Securing Large-scale Data Transfers in Campus Networks: Experiences, Issues, and Challenges

(Invited Paper)

Deepak Nadig

Dept. of Computer Science and Engineering
University of Nebraska-Lincoln
Lincoln, NE, USA
deepaknadig@cse.unl.edu

Byrav Ramamurthy

Dept. of Computer Science and Engineering
University of Nebraska-Lincoln
Lincoln, NE, USA
byrav@cse.unl.edu

ABSTRACT

Increasingly, campus networks manage a multitude of large-scale data transfers. Big data plays a pivotal role in university research and impacts domains such as engineering, agriculture, natural sciences, and humanities. Over the years, numerous solutions have been proposed to manage and secure large-scale data transfers efficiently. Examples consist of the inclusion of security policies at the network edge, optimized middlebox management, and the Science Demilitarized Zone (Science DMZ). These solutions either severely degrade data transfer performance or result in data flows completely bypassing the campus network security controls. In this paper, we present our experience with the design, development, and management of large-scale data transfers using software defined networking (SDN) and network functions virtualization (NFV). We also discuss the issues and challenges associated with securing large-scale data transfers in campus networks.

CCS CONCEPTS

• **Security and privacy** → **Network security**; • **Networks** → **Network security**; *Network architectures*; *Network management*; *Network monitoring*.

KEYWORDS

Data-intensive Science; Security; Application-Awareness; Software Defined Networks; Network Functions Virtualization.

ACM Reference Format:

Deepak Nadig and Byrav Ramamurthy. 2019. Securing Large-scale Data Transfers in Campus Networks: Experiences, Issues, and Challenges: (Invited Paper). In *ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFVSec '19)*, March 27, 2019, Richardson, TX, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3309194.3309444>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SDN-NFVSec '19, March 27, 2019, Richardson, TX, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6179-8/19/03...\$15.00

<https://doi.org/10.1145/3309194.3309444>

1 INTRODUCTION

Typically, campus network infrastructures support multiple network traffic management goals. Commonly supported network traffic types include the following: i) commodity internet traffic to support routine business tasks such as email and web; ii) high-performance networks for tasks associated with the scientific research process such as data retrieval, storage, and analysis from external sources; and iii) specialized (possibly separate) networks to meet specific project/workflow requirements like confidentiality, privacy, and anonymity. These goals often impose conflicting requirements on network design and management. For example, networks designed for routine tasks such as email or web browsing are incapable of supporting large-scale data movement for data-intensive science. In this context, networks optimized and specially engineered for data-intensive tasks are necessary, without which, significant performance degradations would result. Further, many aspects of campus networks are hard to change without impacting normal network operation.

Campus networks, traditionally, are designed to handle a large number of small traffic flows and are ill-suited for high-volume science data transfers. Sharing campus network resources with scientific data transfers has many disadvantages. First, packet loss that is tolerated by campus network LANs (local area networks) generally lead to severe performance degradations with bulk data transfers. Second, the wide area network (WAN) latency effects come into play reducing TCP performance. Lastly, campus network security controls optimize the security of small flows while sacrificing performance and therefore do not meet the stringent requirements for large data transfers over the WAN. SDN and NFV pave the way for novel traffic engineering methods that can overcome the limitations of traditional campus networks. SDN and NFV architectures [13] are used successfully in big data environments for network management [17], QoS provisioning [12], and network design [8].

In this paper, we look at existing architectures for data-intensive science, their implementation, strengths, and weaknesses. Next, we present our experiences and insights on implementing SDN and NFV-based solutions for data-intensive science at a worldwide LHC (large hadron collider) computing grid (WLCG) Tier-2 site. Our solutions span network monitoring, network management, service differentiation, and security for data-intensive science. Lastly, we look at the current issues, emerging challenges, and significant research opportunities arising from scientific data management.

