
E-Center Project

Fermilab Team
Project Review
January, 2012



Outline

- Year 2 Goals
- Year 2 Deliverables
- Major Developed Components
- Data Collection
- Data Retrieval Service (DRS)
- Anomalies Detection Service (ADS)
- Forecasting Data Service (FDS)
- Front-end and UI components
- Site Centric View data service (SCV)
- **Dynamic Circuits Monitoring**
- User's Content
- Year 3 Plan
- Future Ideas



Year 2 Goals

- Production quality E-Center site
- Add Anomalous Detection Service and UI
- Add Forecasting Data Service and UI
- Site Centric Service and UI
- User provided traceroute visualization
- User's generated content and user's engagement



Year 2 deliverables

E-Center now consists of the ***Webservices Ecosystem:***

DRS – Data Retrieval Service – extended with more services

ADS – Anomalous events Detection Service

FDS – Forecasting Data Service

SCV - Site Centric View Data service

Network Data Collection Health Monitoring service – **DRS extension**

E-Center Data Storage was moved to the newly procured hardware

Complete stack of Authentication and Authorization, Single Sign On (Shibboleth) was added



Year 2 deliverables (continued)

Full text search added to the E-Center portal

Network Weather Map for the user's provided **traceroute** results was added

Forecasting front-end UI was implemented and integrated into the E-Center

Anomalous Events Detection front-end UI was implemented and integrated into the E-Center

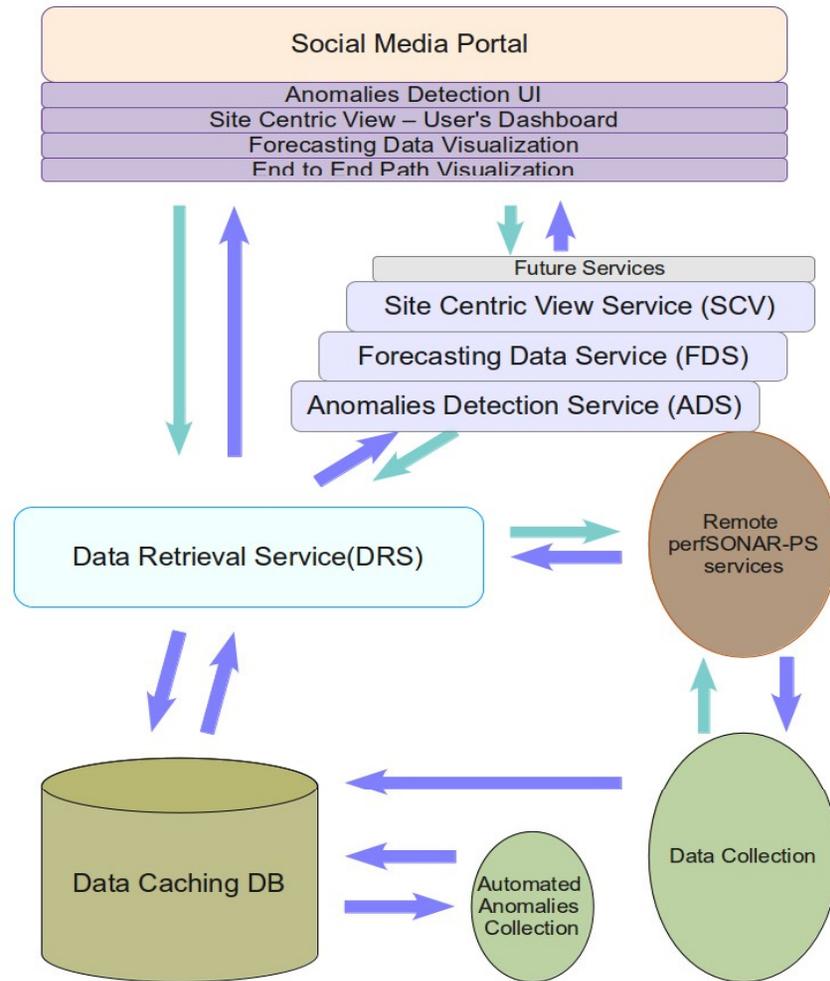
Site Centric View UI was implemented and added to the E-Center

Anonymous user home page and **Dashboard** for the Authenticated user's with **Site Centric View** were added

Started User's engagement via social services provided by the E-Center (area specific groups). Areas for each DOE lab were created to share ideas, concerns.



Major Developed Components



Data Collection

- **perfSONAR-PS** Metadata collected and cached – nightly
 - From ESnet hLS – home Lookup Services
 - From hLSs at DOE Labs
- Data collected and cached – nightly for the past 24 hours
- From all **DOE labs** and **ESnet** measurement mesh:
 - OWAMP (delay, packet loss)
 - BWCTL (throughput)
 - PingER (RTT, packet loss)
- From ESnet centralized services:
 - Interface Utilization, Drops, Errors
 - ESnet Topology
 - Dynamic Circuits Topology and Utilization
 - Traceroutes between ESnet PoPs and DOE lab's



E-Center DRS and Webservices

DRS protocol was updated with new set of functionality (**Site Centric View** and Data for the **User's Provided Traceroute** among them) and now at version 3.8, see -

https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Data_Retrieval_Service%28DRS%29

New services added and protocols documented:

ADS -

https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Anomalies_Detection_Service%28ADS%29

FDS -

https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Forecasting_Data_Service%28FDS%29



