BROCADE

SOLUTION BRIEF

Software-Driven Science DMZ Networks

HIGHLIGHTS

- Enables fast data transfer speeds to support effective collaboration and accelerate research and innovation
- Maximizes available resources to extend scientific and research capabilities
- Protects intellectual property and sensitive research data while facilitating cross-functional research programs

Faster Collaboration and Improved Security: The Benefits of a Software-defined Science DMZ Network

Thanks to advancements in scientific instrumentation, the size of data sets and the volume of data generated by scientific research are increasing exponentially. Research is also becoming more collaborative and competitive, as organizations race to make discoveries that will advance their reputations and yield financial benefits. Labs, meanwhile, are under increasing pressure to find ways to quickly transfer massive amounts of data between researchers worldwide, while protecting intellectual property and ensuring system security.

The traditional campus networks used by these research organizations are not designed to support the fast transfer of such large data volumes. This leads to slowdowns and other issues that impact lab and research productivity, prevent effective collaboration, and jeopardize the integrity of research data. At the same time, large data volumes can overwhelm the infrastructure, hampering day-to-day operations and functions while potentially compromising security systems. Alternative technologies better suited to large data set transfers are often expensive and more than what is needed by most organizations, leading to low utilization rates. A Science DMZ network, however, offers the performance, scale, and security that research organizations require, while Software-Defined Networking (SDN) adds network automation and flexible flow manipulation to the solution.

A Network Architecture for Accelerated Innovation and Analysis

As data volumes surge, research organizations' networks are straining to deliver the required data transfer speeds. The U.S. Department of Energy, Energy Sciences Network (ESnet), estimates that it takes three hours to transfer 1 terabyte (TB) of data over a 1 Gbps network, assuming there are no other traffic flows competing for passage during this time.[1] Consider that the data center at CERN processes approximately 1 petabyte (PB)

[1] ESnet, "Network Requirements and Expectations: Expected Rates to Transfer Data," https://fasterdata.es.net/home/requirements-and-expectations.

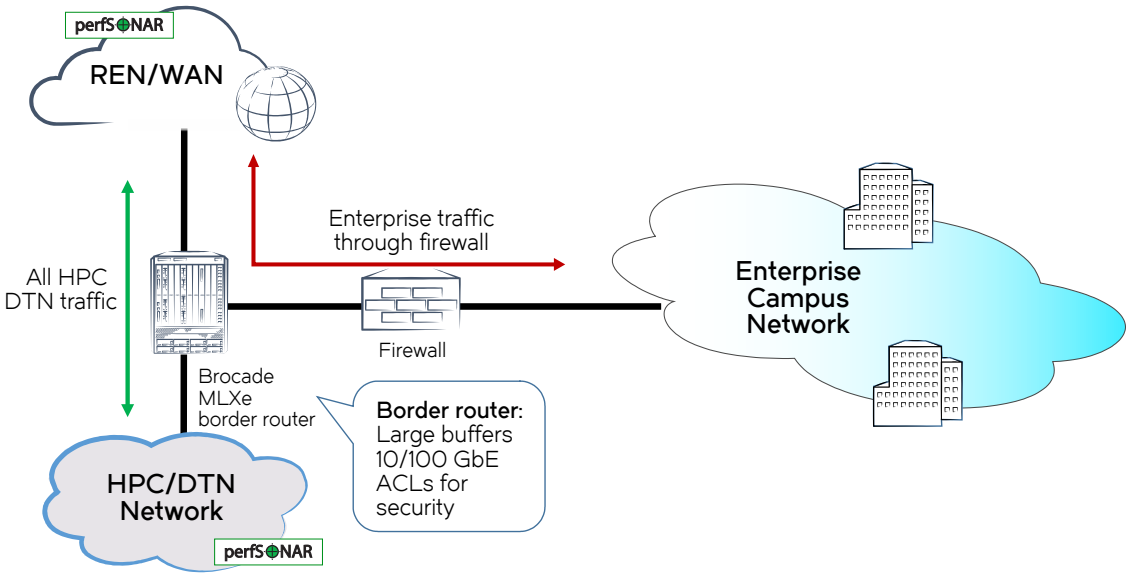

Figure 1: A typical Science DMZ network. The firewall is in-line with operational traffic, while the research data is provided a "fast path."

of data a day, while the Large Hadron Collider can generate 15 PB—or 15,000 TB every year.[2] While most research facilities will never generate data sets that large, this illustrates the challenge researchers face trying to share data or collaborate with researchers at other locations.