The paper is structured as follows: Section 2 provides the context for large-scale data transfers in data-intensive science. In Section 3, we discuss exemplary architectures for data-intensive science, their design goals, impacts, and shortcomings. In Section 4, we briefly outline our experience in managing and securing large data transfers in campus networks. We also discuss the current issues and emerging challenges in securing large data transfers in campus networks in Section 5. Lastly, in Section 6, we conclude our work.

2 DATA-INTENSIVE SCIENCE AND BIG DATA

Data-intensive science relies on massive amounts of data obtained from scientific studies (both experimental studies and simulation). Many science domains are data-intensive and data-driven. They rely on data acquisition, storage, and analysis for advancing fundamental research. Example data-intensive science domains are high energy physics (HEP), climate sciences, biology/genomics, combustion and light sources. In this section, we briefly summarize the big data challenges of different science domains.

High energy physics (HEP) is data-intensive, as advances in this domain require measuring the probability of “interesting events” in a large set of observations (typically 10^{16} or more particle collision observations in a year). The CMS (compact muon solenoid) [7] and ATLAS (A Toroidal LHC ApparatuS) [1] experiments at the large hadron collider (LHC) generate massive amounts of data (~ 10 petabytes per year). The ATLAS experiment currently stores over 100 petabytes of data and it is growing rapidly. The HEP community faces two important challenges namely: i) real-time data reduction, and ii) distributed data processing, analysis, and distribution. The massive volume of analog data generated by the CMS and ATLAS detectors must be reduced in real-time to ensure that it can be stored at an acceptable cost. Furthermore, the data must be shared/distributed worldwide with scientists for further processing and analysis.

Climate science is another example of a data-intensive science domain that critically depends on large amounts of data available from heterogeneous sources worldwide. Sophisticated models of physical processes and interaction between realms (e.g., atmosphere, land, sea) are used by collaborative, multi-disciplinary teams to achieve scientific progress. Example data sources include advanced radiation measurement (ARM) sites and earth observing system (EOS) of satellites. The data from different sources is subject to (near) real-time processing and made available to numerous communities such as scientists, industry, policy-makers, etc.

Other data-intensive domains include the advanced photon source (APS) at Argonne national labs (ANL) and the linac coherent light source (LCLS) facility at SLAC National Accelerator Laboratory. Light sources are used by scientists from many domains ranging from material science to paleontology. Similar to the above domains, there is a strong need for collaborative and high-performance ecosystem at these light source facilities.

Popular services/frameworks for bulk data movement supporting scientific users include: i) The GlobusOnline service built with the GridFTP [3] protocol, and ii) the XROOTD [10] framework. These protocols and services integrate with data transfer nodes (DTNs) and high-speed parallel filesystems such as Lustre or GPFS. The DTNs are purpose-built for bulk WAN data transfers and typically

have high-speed network interfaces (40/100Gbps). In the following section, we describe some exemplary architectures for data-intensive science. These architectures integrate many networking components including SDN, virtual circuits, DTNs, policy frameworks, etc. We also describe the architectural approach for securing high-volume transfers associated with data-intensive science.

3 EXEMPLARY ARCHITECTURES

Many architectures have been proposed to support high-performance data transfers for science applications. These architectures try to meet both operational and security requirements while ensuring simplicity and scalability. Some example architectures are Science DMZ [9], SciPass [6], ScienceSDS [15], Medical Science DMZ [16], etc.

Figure 1: The Science DMZ Architecture (Source: ESNet).

The science DMZ [9] is a scalable network architecture for data-intensive science, built at or near an organization’s local network perimeter. The architecture was developed by ESnet to address the network performance problems associated with high-volume, bulk data transfer needs of science applications. The Science DMZ integrates four key components including a network architecture distinct from general-purpose networks, a dedicated system for data transfers, a performance measurement and network testing system, and policies and enforcement mechanisms for security. The architecture is shown in Figure 1. In this architecture, the DTN is connected directly to the border router. A crucial distinction of this approach is that the offsite data traverses only two devices, the border router, and the Science DMZ switch/router. The DTN is responsible for efficient WAN data movement, and access control lists enforce security policies for the DTN, without the need for a separate firewall.