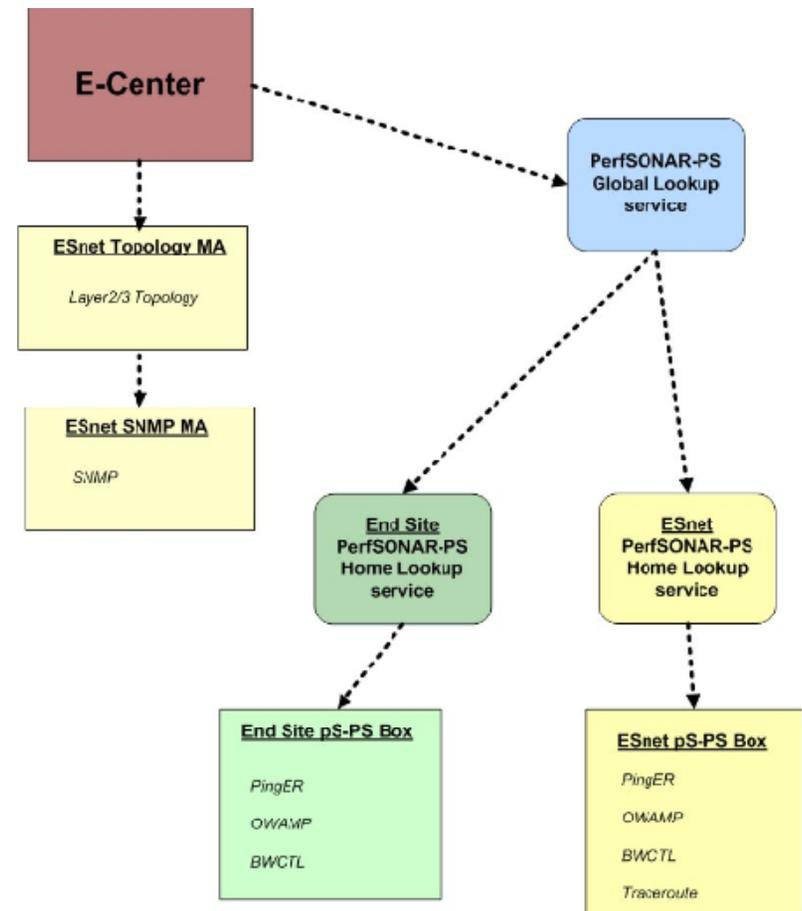
Data Collection: Just Numbers

- Monthly cached data – entries range:
 - OWAMP 60-100M+
 - SNMP 35-50M
 - PingER 10-20M
 - Traceroute 0.6-1.2M
 - BWCTL 60K-120K
- Collecting data from ~300 service endpoints, about 25000 unique measurements (metadata entries).
- Nightly automated anomalies detection for BWCTL and OWAMP data – about 9K metadata entries weekly.



Data Collection: API

- Scheduled pull of all registered pS-PS metadata from ESnet & designated sites
- Scheduled pull of ESnet's Layer2 and Layer3 topology
- Caching of all metadata, ESnet topology elements
- Sharded cache for pS-PS data
- On-fly asynchronous dispatch to remote pS-PS services via E-Center Data API and distributed worker nodes



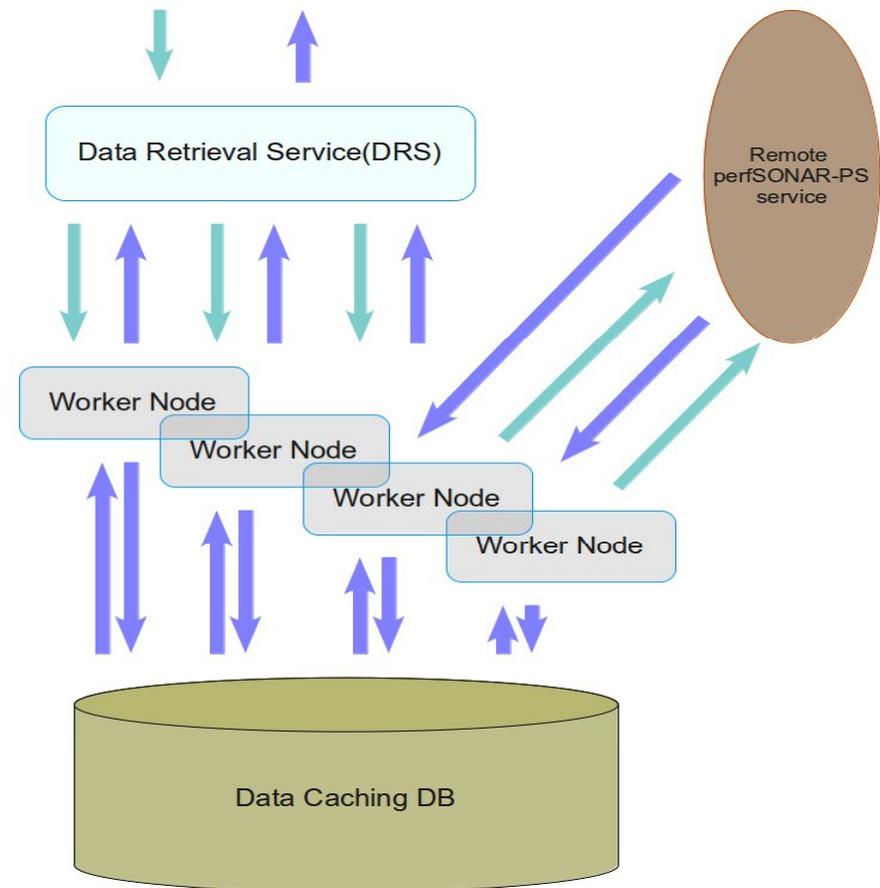
Technical Details

- Choice of the **REST** Data services – simple, flexible
- Data cache - **MySQL DB** engine – scalable, supports replication and clustering for the future extension, ***data sharding*** – slicing each data set in pieces for fast retrieval – we use a ***timestamp*** for monthly tables
- **Gearman** distributed Job submission manager for all backend tasks – to provide desired level of scalability
- **Modern Perl** for the backend and services – stable, 20+ yr old language, great variety of APIs, mature web development framework, compatibility with perfSONAR-PS API
- **PHP, Drupal, and jQuery** for the front-end – popular, well-known open-source development and content management tools with strong support for social functionality.



