

The MyESnet Portal: Making the Network Visible

Authors

Jon M. Dugan, Gopal Vaswani, Gregory Bell, Inder Monga

Energy Sciences Network (ESnet)

Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720-8146, USA

e-mail: jdugan@es.net, gvaswani@es.net, grbell@es.net, imong@es.net

Keywords

User Portals, Network Measurement, Network Visualization, Network Statistics

Abstract

ESnet provides a platform for moving large data sets and accelerating worldwide scientific collaboration. It provides high-bandwidth, reliable connections that link scientists at national laboratories, universities and other research institutions, enabling them to collaborate on some of the world's most important scientific challenges including renewable energy sources, climate science, and the origins of the universe.

ESnet has embarked on a major project to provide substantial visibility into the inner-workings of the network by aggregating diverse data sources, exposing them via web services, and visualizing them with user-centered interfaces. The portal's strategy is driven by understanding the needs and requirements of ESnet's user community and carefully providing interfaces to the data to meet those needs. The '[MyESnet Portal](#)' allows users to monitor, troubleshoot, and understand the real time operations of the network and its associated services.

This paper will describe the MyESnet portal and the process of developing it. The data for the portal comes from a wide variety of sources: homegrown systems, commercial products, and even peer networks. Some visualizations from the portal are presented highlighting some interesting and unusual cases such as power consumption and flow data. Developing effective user interfaces is an iterative process. When a new feature is released, users are both interviewed and observed using the site. From this process valuable insights were found concerning what is important to the users and other features and services they may also want. Open source tools were used to build the portal and the pros and cons of these tools are discussed.

Motivation

Nearly all areas of scientific research have become distributed, collaborative and data-intensive pursuits. The Energy Sciences Network (ESnet) [1] is charged by the US Department of Energy (DOE) to provide the network resources that allow scientists, researchers and staff members at the DOE laboratories to collaborate with scientists all over the world. The collaborations of modern science go far beyond the boundaries of a single institution, agency or even nation and ESnet strives to provide the best connectivity possible between the DOE labs and the rest of the world. To that end, ESnet has operated a dedicated, reliable, high-speed network for over 25 years.

A fast, reliable network is necessary for modern science, but unless scientists can use that network effectively to accelerate their work, ESnet is not fulfilling its mission. ESnet has long been actively involved in ensuring the network is usable end-to-end – even if the other end terminates in another network domain. This means looking beyond the edges of the network and exploring technologies that can help fulfill this mission. ESnet works closely with network engineers and researchers at institutions all over the world to ensure good end-to-end

performance. ESnet has developed technologies such as OSCARS [2] to provide dedicated bandwidth on demand to users. ESnet has also worked with worldwide collaborations such as perfSONAR [3,4] to expose internal network statistics to other organizations to facilitate debugging of network connections end-to-end. The portal is a logical extension of these concepts because it provides an even more detailed view into the ESnet infrastructure.

Providing this level of transparency can be a bit daunting at first. This kind of fear is endemic to all human pursuits – the author who is hesitant to share their manuscript before it is done or the programmer who doesn't want to share their code because it isn't perfect yet. However, the history of collaboration shows that being open allows others to provide feedback that they would not otherwise be able to provide. By exposing more detail about the operation of ESnet to the outside world it will be possible to provide better service to the users and to solve networking problems more quickly.

ESnet constantly collects comprehensive data about the operation of its network. This data comes in many forms, and at first it can seem overwhelming even for those who have extensive networking experience. The MyESnet portal goes beyond simply aggregating this data by providing clean, intuitive interfaces to this data. A great deal of time and effort is spent designing the visualizations. Throughout this design process the user community is consulted to ensure the visualizations show the most useful and relevant data.

Fundamentally, visualizations are a means for understanding and better interpreting raw data. Since many users in the ESnet community are not network data experts, the visualizations have become the heart of the MyESnet portal strategy. While the goals of the portal are to provide clean and intuitive interfaces to the data, the data is very useful in its own right. Having the data well organized and easily accessible greatly simplifies the process of developing new visualizations. Further, some users and researchers may wish to use the data to correlate with their own measurements, some sites may find it useful to incorporate it into their own tools, and others may find new and novel ways to visualize the data. To enable these uses of the data, it has been made accessible in machine-readable formats over the network.

Features and Functions of the MyESnet Portal

This section describes the current features of the MyESnet portal and how they relate to various segments of the ESnet user community. Due to the highly interactive nature of the portal it may be beneficial to visit the portal at <https://my.es.net/>. Screenshots showing various features are included below for ease of reference.

Currently the portal shows information about:

- traffic loads on site facing interfaces
- various breakdowns of network flows involving each site
- monthly and yearly breakdowns of site availability
- scheduled maintenance events
- trouble tickets
- OSCARS circuits
- the progress of new network build outs
- power utilization
- information about demonstrations that ESnet supports.

For each DOE site that ESnet connects, a summary of high-level information is presented to the user. From this summary page a user can delve into more details about the site. Figure 1 shows the site summary page for Fermi National Accelerator Laboratory (Fermilab), a US National Laboratory dedicated to high-energy physics. This summary page highlights some motifs that are repeated many times throughout the site. Most notable is the “apples to apples” comparison between the two graphs. The y-axis range is the same for both graphs making it easy to tell what fraction of the total traffic load is carried by virtual circuits (in the figure it is labeled ‘OSCARS’ traffic). When a user places their mouse over a portion of the graph a tracker is displayed which gives a numerical readout of the values on all graphs at that point in time. In the figure the user has placed their

mouse over a data point, which provides a reading of 26 Gigabits per second (Gbps). From this page a user can navigate to other pages to get more detailed information about the site.

