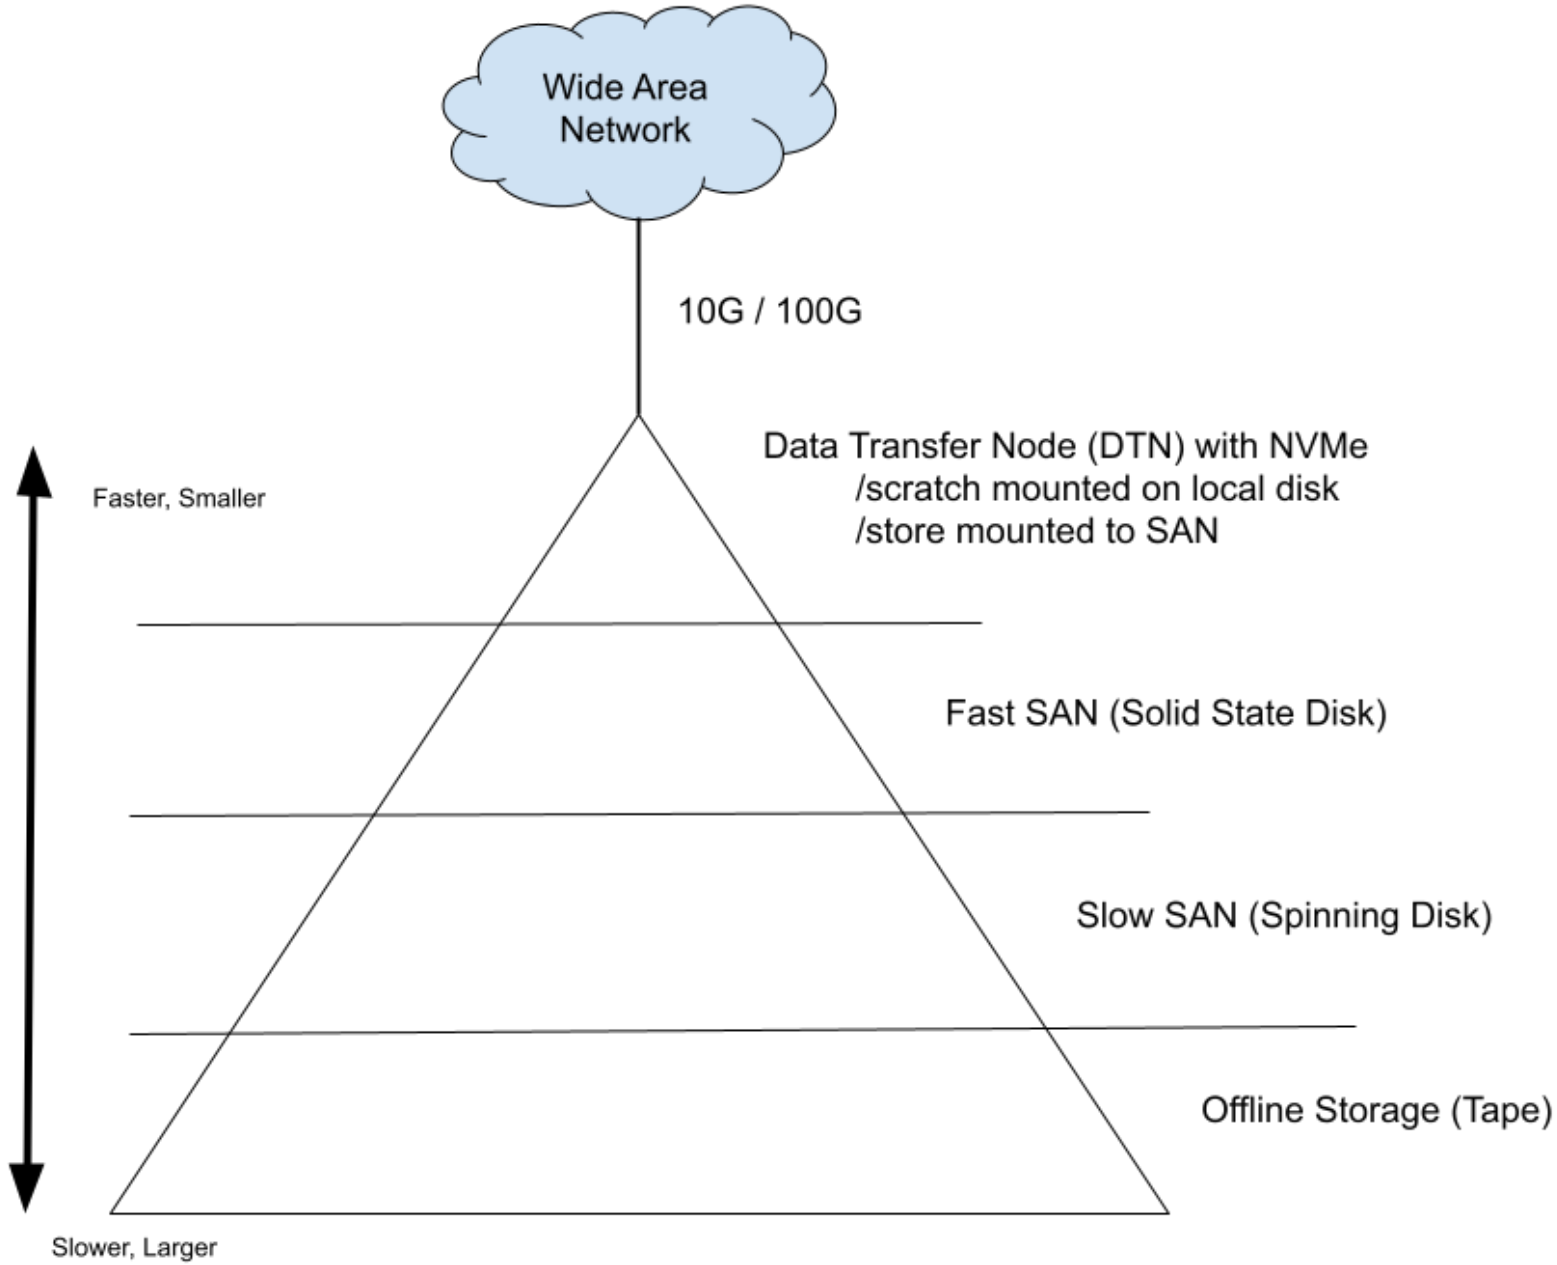

EPOC
Engagement and Performance Operations Center

Solution Space – Data Mobility

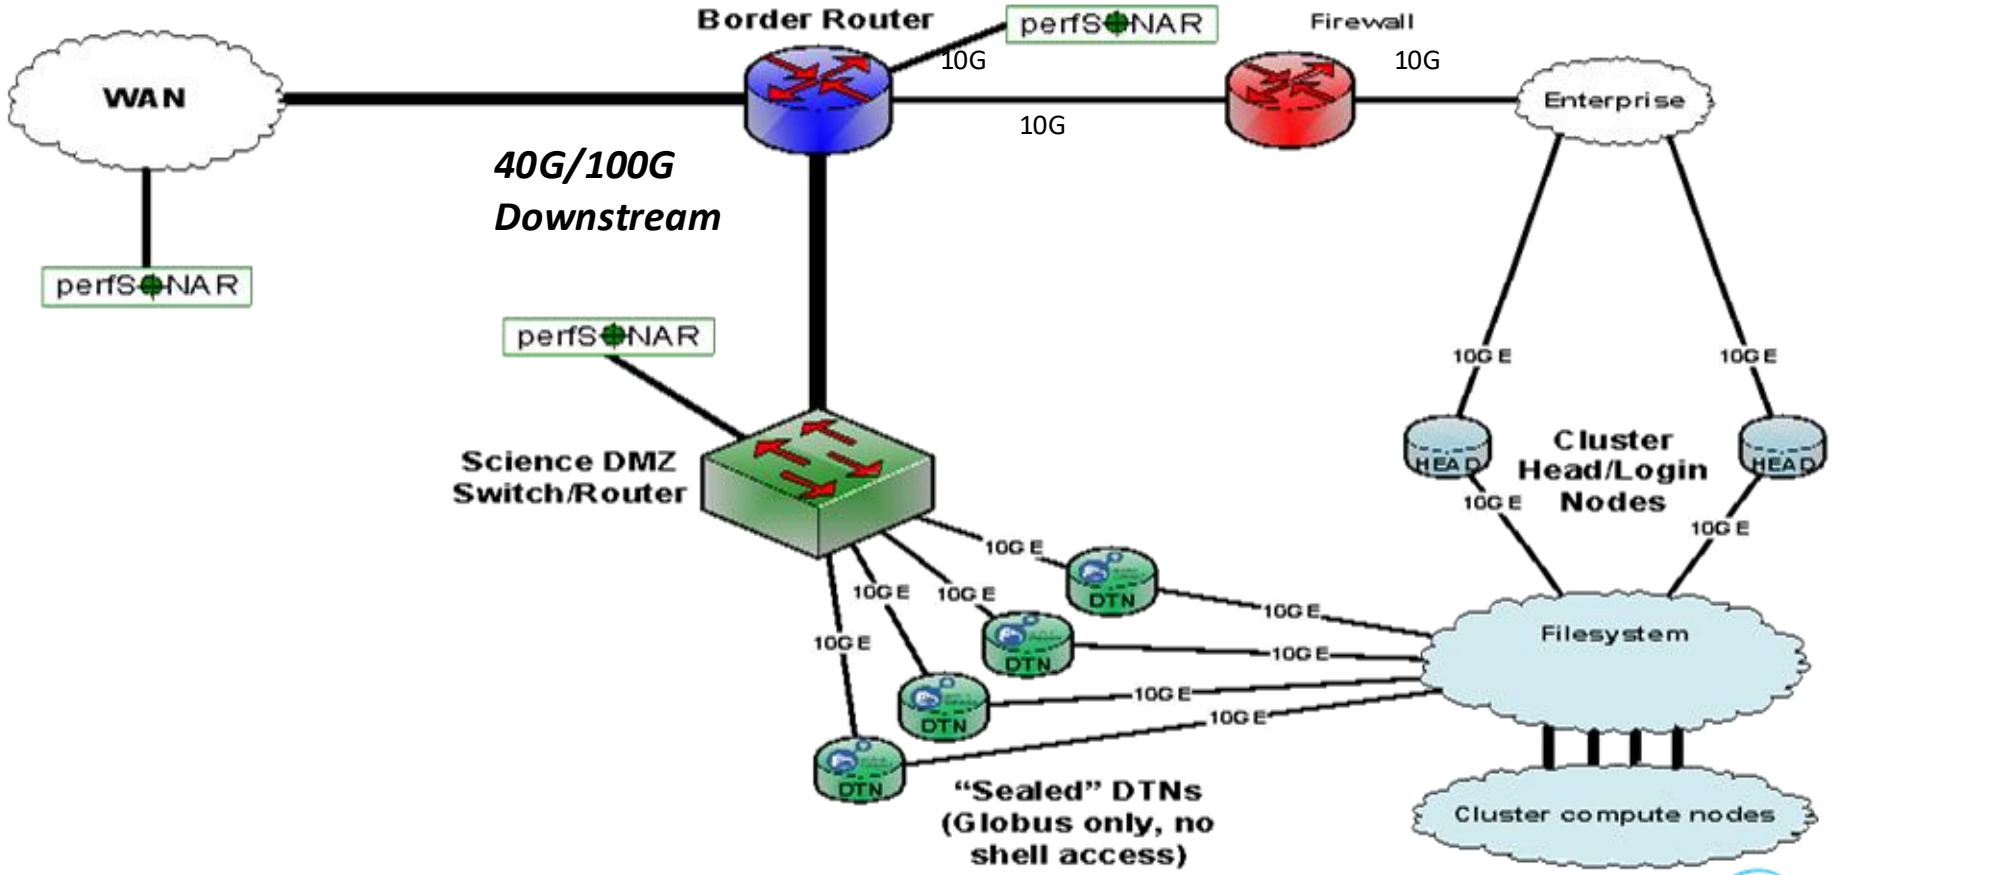

EPOC
Engagement and Performance Operations Center

Solution Space – Data Mobility

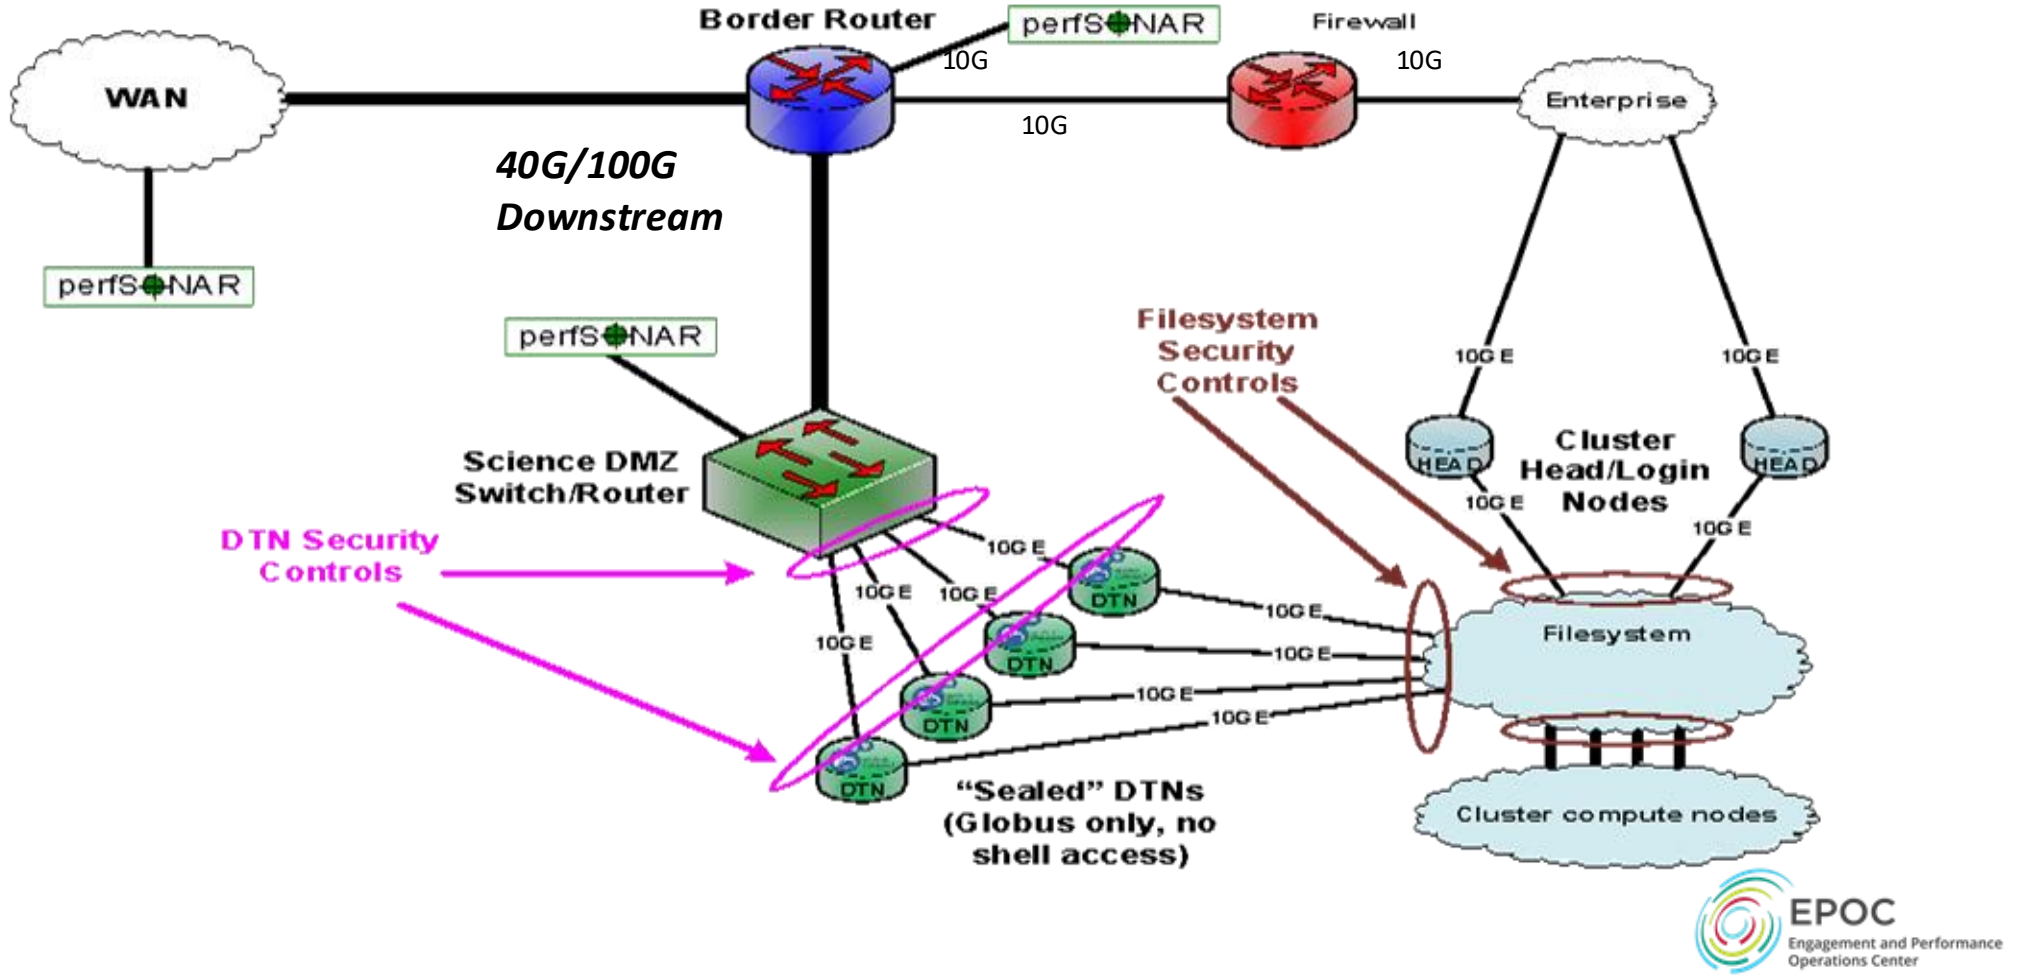

Software – Data Transfer

- Functionality varies
  - Some are command line, some are graphical, some are tied to advanced workflow software
  - All use different protocols (TCP, UDP)
  - All have different port in/out requirements