Additionally, when organizations use the same campus network for operations and for scientific research and collaboration, data transfers must generally pass through firewalls unsuited to such data flows. Large data flows have been found to move up to 20 times more slowly when traversing the firewall. Not only does this create additional delays to the data transfer, but it also can lead to the firewall failing or underperforming, increasing the risk that security will be compromised.

Researchers reliant on a single campus network may subsequently find data transfers take many hours (even days) and are likely to fail and require restarting. This negatively impacts researchers' productivity, ability to meet research schedules, adoption of new research innovations, and ability to collaborate on (and therefore participate in) significant research programs.

The Science DMZ Network

The solution is to build a separate, autonomous network specifically used for scientific research, a Science DMZ (see Figure 1). A Science DMZ network is designed to deliver a "friction free," high-performance, intelligent infrastructure. It houses dedicated Data Transfer Nodes (DTNs) that run software tools for optimizing massive data transfers.

In addition, network administrators can deploy a test and measurement infrastructure, such as perfSONAR, to monitor and ensure network performance.

Ideally, the Science DMZ network should be placed at the perimeter of the Local Area Network (LAN)—with a Science DMZ dedicated switch or router directly connected to a border router—to provide the most efficient and effective system design. A dedicated DTN should be connected directly to the Science DMZ switch or router.

To deliver the performance required to support Big Data transfers, the Science DMZ switch and border routers must meet some key criteria. Buffer space and intelligent optimization of buffering, combined with high-density (100 GbE/10 GbE) port connections,

[2] CERN, "The Worldwide LHC Computing Grid," http://home.cern/about/computing/worldwide-lhc-computing-grid.

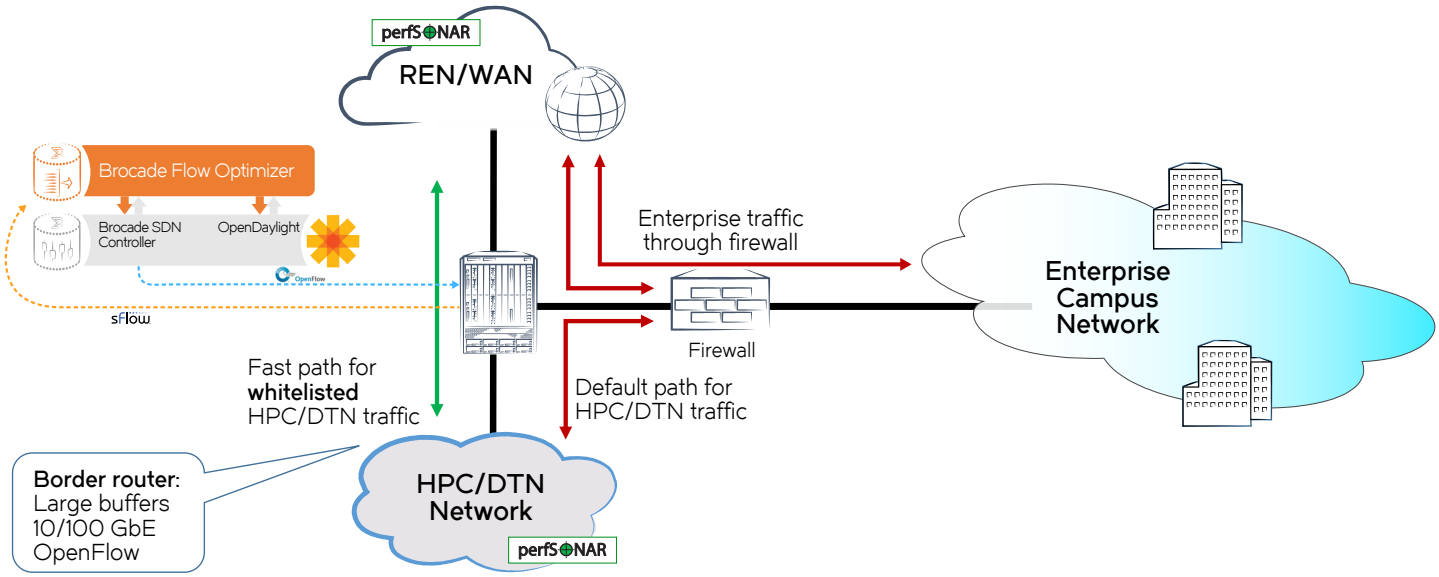

Figure 2: An example of a software-defined Science DMZ network architecture. The firewall protects research-related and operational traffic, while the "whitelisted" traffic is directed around the firewall.

