Scientific Data Movement enabled by the DYNES[*] Instrument

Jason Zurawski
Internet2
zurawski@internet2.edu

Eric Boyd
Internet2
eboyd@internet2.edu

Tom Lehman
ISI East
tlehman@east.isi.edu

Shawn McKee
University of Michigan
smckee@umich.edu

Azher Mughal
California Institute of Technology
azher@hep.caltech.edu

Harvey Newman
California Institute of Technology
newman@hep.caltech.edu

Paul Sheldon
Vanderbilt University
paul.sheldon@vanderbilt.edu

Steve Wolff
Internet2
swolff@internet2.edu

Xi Yang
ISI East
xyang@east.isi.edu

ABSTRACT

Scientific innovation continues to increase requirements for the computing and networking infrastructures of the world. Collaborative partners, instrumentation, storage, and processing facilities are often geographically and topologically separated, thus complicating the problem of end-to-end data management. Networking solutions, provided by R&E focused organizations, often serve as a vital link between these distributed components. Capacity and traffic management are key concerns of these network operators; a delicate balance is required to serve both long-lived, high capacity network flows, as well as more traditional end-user activities. The advent of dynamic circuit services, a technology that enables the creation of variable duration, guaranteed bandwidth networking channels, has afforded operations staff greater control over traffic demands and has increased the overall quality of service for scientific users.

This paper presents the DYNES instrument, an NSF funded cyberinfrastructure project designed to facilitate end-to-end dynamic circuit services. This combination of hardware and software innovation is being deployed across R&E networks in the United States, end sites located at University Campuses. DYNES is peering with international efforts in other countries using similar solutions, and is increasing the reach of this emerging technology. This global data movement solution could be integrated into computing paradigms such as cloud and grid computing platforms, and through the use of APIs can be integrated into existing data movement software.

Categories and Subject Descriptors

J.2 [**Computer Applications**]: Physical Sciences and Engineering

General Terms

Management, Measurement, Performance, Reliability, Standardization

Keywords

ATLAS, CMS, DYNES, FDT, ION, LHC, OSCARS, perfSONAR

1. INTRODUCTION

International scientific collaborations, including High Energy Physics, Astronomy, and biology, note several driving factors that will increase network reliance in the future. Activities that were once localized now are deeply reliant on access to high capacity networking; the technical requirements for different disciplines stand to grow over the coming years [20]:

- **Collaboration Size:** the addition of more researchers, and research facilities, increases the pool of users and resources available to process scientific data sets

[*]NSF Grant 0958998

- **Location of Collaborators:** scientific activity has scaled globally, particularly into non-traditional and remote regions of the planet, forcing the construction of supporting infrastructure
- **Data Collection Rates:** upgrades to the basic instruments of science, e.g. colliders, detectors, telescopes, and genome sequencers, are producing finer grained observations that directly translate into increases in the amount of data to process and store
- **Experimental Expectations:** The time expectation to analyze raw information in search of meaningful results is decreasing, thus pushing technology to be available and responsive to user demands

Keeping these well stated factors in mind, Research and Education (R&E) networking needs to find a way to support scientific demands. This active community has produced many innovations spanning hardware technologies, protocols, software, and services; all of which were directly targeted to domain researchers. To facilitate the bandwidth demands of growing scientific communities, network design has evolved to address two key, yet diametrically opposed, areas:

- **Enterprise requirements**, including the general population of desktops, laptops, and mobile devices on a given network
- **Science requirements**, which encompass the data centers and instruments that have a direct role in the collection, storage, and processing of data sets

These areas, when combined into the cohesive network design strategy informally named the "science DMZ", has gained popularity in the Department of Energy research community as well as within general campus IT infrastructure [21]. In addition to overall network design, the adoption of advanced network services has had a pivotal role in this new paradigm. Monitoring software, such as the perfSONAR framework, can now reliably be used to pinpoint performance trouble spots, or predict potential bottlenecks with the aide of historical measurements [19]. Hybrid networks, including ESnet SDN, GÉANT AutoBAHN, and the Internet2 ION service, offer facilities that enable provisioning of an end-to-end circuits (i.e. a protected Layer 2 VLAN), thus allowing applications to avoid overly congested Layer 3 paths and preventing overuse of general purpose infrastructure [22, 3, 11]. To date, many of the aforementioned technologies were targeted toward the largest component of the worldwide R&E networking infrastructure: Backbone and Regional providers. These far reaching networks, designed to offer large amounts of available capacity, may become congested in certain geographical areas that experience high demand [1, 16]. The ability to create a protected path, separate from the general-purpose IP traffic and using existing network resources (i.e. not requiring new investment), remains highly desirable for bulk data movement applications; particularly those relying on TCP and operating on high RTT/high bandwidth paths, including FDT [8]. This new "circuit" technology, while designed with the WAN in mind, has been slow to migrate to smaller scales where users could take full advantage of this new functionality.