SciPass [6] is a 100Gbps capable SDN-based approach to Science DMZ. It combines an adaptive intrusion detection system (IDS) and dynamic traffic filtering to overcome the challenges posed by campus network firewalls. SciPass combines an SDN controller, a cluster of IDS sensors, a PerfSONAR host, and a firewall with a DTN to identify and filter flows. Policy definitions are used to determine whether flows can bypass the campus firewall controls.

Recently, a medical science DMZ was proposed by the authors in [16]. The proposed solution enables secure, high-volume data transfers of sensitive datasets between organizations while adhering to security and privacy regulations such as Health Insurance Portability and Accountability Act (HIPAA) [11]. The medical science DMZ re-engineers the science DMZ for use with restricted data by creating a secure, high-capacity, data-intensive portion within each organization. The trust relationships at each organization are used to ensure secure data movement to the appropriate computing facility.

Although these architectures provide adequate security and improve network performance for large-volume data transfers, they have some limitations. The science DMZ approach requires networks that are free of packet loss and stateful packet processing. While science DMZ performs selective monitoring and inspection of flows, it is not sufficient to detect all attacks. SciPass relies on IDS to identify flows, and can, therefore, be ineffective when end-to-end data transfers are encrypted. In the following section, we share our experiences in solving these shortcomings. Our solutions are deployed for monitoring, managing and securing large-scale data transfers at a U.S. CMS Tier-2 site.

4 EXPERIENCES AT HOLLAND COMPUTING CENTER @ UNL

As described in Section 3, existing architectures for securing data-intensive science rely on stateful packet processing techniques like deep packet inspection (DPI) or intrusion detection system (IDS) to classify packets/flows. These techniques are ineffective and place an undue burden on middleboxes due to the massive amounts of data transferred between end-points. Further, stateful packet processing can result in packet losses and increased detection latencies, thereby reducing the reliability of such techniques.

We take a unique approach to solving the above problems while minimizing the performance degradation effects of network middleboxes. We develop an application-aware architecture to facilitate the secure exchange of application metadata with the network-layer. Our experience with deploying application-aware solutions for network monitoring, management and security have shown not only the benefits but also the reliability of our approach. In the following, we outline our experiences using application-awareness for monitoring, managing and securing data-intensive science workflows.

4.1 Secure and Reliable Flow Identification using Application-Awareness

As relying on stateful packet inspection for identifying large-volume flows is ineffective, we proposed SDN-managed network architecture for GridFTP transfers (SNAG) [5]. SNAG, unlike other approaches, relies on the exchange of application-layer metadata with the network-layer to provide flow intelligence to the network management system. Application-awareness, from the GridFTP servers in this case, is enabled by exposing a secure application program interface (API) from a trusted GridFTP process. The GridFTP server automatically forwards application metadata to the SDN controller. All GridFTP connection metadata including source/destination addresses, port pairs, session information, etc., are used by

the network layer to identify and classify all active flows reliably. This flow identification and classification is vital to the network operator to ensure that appropriate security policies and enforcement mechanisms are defined.

4.2 Large-scale Data Transfer Monitoring, Management and Service Differentiation

SNAG, using an application-aware approach to traffic classification, helped us develop a unique monitoring system for data-intensive science. Our SNAG-enabled monitoring system can monitor, in real-time, CMS and LIGO (laser interferometer gravitational-wave observatory) flows. Application-layer metadata from the GridFTP servers are used to create fine-grained monitoring views. These views provide additional insights into data-intensive science transfers such as user/project, workflow, and file statistics. The monitoring views created using application-awareness *cannot* be obtained using traditional methods that use stateful packet processing. This is due to the encrypted nature of GridFTP connections. At the Holland Computing Center (HCC), UNL, we have used SNAG to classify over *800 million* connections successfully.

Further, we have used the traffic classification information provided by SNAG to achieve service differentiation [4]. For example, flows identified by SNAG can be individually subject to different quality of service (QoS) policies. Through the secure exchange of application metadata with the network layer, intelligent and adaptive network management decision-making is possible. Also, we can exploit application-layer metadata to define security policies for application flows more accurately.