  - Some require shell access to the machine, some are invoked via other known protocols (HTTP/HTTPS), others can be run 3rd party
- Common themes to a ‘good’ tool:
  - Parallelism
  - Checksumming
  - Aggressive (application layer) tuning
  - API that allows for integration into higher-level software

EPOC
Engagement and Performance Operations Center

Software – Data Transfer (2005)

- 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

EPOC
Engagement and Performance Operations Center

Software – Data Transfer (2023)

- Using the right data transfer tool is ***still*** important
- Sample Results: Berkeley, CA to Argonne, IL (near Chicago) RTT = 53 ms, network capacity = 10Gbps.

Tool	Throughput
scp	330 Mbps
wget, Globus, FDT, 1 stream	6 Gbps
Globus and FDT, 4 streams	8 Gbps (disk limited)

- Notes
  - scp is 24x slower than Globus on this path!!
  - Assume host TCP buffers are set correctly for the RTT

EPOC
Engagement and Performance Operations Center

Workflow

- Now that we have discussed the technology, its important to perform the final engagement step – integration with the end users.
- All workflows are different, but many share common components:
  - Data is created/brought-in/manipulated in one location
  - Data is analyzed/processed stored, possibly in different location
  - Data is shared with others that may be in different locations
  - Different layers of security considerations
  - Requirements for a litany of tools (analysis, transfer, etc.)

