

ESnet: Advanced NETWORKING for SCIENCE

Researchers around the world using advanced computing for scientific discovery are connected via the DOE-operated Energy Sciences Network (ESnet). By providing a reliable, high-performance communications infrastructure, ESnet facilitates the large-scale, collaborative science endeavors fundamental to Office of Science missions.

Energy Sciences Network

In many ways, the dramatic achievements of 21st century scientific discovery—often involving enormous data handling and remote collaboration requirements—have been made possible by accompanying accomplishments in high-performance networking. As increasingly advanced supercomputers and experimental research facilities have provided researchers with powerful tools with unprecedented capabilities, advancements in networks connecting scientists to these tools have made these research facilities available to broader communities and helped build greater collaboration within these communities. The DOE Office of Science (SC) operates the Energy Sciences Network (ESnet). Established in 1985, ESnet currently connects tens of thousands of researchers at 27 major DOE research facilities to universities and other research institutions in the U.S. and around the world (figure 1; sidebar “What is ESnet?,” p50).

ESnet’s mission is to provide an interoperable, effective, reliable, high-performance network communications infrastructure, along with selected leading-edge Grid-related and collaboration services in support of large-scale, collabora-

tive science. These include: sharing of massive amounts of data, supporting thousands of collaborators worldwide, distributed data processing and data management, distributed simulation, visualization, and computational steering, and collaboration with the U.S. and international research and education (R&E) communities.

To ensure that ESnet continues to meet the requirements of the major science disciplines a new approach and a new architecture are being developed. This new architecture includes elements supporting multiple, high-speed national backbones with different characteristics—redundancy, quality of service, and circuit oriented services—all the while allowing interoperation of these elements with the other major national and international networks supporting science.

Evolving Science Environments

Large-scale collaborative science—big facilities, massive amounts of data, and thousands of collaborators—is a key element of the DOE SC scientific mission. The science community that participates in DOE’s large collaborations and facilities is almost equally divided between laboratories and universities, and also has a significant

To ensure that ESnet continues to meet the requirements of the major science disciplines a new approach and a new architecture are being developed.

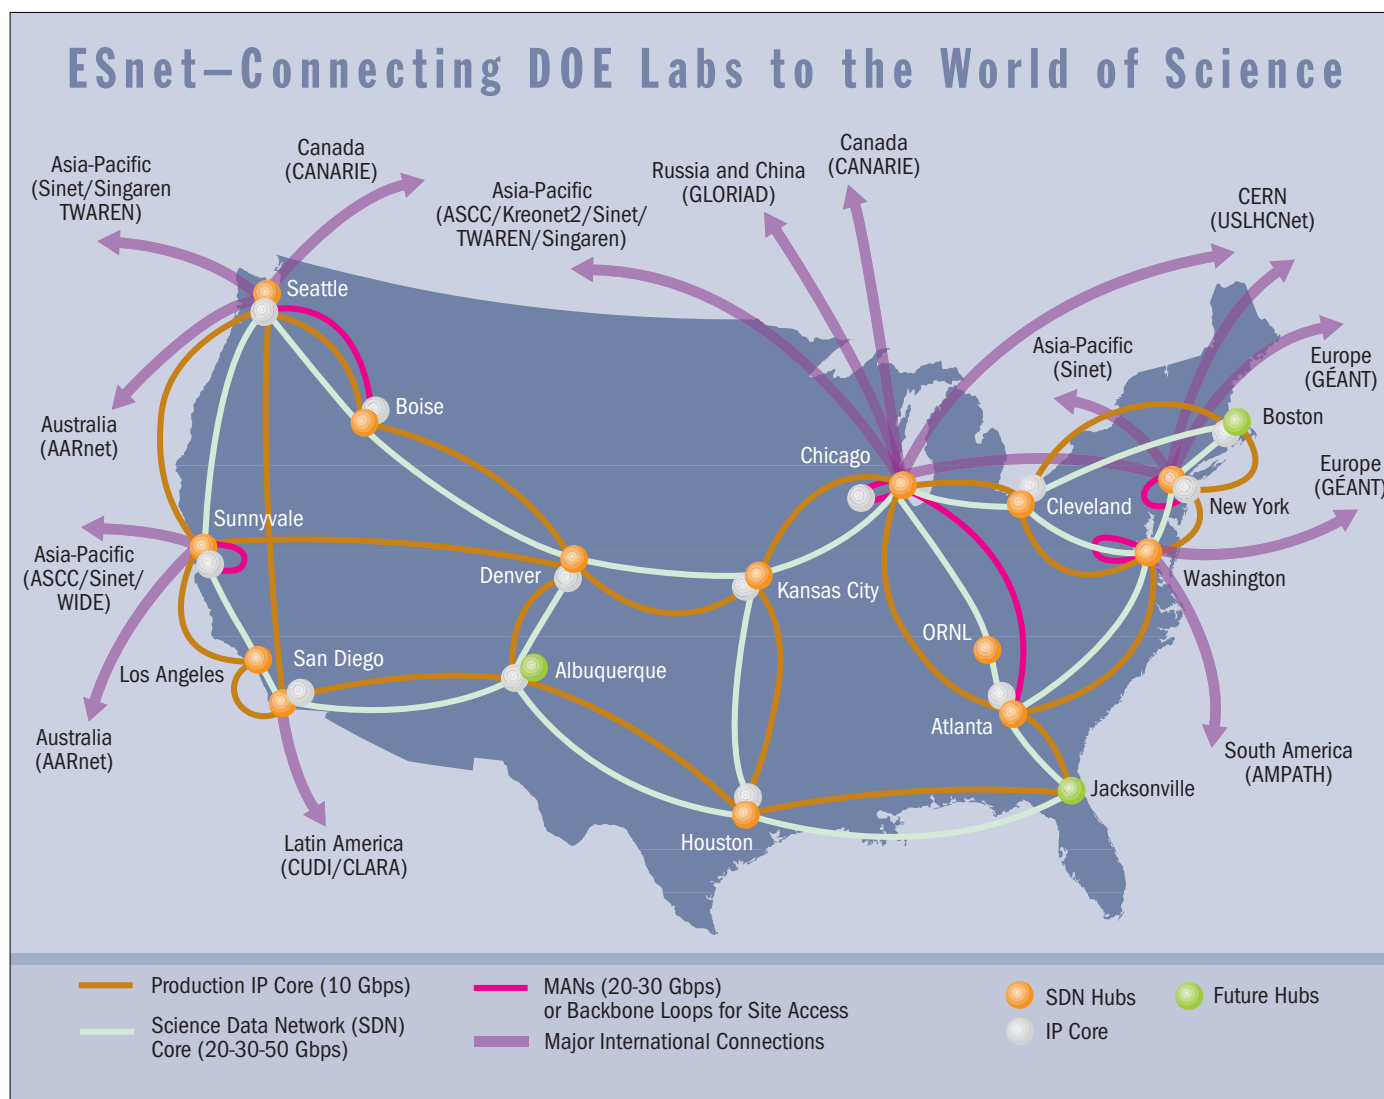


Figure 1. ESnet connects the DOE laboratories to collaborators at research and education institutions around the world.

international component. This component consists of very large international facilities and international teams of collaborators participating in U.S.-based experiments, often requiring the movement of massive amounts of data between the laboratories and these international facilities and collaborators. Distributed computing and storage systems for data analysis, simulations, and instrument operation are becoming common; for data analysis in particular, Grid-style distributed systems predominate.

This science environment is very different from that of a few years ago and places substantial new demands on the network. High-speed, highly reliable connectivity between laboratories and U.S. and international R&E institutions (sidebar “U. S. and International Science,” p51) is required to support the inherently collaborative, global nature of large-scale science. Increased capacity is needed to accommodate a large and steadily increasing amount of data that must traverse the network to get from instruments to scientists and

to analysis, simulation, and storage facilities. High network reliability is required for interconnecting components of distributed large-scale science computing and data systems and to support various modes of remote instrument operation. New network services are needed to provide bandwidth guarantees for data transfer deadlines, remote data analysis, real-time interaction with instruments, and coupled computational simulations.

Requirements from Data, Instruments, Facilities, and Science Practice

There are some 20 major instruments and facilities currently operated or being built by the DOE SC, in addition to the Large Hadron Collider (LHC) at CERN (figure 2, p50) and the ITER fusion research project in France (figure 4, p51). To date, ESnet has characterized 14 of these for their future networking requirements. DOE facilities—such as the Relativistic Heavy Ion Collider (RHIC) at BNL, the Spallation Neutron Source at ORNL,

What Is ESnet?

As the primary DOE network providing production Internet service to almost all of DOE's laboratories and other sites, ESnet serves an estimated 50,000 to 100,000 DOE users, as well as more than 18,000 non-DOE researchers from universities, other government agencies, and private industry, that use SC facilities. The connections with university researchers are critical. As the

single largest supporter of basic research in the physical sciences in the United States, the DOE SC supports 25,500 Ph.D.s, postdoctoral and graduate students, and 21,500 users of facilities, half of whom are at universities.

The goal is to provide the same level of service and connectivity from a DOE laboratory to U.S. and European R&E institutions as the connectivity between any

laboratory-to-laboratory and university-to-university link. The keys to ensuring this are engineering, operations, and constant monitoring. ESnet has worked with the Internet2 and the international R&E community to establish a suite of monitoring tools that can be used to provide a full mesh of paths that continuously check all major interconnection points.



Figure 2. An aerial view of CERN. The Large Hadron Collider (LHC) ring is 27 km in circumference and provides two counter-rotating, 7 TeV proton beams to collide in the middle of the detectors.

the National Energy Research Scientific Computing (NERSC) Center at LBNL, and the Leadership Computing Centers at ORNL and ANL, as well as the LHC at CERN—are typical of the hardware infrastructure for science in the 21st century.

These facilities generate four types of network requirements: bandwidth, connectivity and geographic footprint, reliability, and network services. Bandwidth needs are determined by the quantity of data produced and the need to move the data for remote analysis. Connectivity and

geographic footprint are determined by the location of the instruments and facilities, and the locations of the associated collaborative community, including remote and/or distributed computing and storage used in the analysis systems. These locations also establish requirements for connectivity to the network infrastructure that supports the collaborators. Reliability requirements are driven by how closely coupled the facility is with remote resources. For example, off-line data analysis—where an experiment

U.S. and International Science

ESnet is a network in the traditional sense of the word, connecting end user sites to various other networks. A large fraction of all the national data traffic supporting U.S. science is carried by three networks—ESnet, Internet2, and National Lambda Rail. These three entities represent the architectural scope of science-oriented networks. Internet2 is a backbone network connecting U.S. regional networks to each other and to international networks. National Lambda Rail is a collection of light paths, or lambda channels, that are used to construct specialized R&E networks.

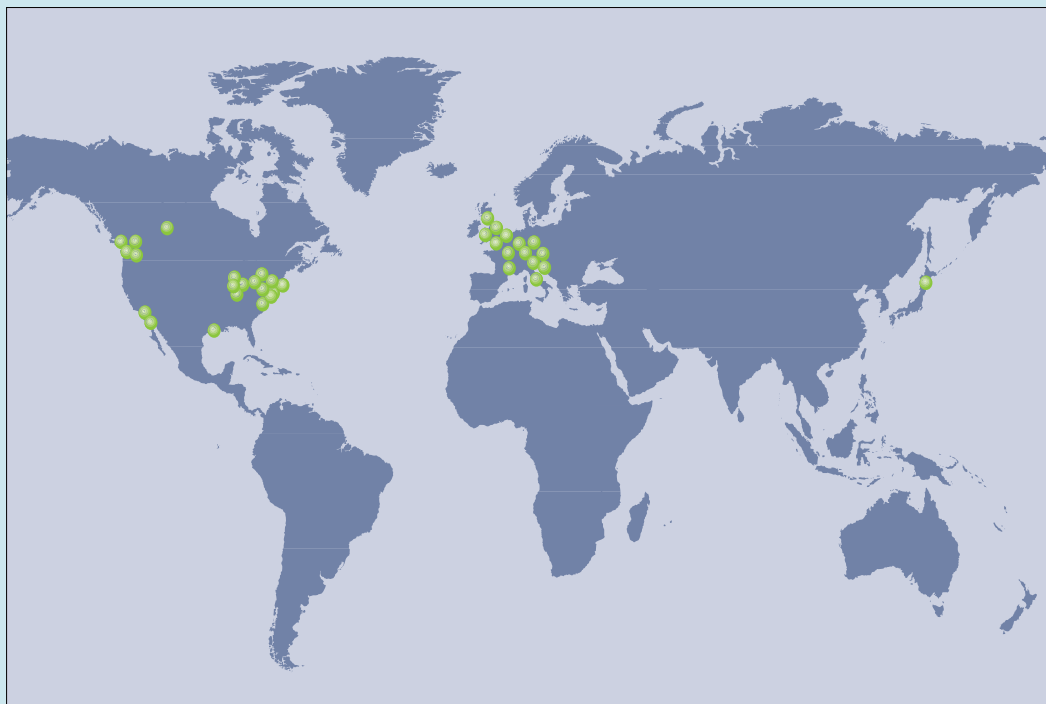


Figure 3. The large-scale data flows in ESnet reflect the scope of DOE SC collaborations. ESnet's top 100 data flows generate 50% of all ESnet traffic (ESnet handles about 3×10^9 flows/month). Of the top 100 flows, 91 are from the DOE Labs (not shown) to other R&E institutions (shown on the map).

runs and generates data, and the data is analyzed after the fact—may be tolerant of some level of network outages. On the other hand, when remote operation or analysis must occur within the operating cycle time of an experiment, or when other critical components depend on the connection, then very little network downtime is acceptable. The reliability issue is critical and drives much of the design of the network. In the past, networks typically provided a single network service—best-effort delivery of data packets—on which are built all of today's higher-level applications (FTP, email, Web, and socket libraries for application-to-application communication), and best-effort Internet Protocol (IP) multicast (where a single outgoing packet is, sometimes unreliably, delivered to multiple receivers). In considering future uses of the network by the science community, several other network services have been identified as requirements, including bandwidth guarantees, traffic isolation, and reliable multicast.

Bandwidth guarantees (sidebar “OSCARs: Guaranteed Bandwidth Service,” p52) are typically needed for online analysis, which always involves time constraints. Another type of application requiring bandwidth guarantees is distributed




Figure 4. An illustration of the ITER project tokamak—a device used in fusion energy research.

OSCARS: Guaranteed Bandwidth Service

DOE has funded the On-demand Secure Circuits and Advance Reservation System (OSCARS) project to develop and deploy the various technologies that provide dynamically provisioned circuits and various qualities of service that can be integrated into a production net environment. Such circuits are called “virtual circuits” because they are defined in software and thus are mutable, as opposed to hardware established circuits.

There are two realms in which OSCARS must operate. The first is intradomain, to

establish a schedulable, guaranteed bandwidth circuit service within the boundary of the ESnet network. The second realm is interdomain, to provide end-to-end services between DOE laboratories and U.S. and European universities.

The OSCARS architecture is based on Web/Grid Services and consists of modules for: user/application request and claiming; the Authentication, Authorization, and Auditing Subsystem (AAAS) to handle access control, enforce policy, and generate usage

records; a Bandwidth Scheduler Subsystem (BSS) that tracks reservations and maps the state of the network (present and future) for managing the allocation of virtual circuits; and a Path Setup Subsystem (PSS) that sets up and tears down the on-demand virtual circuits.

OSCARS, and similar systems in the other networks, is well into development and prototype deployment. A number of OSCARS circuits are currently being tested between DOE labs and European R&E institutions.

From the analysis of historical traffic patterns, several clear trends emerge that result in requirements for the evolution of the network.

workflow systems, such as those used by high energy physics data analysis. The inability of one element, such as a computer, in the workflow system to adequately communicate data to another will ripple through the entire workflow environment, slowing down other participating systems as they wait for required intermediate results, thus reducing the overall effectiveness of the entire system.

Traffic isolation is required because today’s primary transport mechanism, Transmission Control Protocol (TCP), is not ideal for transporting large amounts of data across large distances, such as between continents. While other protocols are better suited to this task, these are not compatible with the fair-sharing of TCP transport in a best-effort network, and are thus typically penalized by the network in ways that reduce their effectiveness. A service that can isolate the bulk data transport protocols from best-effort traffic is needed to address this problem.

Reliable multicast is a service that, while not entirely new, must be enhanced to increase its effectiveness. Multicast provides for delivering a single data stream to multiple destinations without having to replicate the entire stream at the source, as is the case, for example, when using a separate TCP-based connection from the source to each receiver. This is important when the data to be delivered to multiple sites are too voluminous to be replicated at the source and sent to each receiving site individually. Today, IP multicast provides this capability, but in a fragile and limited manner.

Past Traffic Patterns Drive Future Plans

From the analysis of historical traffic patterns, several clear trends emerge that result in requirements for the evolution of the network.

The first and most obvious pattern is the expo-

ponential growth of the total traffic handled by ESnet (figure 5 and figure 6). This traffic trend represents a 10-fold increase every 47 months on average since 1990 (figure 6). ESnet traffic just passed the one petabyte per month level with about 1.5 Gb/s average, steady-state load on the New York–Chicago–San Francisco path. If this trend continues—and all indications are that it will accelerate—the network must be provisioned to handle an average of 15 Gb/s in four years. This implies a minimum backbone bandwidth of 20 Gb/s, because the network peak capacity must be at least 40% higher than the average load in order for today’s protocols to function properly with bursts of traffic, which is normal and expected. In addition, the current traffic trend suggests that 200 Gb/s of core network bandwidth will be required in eight years. This can only be achieved within a reasonable budget by using a network architecture and implementation approach that allows for cost-effective scaling of hub-to-hub circuit bandwidth.

The second major change in traffic is the result of a dramatic increase in the use of parallel file mover applications, such as GridFTP. This has led to the most profound change in traffic patterns in the 21-year history of ESnet. Historically, the peak system-to-system, or “workflow,” bandwidth of the largest network users has increased along with the increases in total network traffic. But over the past 18 months, the peak bandwidth of the largest user systems has been coming down and the number of flows that they generate has been going up, while the total traffic continues to increase exponentially. This reduction in peak workflow bandwidth, together with an overall increase in bandwidth, results from the decomposition of single large flows into many smaller parallel flows. In other

words, the same types of changes that happened in computational algorithms as parallel computing systems became prevalent are now happening in data movement—that is, parallel input/output (I/O) channels operating across the network.

The third clear traffic trend is that over two years the impact of the top few hundred workflows has grown from negligible (before mid-2004) to more than 50% of all traffic in ESnet by mid-2006! This is illustrated in figure 5, where the top part of the traffic bars shows the portion of the total generated by the top 100 hosts.

The fourth significant pattern comes from looking at the sources and destinations of the top data transfer systems, an examination that shows two things. First, the vast majority of the transfers can easily be identified as science traffic since the transfers are between two scientific institutions with systems that are named in ways that reflect the name of the science group. Second, for the past several years the majority of the large data transfers have been between institutions in the U.S., Europe, and Japan, reflecting the strongly international character of large science collaborations organized around large scientific instruments.

Enabling Future Science

Based both on the projections of the science programs and the changes in observed network traffic and patterns over the past few years, it is clear that the network must evolve substantially in order to meet the needs of the DOE SC mission.

The current trend in traffic patterns—the large-scale science projects giving rise to the top 100 data flows that represent about one half of all network traffic—will continue to evolve. As the LHC experiments ramp up in 2006 and 2007, the data to the Tier-1 centers, Fermi National Accelerator Laboratory (FNAL) and BNL, will increase between 200- and 2,000-fold. A comparable amount of data will flow out of the Tier-1 centers to the Tier-2 centers (U.S. universities) for analysis. The DOE National Leadership Class Facility supercomputer at ORNL anticipates a new model of computing in which simulation tasks are distributed between the central facility and a collection of remote “end stations” that will generate substantial network traffic. As climate models achieve the sophistication and accuracy anticipated in the next few years, the amount of climate data that will move into and out of the NERSC Center will increase dramatically. Similarly, the experiment facilities at the new Spallation Neutron Source and Magnetic Fusion Energy facilities will start using the network in ways that require fairly high bandwidth with guaranteed quality of service.

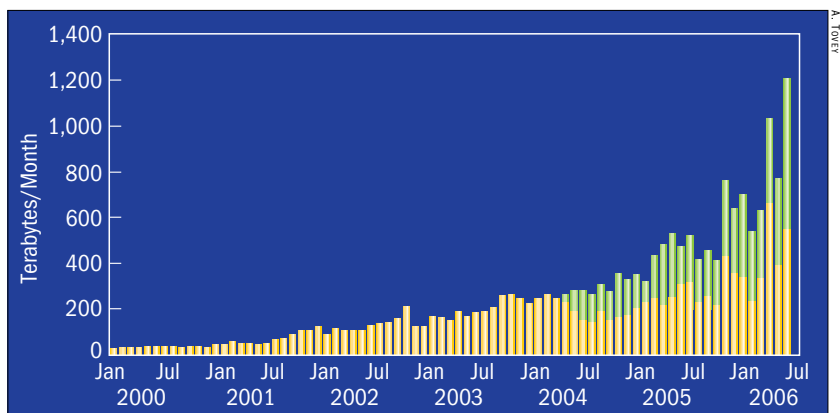


Figure 5. Total ESnet traffic by month, 2000–2006. The segmented bars (from mid-2004 and forward) show that fraction of the total traffic in the top 1,000 data flows, generated from large-scale science facilities. There are typically several billion flows per month in total, most of which are minuscule compared to the top 1,000 flows.

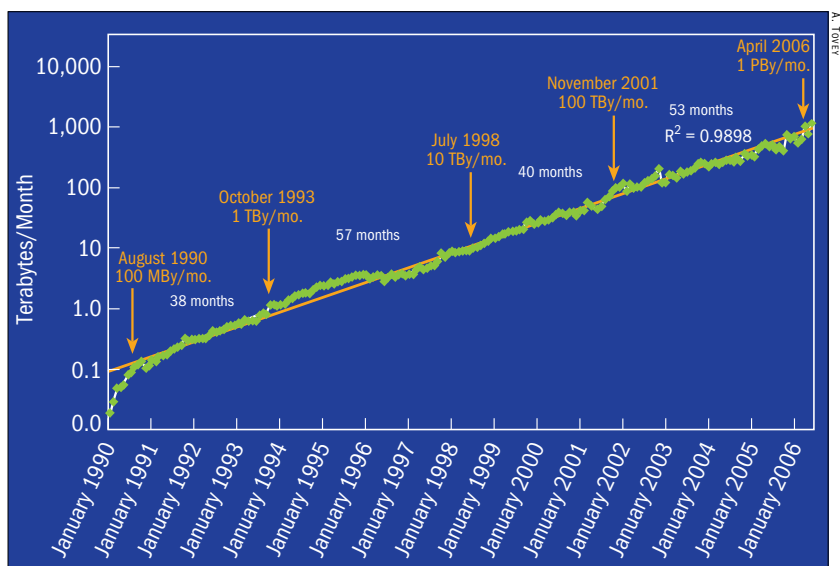


Figure 6. A log plot of ESnet traffic since 1990.