are critical to maintaining line-rate speeds. Organizations such as ESnet recommend using devices that support OpenFlow in order to also enable Software-Defined Networking (SDN). Using a Brocade® MLX® Series router as a border router ensures that these core criteria are met, while offering flexible Access Control List (ACL) functionality for stateless security.

A Software-defined Science DMZ Network

Research organizations must support collaboration between researchers, but also protect intellectual property and sensitive data. They are therefore concerned about moving data to host sites outside the firewall, and how this practice may compromise overall security. One way to address this data transfer challenge is to use ACLs and Policy-Based Routing (PBR) configurations within the Science DMZ network to bypass the firewall, mostly during data transfer. However, this can be a complex, sometimes manual, process with potentially serious security implications, and it requires keeping the ACL up to date at all times.

A software-defined Science DMZ network offers additional layers of flexibility and security to the DMZ, by automating processes and enforcing preset policies related to specific data flows. The software-defined Science DMZ network can be built using an SDN-ready hardware infrastructure, an SDN controller, and an application that works in conjunction with the controller to identify and optimize trusted data transfer flows. The Brocade SDN Controller, an SDN-ready switch (such as the Brocade MLXe), and the Brocade Flow Optimizer application are ideally suited for the software-defined Science DMZ network. This integrated solution offers ease of use and flexibility that has been fully tested in high-performance environments around the world.

In a software-defined network environment, research support staff can pre-program policies, which the switch will then enforce with OpenFlow via the border router. These policies determine which data transfer flows can be directed around the firewall, increasing transfer speed and preventing the firewall from becoming overwhelmed and compromised. Policy management and enforcement are thus largely automated, creating trusted flows across a predefined path, so researchers can transfer large data volumes without requiring additional IT resources or support.

Once deployed, the Science DMZ network can increase data transfer speeds by a factor of 20,[3] improving collaboration and productivity. Resources, instruments, and the network can be utilized more efficiently because transfers do not have to be preplanned or limited to running during specific periods (such as overnight).

The Science DMZ network and campus network remain protected behind the firewall, reducing the risk that security will be compromised or challenged. Moreover, the firewall itself remains protected and optimized to focus on smaller, general flows. At the same time, this approach allows for early adoption and testing of SDN, which can provide benefits to the entire IT organization.

The best way to deploy a software-defined Science DMZ depends on an organization's network topology (see Figure 2). A hybrid port for the SDN-ready switch is required when the SDN DMZ is not deployed on an OpenFlow-only network.

Conclusion

Growing numbers of labs and research institutions worldwide are choosing to deploy Science DMZ networks to meet their data transfer performance, scale, and security requirements. Many of the early adopters have been larger facilities, which have the financial resources to invest in infrastructure to support their evolving science and research programs. The next wave of adopters will likely consist of midsized facilities, where collaboration among different research groups is critical.

Software-defined Science DMZ networks automate many of the tasks required for secure and efficient data transfer, reducing the resource burden on IT and enabling researchers to collaborate and access results and analyses faster. Security is maintained by protecting and optimizing standard-use firewalls. By replacing some network functions with software-based or virtual solutions, research institutions can benefit from a lower total cost of ownership and maximize their infrastructure investments.

Learn More

Brocade partners with companies of all sizes to deliver innovative solutions that help organizations maximize the value of their most critical information. To learn more, visit www.brocade.com.

About Brocade

Brocade networking solutions help organizations transition smoothly to a world where applications and information reside anywhere. Innovative Ethernet and storage networking solutions for data center, campus, and service provider networks help reduce complexity and cost while enabling virtualization and cloud computing to increase business agility. Learn more at www.brocade.com.

[3] ESnet, "ESnet Fasterdata Knowledge Base," https://fasterdata.es.net.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE