ESnet

ENERGY SCIENCES NETWORK

The Science DMZ

Eli Dart, Network Engineer

ESnet Science Engagement

Lawrence Berkeley National Laboratory

Building the Modern Research Data Portal

GlobusWorld

Chicago, IL

April 20, 2016

Outline

- Science DMZ in brief
- Context – Science DMZ in the community
- Science DMZ and Data Portals

- This assumes you already have a Science DMZ
  - If you don't have one, we can chat about how you might build one
  - If it would be helpful, I can talk to your systems and networking folks
  - Or check out the fasterdata knowledgebase:
    - http://fasterdata.es.net/science-dmz/

ESnet

Science DMZ Design Pattern (Abstract)

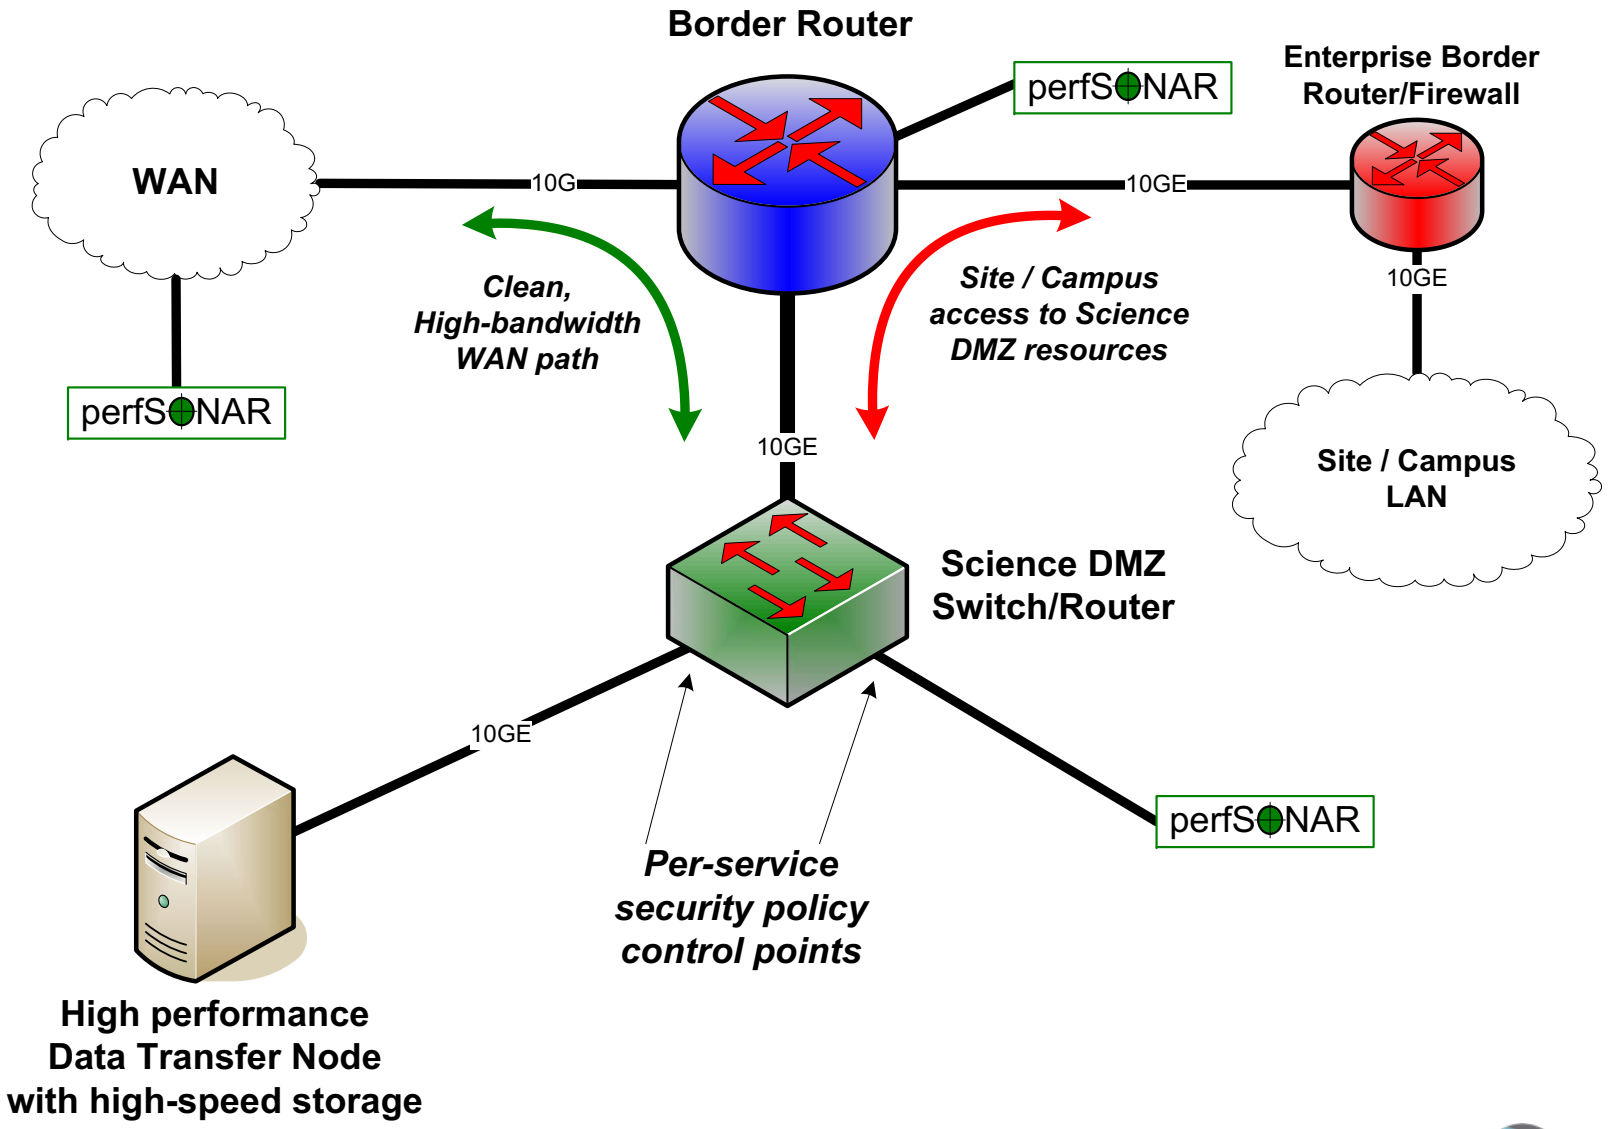

Supercomputer Center Deployment

- 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

HPC Center

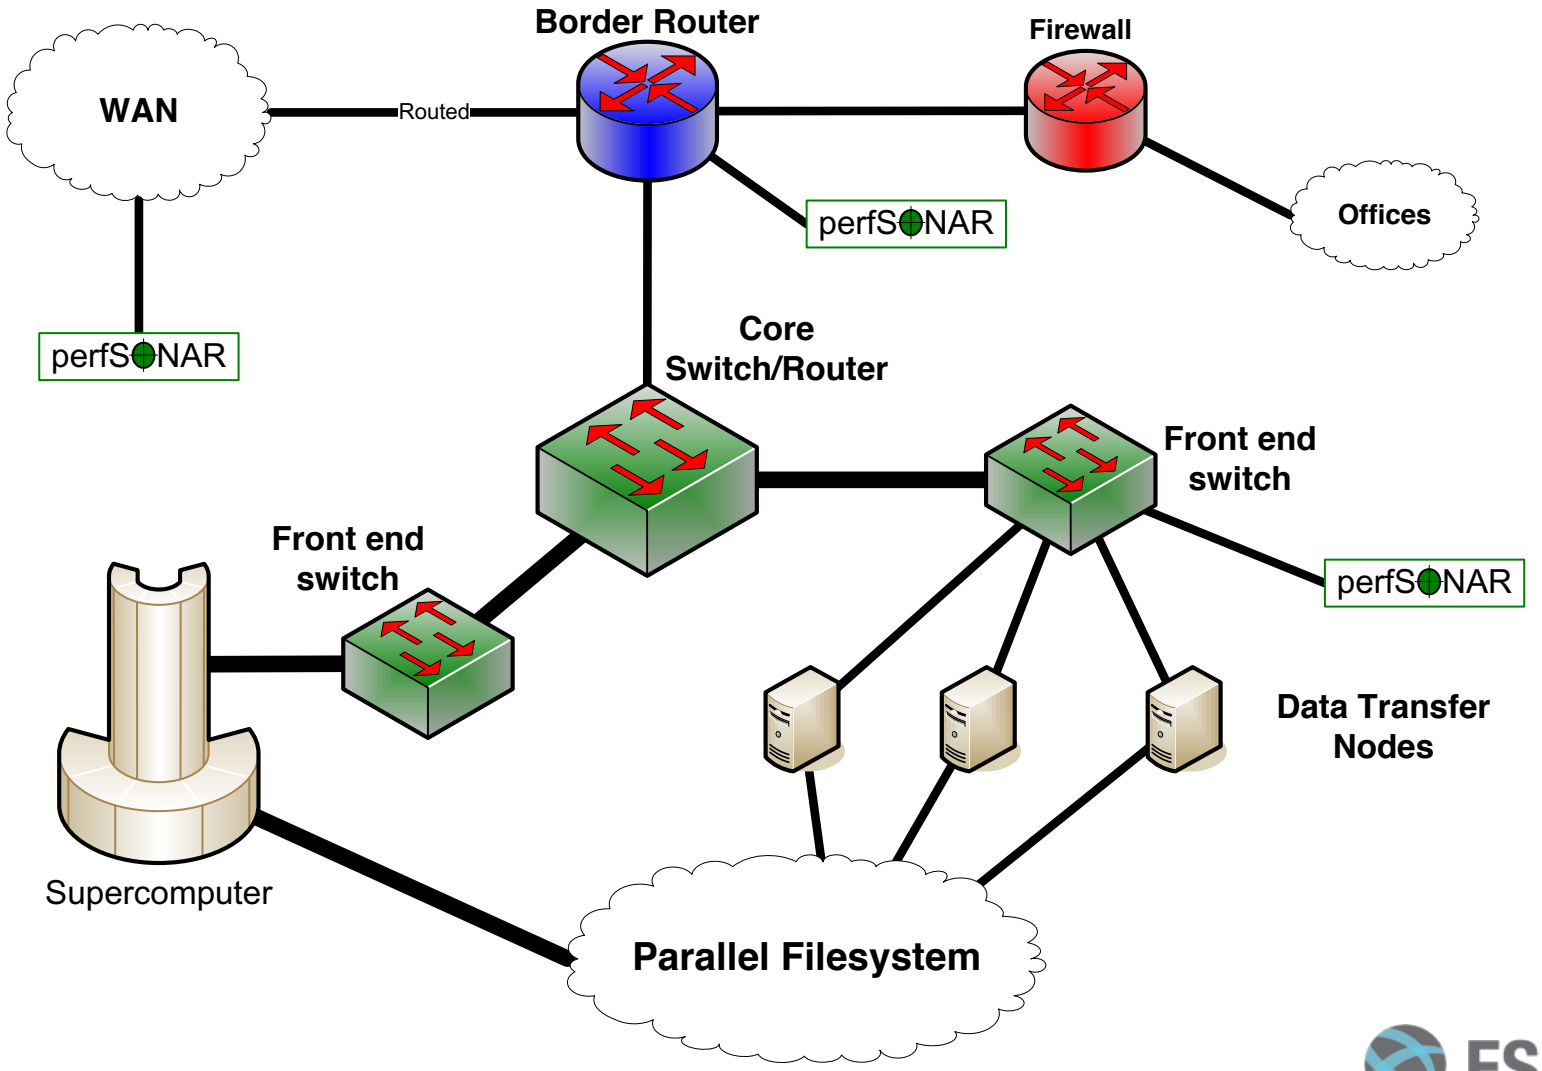

HPC Center Data Path

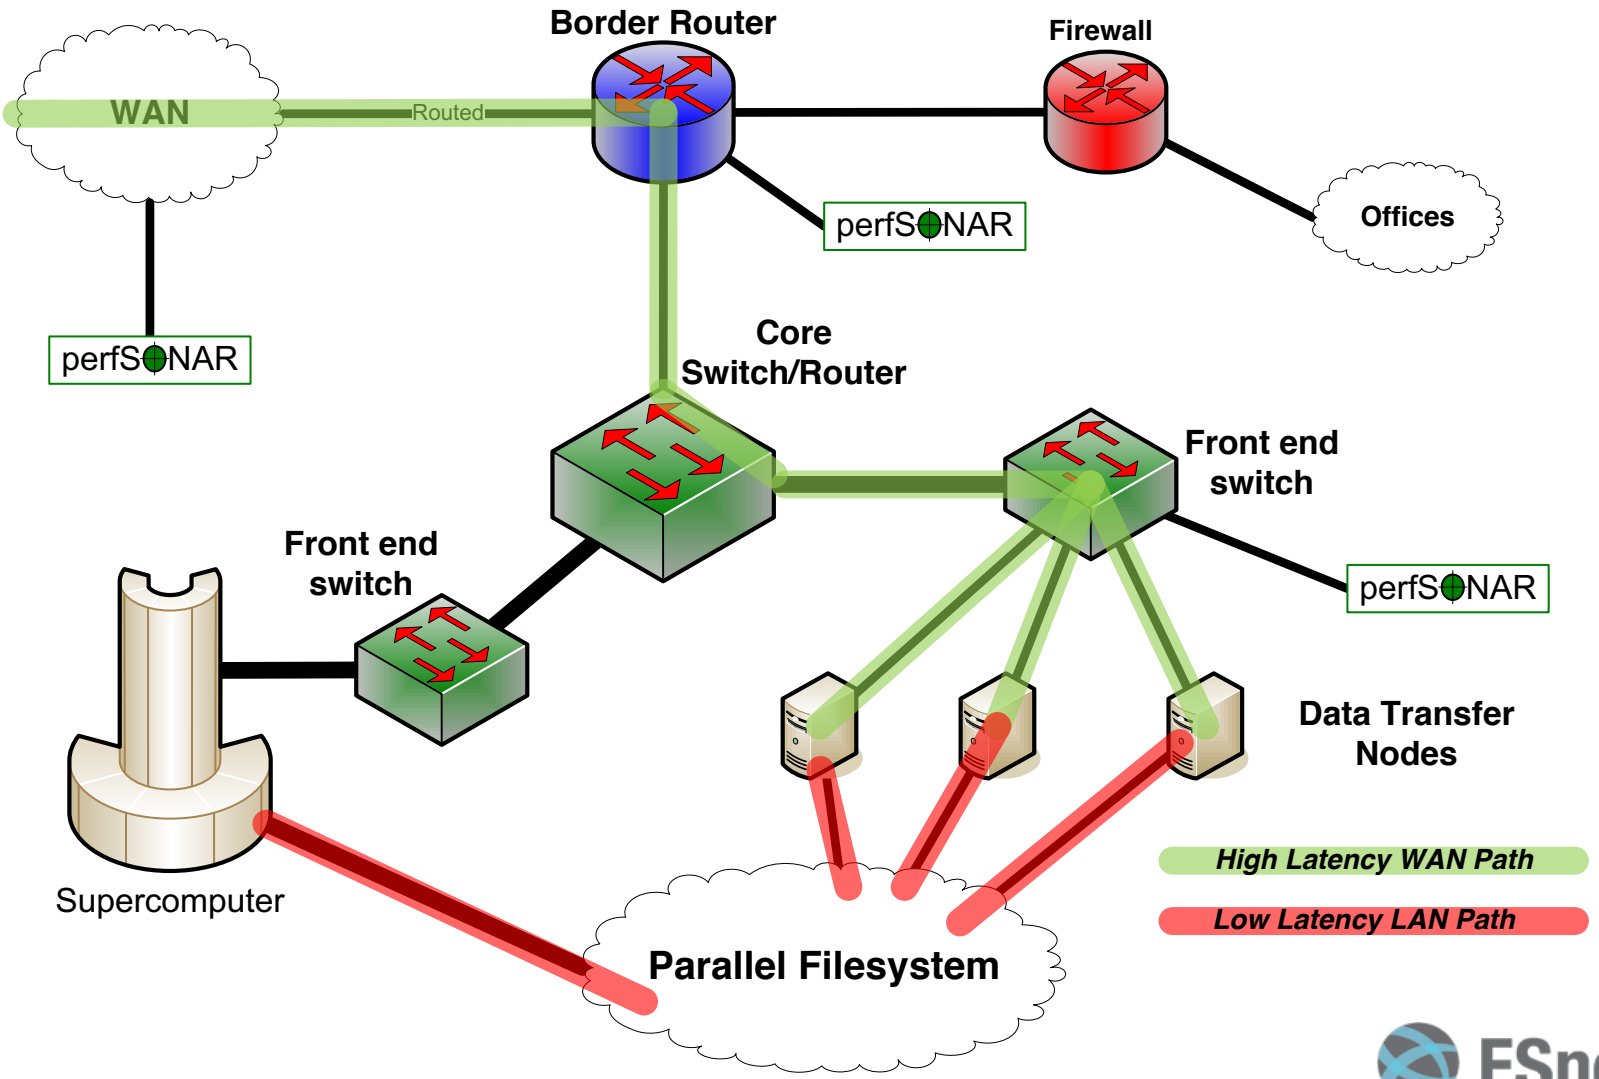

ESnet

Context: Science DMZ Adoption

- DOE National Laboratories
  - HPC centers, LHC sites, experimental facilities
  - Both large and small sites
- NSF CC* programs have funded many Science DMZs
  - Significant investments across the US university complex
  - Big shoutout to the NSF – these programs are critically important
- Other US agencies
  - NIH
  - USDA Agricultural Research Service
- International
  - Australia https://www.rdsi.edu.au/dashnet
  - Brazil
  - UK

ESnet

Strategic Impacts

- What does this mean?
  - We are in the midst of a significant cyberinfrastructure upgrade
  - Enterprise networks need not be unduly perturbed ☺
- Significantly enhanced capabilities compared to 3 years ago
  - Terabyte-scale data movement is much easier
  - Petabyte-scale data movement possible outside the LHC experiments
    - ~3.1Gbps = 1PB/month
    - ~14Gbps = 1PB/week
  - Widely-deployed tools are much better (e.g. Globus)
- Metcalfe's Law of Network Utility
  - Value of Science DMZ proportional to the number of DMZs
    - n^2 or $n(\log_n)$ doesn't matter – the effect is real
  - Cyberinfrastructure value increases as we all upgrade

ESnet

Next Steps – Building On The Science DMZ

- Enhanced cyberinfrastructure substrate now exists
  - Wide area networks (ESnet, GEANT, 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
- ***The underlying cyberinfrastructure components (networks, Science DMZs, DTNs, etc.) are part of the instrument of discovery***
- Large-scale data-intensive science requires that we build larger structures on top of those components

ESnet

Science Data Portals

- Large repositories of scientific data
  - Climate data
  - Sky surveys (astronomy, cosmology)
  - Many others
  - Data search, browsing, access
- Many scientific data portals were designed 15+ years ago
  - Single-web-server design
  - Data browse/search, data access, user awareness all in a single system
  - All the data goes through the portal server
    - In many cases by design
    - E.g. embargo before publication (enforce access control)