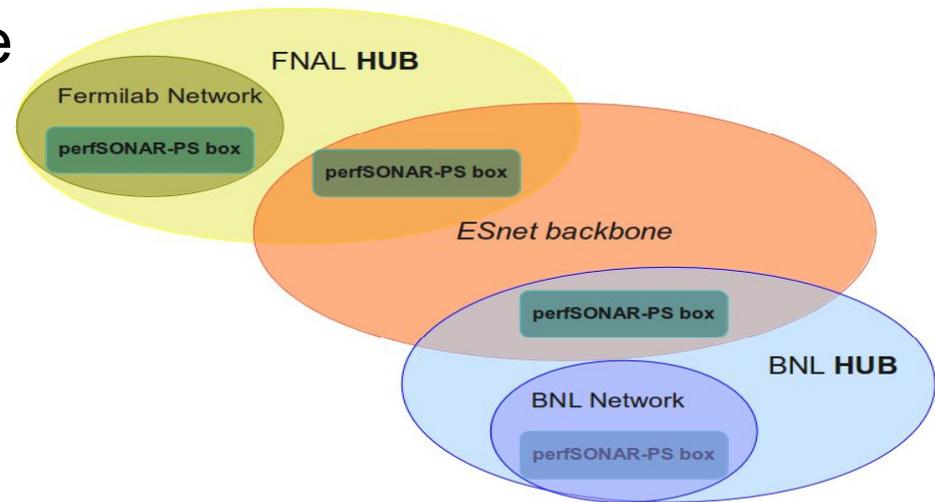
DRS: Data Retrieval Service

- Scalable, extendable web service
- Returns cached data or data from remote pS-PS services
- End to End tomography via SNMP & Traceroute
- Parameters:
 - Time period
 - Source/Destination IPs
 - Source/Destination HUBs
 - Data type
 - User's provided Traceroute
 - Resolution - number of aggregated data points



DRS: What is the HUB ?

- Logical aggregation of the End-to-End pS-PS monitoring endpoints - End Site located with ESnet edge located
- Allows one to see all available data for WAN network path
- Every pS-PS service and as result every measurement is mapped to specific HUB upon discovery



ADS: Anomalies Detection Service

- **REST** based web service, can be located anywhere
- Designed to work with any E-Center supported data sample or may send request to the DRS for the time series dataset
- Implements Static (**SPD**) or Adaptive Plateau Detection Algorithms (**APD**):
 - developed at Ohio Supercomputing Center (Prasad Calyam)
- Protocol is documented at:
[https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Anomalies_Detection_Service\(ADS\)](https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Anomalies_Detection_Service(ADS))

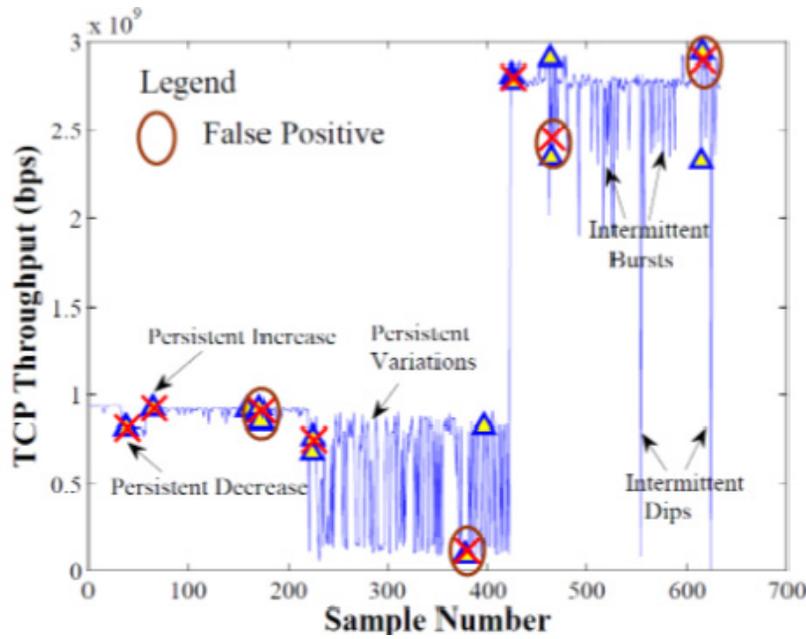


ADS: Plateau Anomaly Detection

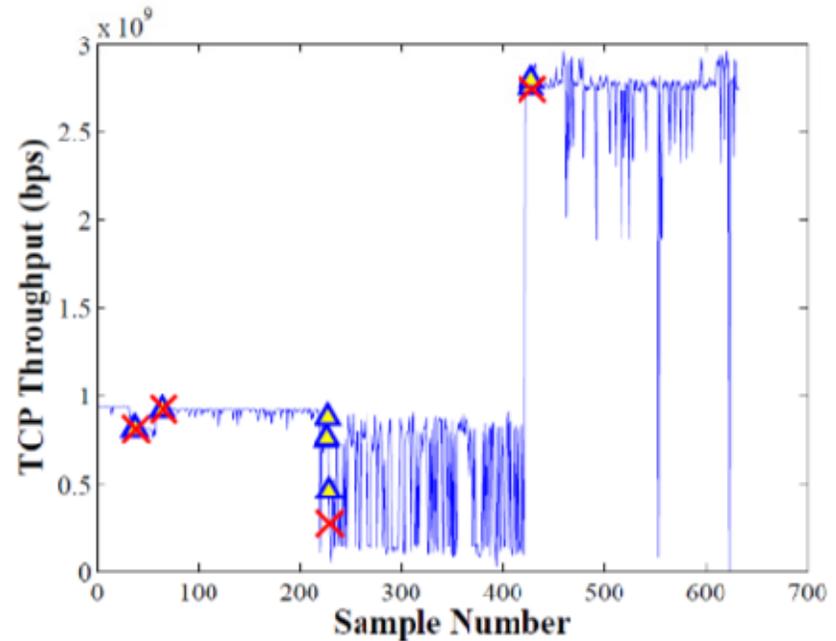
- Enhanced mean \pm standard deviation (MSD) algorithm
- Plateau detector uses two salient thresholds
 - Sensitivity (s): specifies magnitude of plateau change that may result in anomaly
 - Trigger duration (td): specifies duration of anomaly event before trigger is signaled
- Network health norm determined by calculating mean for set of measurements sampled recently into “summary buffer”
 - Number of samples in “summary buffer” is user defined, and labeled ‘summary window count’ (swc)