Innovation occurs within facilities of all sizes. Federal laboratories may features 10s to 100s Gigabits of available connectivity; institutions with a smaller population of researchers may be equipped to deal with traffic of an order of magnitude less than these larger facilities. In either case, simple steps to alter the network architecture or enable innovative networking technologies, could lead to significant gains in productivity for scientific users and traffic management ability for network operations staff. The DYnamic NEtwork System (DYNES) is an effort designed to assist in these two key areas by providing efficient hardware, and accompanying advanced network services, to address the needs of the R&E community [7]. Funded by the NSF through the MRI program, DYNES is developing a nationwide "cyber-instrument"; enabling dynamic circuit capability, fast data transfer, and network monitoring at numerous U.S. based universities and regional networks.

This paper will proceed as follows: Section 2 will discuss some of the agile networking considerations for data intensive science. Section 3 introduce the DYNES solution, focusing on the hardware and software interactions of this emerging framework. Section 4 will present some performance numbers used to validate the basic use case of DYNES. Section 5 will present an example of DYNES interactions on a global scale, and demonstrate scalability to similar use cases. We will conclude and present future work in Section 6.

2. AGILE R&E NETWORKING

Traditional LAN designs favor the important notions of "protection" and "availability"; campus environments will place emphasis on deploying firewalls to protect the client machines from malicious attacks, and packet shaping technology to silently reduce the overall consumption of bandwidth. Figure 1 shows a simplified campus LAN, highlighting the "complete protection" afforded by a single device positioned on the border. This design affords the same treatment to all users — students in dormitories surfing the web as well as the supercomputing center with sophisticated instrumentation. While effective in protecting users and network resources, data intensive science and remote collaboration may suffer unexpected consequences due to the overall architectural considerations.

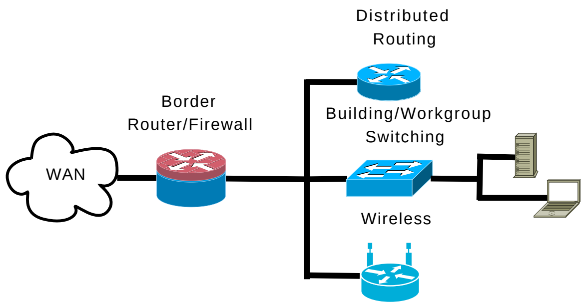

Figure 1: Traditional Campus Network

Applications designed to move large amounts of data are often based on TCP. FDT, a tool used as the basis for data movement in the LHC Virtual Organization (VO), utilizes multiple streams of TCP traffic to transmit data. While this is an effective way to consume more of the available resources, it is unfair to other users, and simply creates an "arms race" for use of the network. Firewalls and packet shaping devices have a profound effect on TCP performance, particularly as the RTT between locations increases [15]. These devices often have small memory buffers, which hampers the ability to handle a single large flow, not to mention multiple flows of various sizes. Architectural changes, including the concept of a "science DMZ", offer a vast improvement to data movement activities by reducing the number of disruptive devices on the local portion of the path, and allow storage and processing a more direct access method to the WAN. Enterprise protection requirements can still be met, as demonstrated in Figure 2. Additionally, the use of access control lists on routing devices and per-host and per-service security settings can offer similar functionality to a traditional firewall device at a much lower overhead.

These simple architectural measures have had a profound effect on functionality of the network for research use. Additional support, provided by advanced software and services, can further improve operational concerns on R&E networks.