ESnet

Legacy Portal Design

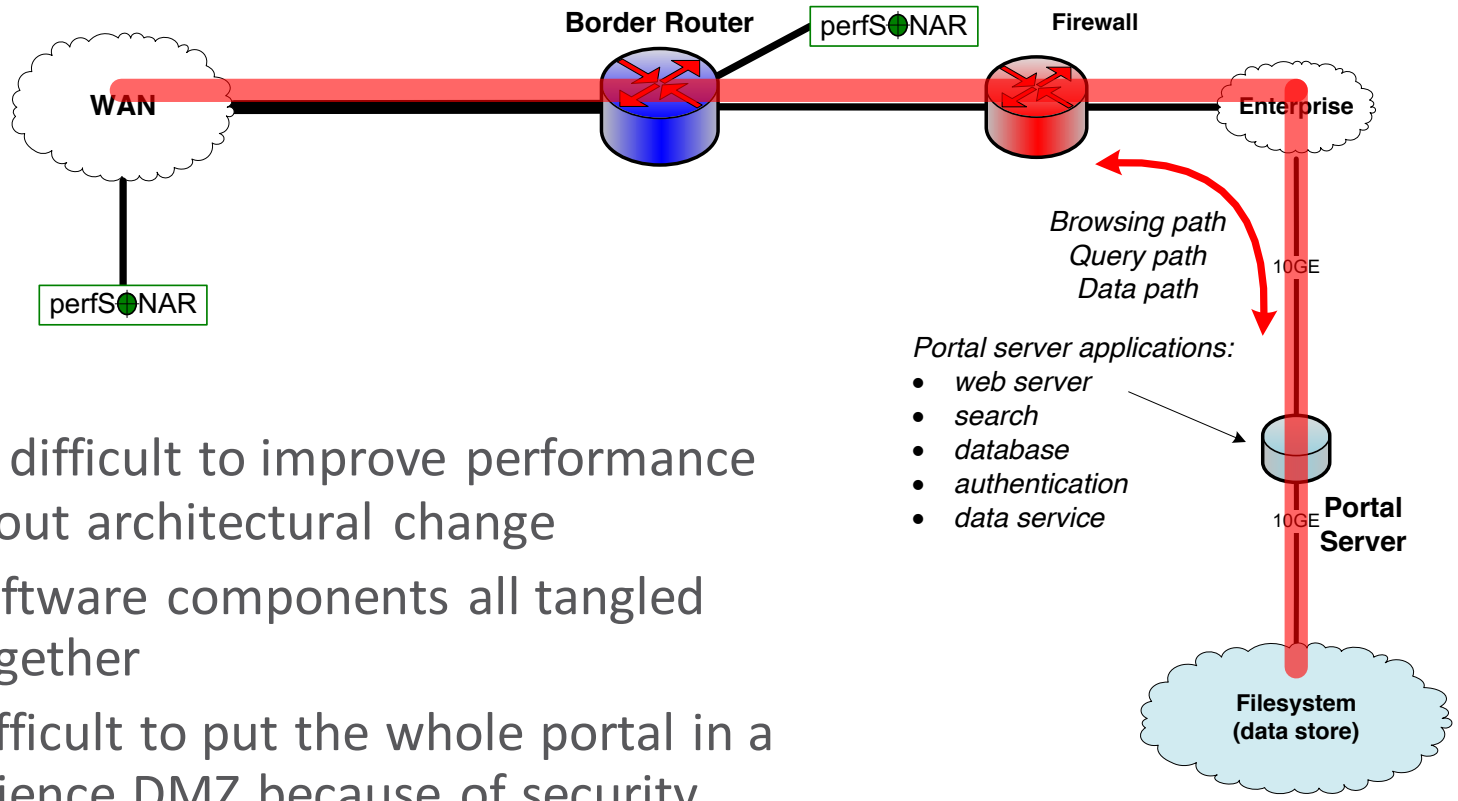

- Very difficult to improve performance without architectural change
  - Software components all tangled together
  - Difficult to put the whole portal in a Science DMZ because of security
  - Even if you could put it in a DMZ, many components aren't scalable
- What does architectural change mean?

ESnet

Example of Architectural Change – CDN

- Let's look at what Content Delivery Networks did for web applications
- CDNs are a well-deployed design pattern
  - Akamai and friends
  - Entire industry in CDNs
  - Assumed part of today's Internet architecture
- What does a CDN do?
  - Store static content in a separate location from dynamic content
    - Complexity isn't in the static content – it's in the application dynamics
    - Web applications are complex, full-featured, and slow
      - Databases, user awareness, etc.
      - Lots of integrated pieces
    - Data service for static content is simple by comparison
  - Separation of application and data service allows each to be optimized

ESnet

Classical Web Server Model

- Web browser fetches pages from web server
  - All content stored on the web server
  - Web applications run on the web server
    - Web server may call out to local database
    - Fundamentally all processing is local to the web server
  - Web server sends data to client browser over the network
- Perceived client performance changes with network conditions
  - Several problems in the general case
  - Latency increases time to page render
  - Packet loss + latency cause problems for large static objects

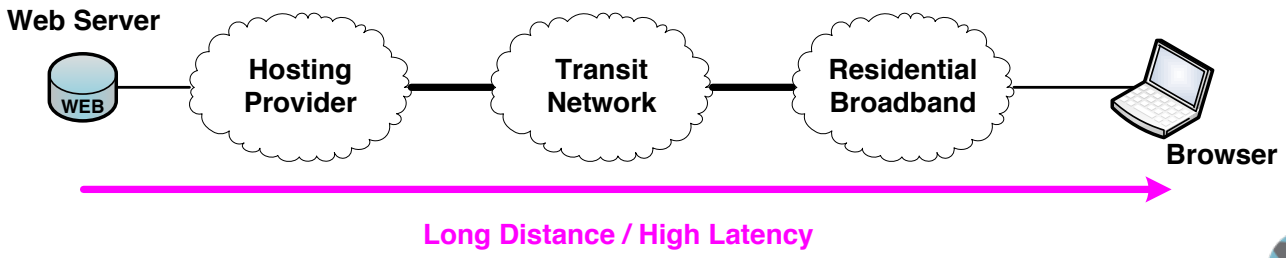

ESnet

Solution: Place Large Static Objects Near Client

- CDN provides static content "close" to client
  - Latency goes down
    - Time to page render goes down
    - Static content performance goes up
  - Load on web server goes down (no need to serve static content)
  - Web server still manages complex behavior
    - Local reasoning / fast changes for application owner
- Significant win for web application performance