ADS: SPD and APD samples



Static Sensitivity in SPD Scheme

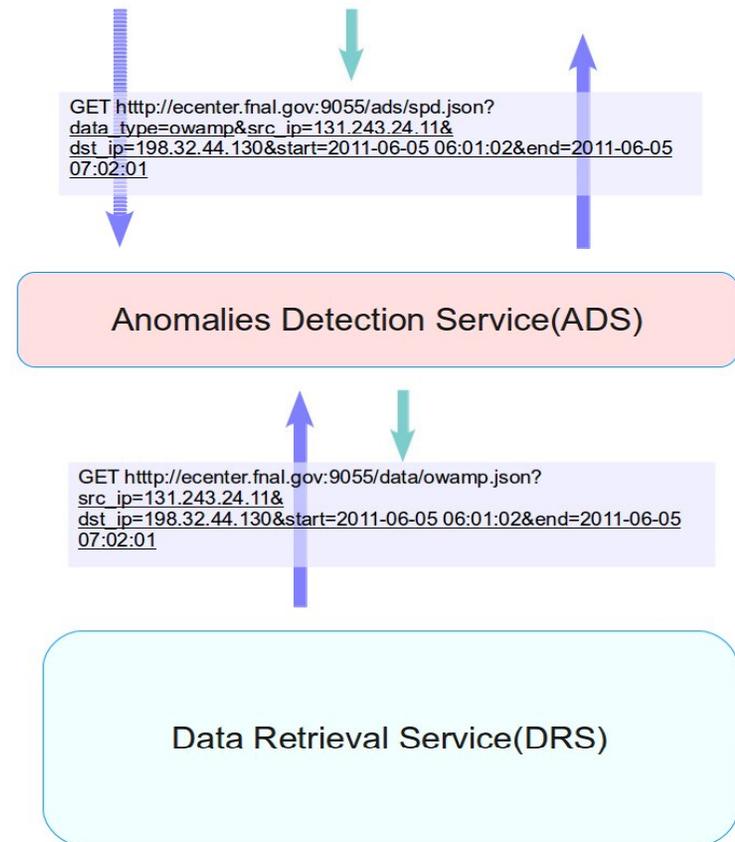


Dynamic Sensitivity in APD Scheme



ADS: Integration with DRS

- ADS is standalone web service
- Can return results for supplied data sample
- OR, can send request to DRS for specific data type for time period
- Returns all warnings & critical anomalies found according to plateau detection parameters



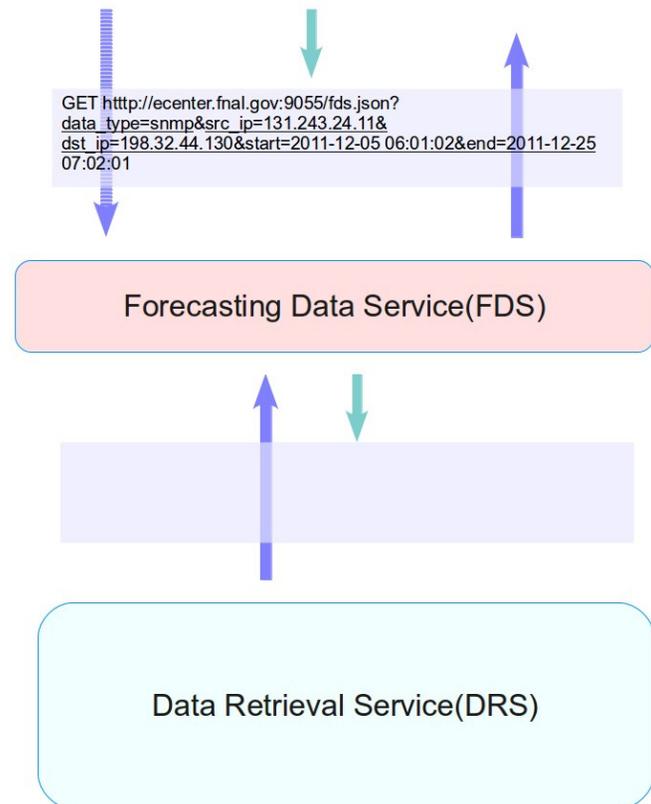
FDS: Forecasting Data Service

- **REST** based web service, can be located anywhere
- Designed to work with any E-Center supported data sample or may send request to the DRS for the time series dataset
- Based on Forecasting algorithm developed by Michael Frey at Bucknell University
- Protocol is documented at:
[https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Forecasting_Data_Service\(FDS\)](https://cdcvs.fnal.gov/redmine/projects/ecenter/wiki/Forecasting_Data_Service(FDS))