This evolution in traffic patterns and volume will result in the top 100 to 1,000 flows accounting for a very large fraction of the traffic in the network, even as total ESnet traffic volume grows. This means that the large-scale science data flows will overwhelm everything else on the network.

The current few Gb/s of average traffic on the backbone will increase to 40 Gb/s (LHC traffic; sidebar “Data Analysis for the Large Hadron Collider,” p54) and then increase to probably double that amount as the other science disciplines move into a collaborative production, simulation, and data analysis mode on a scale similar to the LHC. This will get the backbone traffic to 100 Gb/s as predicted by the science requirements analysis three years ago.

ESnet’s Evolution

The architecture of the network must change to accommodate this growth and the change in the

Based both on the projections of the science programs and the changes in observed network traffic and patterns over the past few years, it is clear that the network must evolve substantially in order to meet the needs of the DOE SC mission.

Data Analysis for the Large Hadron Collider

The major high energy physics (HEP) experiments of the next 20 years will break new ground in our understanding of the fundamental interactions, structures and symmetries that govern the nature of matter and space-time. Among the principal goals are to find the mechanism responsible for mass in the Universe, and the Higgs particles associated with mass generation, as well as the fundamental mechanism that led to the predominance of matter over antimatter in the observable cosmos.

The largest collaborations today, such as those building experiments for CERN's Large Hadron Collider (LHC; figures 2, 7–9) program, each encompass some 2,000 physicists from 150 institutions in more than 30 countries. The current generation of operational experiments—BaBar at the Stanford Linear Accelerator Center (SLAC) and D0 and CDF at FNAL, as well as the experiments at the RHIC program at BNL—face similar challenges. BaBar, for example, has already accumulated datasets approaching a petabyte.

Collaborations on this global scale would not have been attempted if the physicists could not plan on high capacity networks to interconnect the physics groups throughout the lifecycle of the experiment, and to make possible the construction of Data Grids capable of providing

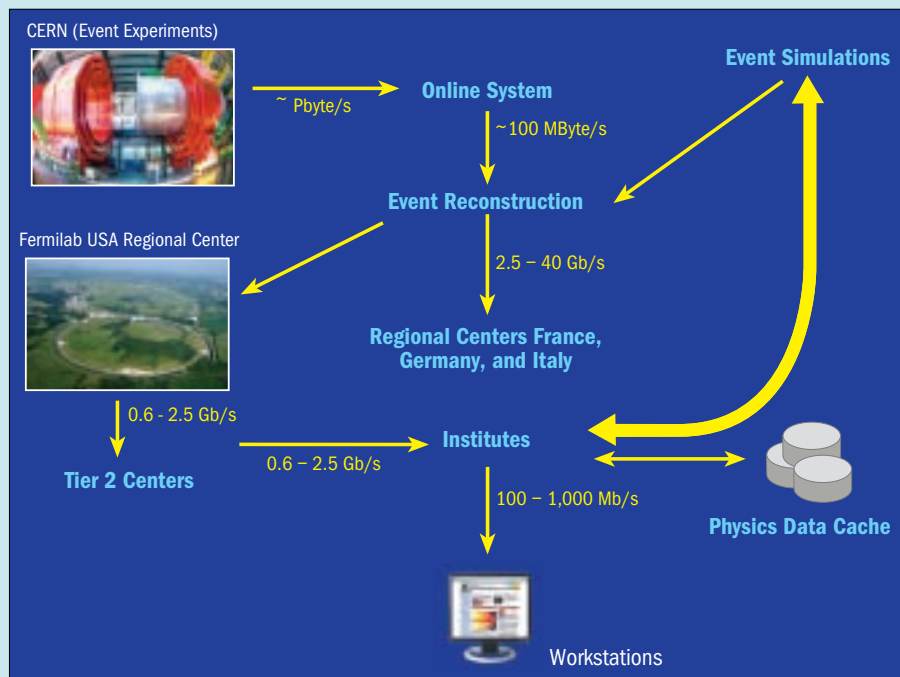


Figure 7. This chart depicts an example of how information flows from HEP facilities, such as CERN, to researchers at centers and institutes around the world. Such HEP experiments will be collecting, cataloguing, and analyzing several petabytes of data per year by 2009.

access, processing, and analysis of massive datasets. These datasets will increase in size from petabytes to exabytes (10^{18} bytes) within the next decade. Equally important is highly capable middleware—the Grid data

management and underlying resource access and management services—to facilitate the management of worldwide computing and data resources that must all be brought to bear on the data analysis problem of HEP.