4.3 An SDN/NFV Approach to Securing Data-Intensive Science Transfers

We can also use the SNAG application-aware architecture to secure data-intensive science workflows. Using the traffic classification information, we developed ScienceSDS [15]. ScienceSDS is a software-defined security framework that employs a virtualized security services infrastructure to provide security-as-a-service for data-intensive science. Here, we use the traffic classification information to redirect only flows of interest intelligently to virtualized security services. Unlike traditional middleboxes and security appliances, our approach using virtualized services is highly scalable. The SNAG application-aware approach also forms an important component of our GridFTP transfer anomaly detection system proposed in the work in [14].

5 ISSUES AND EMERGING CHALLENGES

Although existing network architectures described in this paper provide adequate security and balance the stringent requirements of high-volume science transfers, many issues and challenges remain. Firstly, these architectures represent a radically different design philosophy from the traditional campus networks, inevitably leading to the creation and management of multiple networks. With campus networks managing separate commodity internets and research networks, operators are faced with increased complexity in managing, provisioning, and securing the network. In this context, we have to analyze carefully, the impacts of the resulting network

split such as latency/loading issues, under-utilization, and associated technical challenges. Challenges also exist in evaluating novel solutions for data-intensive science; the complex nature of these networks make them less amenable to at-scale testing and experimentation. Further, for architectures designed to manage restricted and privacy-sensitive information such as the medical science DMZ, their “risk-management” abilities must be carefully evaluated.

Numerous research opportunities exist in developing novel approaches to securing data-intensive science. As more and more science data transfers utilize high-speed optical transports, novel approaches to both optimizing and securing multi-layer software-defined optical networks (SDONs) become essential. Further, numerous opportunities exist in extending and evaluating the suitability of future internet architectures for data-intensive science. Specifically, content-centric or information-centric networking architectures (such as Named Data Networking (NDN) [2, 18]) may be well-suited for this purpose.

6 CONCLUSIONS

In this paper, we have presented our experiences, issues, and challenges in securing large-scale data transfers in campus networks. Data-intensive science continues to play an essential role in university research impacting multiple domains. We presented an overview of the most active data-intensive science domains, along with the network architectures designed to secure science data transfers. Our experience at Holland Computing Center, UNL, a U.S. CMS Tier-2 site, has shown that application-aware approaches are beneficial to network management and security. However, limitations remain, and the data-intensive science community has to overcome many challenges across computing, storage, data movement, and analysis as it prepares for exascale computing.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant Numbers OAC-1541442 and CNS-1817105. This work was completed using the Holland Computing Center of the University of Nebraska, which receives support from the Nebraska Research Initiative. The authors would like to thank David Swanson, Brian Bockelman, and Garhan Attebury, Holland Computing Center at UNL for their valuable support and discussions.