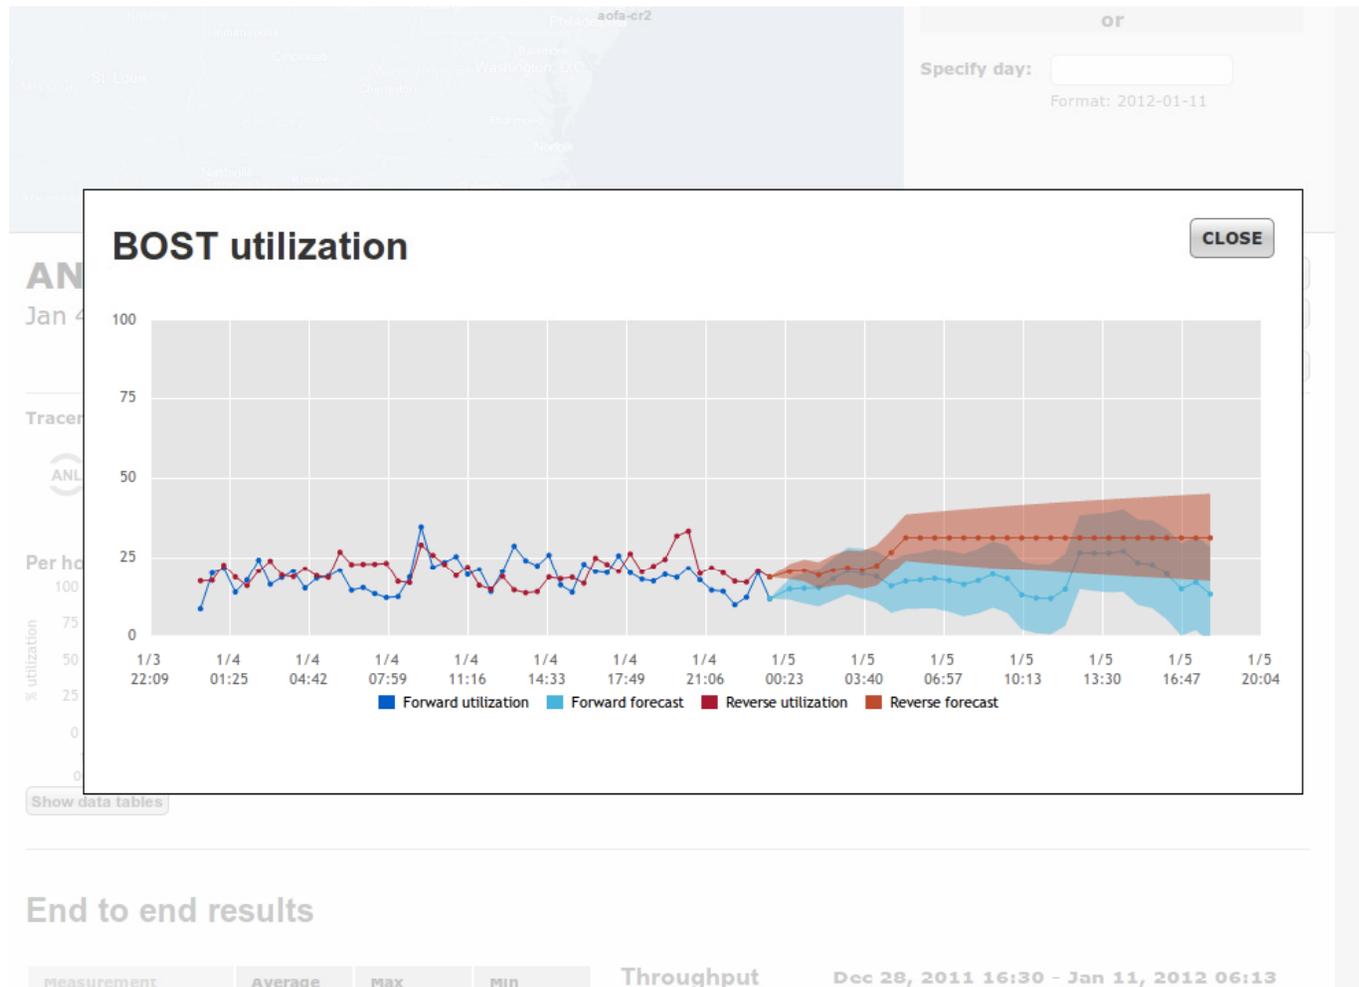


FDS: Integration with DRS

- FDS is standalone web service
- It sends request to DRS for specific data type for some time period but for not less 200 data points – analysis limit
- Returns forecasted values with error bands



FDS: In Action



Front-end: Forecasting UI

or

Specify day: Format: 2012-01-11

ANL to BNL

Jan 4, 2012 00:00 to Jan 5, 2012 00:00

<http://localhost/center/ixTa> [Link](#)

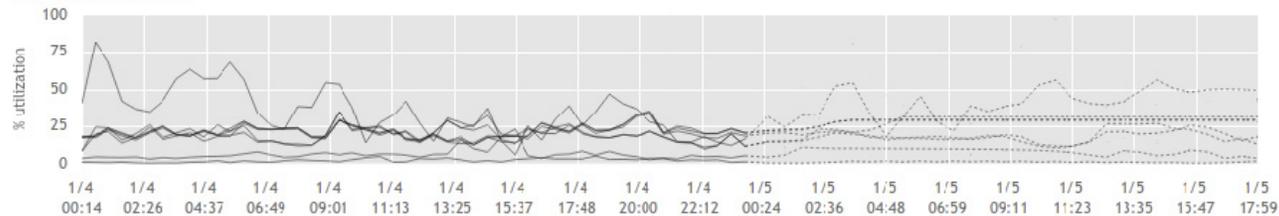
Timezone: America/Chicago

Type: SNMP [Get forecast](#) [Anomaly settings](#) [Detect anomalies](#)

Traceroute (logical) ■ Forward ■ Reverse [Click nodes to see detail](#)



Per hop utilization



[Show data tables](#)

End to end results

Measurement	Average	Max	Min
Forward Throughput	4080mbps	4606mbps	2209mbps

Throughput Dec 28, 2011 16:30 - Jan 11, 2012 06:13



Front-end

- Powered by Drupal Content Management System
- Popular and proven: Great community, stable, widely-known. Strong support for social networking features (issue tracking, blogs, knowledge base, tag cloud)
- Widely adopted by government: Challenge.gov, Whitehouse.gov, Energy.gov, Recovery.gov
- Key social features
 - Data conversations: Issue & data notification system allows conversation & collaboration across DOE networking community.
 - Knowledge base: A powerful WIKI with an easy editor for generating and sharing documentation.
 - Groups: Lightweight system for sharing conversations & documentations with a task force or working group
 - Activity stream: A birds' eye view of important conversations, anomalies, and knowledge base articles customized for each user



Front-end: Dashboard

Dynamic, per-user dashboard:

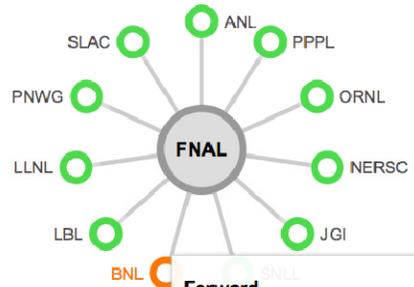
- ❑ encourages users to interact with site and each other
- ❑ user-selectable site-centric view and other data tools to give user instant snapshot of network status
- ❑ site-wide activities using familiar "activity stream" interface
- ❑ Activity stream includes detected anomalies, new content, comments, and more, based on user's group membership



Dashboard

FNAL status

2011-09-08 04:00:00
to
2011-09-08 10:00:00



Tools

[Weathermap](#)
[Create issue](#)

[Network health](#)
[Create wiki page](#)

Hot conversations

The most active public conversations in the past 2 weeks

[Hello, conference attendees!](#)

by [eads](#) | 6 recent comments

Last comment: 2 weeks 5 days ago

Most popular content

The most popular content in the past month

[FNAL to BNL, 2011-06-06 05:00:00 - 2011-06-07 05:00:00 \(UTC\)](#)

by [maximgrp](#) | Issue

Last updated: 1:34pm 9-07-2011

[Hello, conference attendees!](#)

by [eads](#) | Issue

Last updated: 4:20pm 7-14-2011

Forward

Measurement	Value	Interface
Utilization	89.22	134.55.221.138
Errors	0.00	NA
Drops	0.00	NA

Reverse

Measurement	Value	Interface
Utilization	85.64	134.55.41.121
Errors	0.00	NA
Drops	0.00	NA

Site activity

Activity on the site and in your groups



Issue [FNAL to BNL, 2011-06-06 05:00:00 - 2011-06-07 05:00:00 \(UTC\)](#) updated by [maximgrp](#)

[maximgrp](#) wrote: *There is a higher than normal one way*



[admin](#) commented on [Hello, conference attendees!](#) 3 weeks, 1 day ago

[admin](#) wrote: *Replied from Outlook web express.*



Wiki page [Data Retrieval Service Specification](#) updated by [admin](#) Jul 14, 2011

Log message: *Current version is 3.3*



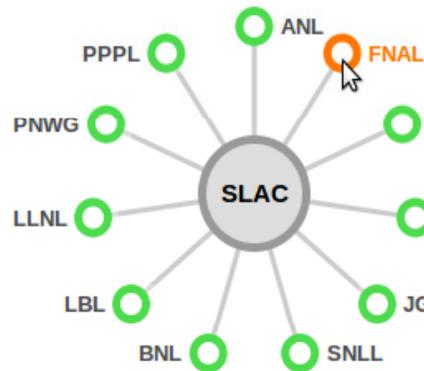
Front-end: Site-centric view

- User-customizable view of connections to other DOE sites.
- 24-hour snapshot to find where problems exist
- Cached hourly for speedy display on dashboard

E-Center dashboard

SLAC status:

2011-07-05 04:00:00 -
2011-07-05 16:00:00



Tools

Forward

Measurement	Value	Interface
Utilization	87.12	134.55.219.122
Errors	0.00	NA
Drops	0.00	NA

Reverse

Measurement	Value	Interface
Utilization	55.51	134.55.209.45
Errors	0.00	NA
Drops	0.00	NA

8 recent comments

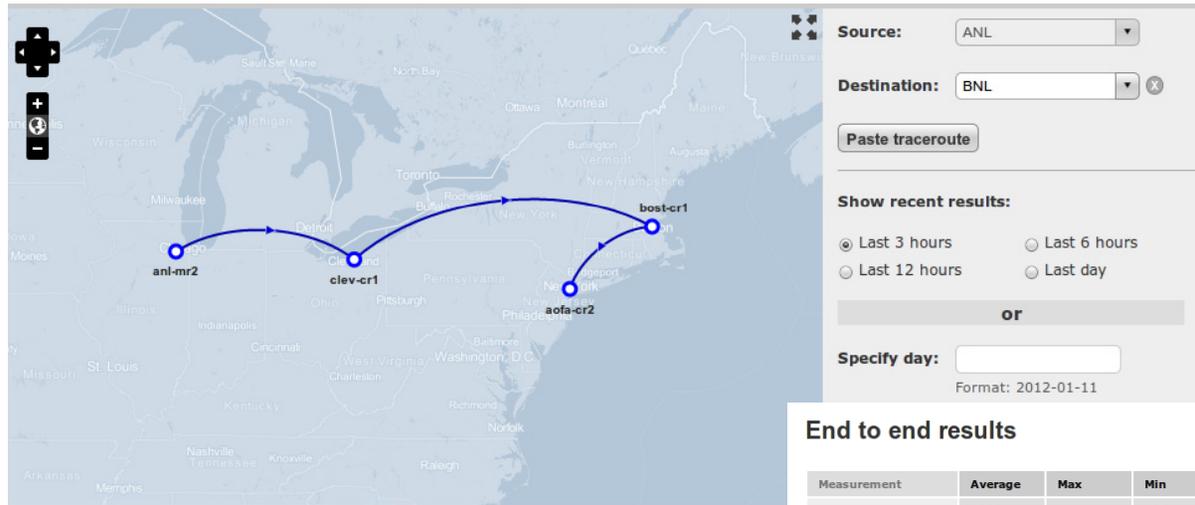
Last: 4 weeks 6 hours ago

Site activity

Activity on the site and in your groups



Front-end: E2E UI & data visualization



ANL to BNL

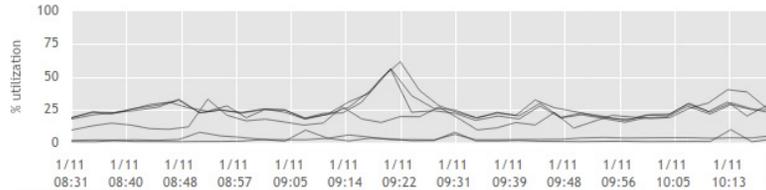
Jan 11, 2012 08:30 to Jan 11, 2012 11:30

Type: **SNMP** [Get forecast](#)