The strategy for the next-generation ESnet is based on a set of architectural principles that lead to four major network elements and a new network service for managing large data flows.

types of traffic. The general requirements for the new architecture are that it provide: high-speed, scalable, and reliable production IP networking; connectivity for university and international collaboration; highly reliable site connectivity to support lab operations as well as science; global Internet connectivity; support for the high-bandwidth data flows of large-scale science, including scalable, reliable and very high-speed network connectivity to DOE labs; and dynamically provisioned, virtual circuits with guaranteed quality of service (for dedicated bandwidth and for traffic isolation).

In order to meet these requirements, the capacity and connectivity of the network must increase to include fully redundant connectivity for every site, high-speed access to the core for every site (at least 20 Gb/s, generally, and 40–100 Gb/s for some sites) and a 100 Gb/s national core/backbone bandwidth by 2008 in two independent backbones.


Meeting the Science Requirements

The strategy for the next-generation ESnet is based on a set of architectural principles that lead to four major network elements and a new network service for managing large data flows. One of the architectural principles involves using ring topologies for path redundancy in every part of the network, rather than just in the core. Another principle provides multiple independent connections everywhere to guard against hardware and fiber failures. A third principle provisions one core network—the IP network—specialized for handling the huge number (3×10^9 per month) of small data flows (hundreds to thousands of bytes each) of the general IP traffic. Provisioning a second core network, the Science Data Network (SDN), is the last architectural principle. The SDN is specialized for the relatively small number (hundreds to thousands) of massive data flows (gigabytes to terabytes each) of large-scale science, which by volume already accounts for 50%



CERN

Figure 8. Several images from the CERN high energy physics facility. At the upper left, a view of the LHC tunnel with a worker inside; at the upper right, detail of the sensor from the first half tracker inner barrel (TIB); the two lower panels are views of the ATLAS Experiment detector. Physicists depend on ESnet to transport data from HEP experiments to researchers around the world.



CERN

Figure 9. A simulated event of the collision of two protons in the ATLAS experiment viewed along the beam pipe. The colors of the tracks emanating from the center show the different types of particles emerging from the collision.

“We must build the second generation of the Internet so that our leading universities and national laboratories can communicate in speeds 1,000 times faster than today, to develop new medical treatments, new sources of energy, new ways of working together ...”

PRESIDENT WILLIAM J. CLINTON
1997 State of the Union Address

Science Services

Authentication and Trust Federation

Cross-site identity authentication and identity federation is critical for distributed, collaborative science in order to enable routine sharing of computing and data resources, as well as other Grid services. ESnet provides a comprehensive service to support secure authentication. Managing cross-site trust agreements among many organizations is crucial for authorization in collaborative environments. ESnet assists in negotiating and managing the cross-site, cross-organization, and international trust relationships to provide policies that are tailored to collaborative science.

Public Key Infrastructure

Grid computing and data analysis systems rely on Public Key Infrastructure (PKI) for their security. ESnet provides PKI and X.509 identity certificates that are the basis of secure, cross-site authentication of people and Grid systems. The ESnet PKI is focused entirely on enabling science community resource

sharing and its policies are driven entirely by the science communities that it serves. That is, the trust requirements of the science communities are formally negotiated and encoded in the Certification Policy and Certification Practice Statement of the Certification Authorities.

Voice, Video, and Data Tele-Collaboration Service

The human communication aspect of collaboration, especially in geographically dispersed scientific collaborations, represents an important and highly successful ESnet Science Service. This service provides audio, video, and data teleconferencing with the central scheduling essential for global collaborations. The ESnet collaboration service supports more than a thousand DOE researchers and collaborators worldwide with video conferences, audio conferencing, and data conferencing. Web-based, automated registration and scheduling is provided for all of these services.

of all ESnet traffic and will completely dominate it in the near future.

These architecture principles lead to four major elements for building the new network.

The first element is a high-reliability IP core network based on high-speed, highly capable IP routers to support internet access for both science and lab operational traffic and some backup for the science data carried by SDN, science collaboration services, and peering with all of the networks needed for reliable access to the global Internet. The second element involves an SDN core network based on layer 2 (Ethernet) and/or layer 1 (optical) switches for: multiple 10 Gb/s circuits with a rich topology for very high total bandwidth to support large-scale science traffic and for the redundancy needed to ensure high reliability; dynamically provisioned guaranteed bandwidth circuits to manage large, high-speed science data flows; dynamic sharing of some optical paths with the R&E community for managing peak traffic situations and for providing specialized services such as all-optical, end-to-end paths for uses that do not yet have encapsulation interfaces (such as InfiniBand); and an alternate path for production IP traffic. The third element, Metropolitan Area Network (MAN) rings, connect labs to the core(s) to provide more reliable (ring) and higher bandwidth (multiple 10 Gb/s circuits) site-to-core connectivity, support for both pro-

duction IP and large-scale science traffic, and multiple connections between the SDN core, the IP core, and the sites. The fourth element is composed of loops off the core rings to provide for dual connections to remote sites where MANs are not practical.