REFERENCES

- [1] Georges Aad, JM Butterworth, J Thion, U Bratzler, PN Ratoff, RB Nickerson, JM Seixas, I Grabowska-Bold, F Meisel, S Lokwicz, et al. 2008. The ATLAS experiment at the CERN large hadron collider. *Jinst* 3 (2008), S08003.
- [2] M. Alhowaidi, B. Ramamurthy, B. Bockelman, and D. Swanson. 2017. The Case for Using Content-Centric Networking for Distributing High-Energy Physics Software. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. 2571–2572. <https://doi.org/10.1109/ICDCS.2017.295>
- [3] W. Allcock, J. Bresnahan, R. Kettimuthu, et al. 2005. The Globus Striped GridFTP Framework and Server. In *Supercomputing, 2005. Proc of the ACM/IEEE SC 2005 Conference*. 54–54. <https://doi.org/10.1109/SC.2005.72>
- [4] D. N. Anantha, B. Ramamurthy, B. Bockelman, et al. 2017. Differentiated network services for data-intensive science using application-aware SDN. In *2017 IEEE ANTS*. 1–6. <https://doi.org/10.1109/ANTS.2017.8384105>
- [5] Deepak Nadig Anantha, Zhe Zhang, Byrav Ramamurthy, et al. 2016. SNAG: SDN-managed Network Architecture for GridFTP Transfers. In *3rd Innovating the Network for Data-Intensive Science (INDIS) '16*, SC16. 10.
- [6] Edward Balas and A Ragusa. 2014. SciPass: a 100Gbps capable secure Science DMZ using OpenFlow and Bro. In *Supercomputing 2014 conference (SC14)*.
- [7] D. Bonacorsi. 2007. The CMS Computing Model. *Nuclear Physics B - Proceedings Supplements* 172 (2007), 53 – 56. <https://doi.org/10.1016/j.nuclphysbps.2007.07.030>
- [8] L. Cui, F. R. Yu, and Q. Yan. 2016. When big data meets software-defined networking: SDN for big data and big data for SDN. *IEEE Network* 30, 1 (January 2016), 58–65. <https://doi.org/10.1109/MNET.2016.7389832>
- [9] Eli Dart, Lauren Rotman, Brian Tierney, et al. 2014. The Science DMZ: A network design pattern for data-intensive science. *Scientific Programming* 22, 2 (2014), 173–185.
- [10] Alvise Dorigo, Peter Elmer, Fabrizio Furano, and Andrew Hanushevsky. 2005. XROOTD-A Highly scalable architecture for data access. *WSEAS Transactions on Computers* 1, 4.3 (2005).
- [11] Centers for Disease Control, Prevention, et al. 2003. HIPAA privacy rule and public health. Guidance from CDC and the US Department of Health and Human Services. *MMWR: Morbidity and mortality weekly report* 52, Suppl. 1 (2003), 1–17.
- [12] S. Jain, M. Khandelwal, A. Katkar, et al. 2016. Applying big data technologies to manage QoS in an SDN. In *2016 12th Conference on Network and Service Management (CNSM)*. 302–306. <https://doi.org/10.1109/CNSM.2016.7818437>
- [13] I. Monga, E. Pouyoul, and C. Guok. 2012. Software-Defined Networking for Big-Data Science - Architectural Models from Campus to the WAN. In *2012 SC Companion: High Performance Computing, Networking Storage and Analysis*. 1629–1635. <https://doi.org/10.1109/SC.Companion.2012.341>
- [14] Deepak Nadig, Byrav Ramamurthy, Brian Bockelman, and David Swanson. 2018. Identifying Anomalies in GridFTP Transfers for Data-Intensive Science Through Application-Awareness. In *Proceedings of the 2018 ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization (SDN-NFVSec'18)*. ACM, New York, NY, USA, 7–12. <https://doi.org/10.1145/3180465.3180469>
- [15] Deepak Nadig Anantha and Byrav Ramamurthy. 2017. ScienceSDS: A Novel Software Defined Security Framework for Large-scale Data-intensive Science. In *ACM Security in Software Defined Networks & Network Function Virtualization (SDN-NFVSec '17)*. ACM, 13–18. <https://doi.org/10.1145/3040992.3040999>
- [16] Sean Peisert, Eli Dart, William Barnett, Edward Balas, James Cuff, Robert L. Grossman, Ari Berman, Anurag Shankar, and Brian Tierney. 2017. The medical science DMZ: a network design pattern for data-intensive medical science. *Journal of the American Medical Informatics Association* 25, 3 (10 2017), 267–274. <https://doi.org/10.1093/jamia/ocx104>
- [17] Guohui Wang, T.S. Eugene Ng, and Anees Shaikh. 2012. Programming Your Network at Run-time for Big Data Applications. In *Hot Topics in Software Defined Networks (HotSDN '12)*. ACM, 103–108. <https://doi.org/10.1145/2342441.2342462>
- [18] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, et al. 2014. Named Data Networking. *SIGCOMM Comput. Commun. Rev.* 44, 3 (July 2014), 66–73. <https://doi.org/10.1145/2656877.2656887>