Traceroute (logical) ■ Forward ■ Reverse [Click nodes to see detail](#)



Per hop utilization

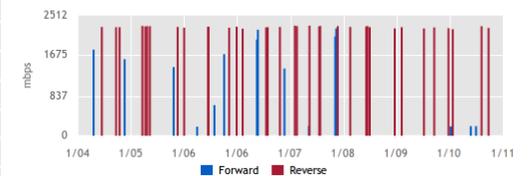


[Show data tables](#)

End to end results

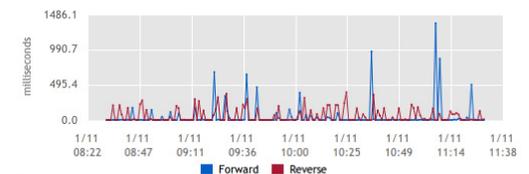
Measurement	Average	Max	Min
Forward Throughput ?	1202mbps	2228mbps	186mbps
Reverse Throughput ?	2257mbps	2284mbps	2213mbps
Forward Duplicates ?	0	0	0
Reverse Duplicates ?	0	0	0
Forward Loss ?	0.0	0.0	0.0
Reverse Loss ?	0.0	0.0	0.0
Forward Max delay ?	58ms	1351ms	13ms
Reverse Max delay ?	67ms	396ms	13ms
Forward Min delay ?	12ms	13ms	12ms
Reverse Min delay ?	13ms	13ms	12ms

Throughput Jan 4, 2012 10:38 - Jan 11, 2012 10:19



[Popup chart](#) [Show data tables](#)

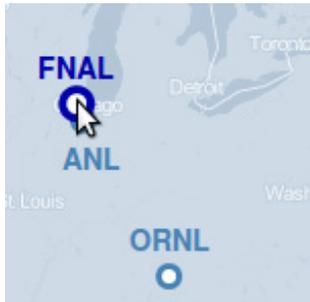
Max delay Jan 11, 2012 08:31 - Jan 11, 2012 11:29



[Popup chart](#) [Show data tables](#)



Front-End: Weathermap UI usage



Select source and destination on map

A form with two dropdown menus. The first is labeled 'Source:' and has 'FNAL' selected. The second is labeled 'Destination:' and has 'A' selected. Below the second dropdown is a list of suggestions: 'ANL' (highlighted in yellow) and 'SLAC'. There is a 'Show recent' link to the left of the list.

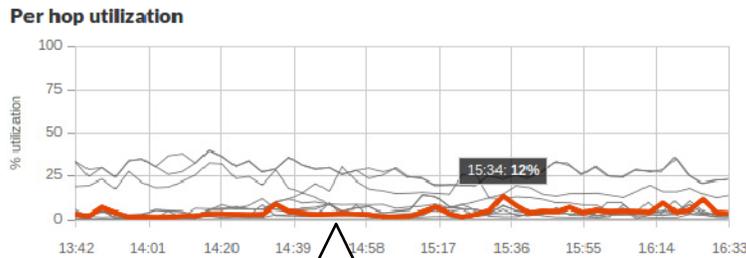
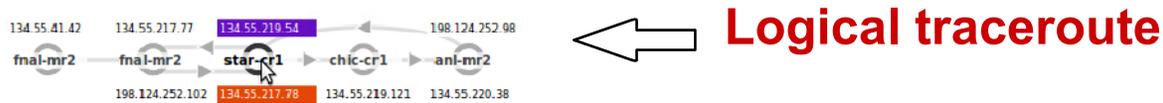
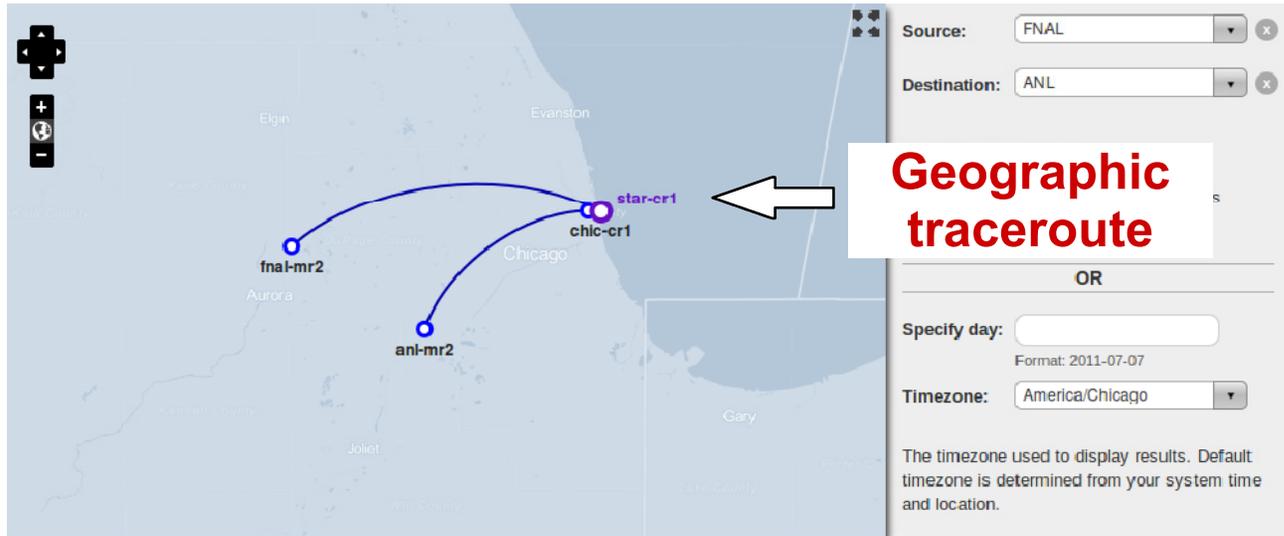
...or by typing to search or selecting from list of sites

A form with four radio buttons for date ranges: 'Last 3 hours', 'Last 6 hours', 'Last 12 hours', and 'Last day'. The 'Last day' option is selected. Below the radio buttons is a horizontal line with the word 'OR' centered. Underneath is a text input field labeled 'Specify day:' with the format '2011-07-07' shown below it. Below that is a dropdown menu labeled 'Timezone:' with 'America/Chicago' selected. At the bottom, there is a paragraph of text: 'The timezone used to display results. Default timezone is determined from your system time and location.'