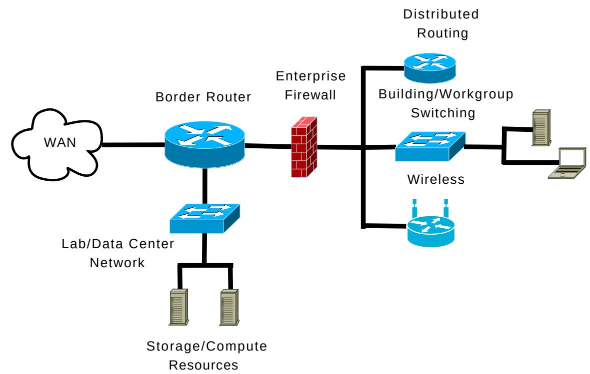

Figure 2: Network Featuring a Science DMZ

3. DYNES

The effective design of networks to support scientific activities addresses many key issues, namely by encouraging the development of a physical separation of traffic classes and broader hardware management strategies. These steps alone will often lead to performance gains in a local setting; scientific innovation, however, is no longer a locally based concern. Global collaborations rely on the existence of worldwide R&E networking infrastructure as well as the assurance of end-to-end capacity and performance. This broad view applies to all networks - from the largest providers to the smallest consumers. Backbone networks are designed to transit traffic from regional aggregation points, each serving countless facilities. These far reaching networks, designed to offer large amounts of available capacity, may become congested in certain geographical areas experiencing high demand. International links, funded by parties including the DOE and NSF, link continents and countries to enable distributed activities [24, 12]. The sum of available networking resources globally remains high, and continues to grow. Effective use of the available resources continues to be a struggle that is a actively being addressed in the network research community.

The ability to create a protected path, separate from the general-purpose IP traffic and using existing network resources, remains highly desirable for bulk data movement applications; particularly those relying on the TCP and operating on high RTT and high bandwidth paths. This new "circuit" based technology, called as such due to the use of end-to-end Virtual Local Area Neworks (VLANs), was designed with the WAN in mind and has been slow to migrate to smaller scales. The DYnamic NEtwork System (DYNES) addresses this technology gap by providing hardware and software solutions to regional and campus networks — the overarching goal being to extend the technology already available on many backbone networks. This NSF sponsored project is developing and deploying a nationwide "cyber-instrument", designed to span approximately 40 US universities and 11 regional networks. DYNES was awarded to a collaborative team including Internet2, the California Institute of Technology, the University of Michigan, and Vanderbilt University in 2010. The DYNES team will partner with the LHC and astrophysics communities, the OSG, and Worldwide LHC Computing Grid (WLCG) to deliver these capabilities to the LHC experiment as well as others such as LIGO, and the SDSS/LSST astronomy programs, broadening existing Grid computing systems by promoting the network to a reliable, high performance, actively managed component [13, 14, 23, 18, 25].

By integrating existing and emerging protocols, software for dynamic circuit provisioning and scheduling, in-depth end-to-end network path and end-system monitoring, and higher level services for management on a national scale, DYNES will allocate and schedule channels with bandwidth guarantees to several classes of prioritized data flows with known bandwidth requirements, and to the largest high priority data flows, enabling scientists to utilize and share network resources effectively. DYNES is designed to support many data transfers which require aggregate network throughputs between sites of 1-20 Gbps, rising to the 40-100 Gbps range as the underlying network technology is upgraded. This capacity will enhance researchers' ability to distribute, process, access, and collaboratively analyze 1 to 100 TB datasets at university-based Tier2 and Tier3 centers now, and PB-scale datasets in the future.

DYNES is based on a "hybrid" packet and circuit architecture composed of Internet2's ION service and extensions over regional and state networks to US campuses. It will connect with transoceanic (IRNC, USLHCNet), European (GÉANT), Asian (SINET3) and Latin American (RNP and ANSP) R&E networks through the aid of related efforts including IRNC DyGIR [2]. DYNES will build on existing key open source software components that have already been individually field-tested and hardened in part by the PIs: DCN Software Suite (OSCARS/DRAGON), perfSONAR, UltraLight Linux kernel, and FDT.

3.1 DYNES Hardware

The DYNES framework requires 3 key pieces of hardware:

- A network device (e.g. a switch) that can be dynamically controlled
- A controller to integrate with the OSCARS control plane
- A data movement server, capable of storing large amounts of data and utilizing dual network connections

Figure 3 shows these components in a block diagram. The OSCARS software, described in Section 3.2, functions as the software "glue", linking together the networks participating in dynamic control, and allowing the creation of dynamic end-to-end circuits. FDT, described in Section 3.3, functions end to end — allowing the data movement servers provided by DYNES (as well as other existing compute and storage resources) to stage and migrate information over short term circuits. perfSONAR, described in Section **??** is available to monitor the status and health of the network participants.

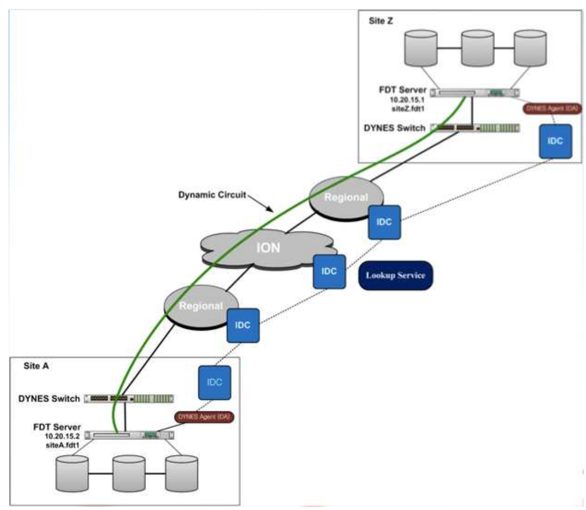

Figure 3: DYNES Hardware

3.2 OSCARS

The On-Demand Secure Circuits and Advance Reservation System, or OSCARS, provides multi-domain, high-bandwidth virtual circuits that guarantee end-to-end network data transfer performance [17]. Originally a research concept, OSCARS has grown into a robust delivery mechanism in the form of ESnet Science and Data Network (SDN) and the Internet2 ION Service. OSCARS virtual circuits carry fifty percent of ESnets annual 60 petabytes of traffic. This segmentation of network resources has had a profound effect on research innovation — users are free to worry about science instead of becoming experts at network design. The dynamically provisioned, multi-domain, guaranteed bandwidth circuits provided by OSCARS extend Layer 2 VLAN concepts across a complete end-to-end path. Muti-domain circuit operation is accomplished through a series of protocols designed to negotiate on various resources. Each domain maintains local control over all components, and has the ability to set policy regarding use; this is shown in Figure 4. The concept of virtual circuits integrates cleanly with existing networking solutions, and often co-exists on the same physical infrastructure. Thus it becomes possible to make a conscious decision to switch a target machine's connect between a dynamically provisioned circuit or the Layer 3 infrastructure with simple networking commands. APIs and web-based toolkits are also available to ease adoption further; applications can be modified to use either functionality [8]. Prior demonstrations of this technology have shown use cases both as a way to manually dedicate resources in a traffic management capacity, but also dynamically to support applications including data transfer and high definition video [10, 6].

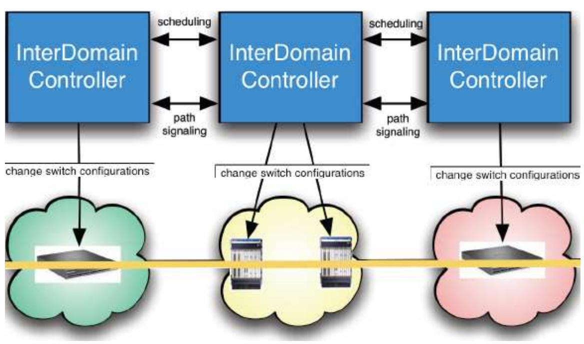

Figure 4: OSCARS in Action: Controlling Switch Operation

3.3 FDT

FDT is an application designed to enable efficient data transfers; performance emphasis is placed upon the interaction between the hosts, their 'internal' components (e.g. disk, bus, processor, and network card) and the wide area network. FDT is written in the JAVA programming language, which has enabled deployment on many major operating system and architectural platforms.

Figure 5 demonstrates the typical use case for FDT. FDT is based on an asynchronous, flexible multi-threaded system and includes the following features:

- Streams a complete dataset (e.g. a list of files) continuously, using a managed pool of buffers through one or more TCP sockets on the host machine

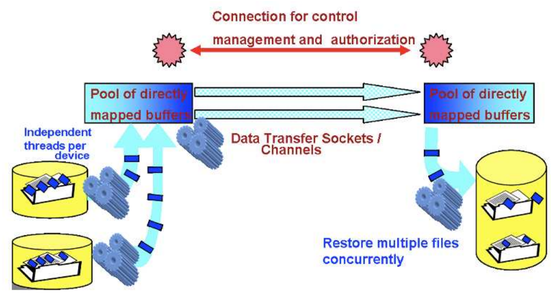

Figure 5: FDT Operation

- Uses independent threads to read and write on each physical device
- Can transfer data in parallel using multiple TCP streams
- Optimizes the interaction between network and disk I/O