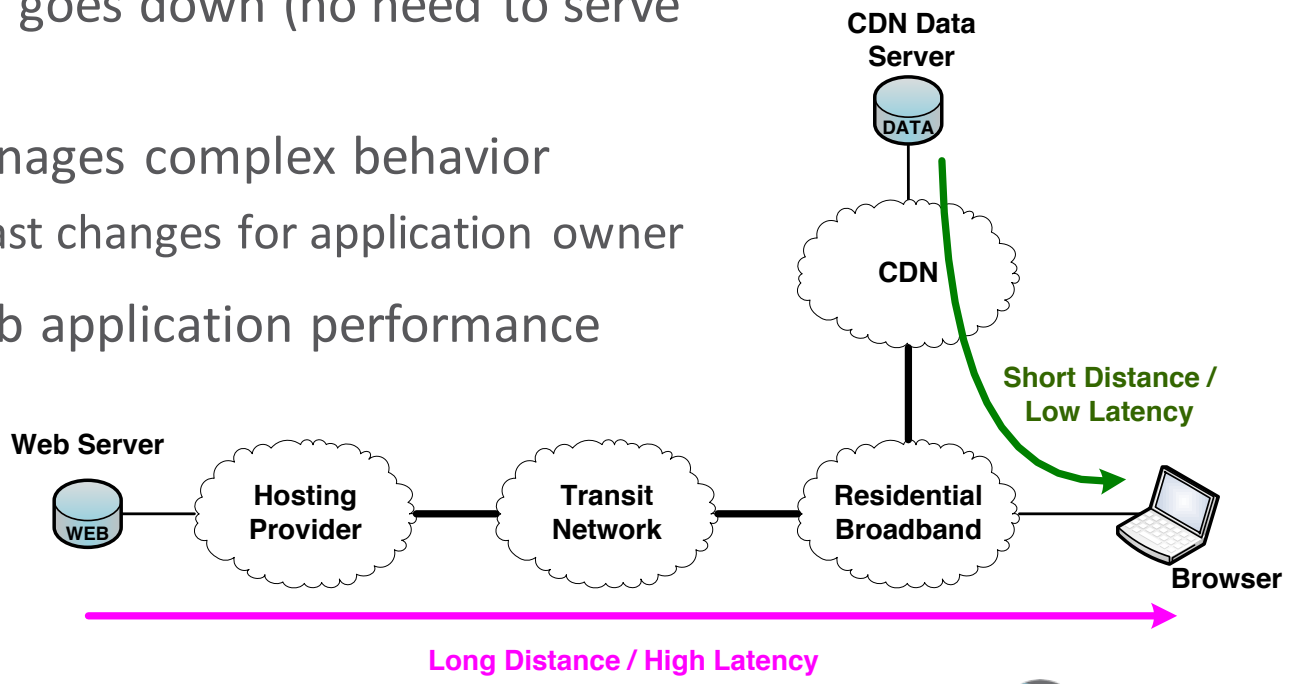

ESnet

Client Simply Sees Increased Performance

- Client doesn't see the CDN as a separate thing
  - Web content is all still viewed in a browser
    - Browser fetches what the page tells it to fetch
    - Different content comes from different places
    - User doesn't know/care
- CDNs provide an architectural solution to a performance problem
  - Not brute-force
  - Work smarter, not harder

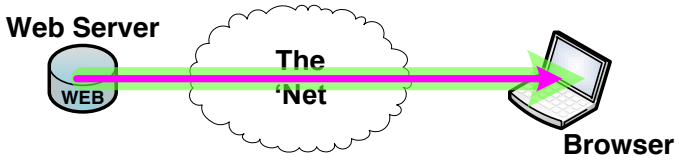

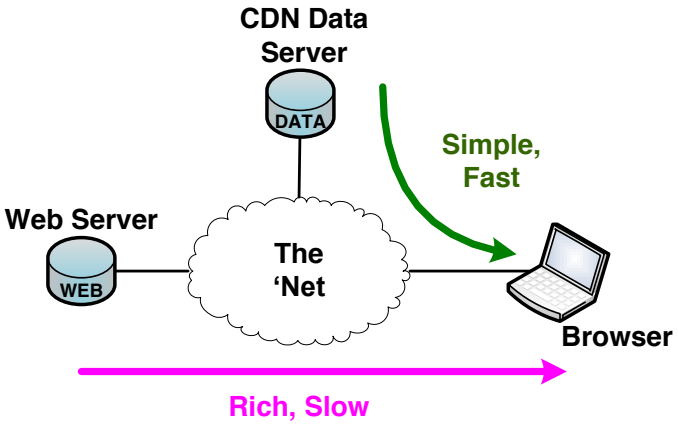

ESnet

Architectural Examination of Data Portals

- Common data portal functions (most portals have these)
  - Search/query/discovery
  - Data download method for data access
  - GUI for browsing by humans
  - API for machine access – ideally incorporates search/query + download
- Performance pain is primarily in the data handling piece
  - Rapid increase in data scale eclipsed legacy software stack capabilities
  - Portal servers often stuck in enterprise network
- Can we "disassemble" the portal and put the pieces back together better?
  - Use Science DMZ as a platform for the data piece
  - Avoid placing complex software in the Science DMZ

ESnet

Legacy Portal Design

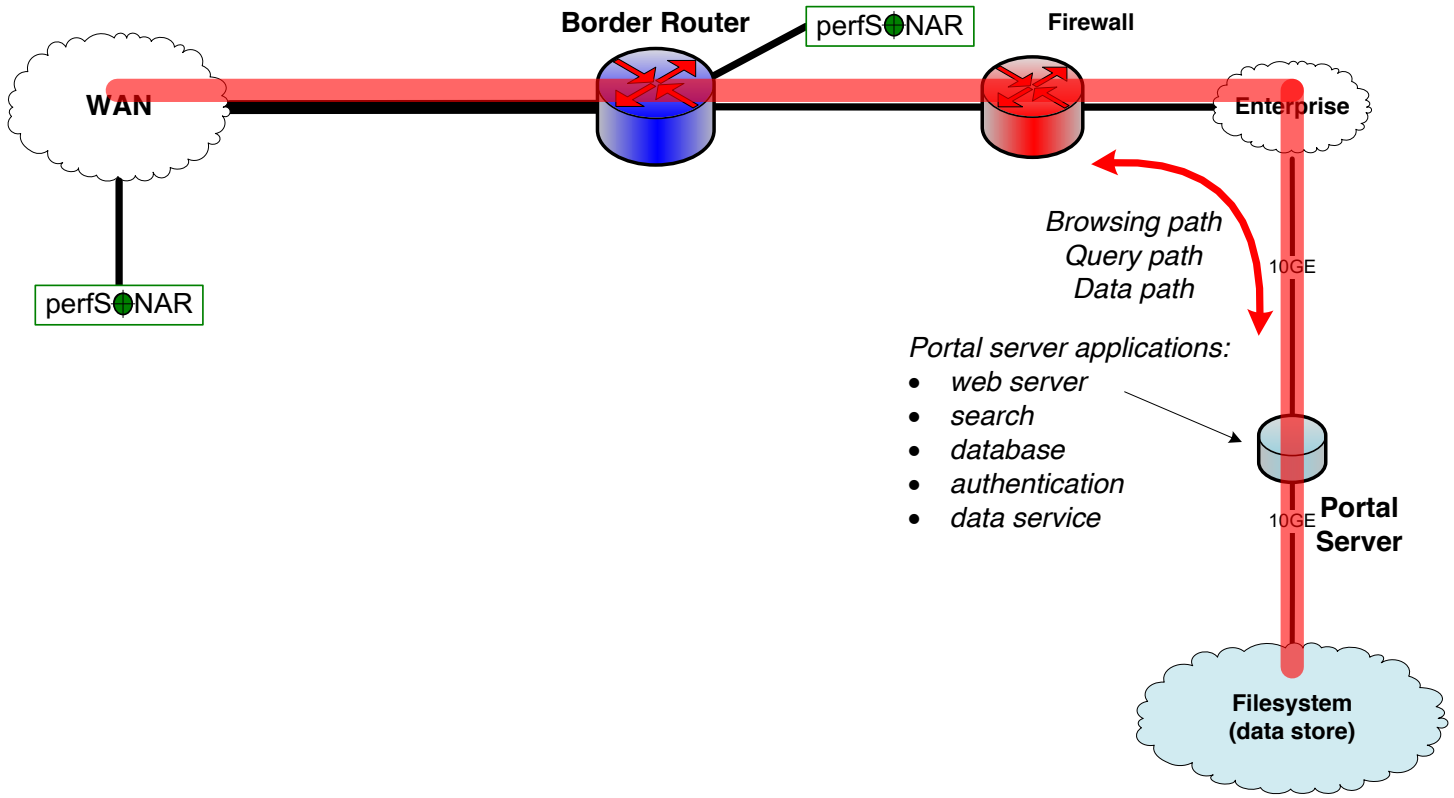

Next-Generation Portal Leverages Science DMZ

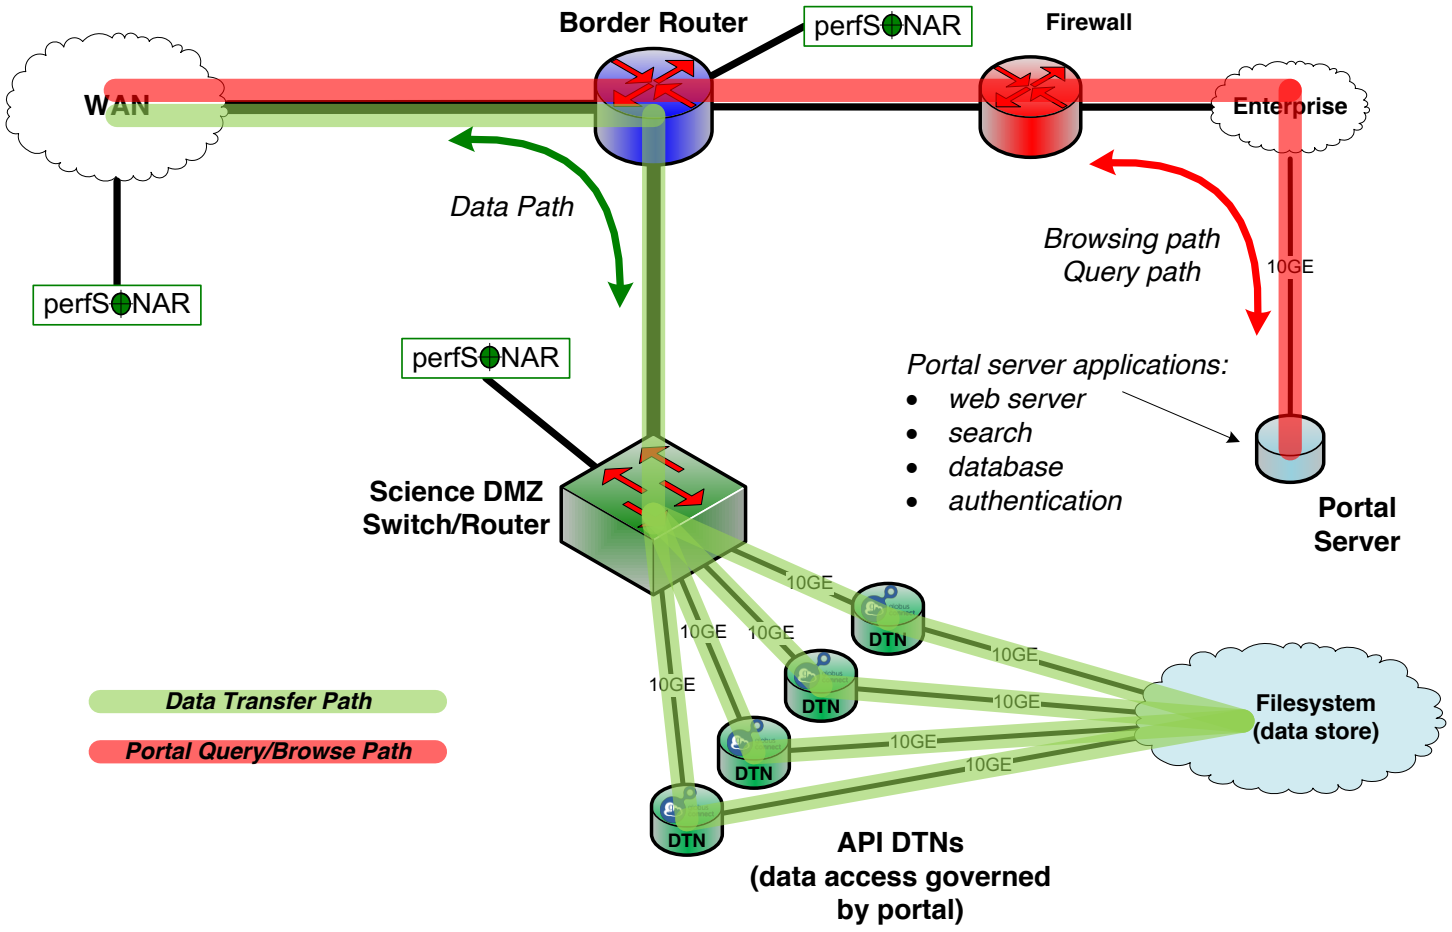

Put The Data On Dedicated Infrastructure

- We have separated the data handling from the portal logic
- Portal is still its normal self, but enhanced
  - Portal GUI, database, search, etc. all function as they did before
  - Query returns pointers to data objects in the Science DMZ
  - Portal is now freed from ties to the data servers (run it on Amazon if you want!)
- Data handling is separate, and scalable
  - High-performance DTNs in the Science DMZ
  - Scale as much as you need to without modifying the portal software
- Outsource data handling to computing centers
  - Computing centers are set up for large-scale data
  - Let them handle the large-scale data, and let the portal do the orchestration of data placement

ESnet

Ecosystem Is Ready For This

- Science DMZs are deployed at Labs, Universities, and computing centers
  - XSEDE sites
  - DOE HPC facilities
  - Many campus clusters
- Globus DTNs are present in many of those Science DMZs
  - XSEDE sites
  - DOE HPC facilities
  - Many campus clusters
- Architectural change allows data placement at scale
  - Submit a query to the portal, Globus places the data at an HPC facility
  - Run the analysis at the HPC facility
  - The results are the only thing that ends up on a laptop or workstation

ESnet

Links and Lists

- ESnet fasterdata knowledge base
  - http://fasterdata.es.net/
- Science DMZ paper
  - http://www.es.net/assets/pubs_presos/sc13sciDMZ-final.pdf
- Science DMZ email list
  - Send mail to sympa@lists.lbl.gov with subject "subscribe esnet-sciencedmz"
- perfSONAR
  - http://fasterdata.es.net/performance-testing/perfsonar/
  - http://www.perfsonar.net
- Globus
  - https://www.globus.org/

ESnet

ENERGY SCIENCES NETWORK

Thanks!

Eli Dart dart@es.net

Energy Sciences Network (ESnet)

Lawrence Berkeley National Laboratory

http://fasterdata.es.net/

http://my.es.net/

http://www.es.net/

Extra Slides

DTN Cluster Detail

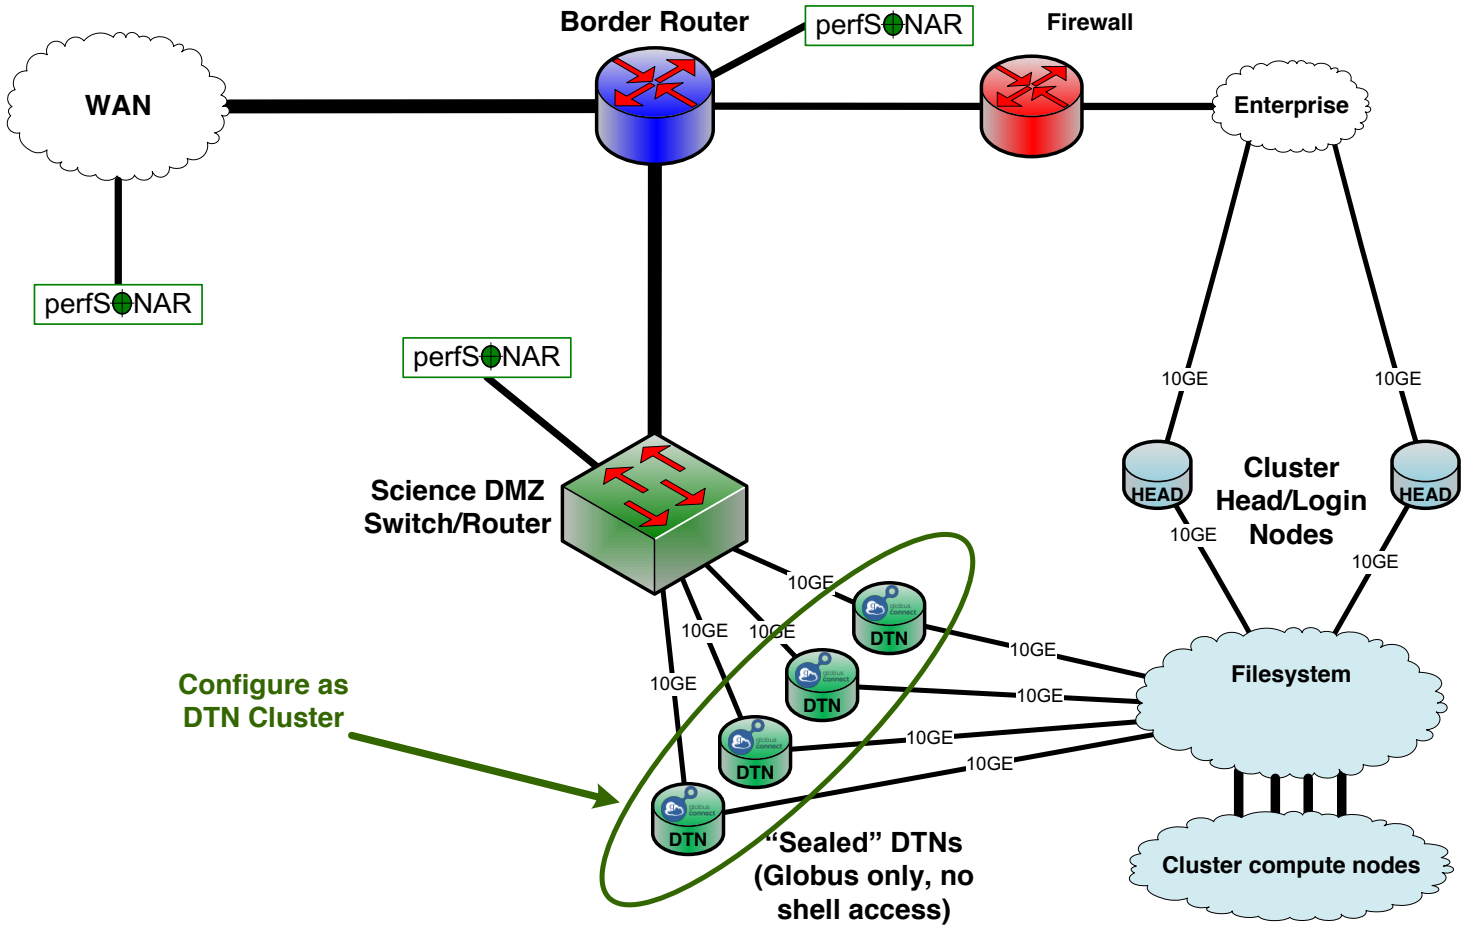

DTN Cluster Design

- Configure all four DTNs as a single Globus endpoint
  - Globus has docs on how to do this
  - https://support.globus.org/entries/71011547-How-do-I-add-multiple-I-O-nodes-to-a-Globus-endpoint-
- Recent options for increased performance
  - Use additional parallel connections
  - Distribute transfers across multiple DTNs (Globus I/O Nodes)
  - Critical – only do this when all DTNs in the endpoint mount the same shared filesystem
- Use the Globus CLI command **`endpoint-modify`**
  - Use the --network-use option
  - Adjusts concurrency and parallelism
  - More info at globus.org (http://dev.globus.org/cli/reference/endpoint-modify/)

ESnet

Security Footprint of a Globus Transfer

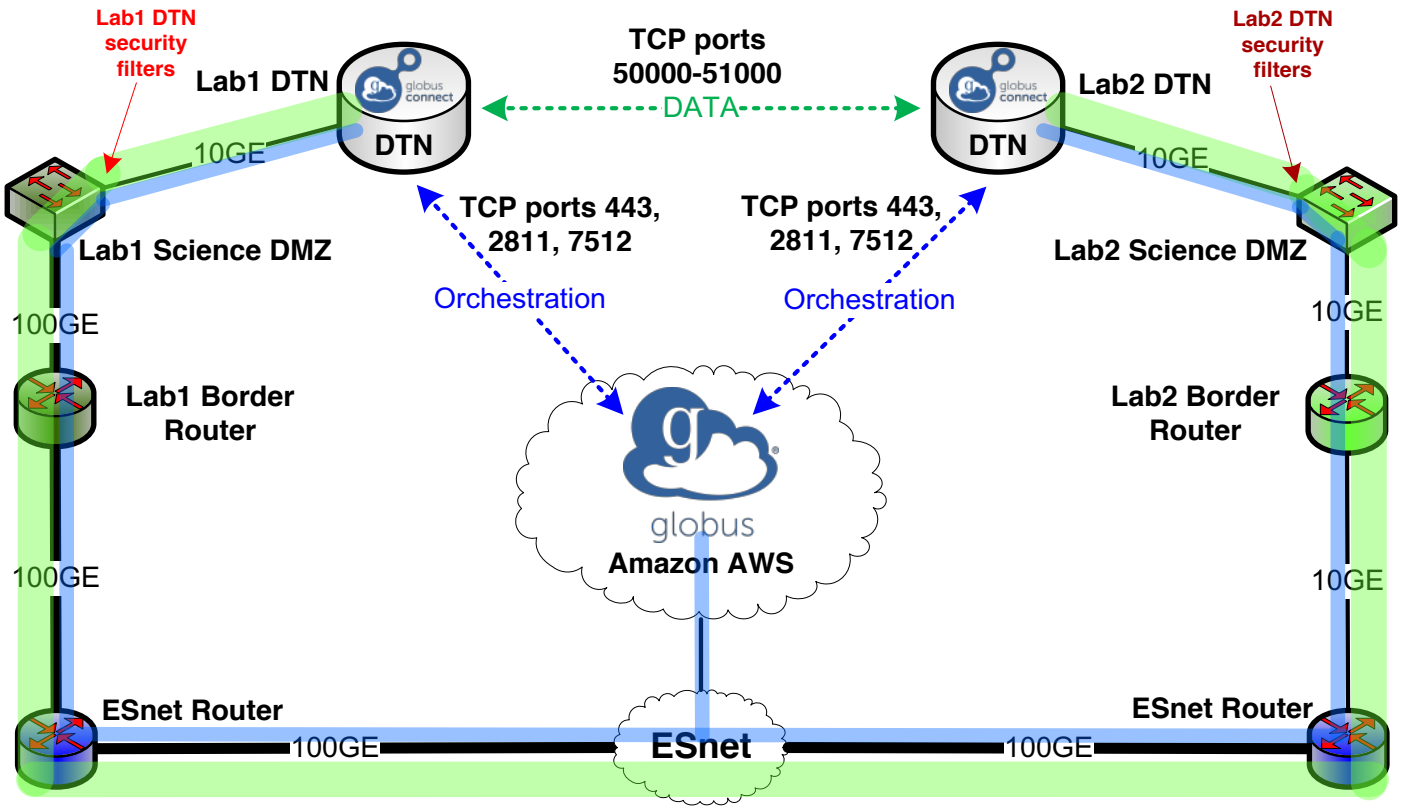

Logical data path

Physical data path

Logical control path

Physical control path

Lab1 DTN security filters

Src Address	Src Port	Dst Address	Dst Port
Lab1 DTN	TCP 50000-51000	Lab2 DTN	TCP 50000-51000
Lab1 DTN	TCP 443, 2811, 7512	Globus Cloud	TCP unprivileged
Lab2 DTN	TCP 50000-51000	Lab1 DTN	TCP 50000-51000
Globus Cloud	TCP unprivileged	Lab1 DTN	TCP 443, 2811, 7512

Lab2 DTN security filters

Src Address	Src Port	Dst Address	Dst Port
Lab2 DTN	TCP 50000-51000	Lab1 DTN	TCP 50000-51000
Lab2 DTN	TCP 443, 2811, 7512	Globus Cloud	TCP unprivileged
Lab1 DTN	TCP 50000-51000	Lab2 DTN	TCP 50000-51000
Globus Cloud	TCP unprivileged	Lab2 DTN	TCP 443, 2811, 7512

ESnet

Security Footprint of a Globus DTN

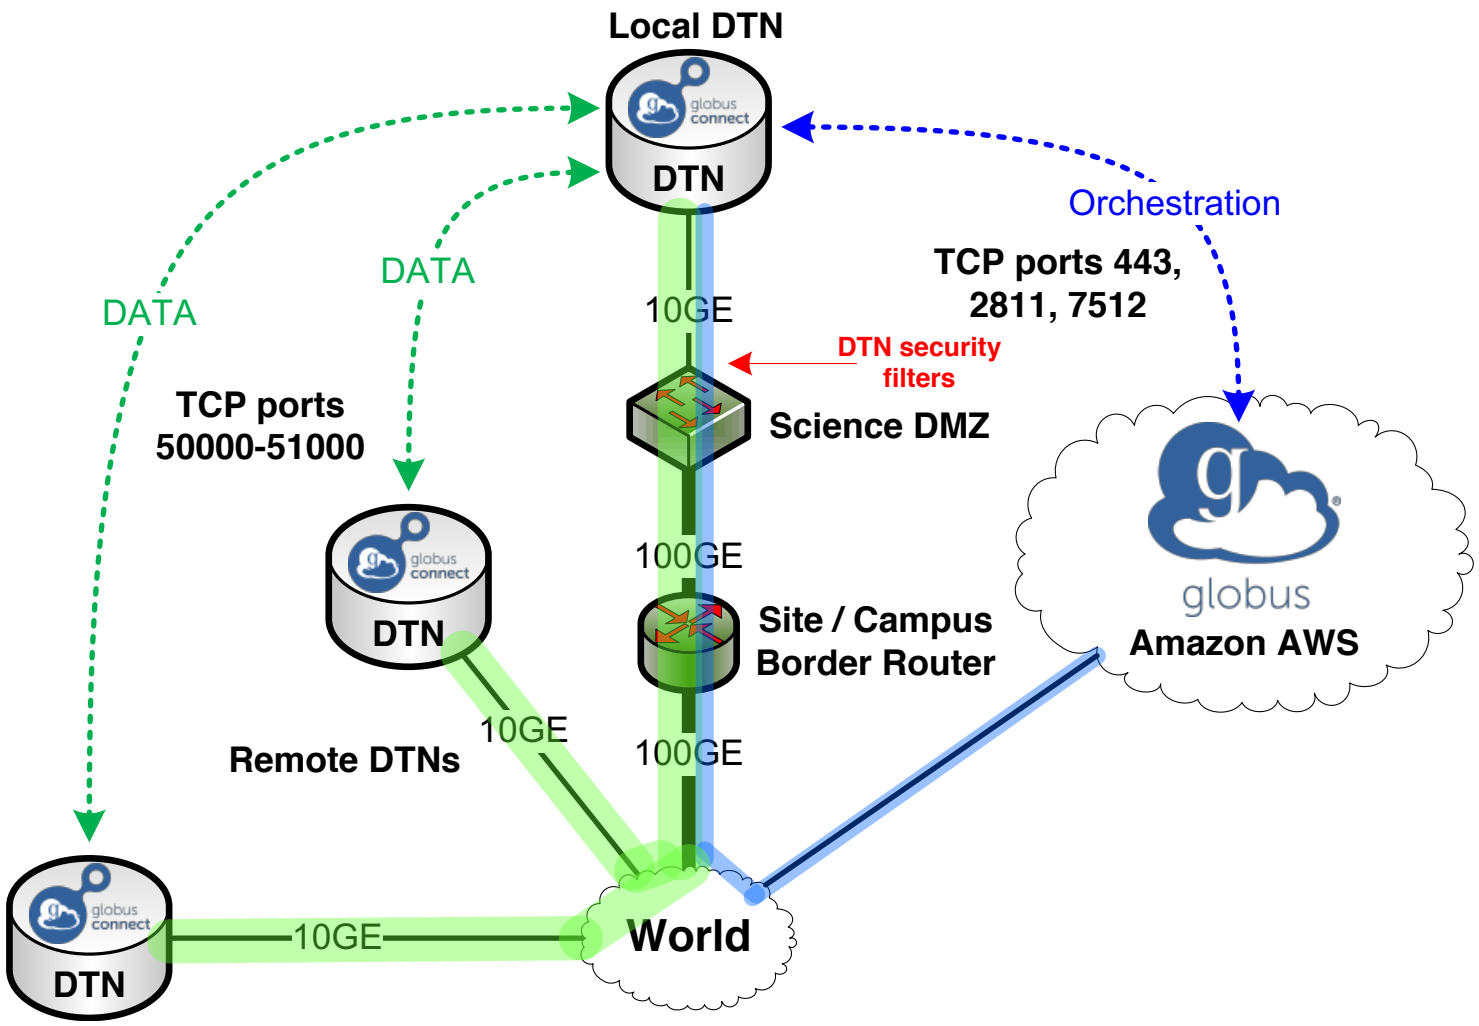

Src Address	Src Port	Dst Address	Dst Port
Local DTN	TCP 50000-51000	World	TCP 50000-51000
Local DTN	TCP 443, 2811, 7512	Globus Cloud	TCP unprivileged
World	TCP 50000-51000	Local DTN	TCP 50000-51000
Globus Cloud	TCP unprivileged	Local DTN	TCP 443, 2811, 7512