These elements are structured to provide a network with fully redundant paths for all of the SC Labs. The IP and SDN cores are independent of each other and both are ring-structured for resiliency. These two national cores are interconnected at several locations with ring-structured MANs that also incorporate the DOE labs into the ring. This will eliminate all single points of failure except where multiple fibers may be in the same conduit, as is frequently the case between metropolitan area points of presence and the physical sites. In the places where metropolitan rings are not practical, such as for the geographically isolated labs, resiliency is obtained with dual connections to one of the core rings.

Services Supporting Science

The evolution of ESnet is being guided by the results of several studies to determine the key requirements of the DOE research community. These studies identified various middleware services that, in addition to the network and its services, need to be in place in order to provide an effective distributed science environment.

History of ESnet

ESnet has been supporting the scientific missions of DOE for nearly 20 years. ESnet was originally developed to serve the needs of the HEP and Magnetic Fusion Energy (MFE) research communities, which had been dependent upon their own networks, HEPnet and MFEnet, respectively. In the mid-1980s, HEP and MFE researchers recognized the need for substantially improved computer network facilities. In response, Dr. Alvin Trivelpiece,

then Director of Energy Research, recommended that HEPnet and MFEnet initiatives be combined in order to provide a more functional and efficient network to serve the Energy Research community, in general. The Energy Sciences Network Steering Committee was formed, and an evolutionary model—endorsing a phased approach aimed at long-term goals—for the development of the new network was proposed.

ESnet was initially constructed to link U.S. centers of Energy Research, but then also began providing international connectivity in collaboration with the NSF, NASA, and the Defense Advanced Research Projects Agency (DARPA). Now an integral part of the Internet, ESnet is still expanding, currently with the upgrade to ESnet4. More details about the history of ESnet are available online, at: <http://www.es.net/hypertext/esnet-history.html>

These services are called “science services,” and are simply services that support the practice of science. Examples include trust management for collaborative science, cross-site trust policies negotiation, long-term key and proxy credential management, human collaboration communication, end-to-end monitoring for Grid/distributed application debugging and tuning, and persistent hierarchy roots for metadata and knowledge management systems.

Because of its established characteristics, ESnet is a natural provider for a number of these services. For example, ESnet is trusted, persistent, and has a large (nearly comprehensive within DOE) user base. ESnet also has the facilities to provide reliable access and high availability of services through assured network access to replicated services at geographically diverse locations.

However, given the small staff of an organization like ESnet, a constraint on the scope of such services is that they must be scalable in the sense that, as the service user base grows, ESnet interaction with the users does not grow.

There are three such services that ESnet provides to the DOE and/or its collaborators (sidebar “Science Services”). Federated trust is policy established by the international science collaboration community to meet its needs. Public Key Infrastructure certificates provide remote, multi-institutional identity authentication. Human collaboration services involve technologies such as video, audio, and data conferencing.

Conclusions

ESnet can trace its origins to a dialup modem service provided to users of the Magnetic Fusion Energy Computer Center, known today as the NERSC Center. Over the years, remote access terminals replaced the dialups, and leased telephone lines were deployed and then replaced with satellite connections. In 1985, DOE responded to the growing demand for networking by combining separate

fusion energy and high energy physics networking initiatives to lay the foundation for ESnet.

Today, as an integral part of the DOE SC, ESnet provides seamless, multiprotocol connectivity among a variety of scientific facilities and computing resources in support of collaborative research, both nationwide and internationally. High-performance computing has now become a critical tool for scientific and engineering research. In many fields of research, computational science and engineering have become as important as the more traditional methods of theory and experiment. Additionally, the construction of large experimental facilities used by international collaborations has driven requirements for large-scale data transfer, often to multiple sites at the same time. Progress and productivity in such fields depend on interactions between people and machines located at widely dispersed sites, interactions that can only occur rapidly enough via high-performance computer networks. The ubiquity of networks has provided researchers with unexpected capabilities and unique opportunities for collaborations.

These benefits have only whetted the scientific community’s appetite for still higher levels of network performance to support wider network usage, the transmission of ever-greater volumes of information at faster rates, and the use of more sophisticated applications. The scientific community is also increasingly sensitive to the importance of securely protecting privacy and intellectual property. The mission of ESnet is to satisfy these needs as fully as possible for DOE researchers and, with its new architecture, ESnet is working hard to meet these needs for years to come.

Contributor: Dr. William E. Johnston is the ESnet Department Head and Senior Scientist at LBNL

Further Reading
<http://www.es.net/>

Today, as an integral part of the DOE SC, ESnet provides seamless, multiprotocol connectivity among a variety of scientific facilities and computing resources in support of collaborative research, both nationwide and internationally.