EPOC
Engagement and Performance Operations Center

Science Workflow Consultation

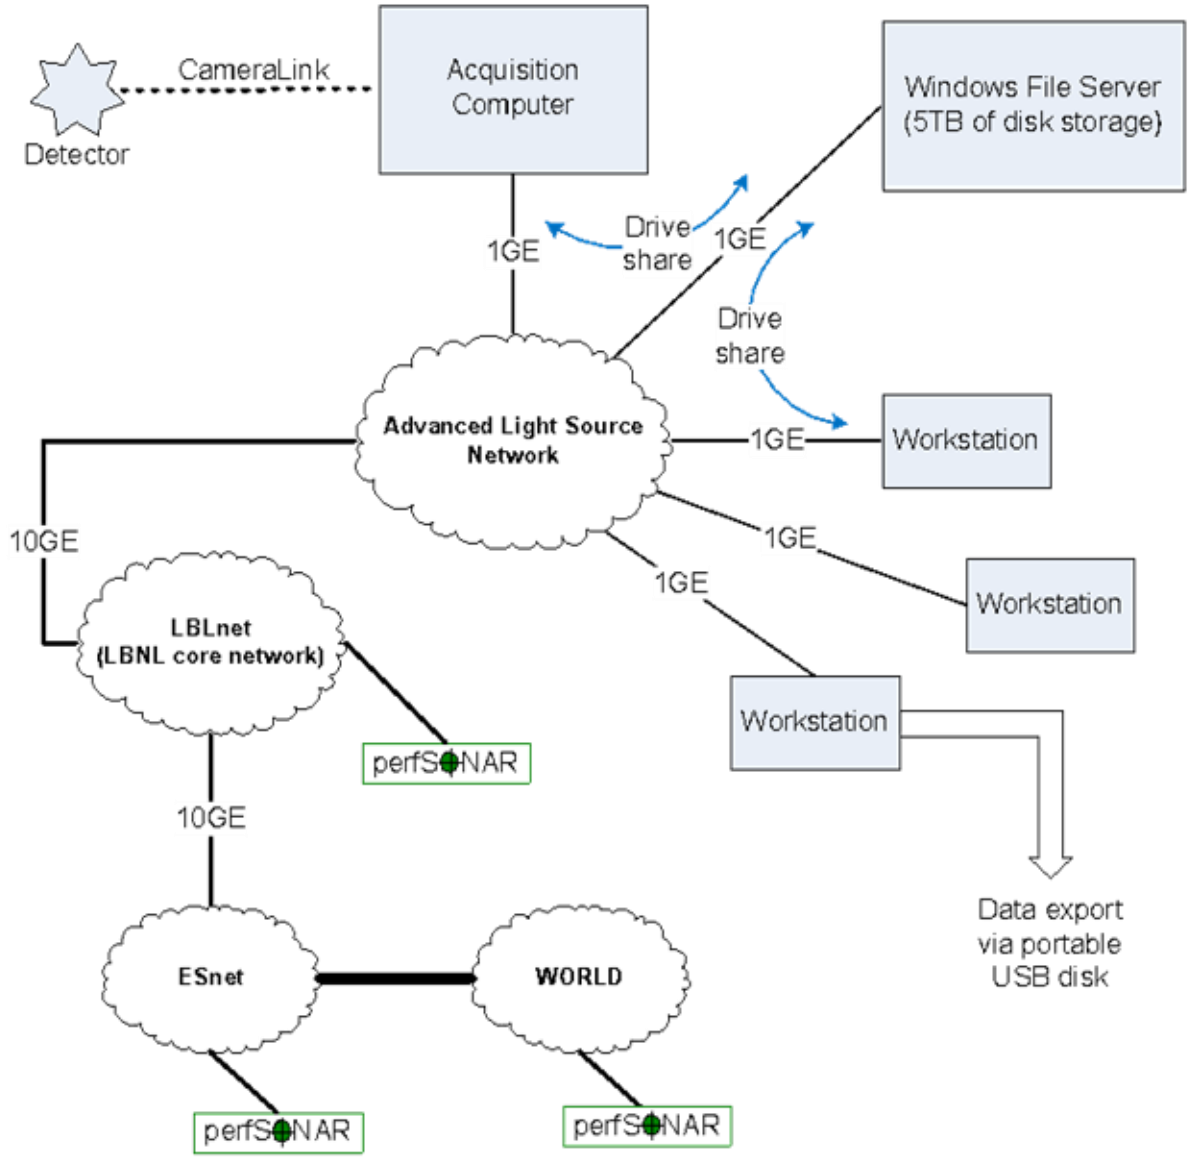

EPOC
Engagement and Performance Operations Center

Improved Workflow Infrastructure

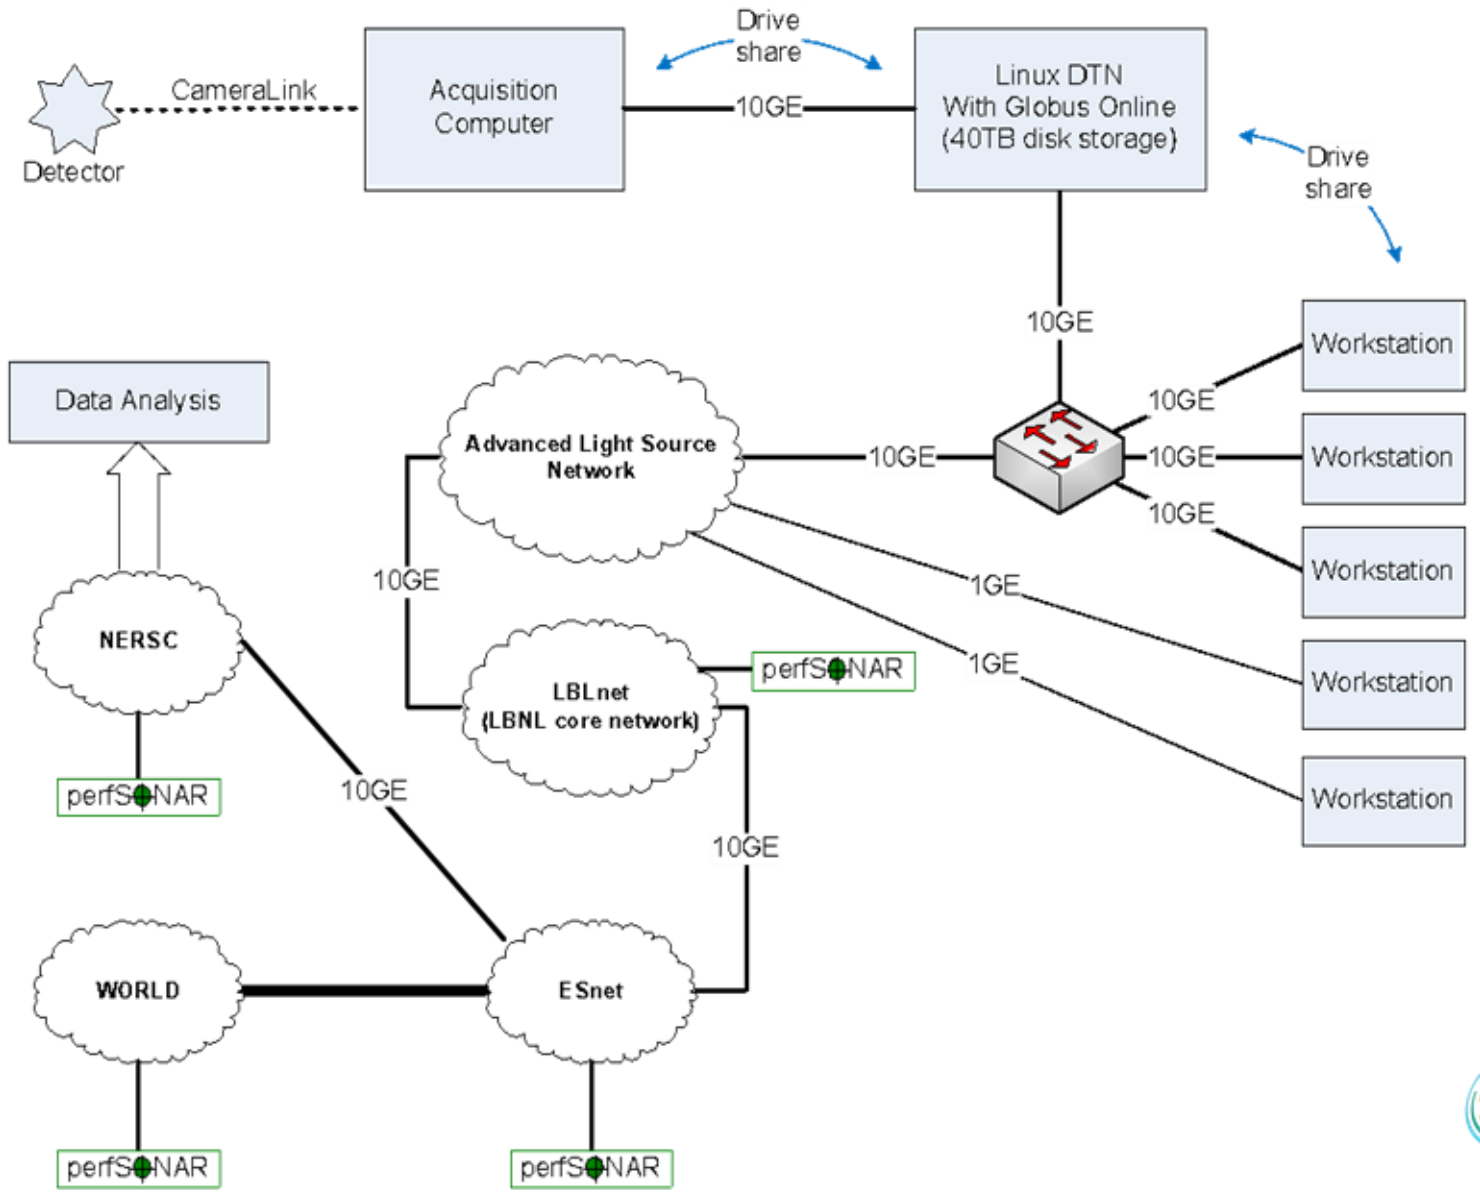

EPOC
Engagement and Performance Operations Center

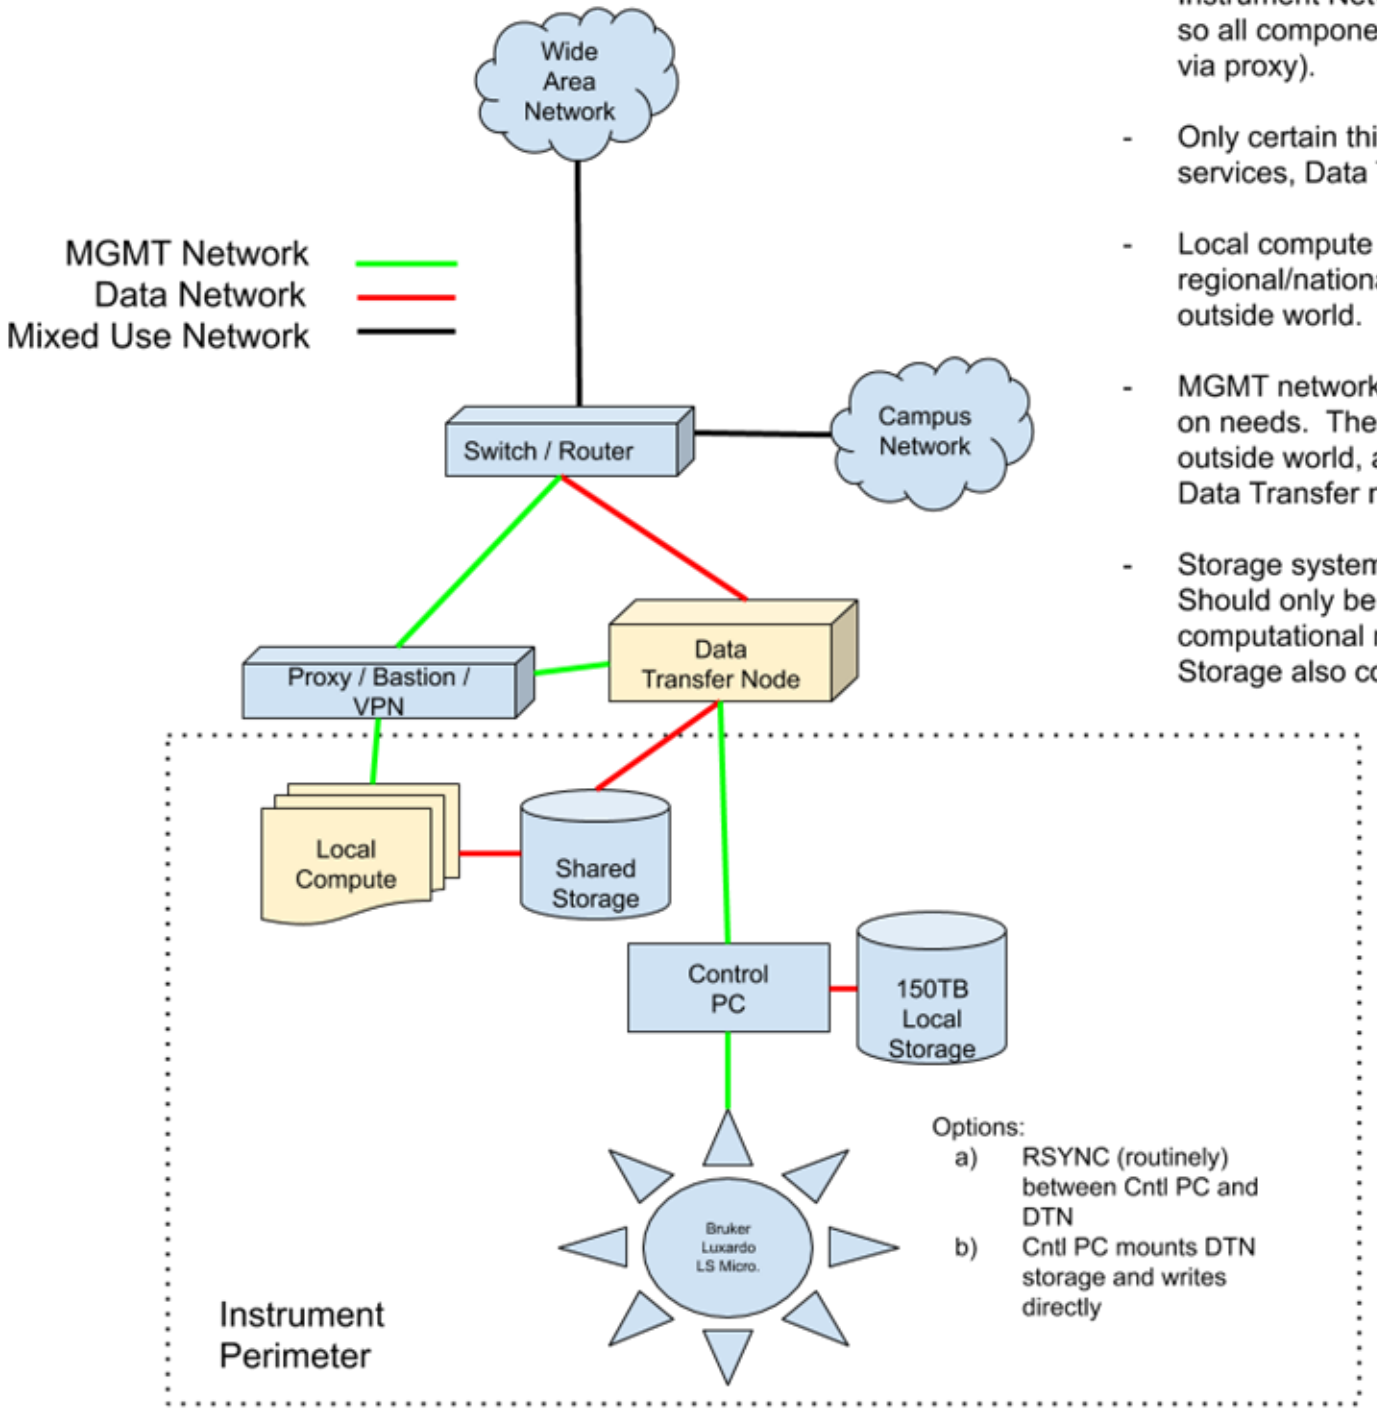

- Instrument Network can features static internal addressing scheme, so all components can function without external networking (except via proxy).
- Only certain things exposed with external address: Proxy/internet services, Data Transfer Node, Bastion/VPN.
- Local compute can be bolted on to complete analysis. Can also use regional/national compute, and use Data Transfer node to send to outside world.
- MGMT network could have connections to multiple things - depends on needs. The idea here is that the control PC is isolated from the outside world, and has to Proxy through either the VPN/Bastion or Data Transfer node.
- Storage system is meant to be protected from external access. Should only be accessible by instrument, data transfer, and computational resources (e.g. establish a 'data VLAN' for access). Storage also could just be inside of the data transfer node.