FDT has embraced the concept of dynamic circuit networking by integrating with the OSCARS control framework through the use of programatic APIs. This mode of operation allows the FDT agent on a host to transmit data using traditional Layer3 infrastructure, or invoke an on-demand circuit to perform a point to point transfer for short periods of time. FDT is capable of requesting bandwidth, and managing the nuances of the circuit, freeing the user to ignore details about the underlying network.

3.4 perfSONAR

The PERFormance Service Oriented Network monitoring Architecture, or perfSONAR, is a federated network monitoring framework, that facilitate end-to-end sharing of performance measurement data. This component based suite of tools decouples the tasks of measuring network performance from mechanisms used to store data, share and visualize results, and authenticate user permissions [9]. perfSONAR has been widely adopted by the R&E education community, and is used to expose active and passive network measurements on networks of all sizes, as well as for resources managed at end sites including scientiMeasurement data, provided through a mechanism such as perfSONAR, is vital to network aware applications including remote collaboration tools and software designed for the task of data movement. Knowledge of current performance, as well as historical trends, can be used to making routing decisions, invoke dynamic connectivity, or schedule future data movement activities [26].

4. PERFORMANCE EVALUATION

The DYNES solution depends on clean interactions between the data movement software, the available hardware, and the software designed to shape and control the flow of networking traffic. There are two concerns for effective operation:

- Interaction between the storage, networking components, and operating system that may limit the overall throughput of FDT
- Interaction between network devices (e.g. switches, hosts) and the effect of an intensive TCP flow

4.1 Disk Performance

Network performance can often be limited by the realities of the hosts involved, namely the spinning of a physical disk, use of RAID techniques, file system interactions, and operating system configuration. The DYNES team performed several tests to evaluate the choice of hardware and software, and the expected performance that could be achieved from the data movement server using the FDT tool. The DYNES servers are Intel x86 64 bit architecture, and are running the CentOS 5.6 operating system [5]. Two operating system kernels were tested: the default CentOS kernel as well as the latest version from the mainline kernel development project. Typically operating system vendors use a stable version of the kernel as they evaluate new features and functionality — these tests were to reveal any potential performance gains that may be available in the recent development.

In addition to these factors, the DYNES project tested different I/O scheduling routines (NOOP vs Deadline), varied the strip size used in the RAID 5 configuration (128K vs 256K), varied the caching mechanism ("Write Through" vs "Write Back"), and configured the underlying filesystem to be XFS. The Bonnie++ tool was used to help evaluate these tests [4]. Table 4.1 shows the tests using the kernels and the RAID 128K strip size. Note that the table is rather sparse due to complications with the default CentOS kernel, the filesystem, and the configuration choices which did not allow collection of results due to failures.

128K Strip Size				
	Write Through		Write Back	
	Deadline	NOOP	Deadline	NOOP
Kernel				
2.6.18.238.9.1	N/A	N/A	N/A	N/A
2.6.38	9.26	8.98	N/A	N/A

Table 1: Disk Performance (in Gbps) with various cache and schedulers

Table 4.1 shows a cleaner view of the kernel testing. The aforementioned problems with the CentOS kernel, filesystem, and RAID configuration prevented some results from being collected in this series of tests as well.

256K Strip Size				
	Write Through		Write Back	
	Deadline	NOOP	Deadline	NOOP
Kernel				
2.6.18.238.9.1	N/A	N/A	5.60	5.60
2.6.38	9.05	7.88	9.19	8.58

Table 2: Disk Performance (in Gbps) with various cache and schedulers

Given the limited amount of flexibility provided by the CentOS default kernel, the DYNES project has pursued using the latest versions from the kernel development project in an effort to increase the overall performance on the data movement host. Further testing has revealed some of the nuances of disk performance for both read and write operations. The Deadline scheduler demonstrates balanced performance for both operations. The NOOP scheduler, shown in Figure 6 produces results that favor read operations over writes, reducing the overall performance.

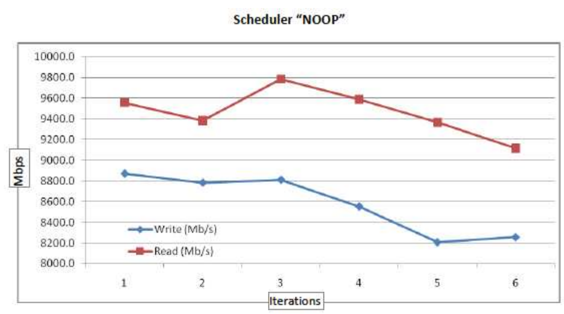

Figure 6: 2.6.38 Kernel, 256K Strip Size, NOOP Scheduling

These tests form one aspect that DYNES wishes to address, but the issue of network performance is crucial for the overall success of the framework.

4.2 Network Performance

The typical use case for a data movement application, such as FDT, is WAN transfers crossing multiple domains. DYNES equipment includes a Layer2/Layer3 capable device that serves as a first entry point onto the network; this device can be dynamically controlled by the OSCARS component to create end-to-end dyanmic circuits. While these paths are configurable in terms of bandwidth and free from congestion that may otherwise alter the overall performance, the hardware itself must be evaluated to ensure that performance matches user expectations and the host capabilities described in Section 4.2. For this test two data movement servers were connected through a common switch at a rate of 10Gbps. A simple "memory to memory" test can be used to set a baseline, and does not use the host's physical disk. The performance was shown to be close to line rate in this case, which is expected.

The second test, shown in Figure 7, uses the host configuration from Section 4.2 (The XFS file system, Deadline Scheduler and a RAID 5 with a strip size of 256K). This test shows a "memory to disk" transfer, and was designed to illustrate disk "write" performance (a typical bottleneck for data movement). The results show performance approaching 75% link utilization for the length of the test.

Figure 7: 10G Memory to Disk Test through DYNES Switch

These results have verified the performance of the hardware and software, albeit in a local setting, as a baseline to set expectations as the DYNES infrastructure is deployed around the country. Given the complications of WAN transfers, e.g. the effect that latency has on the bandwidth delay product, it is expected that further tuning will be required through operational experience. This may result in configuration changes, software updates, and eventual hardware refreshes by the DYNES project over the life of the deployment.

5. GLIF 2011 DEMONSTRATION

The partners of the DYNES project, in collaboration with the AMPATH international exchange point, Rede Nacional de Ensino e Pesquisa (RNP), São Paulo State University (UNESP), São Paulo Research and Analysis Center (SPRACE), Rio de Janeiro State University (UERJ), and Tier2 HEPGrid Brazil are collaborating on a demonstration of scientific data movement through dynamic circuit capabilities for the 2011 Global Lambda Integrated Facility (GLIF) meeting. This interoperability exercise is a mixture of cutting edge research capabilities and production quality infrastructure . All resources are individually managed, and rely on advanced network services to perform the task of creating, managing, and utilizing the networking infrastructure.

The DYNES infrastructure, consisting of end sites located at the California Institute of Technology, the University of Michigan, Vanderbilt University, and the Internet2 ION service, will offer the ability to create multi-gigabit dynamic circuits within the United States. DYNES has been integrated with the AmLight IRNC link via the DyGIR grant — thus enabling dynamic capabilities across an international connection. The Brazilian R&E backbone, RNP, has enabled OSCARS capabilities within their network; this investment will allow end sites within Brazil, including UNESP and UERJ, to connect into the multi-domain capabilities stretching far in the United States. The SPRACE and HEPGrid Tier2 centers, active participants in the LHC collaboration, may connect to the Tier2 centers in the U.S. and utilize the FDT tool to receive scientific data.

6. CONCLUSION AND FUTURE WORK

The technological realities of scientific collaboration are trending toward increased use of network resources. The R&E networking community has responded with several solutions that work in concert:

- Increased capacity on the backbone, regional and campus networks
- Suggestions to design dual-use network architectures to serve scientific users as well as the general population
- Innovative technologies to ensure end-to-end performance for high capacity network use cases

DYNES, an infrastructure designed to deliver a hardware and software based solution to end users and regional networks, is in the process of deploying a nationwide instrument to address the needs of the scientific community. This infrastructure integrates with similar efforts already in place on backbone networks, and will inter-operate with similar efforts designed to address international connectivity. Recent activities to demonstrate this capability have shown that this technology is ready for production networks, and capable of meeting the needs of scientists.