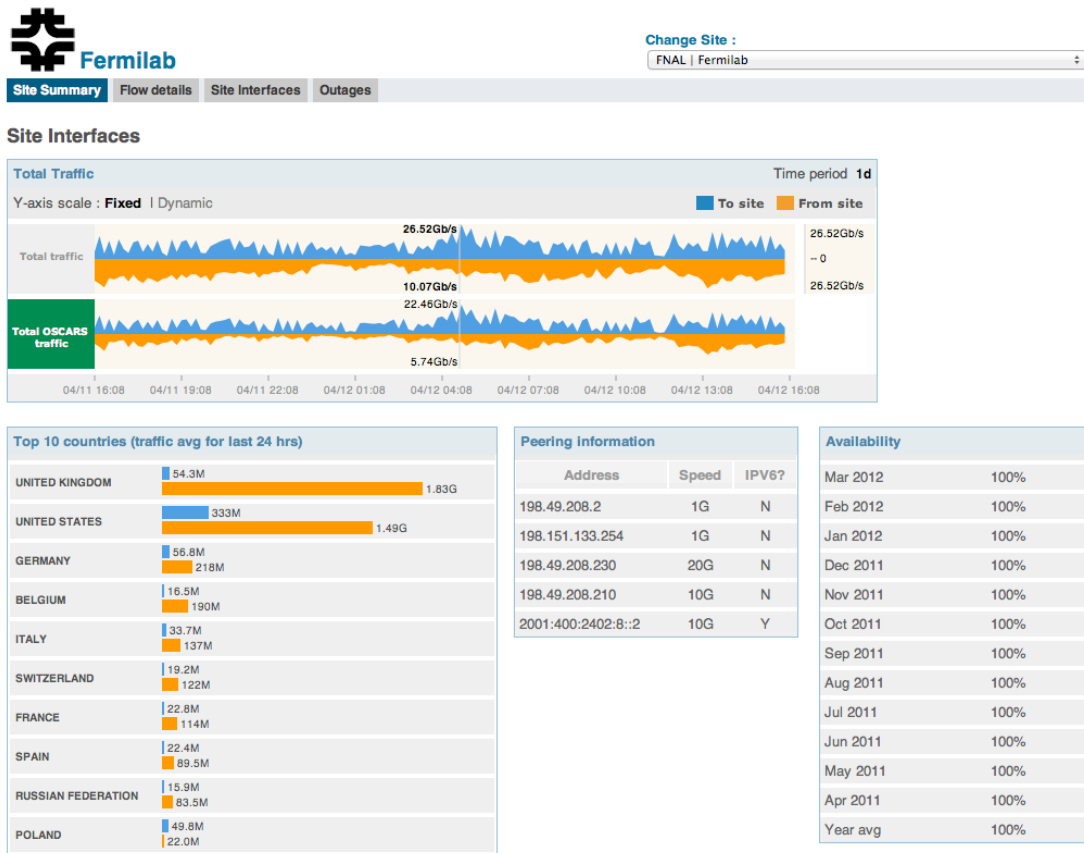



Figure 1 – Site Summary Page

Network flow data is broken down by protocol, application, source autonomous system, next hop autonomous system and top talkers. The network flow information is retrieved from an Arbor PeakFlow SP [5] system using its web services interface. For users affiliated with the site there are two additional traffic breakdowns not available to the public. The breakdowns show the top talkers into and out of the site. Due to the potentially sensitive nature of this information the last octet of the IP address is redacted when shown to the general public.

Flow Details

From : Apr 10 14:41 To : Apr 11 14:41

To site From site


Y-axis Scale

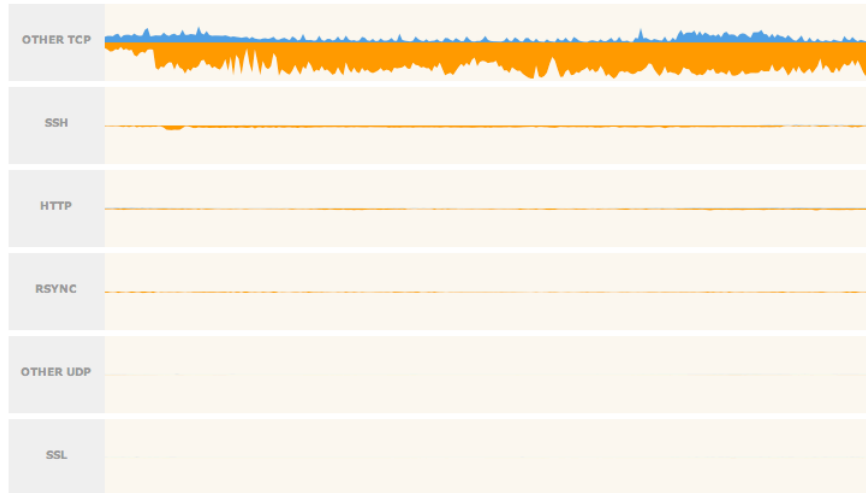
Fixed | Dynamic

Time range

Period: 1d | 7d | 30d | custom

Traffic split by : 'Application'

Apr 11 00:10


Split Traffic by :

Protocol

Application

Autonomous System (origin)

Autonomous System (peer)

Country

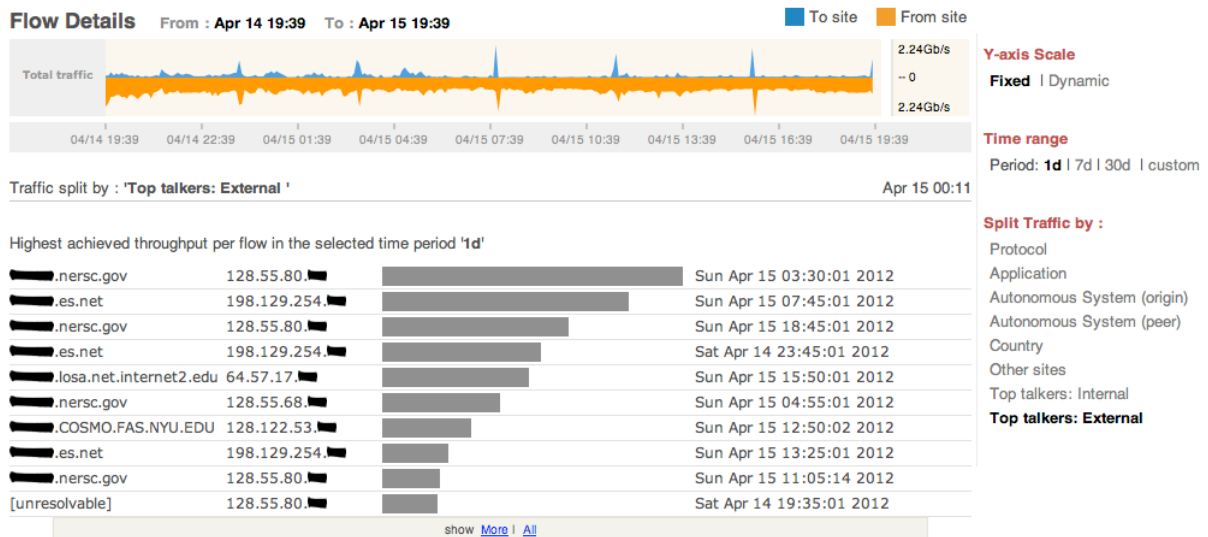
Other sites

Top talkers: Internal

Top talkers: External

Figure 2 – Flow details page

Figure 2 shows the application breakdown for Fermilab and Figure 3 shows the top talkers visualization. Once again in this sample screenshot it is clear that *Other TCP* traffic dominates the traffic mix with SSH, HTTP and RSYNC making up a smaller percentage of the traffic. Just like the site summary page, the graphs use the same range on the y-axis and have the tracker feature. To the right of the graphs there are various options for changing what is displayed including selecting which traffic breakdown is displayed, changing the time range of the graph, and giving each graph a y-axis with a dynamic range. By selecting the dynamic y-axis scale the user can see finer details in the traffic pattern. This mode, however, limits the ability to do comparisons of magnitude amongst the graphs.


Figure 3 – Top talkers page

The traffic load to and from the sites is derived from SNMP counter data. The top of the page shows a summary graph for the total traffic load into and out of a given site. A second graph is also displayed showing the fraction of the total traffic that is delivered via OSCARS. Once again these graphs have the tracker feature and by default use the same range on the y-axis to allow for easy comparisons. As with the flow page there are various filtering options to change what is displayed on the graphs. Due to the similarity to the flow page no screenshot is shown, however the page can be seen online at <https://my.es.net/site/fnal/interface/>.

Another page shows details about upcoming and past outages both planned and unplanned. This page derives its data from an in-house planned maintenance calendar [6] which collects and tracks information about outage events and their impact on the network. There is linkage between these outage and maintenance events and the trouble tickets associated with them.

Many sections of the portal show visualizations of the whole of ESnet as opposed to a single site. These sections include a matrix showing the overall availability of all of the sites within ESnet, a list of all currently active OSCARS circuits along with a graphical display of their path, details about network build outs, power utilization, a matrix of perfSONAR derived measurements to critical collaborators and details about demonstrations that ESnet is supporting or has supported.

ANI Prototype Network, Phase I

The network in this phase will fulfill the base requirements of the ANI project and additionally deliver 100G to Supercomputing 2011.

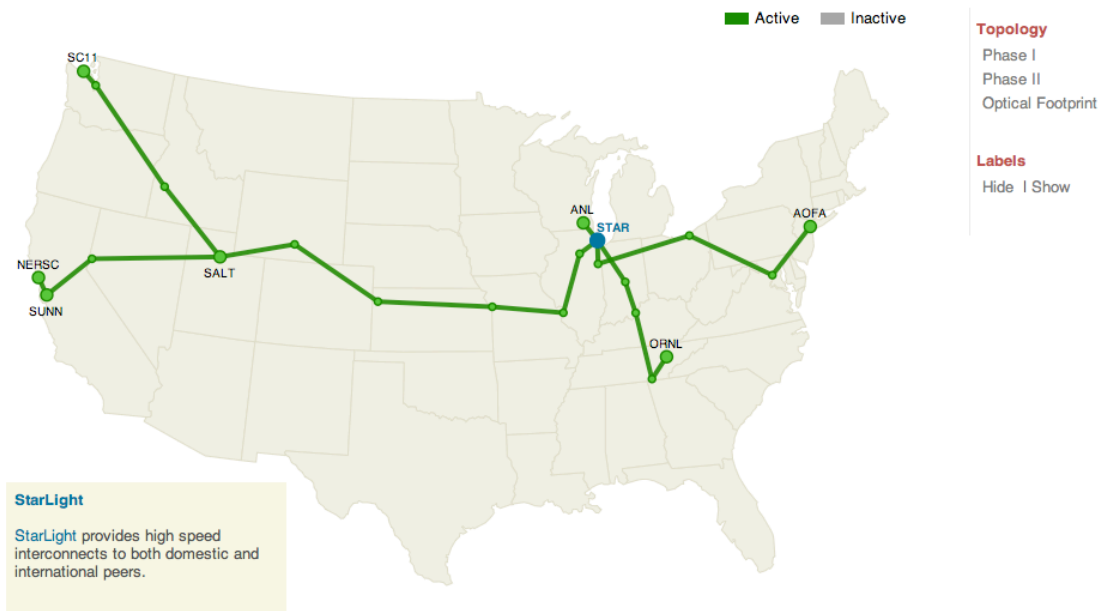


Figure 4 – Network build out status for ANI Phase I

In 2009, DOE launched Advanced Networking Initiative (ANI) [17] through a \$62 million investment as part of the U.S. American Recovery and Reinvestment Act (ARRA). ANI has enabled ESnet to build a brand new, state of the art network platform for accelerating scientific discovery in this new age of data intensive research. ANI is providing the funds for ESnet to upgrade its national network to 100 Gigabit per second in capacity as well as create a separate national test bed for network research experimentation. These resources will be built on top of a national fiber optic footprint that will be for the first time owned and operated by DOE.

To help the ESnet community keep track of the new network build, the “Network Build Out” page depicts the footprint of the new ESnet network on the continental United States. The maps are data driven and can be updated dynamically. Making the maps dynamic has allows the highlighting of which sections are built and which are not yet built. It also allows for extended interactivity and makes changes relatively straightforward. As the ANI project has progressed both the optical and logical progression of the network has been displayed. Figure 4 shows an example of the ANI Phase I build out including an additional link that was created to connect the Supercomputing 2011 conference held in Seattle, Washington in November 2011.

The power utilization visualization is a work in progress. Currently, it shows the power consumption of ESnet’s routers within the first phase of the ANI network build. As the rest of the new ESnet infrastructure is deployed, facilities for collecting power data for all devices both at Layer 3 and at the optical layer will be installed. The current visualization shows the live power used and also shows the number of joules per bit. The joules per bit calculation is an attempt to characterize energy spent per unit of useful work. This visualization makes it very clear that network equipment vendors have a lot of room to improve in terms of power proportionality. This type of data has not been published for any other network; as a result this type of visualization has not been attempted before. It presents an interesting baseline for potential research in this space.

es.net-817 | FNAL - MIT, VLAN 3001 | 02-25-2009 To 02-25-2019




Figure 5 – OSCARS details page

Figure 5 shows the visualization of an OSCARS circuit including a diagram of which routers and interfaces carry the traffic for a particular virtual circuit and a graph of the utilization of the circuit. Presently only currently active circuits can be visualized but there are plans to add support for viewing information on circuit reservations that have been completed. This OSCARS visualization took a great deal of inspiration from the “subway map” views used in the E-Center [7] project to visualize network

paths

PerfSONAR Tests

Select Test esnet geant tests

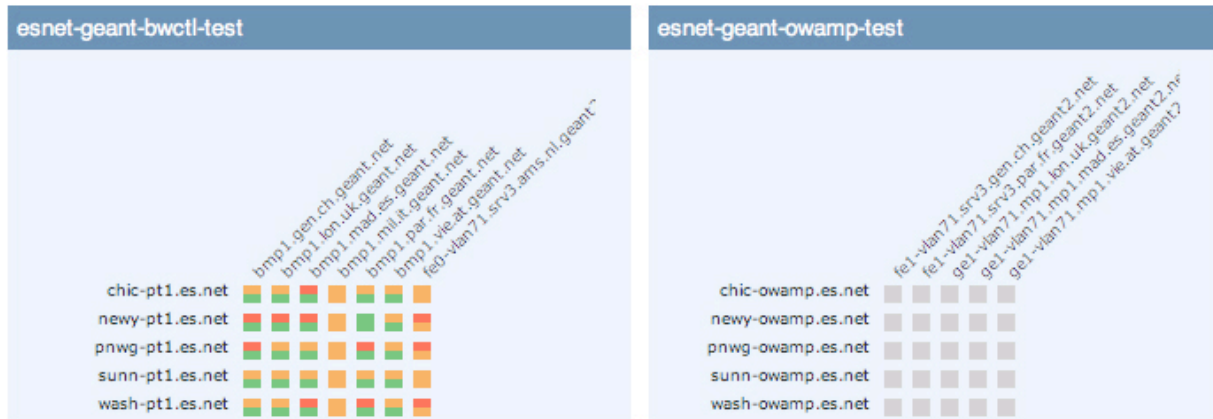


Figure 6 – Matrix showing test results to GÉANT

Figure 6 shows a matrix of BWCTL [21] and OWAMP [22] measurements between ESnet and GÉANT test hosts. This data is extracted from perfSONAR and is an example of how the portal highlights key collaborations and peers. In the diagram above each cell is split in half and each half shows the results for one direction of a bidirectional measurement. Future work on the portal will add many more capabilities for displaying relevant data about the communities and collaborations in which ESnet participates. Eventually it will be possible for a user outside of ESnet to compose a page for a specific community and share it via the portal.

Scaling the Earth System Grid to 100Gbps Networks (LBNL)


Presenters: Alex Sim and Mehmet Balman (LBNL Computational Research Division)

Venue: Booth 512


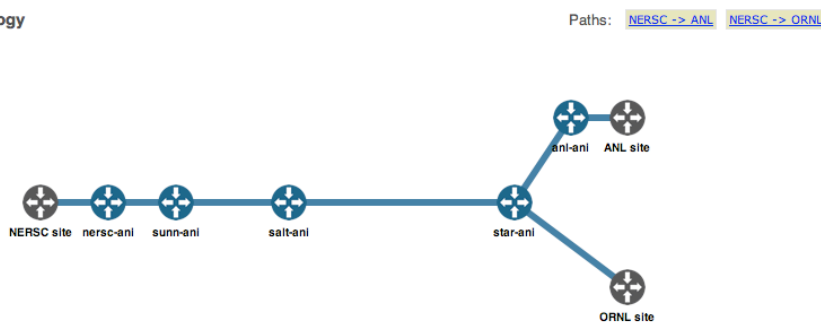
Climate change research is one of the critical data intensive sciences, and the amount of data is continuously growing. Climate simulation data is geographically distributed over the world, and it needs to be accessed from many sources for fast and efficient analysis and intercomparison of simulations. This demonstration will leverage a 100 Gbps link connecting the National Energy Research Scientific

[more](#)

Live demo stats



Topology



SC11

- Overview
- SC Traffic
- Demonstrations

Demo details

Scaling the Earth System Grid to 100Gbps Networks (LBNL)

Visualizing the Universe at 100 Gbps (LBNL)

Global Range 40Gbps InfiniBand Demo (IBTA)

Demonstration of High Performance 100 Gbps Networking for Petascale Science (NASA)

The new HPC cloud: Data-intensive scientific applications, Lustre-WAN, and OpenFlow over 100Gb Ethernet (Indiana)

RFTP: RDMA-based file transfer application for ultra-high speed networks (BNL)

End-to-End Virtualization: Campus, WAN and Data Center (ESnet/UDel)

Just in Time Data Placement for LHC Analysis (FNAL/UCSD)

40G LHC Data Transfer Between SC11 and Univ Michigan (Caltech)

Dynamic, On-demand Resources and Provisioning (RENCI)

Figure 7 – SC11 Demo Page

Part of supporting the DOE science community involves technology demonstrations. Supercomputing – the International Conference on High Performance Computing, Networking, Storage and Analytics (usually referred to as SC) is one of the most prominent showcases for the work of the scientists and researchers in the DOE complex. As a result ESnet has a large presence at the conference and we use the portal to showcase our efforts at SC. Figure 7 shows the details of a demonstration that ESnet supported for SC11. This visualization shows live traffic as well as a trend line throughout the week of SC. (The demos had limited time windows in which to use the shared 100G circuit.) It also gives a textual description and a diagram of the network topology of the demo.

User Community

Up to this point the portal has been primarily designed with the ESnet site coordinators (the primary points of contact between the sites and ESnet) in mind. Site coordinators have a high degree of technical sophistication, especially as it pertains to networking. The site coordinators will always be one of ESnet’s most important communities but we aspire to make the portal usable to scientific end users of ESnet who have little or no network expertise.

In order to be usable to a broader audience nearly all of the content on the portal is visible without using a login. The only areas that we have limited access to are those that have been deemed to be potentially sensitive because they pose security concerns. Creating a login on the portal does however give the user some ability to customize their experience. In the future, users will have greater ability to customize their experience. In order to be open and participate fully in the R&E community we have added federated logins via Shibboleth.

Relationship with perfSONAR

The portal is a complementary, not a competing effort to perfSONAR. perfSONAR focuses on sharing a specific, crucial set of measurements across administrative boundaries. The portal– on the other hand– is focused on providing a wide variety of data and information about ESnet to make ESnet more transparent. Some data already comes from perfSONAR and much of the data the portal makes use of is also distributed via perfSONAR. The visualizations used in the portal deal with network data making it very easy to use them with network data obtained from perfSONAR. The experience gained by developing the portal will also be used to inform suggestions for improvements to perfSONAR in the future.

Presently the portal is making only limited use of perfSONAR data and would greatly benefit from more integration with perfSONAR. A key impediment to using more perfSONAR-collected data is the lack of language bindings for Python, the language used to develop the portal. There are, however, a few approaches that may help to solve this issue. First, the DRS service developed as part of the E-Center project harvests and aggregates perfSONAR data and makes it accessible via a REST interface. A discussion with the E-Center developers is in progress to determine the viability of this approach. Secondly, multiple groups are planning to add REST interfaces to perfSONAR which would support greater integration. These various approaches will help the portal make better use of perfSONAR data.

It takes time to add new kinds of data to perfSONAR since it is developed by the consensus of a large community. In order to provide access to data that is not yet available via perfSONAR and to users that find it difficult to get data using perfSONAR, ESnet has undertaken the development of the ESnet API.

The ESnet API: api.es.net

Well-defined interfaces to data help enormously when building tools to visualize data. The earliest data sets visualized by the portal were SNMP and flow data because the backend systems had usable APIs. Over time it became clear that having easy access to all of the relevant data within ESnet would be crucial to building future visualizations and to keeping the data itself fresh. This led to the development of the ESnet API which will be referred to as *api.es.net* below.

Presently data is collected together for consumption in the portal using multiple approaches. In some cases local copies of the data are kept in a portal specific database. In other cases data is proxied from backend services and cached by the portal. While a clever user could access the raw data using these internal APIs, they are not well structured and have a distinctly portal centric viewpoint. `api.es.net` provides a cleaner and broader view of the data set. Ultimately the portal will use `api.es.net` as its source of truth. It remains to be seen if for performance reasons if it will be necessary for the portal to cache data it consumes from `api.es.net`.

In some cases within ESnet there were multiple versions of manually maintained data. As a result it was easy for there to be inconsistent or missing data in some versions. This caused a great deal of confusion and lost time. Many of ESnet's tools coped with this by keeping a local mapping of one set of names to another which only compounded the problem. Both for the portal and for other internal tools, having access to a consistent view of the data is crucial.

The list of ESnet sites is a good example of how the API can deliver a consistent set of data. It may seem trivial but this set of names (and more specifically the list of abbreviations) is a key that allows many data sources within ESnet to be correlated with each other. Historically this had been maintained manually in several databases. With the API in place all tools have access to the same list of sites. This list only needs to be updated at the source of truth: `api.es.net`. As systems are automated it is extremely useful to have a naming convention that is consistently applied; this allows for machine correlation without the aid of human intuition. This list of sites is used in the portal, the flow collection system, monthly statistics reporting, outage and preventive maintenance tracking, incident management, and reporting to our funding agency. Centralizing this in the ESnet API is just one example of how having a central source of truth is a big win.

All access to the ESnet API is through a RESTful interface to our data. REST is a concept that came out of Roy Fielding's dissertation [8]. There has been much debate on how to apply and interpret the REST principles. `api.es.net` has been developed with a pragmatic rather than dogmatic approach to REST. RESTful API Design, Second Edition [9] provides a detailed and pragmatic overview of how to design RESTful APIs. For the purposes of this paper REST means the exposing of resources over HTTP so that they can be queried and when appropriate modified. `api.es.net` uses JSON [20] representations of the resources. The REST framework used by the portal supports alternative representations such as XML but JSON is the primary representation.

Why HTTP and JSON? HTTP and JSON are usable from nearly any environment and have become a de facto standard. XML has a broad base of support but is much more cumbersome to work with and has been losing mindshare relative to JSON. One great strength of JSON is that it deals in basic data types and collections which are familiar to all programmers. Originally JSON was intended to make data easy to work with in JavaScript but JSON libraries provide a very natural way to work with data in other languages. Additionally the rise of client side development using JavaScript in the browser has made JSON a de-facto standard.

Demanding that every tool use a particular library or language is not practical, especially in the ESnet environment where there are tools written in several programming languages including Python, Perl and Java. This allows each tool to be written in the language best suited for the task because nearly every programming language has libraries that support HTTP and JSON .

The API is read-only for most resources but there are some cases where it may make sense to allow external users to make changes to the data. The most obvious case for this is allowing user communities to define sites of interest to build a dashboard for their community. This would require managing permissions and access. OAuth [23] will be used to manage authorization in conjunction with local accounts, OpenID [24] or Shibboleth [25] for authentication.

Design Principles

The portal is guided by a few principles, including:

- users come first

- simplicity is key
- develop new features rapidly for quick use by the community
- feedback from the community is critical for guiding future development

The Agile programming movement [18] has greatly influenced the way the portal and api.es.net are developed, however the principles are applied pragmatically rather than dogmatically.

We have strived to follow ‘User centered design’ principles to make the portal useful and easy to use for our end users. For any feature in our portal we start with a couple of fundamental questions: “Who are the users?” and “What are their needs?” We try to understand the daily activities and goals of our users and use various methods like surveys and face-to-face interviews to record their needs and requirements. In the initial phase, we have focused on the most core community of ESnet, site coordinators. These are the people at each site that act as a liaison between the site and ESnet. Even though the site coordinator is the primary persona at present, most features are useful to a much broader user base which is regularly solicited for feedback.

Once an initial set of user requirements are gathered, multiple low fidelity prototypes are developed which are used as basis for discussion. Multiple iterations are made to improve these prototypes based on internal and external feedback.

The next step involves charting out the detailed design and development tasks for the feature. This also involves identifying the data sources, possible adjustment of the data model and exposing it through standard API interfaces.

After a feature is developed, it is presented to test users and feedback is gathered which is used to refine the features for production. The approach for gathering feedback is both asking for input and watching the users interact with the site. Often times the observations as to how users use the site are just as fruitful if not more so. In particular, many times it has become very clear that things which were obvious to the designers were not at all obvious to the users. Because of the tremendous value in both interviewing and observing the users, significant effort is made to meet with as many users as possible.

The ESnet site coordinators meet twice annually to share best practices, discuss challenges and develop common strategies to tackle these challenges. At the upcoming site coordinators meeting in July 2012, a session will be held to gather feedback on the portal and its development thus far.

Furthermore, the broader structure of the portal is revisited every few months. The entire information architecture of the portal is analyzed through the lens of the new features added. Features are organized into logical groupings based on user goals and networking domain. By doing this it is possible to ensure the usability of the portal as it grows in size and scope.

In order to illustrate the process used the details of implementing the OSCARS circuit visualization are detailed below.

The portal developers approached OSCARS team to understand which important features should be highlighted in this visualization. Through a series of interviews (and based on prior knowledge) the key tasks for monitoring the circuit were identified. A lightweight mockup of the visualization was created in Adobe Illustrator. This allowed the rapid exploration of different options based on the requirements identified earlier.

This document along with detailed annotations became the “visual requirements document” which served as the basis of feature design. This prototype was presented to the OSCARS team and through a series of iterations it was refined based on their feedback. This lightweight mockup, essentially a detailed drawing, allowed the exploration of all the possible options without the constraints imposed by the browser environment.

Once the team was satisfied with the mockup, the feasibility of the designs was evaluated based on the OSCARS circuit data available and the deadline. In the spirit of “Agile development” it was decided that it would be best to roll the feature out in a few discrete steps. The first step involved visualizing only circuits for

which both end-points within ESnet; the second step will extend the visualization to work with circuits that have only a single end-point within ESnet.

At this point effort was split into two parallel lines of execution. One was to focus on the visualization development using the sample data. The other was to create an internal interface to provide the data about all available circuits. This allowed each developer to work at their own pace. Once both portions of the implementation were finished, they were merged to form the working prototype which was released to the OSCARS team for testing.

At this point there was a series of iterations, during which various aspects of the user interface and visualization were refined and readied for production. Once the visualization was deemed complete it was deployed to the public instance of the portal. The visualizations used in this feature were packaged as configurable widgets for later reuse and sharing with the wider community.

Data Sources

While there is a great deal of data available about the network, it is often hard to interpret in its raw form and it is often not easily accessible. Making the data easier to interpret is covered elsewhere in this paper. In this section the latter challenges are addressed particularly the opportunities presented by cataloging and organizing the data within ESnet.

When development of the portal began a significant amount of time was dedicated to creating a catalog of all the data sources available within ESnet – referred to as the ESnet Data Atlas. It captures each data source within ESnet including the kind of data, how the data is collected or produced and how to access it.

Data	Description	Access Method
SNMP	SNMP Measurements	REST, SOAP (perfSONAR)
Flow data	Netflow from Arbor and homegrown system	REST, custom
Topology	Network topology information	custom (flat files)
BGP Peering	Information about BGP peers	REST, custom (flat files)
OSCARS	OSCARS Reservation database	SOAP
Contact Info	Contact information for sites, peers and vendors	MySQL
Maintenance/Outage	Planned maintenance and unplanned outage descriptions	custom (flat files)
Status Monitoring	Status of routers and hosts	custom, SOAP, REST
syslog	logs from routers and hosts	custom (flat files)
Trouble tickets	Incident reports as reported to the helpdesk	custom
Security data	Intrusion detection, network scanners	custom
DNS	DNS entries	flat files, proprietary database

Table 1 – Partial list of ESnet data sources

Table 1 provides a sampling of the kinds of entries that are documented in the ESnet Data Atlas. Compiling this catalog provided a starting point for determining what could be visualized in the portal. It also helped identify which pieces might be more easily visualized because they already had a well-defined interface for accessing them.

Development Cycle

The creation of the portal and api.es.net has over time suggested a pattern or cycle of development (Figure 8). While not a particularly original pattern, this defined cycle has helped the development process remain focused and forward moving. The team endeavors to stay on a two-week cadence, resulting in the release of a new feature or enhancement every other week.

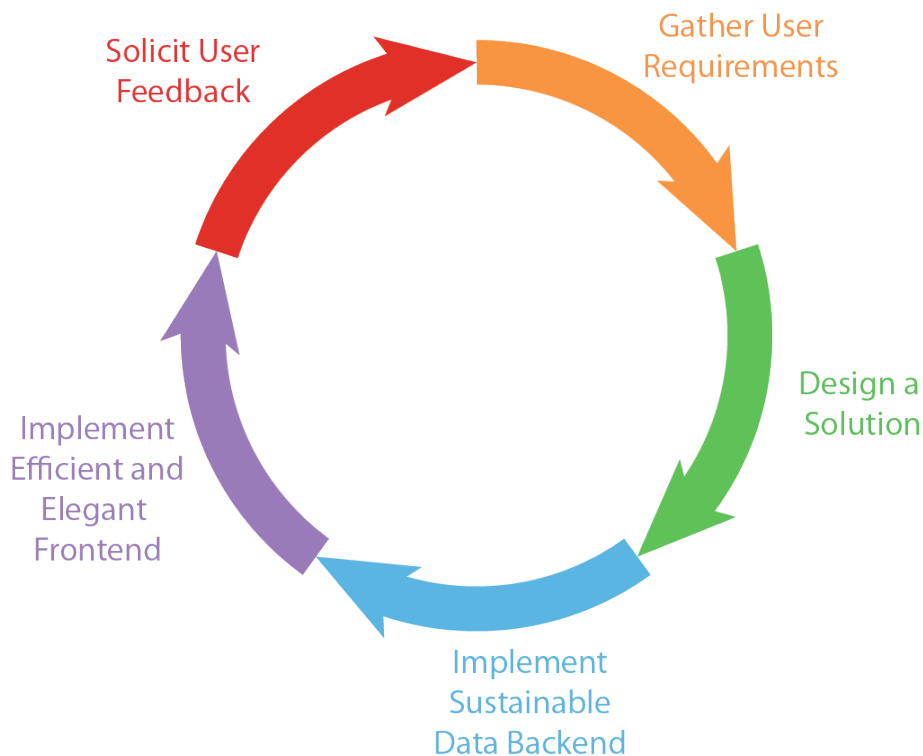


Figure 8 – Portal development cycle

Following the principle of putting the user first, development begins by gathering requirements. Initially all of the use cases and potential requirements are considered and recorded. Once all of the requirements are on the table they are prioritized. Some features are manageable in a single step and some get broken into two or more steps. It is better to make regular incremental progress rather than trying to take large quantum leaps forward.

Once a clear set of requirements is defined, the team works on designing the feature both from a data and a front-end perspective. At this stage mockups and example data sets are used to try to understand how each of the pieces fits together.

Once a solution has been designed we move into implementation. The implementation generally is broken into two distinct steps: the data backend and the user frontend.

The early versions of our data backend was fairly ad hoc and inflexible which in turn made other parts of the code complex and confusing. A commitment has been made to take the time to implement a backend that works well. api.es.net is a direct result of the commitment to a quality data backend. Building a quality data backend requires keeping data up to date, delivering the data quickly to the client browser and properly organizing the data. These issues are explored in more detail in the section about the ESnet API.

The front-end development process has also been very iterative. As visualizations are built there is an effort to use consistent user interfaces and widgets. The observation of users interacting with the portal combined with conceptual feedback has provided valuable insights for improving the user experience. The principles of good user interface design and user experience are applied at each iteration.

Finally once a feature has been implemented it is rolled out to internal users for initial, which often leads to some changes – sometimes they are minor and in other cases they are quite extensive. After further revision and testing, the new feature is deployed to the main portal site. The feature, however, continues to be refined based on user feedback and experience. Changes are also made to reflect updates in the toolkits we use, changes in the data backend and other improvements which are discovered along the way.

Technologies Used in the Portal

While the approach used and the guiding principles are not tied to any particular technology it is worth taking a small amount of space to describe the technologies used. On the server side the Django [10] web framework is used with PostgreSQL [11] as the backend database engine. On the client side the primary toolkit is d3.js [12] with a little bit of help from jQuery [13].

Django was chosen because the developers were familiar with it and comfortable working in Python. It has proven to be a very capable framework and it has been quite easy to work with. While no framework is perfect there is little to complain about with Django. Because Django makes it very easy to break the pieces of the site into modules (“apps” in Django parlance) we initially had many small modules but their division was not well thought out. Over time we have consolidated the code into a smaller number of modules. Due to its popularity there is a plethora of third party add on modules which have proven very useful. In particular the Tasty Pie [14] module is very flexible and convenient for building REST interfaces. It is currently used by api.es.net and it may soon be used for the RESTful parts of the portal. There is a Shibboleth module [15] for Django as well. While it was a bit dated it was sufficient as a basis for the Shibboleth integration needed for the portal.

d3.js is a very flexible and powerful framework for visualizing data in web browsers. The initial implementation of the portal used protovis [16] but quickly moved to d3.js when it was clear that protovis development was ceasing. The choice of protovis and later d3.js was guided by the belief that there is tremendous potential to visualize network data in many novel and rich ways which have not yet been explored. To effectively explore this space the developers felt it was important to use tools that allowed for the flexibility to try new types of visualizations. d3.js is not a conventional charting library which simply exposes a set of predefined charts. Rather it is framework that allows for binding any data to visual elements and efficient manipulation those elements based on the data. d3.js allows the developer to take advantage of the full capabilities of underlying technologies such as CSS3, HTML5 and SVG.

The network visualization effort put into the portal is of use to the wider visualization community. One primary goal in the development of the visualization engine for the portal has been modularity and reusability. Many of the visualizations used in the portal have been turned into configurable widgets. These widgets will be made into a separate package that will be released as open source in the near future.

Future Directions

The portal today is quite a bit different than it was the day it was launched and it continues to evolve on a weekly basis. Throughout the course of the next year many new features will be added to the portal. A major thrust is to rework the information architecture to make it more consistent and coherent. Part of this information architecture will include a concept centered on our collaborators and the communities that ESnet and ESnet sites participate in. This concept will lead to features that highlight relationships with peer networks and key communities that use ESnet. A tool to help users diagnose problems between a pair of end systems is also under development. This tool will use ESnet local data as well as perfSONAR data. A large focus area will be research and development of a visualization which demonstrates the science impact of ESnet. This visualization will explore the available data to find the ways to demonstration the ways in which ESnet has helped scientific progress.

Other near term improvements include:

- A visualization showing historic and present day monthly traffic volumes
- More details about our connections to our sites and peers
- A visualization highlighting where ESnet is IPv6 enabled for World IPv6 day
- ServiceNow [19] incident management integration
- Extending the OSCARS visualization to show virtual circuits that extend from ESnet into one or more peer networks
- Power visualization will be greatly expanded to include data both from new Layer 3 elements and from optical elements

ESnet will be transitioning from ESnet 4, which is based on 10G circuits, to a new architecture, ESnet 5, which is based on 100G circuits. As this transition occurs the portal will track the status of the build out. ESnet will be supporting many demonstrations at SC12 and the portal will be extended to showcase those demonstrations.

One further future development will not be visible to the end users but instead only to ESnet staff. The portal has become a nexus of information within ESnet and so it is only natural that some views and visualizations which are only appropriate for use by ESnet staff be integrated into the portal. For example recent configuration changes might be displayed along with the publicly visible information for a site. Other data that would be visible to staff but not normal users includes detailed technical drawings, physical addresses and private contact information. This is strong evidence that the portal is useful not only for making the network visible to external users but also for giving internal staff increased visibility into the network.

There has been some exploration of how to resolve issues with regard to accessing data sets more quickly to ensure the visualizations load more rapidly. The developers strive to continuously improve the code; as suboptimal ways of doing things are discovered they are replaced with better solutions.

Conclusion

The portal has been well received by users, project management and by other members of the community. All of the interest that has been shown in the portal is very humbling but also very encouraging. Because of the increasing complexity and interdependency of scientific collaborations and the networks that enable that collaboration, it is imperative that users from all areas have better visibility into the network. Between the user focused portal and the data access afforded by the API ESnet is leading the way to a more visible network.

Acknowledgements

The authors would like to thank the our program managers at the US Department of Energy, particularly Vincent Dattoria and Rich Carlson and the staff at ESnet for support and tireless feedback. Special thanks also to Linda Winkler, Ted Sopher and Phil DeMar for their invaluable feedback during the design and implementation of the portal. Thanks to the E-Center project their work in making the network visible. This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 with the University of California.

References

- [1] See the ESnet website, <http://www.es.net/> for more details about ESnet.
- [2] “Intra and Interdomain Circuit Provisioning Using the OSCARS Reservation System.” Chin Guok; Robertson, D.; Thompson, M.; Lee, J.; Tierney, B.; Johnston, W., Energy Sci. Network, Lawrence Berkeley National Laboratory. In BROADNETS 2006: 3rd International Conference on Broadband Communications, Networks and Systems, 2006 – IEEE. 1-5 Oct. 2006. Available at <http://es.net/news-and-publications/publications-and-presentations/> also see “Network Services for High Performance Distributed Computing and Data Management,” W. E. Johnston, C. Guok, J. Metzger, and B. Tierney, ESnet and Lawrence Berkeley National Laboratory, Berkeley California, U.S.A. The Second International Conference on Parallel,

Distributed, Grid and Cloud Computing for Engineering, 12-15 April 2011, Ajaccio - Corsica – France.

Available <http://es.net/news-and-publications/publications-and-presentations/>

- [3] See “perfSONAR Technical Overview,” <http://www.perfsonar.net/overview.html>
- [4] See “OGF Network Measurement Working Group (NMWG),” <http://nmwg.internet2.edu/>
- [5] See “Peakflow SP: Traffic Anomaly Detection,” <http://www.arbornetworks.com/peakflowsp>
- [6] O’Connor, Michael. “ESnet Planned Maintenance Calendar”. Available at: <http://es.net/news-and-publications/publications-and-presentations/>
- [7] E-Center, <https://ecenter.fnal.gov/>.
- [8] Fielding, Roy Thomas (2000), *Architectural Styles and the Design of Network-based Software Architectures*, Doctoral dissertation, University of California, Irvine. See also: Fielding, Roy T.; Taylor, Richard N. (2002-05), "Principled Design of the Modern Web Architecture" (PDF), *ACM Transactions on Internet Technology (TOIT)* (New York: Association for Computing Machinery) 2 (2): 115–150,
- [9] Mulloy, Bryan, RESTful API Design, Second Edition. Available online at: <http://www.slideshare.net/apigee/restful-api-design-second-edition>
- [10] The Django Web Framework, <https://www.djangoproject.com/>
- [11] The PostgreSQL Database, <http://www.postgresql.org/>
- [12] d3.js, <http://mbostock.github.com/d3/>
- [13] jQuery, <http://jquery.com/>
- [14] django-tastypie, <https://github.com/toastdriven/django-tastypie>
- [15] django-shibboleth, <http://code.arcs.org.au/gitorious/django/django-shibboleth>
- [16] protovis, <http://mbostock.github.com/protovis/>
- [17] Advanced Networking Initiative, <http://www.es.net/RandD/advanced-networking-initiative/>
- [18] Agile Manifesto, <http://agilemanifesto.org/principles.html>
- [19] ServiceNow, <http://www.service-now.com/>
- [20] Crockford, Douglas. RFC 4627: *The application/json Media Type for JavaScript Object Notation (JSON)*. July 2006. <http://www.ietf.org/rfc/rfc4627.txt>
- [21] Boote, Jeff and Brown, Aaron. *Bandwidth Test Controller (BWCTL)*. <http://www.internet2.edu/performance/bwctl/>
- [22] Shalunov, S. et al. RFC 4656: *A One-way Active Measurement Protocol (OWAMP)*. September 2006. <http://www.rfc-editor.org/rfc/rfc4656.txt>
- [23] Recordon, D. and Hardt D. *The OAuth 2.0 Authorization Protocol*. Internet draft available at: <http://tools.ietf.org/html/draft-ietf-oauth-v2-25>
- [24] OpenID, http://openid.net/specs/openid-authentication-2_0.html
- [25] Shibboleth, <http://shibboleth.internet2.edu/>

Biographies

Jon M. Dugan is a network engineer and software developer at ESnet. His professional interests include high performance networking, network monitoring and visualization, network configuration automation/auditing and network performance tuning. Dugan is also keen to see ESnet used to enable science for all kinds and sizes of experiments and researchers. He received a Bachelor of Science degree from the University of Illinois at Urbana-Champaign. Dugan spent the early part of his career as a network engineer at the National Center for Supercomputing Applications (NCSA) and is very active in the joint IEEE/ACM [Supercomputing](#) conference series.

Gopal Vaswani is a computer systems engineer in ESnet's Network Engineering Group. His responsibilities include conceptualizing and developing tools and services for the network’s infrastructure. Vaswani has a master's degree in Information Management and Systems from the University of California, Berkeley. After graduation, Vaswani worked at two Internet startup companies before coming to ESnet—Virtual Vehicle LLC, a company that he co-founded with a U.C. Berkeley classmate, and ifeelgoods Inc., an early phase social media startup. Vaswani is interested in how technology shapes our perceptions and society and has a strong background in user interface design. Prior to graduate school, Vaswani was a user experience researcher at Rediff.com and Honeywell Technology Labs in Bangalore, India. Originally from India, Vaswani earned a bachelor's degree in design at the Indian Institute of Technology.

Gregory Bell is Acting Division Director for the Scientific Networking Division at Lawrence Berkeley National Laboratory. He previously served as Chief Information Strategist and Area Lead for ESnet, and as Chief Technology Architect for the Office of the CIO at LBNL. Bell has a B.A. from Harvard and a Ph.D. from UC Berkeley.

Inder Monga is ESnet's Chief Technologist and advances networking services for collaborative and distributed science at ESnet, as well as contributes to ongoing research projects including OSCARS, perfSONAR, as well as the Advanced Network Initiative. Monga's research interests include network virtualization, grid/cloud computing and sensor networking. His work experience in the private sector has included network engineering for Wellfleet Communications and Canadian telecom company Nortel, where he focused on application and network convergence. Originally from India, Monga spent about 19 years in the greater Boston area. He earned his undergraduate degree in electrical/electronics engineering from Indian Institute of Technology in Kanpur, India, before coming to the U.S. to pursue his graduate studies in Boston University's Electrical and Electronics Communication Systems Department.