EPOC
Engagement and Performance Operations Center

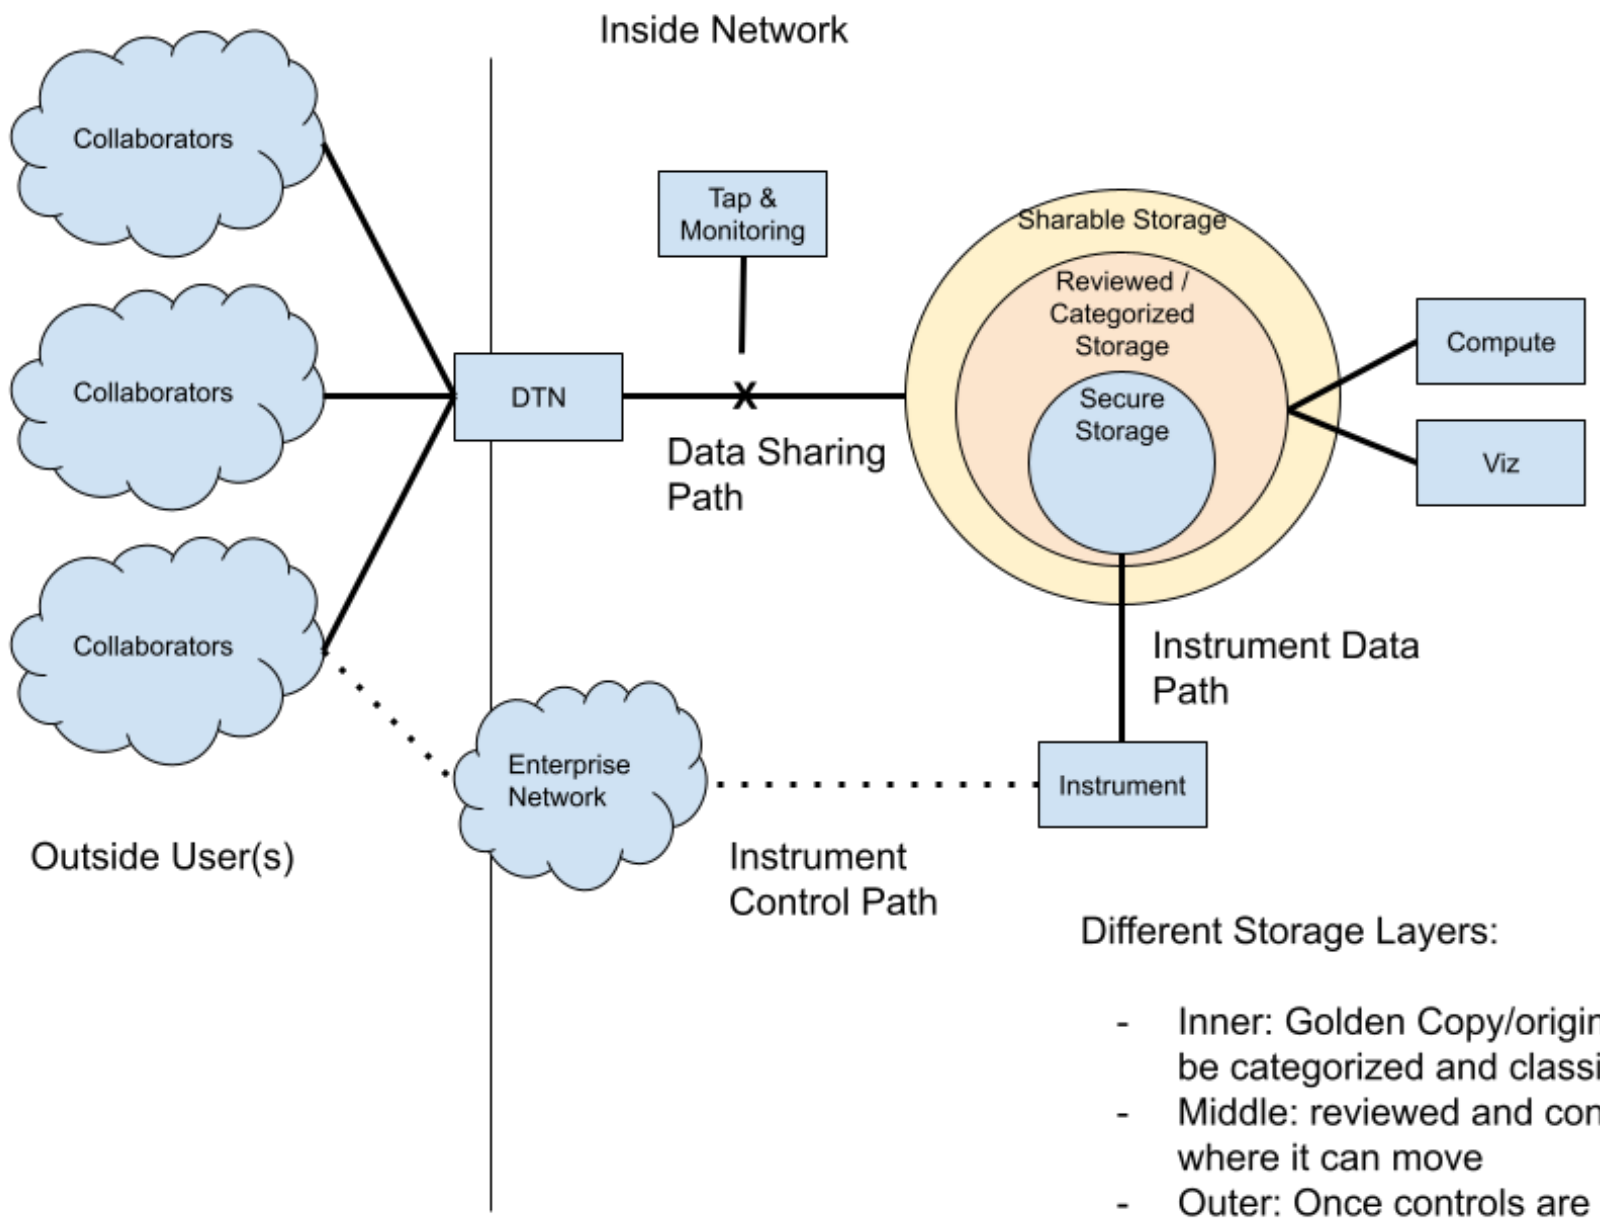

Different Storage Layers:

- Inner: Golden Copy/origin until it can be categorized and classified
- Middle: reviewed and controls placed where it can move
- Outer: Once controls are in place, it can be sent to different use cases, maybe several of these (for internal or external use)

EPOC
Engagement and Performance Operations Center

To Reiterate:

- Data movement is hard to get right
- Lots of moving parts
  - Software, Servers, Networks, and People
- Testing will reveal that it may not be ideal
- Shared experience around the community – lift all the boats, share all the knowledge, etc.

EPOC
Engagement and Performance Operations Center

Outline

- Introduction
- Solution Space
- ***Conclusions / QA***

EPOC
Engagement and Performance Operations Center

Questions?

zurawski@es.net

EPOC
Engagement and Performance Operations Center

EPOC
Engagement and Performance Operations Center

Science DMZ Overview: Practical Designs and Use Cases

Jason Zurawski
zurawski@es.net
ESnet / Lawrence Berkeley National Laboratory

Empowering Secure Data-Driven Research
January 14, 2026

TACC TEXAS ADVANCED COMPUTING CENTER
ESnet ENERGY SCIENCES NETWORK