7. ACKNOWLEDGEMENTS

The authors would like to thank staff members from the California Institute of Technology, Internet2, the University of Michigan, and Vanderbilt University for their assistance in the operation and construction of DYNES. The DYNES project is also grateful for the contributions of the FDT, OSCARS, and perfSONAR-PS projects — the underlying software set that controls DYNES. Lastly, the authors would like to acknowledge Dell Computers for assisting the DYNES team in our evaluation and purchase of hardware for this effort. These contributions have been crucial to the overall success in designing and implementing the DYNES framework.

The authors would also like to acknowledge funding bodies that support this ongoing work. The DYNES project is supported by the NSF grant **0958998**. Related work for the DISUN, Ultralight, and PlaNetS projects is sponsored by NSF grants **PHY-0533280**, **PHY-0427110**, and **PHY-0622423**. The Department of Energy (DOE) sponsors US LHCNet under grant **DE-FG02-08ER41559**.

8. REFERENCES

[1] Internet2 ION Case Study: Large Hadron Collider (LHC). Case study, Internet2, September 2009. http://www.internet2.edu/pubs/200909-CS-ION-LHC.pdf.

[2] National Science Foundation International Research Network Connections (IRNC) Awards. Case study, Internet2, October 2010. http://www.internet2.edu/pubs/201010-IS-IRNC.pdf.

[3] GEANT2 AutoBAHN. `http://www.geant2.net/server/show/ConWebDoc.2544`.

[4] Bonnie++. `http://www.coker.com.au/bonnie++/`.

[5] CENTOS Linux. `http://www.centos.org/`.

[6] N. Charbonneaum, V. Vokkarane, C. Guok, and I. Monga. Advance reservation frameworks in hybrid ip-wdm networks. *IEEE Communications*, pages 132–139, May 2011.

[7] MRI-R2 Consortium: Development of Dynamic Network System (DYNES). `http://www.internet2.edu/ion/dynes.html`.

[8] Fast Data Transfer (FDT). `http://monalisa.cern.ch/FDT`.

[9] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, M. Swany, S. Trocha, and J. Zurawski. Perfsonar: A service oriented architecture for multi-domain network monitoring. In *Third International Conference on Service Oriented Computing - ICSOC 2005, LNCS 3826, Springer Verlag*, pages 241–254, Amsterdam, The Netherlands, December 2005.

[10] A. Hutanu, J. Ge, C. Toole, R. Paruchuri, A. Yates, and G. Allen. Distributed Visualization Using Optical Networks: Demonstration at Super-computing 2008. Technical report, LSU CCT, October 2008. http://www.cct.lsu.edu/CCT-TR/CCT-TR-2008-10.

[11] Internet2 ION. `http://www.internet2.edu/ion/`.

[12] International Research Network Connections. `http://irnclinks.net/`.

[13] Laser Interferometer Gravitational Wave Observatory. `http://www.ligo.caltech.edu/`.

[14] Large Synoptic Survey Telescope. `http://www.lsst.org`.

[15] M. Mathis. Pushing up Performance for Everyone. Presentation, Internet2, December 1999. http://staff.psc.edu/mathis/papers/.

[16] I. Monga, C. Guok, W. Johnston, and B. Tierney. Hybrid networks: Lessons learned and future challenges based on esnet4 experience. *IEEE Communications*, pages 114–121, May 2011.

[17] ESnet OSCARS. `http://www.es.net/services/virtual-circuits-oscars/`.

[18] Open Science Grid. `http://www.opensciencegrid.org/`.

[19] perfSONAR-PS. `http://psps.perfsonar.net`.

[20] ESnet Network and Science Requirement Workshops Reports. `http://www.es.net/about/science-requirements/reports/`.

[21] Science DMZ. `http://fasterdata.es.net/fasterdata/science-dmz/`.

[22] ESnet SDN. `http://www.es.net/hypertext/network.html`.

[23] Sloan Digital Sky Survey. `http://www.sdss.org/`.

[24] USLHCNet. `http://lhcnet.caltech.edu/`.

[25] Worldwide LHC Computing Grid. `http://lcg.web.cern.ch/lcg/`.

[26] R. Wolski. Dynamically forecasting network performance using the network weather service. *Journal of Cluster Computing*, pages 119–132, January 1998.