Select date range and timezone for query



Front-end: Weathermap Results



Path utilization chart

FNAL (fna1-mr2) to ANL (an1-mr2)

Jul 7, 2011 08:40 - Jul 7, 2011 11:40

Create issue from these results

Permanent link (bookmark and share results)

<http://131.225.82.131/center/network?src=hub%3AFNAL&dst=>

Create issue from query

Permanent link to query



Front-end: Weathermap Results (II)

End to end results

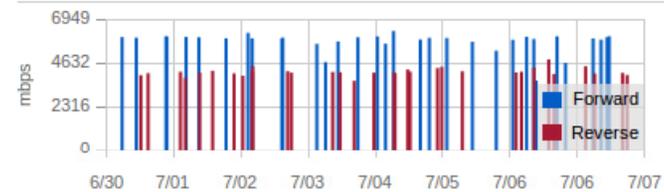
▸ Anomaly detection settings

Detect anomalies

Measurement	Average	Max	Min
Forward Throughput	5778mbps	6317mbps	3690mbps
Reverse Throughput	4157mbps	4810mbps	3678mbps
Forward Duplicates	0	0	0
Reverse Duplicates	0	0	0
Forward Loss	0.0	0.0	0.0
Reverse Loss	0.0	0.0	0.0
Forward Max delay	85ms	531ms	2ms
Reverse Max delay	88ms	476ms	2ms
Forward Min delay	1ms	2ms	1ms
Reverse Min delay	2ms	2ms	1ms

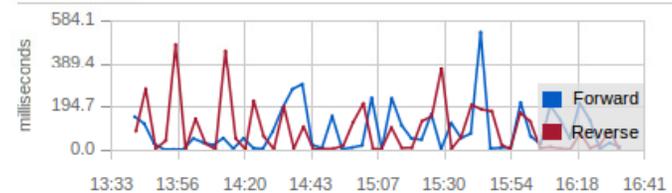
Throughput

Jun 30, 2011 21:26 - Jul 7, 2011 13:24



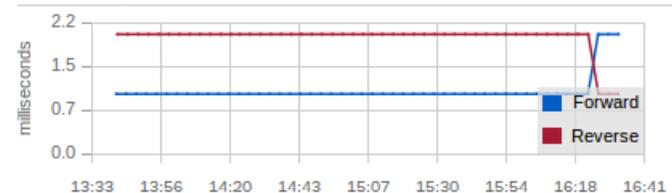
Max delay

Jul 7, 2011 13:42 - Jul 7, 2011 16:32



Min delay

Jul 7, 2011 13:42 - Jul 7, 2011 16:32



Front-end: Weathermap Results

▼ Anomaly detection settings

Algorithm:

Sensitivity:

Elevation 1:

Elevation 2:

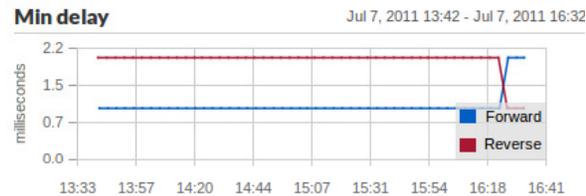
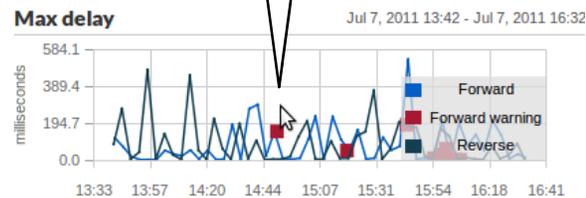
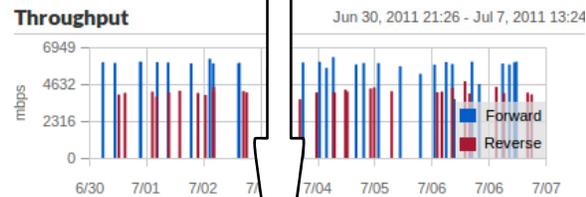
Window count:

Detect anomalies

← Experimental anomaly detection

Anomalies shown on charts

Measurement	Average	Max	Min
Forward Throughput	5778mbps	6317mbps	3690mbps
Reverse Throughput	4157mbps	4810mbps	3678mbps
Forward Duplicates	0	0	0
Reverse Duplicates	0	0	0
Forward Loss	0.0	0.0	0.0
Reverse Loss	0.0	0.0	0.0
Forward Max delay	82ms	531ms	2ms
Reverse Max delay	88ms	476ms	2ms
Forward Min delay	1ms	2ms	1ms
Reverse Min delay	2ms	2ms	1ms



Authentication & User's Registration

- Make signup/login easy AND safe
- Support for Single Sign-On AA systems deployed at DOE: **OpenID**, **Shibboleth** (experimental)
- Supports common OpenID providers (Google, Facebook, etc.) as well
- Fallback to domain-based white-lists and manual user management

User account

Click your account provider:



Username: *

Enter your E-Center username.

Password: *

Enter the password that accompanies your username.

Log in



Front-end: Easy Content Creation

- Rich text editor
- WYSIWYG and Markdown modes
- Embedded media support (YouTube, Vimeo)
- Code/syntax highlighting
- Powerful, secure, nice to use
- Email integration: custom notifications, replied by-email functionality



Easy Content Creation (cont)

Body:

The screenshot shows a rich text editor with a toolbar at the top containing icons for undo, redo, bold, italic, bulleted list, numbered list, link, unlink, image, link icon, unlink icon, and a 'Heading 2' dropdown menu. The main editing area contains:

```
{syntaxhighlighter class="brush: php;"}  
/* PHP code */  
function hello_world() {  
    return 'hello world';  
}  
{/syntaxhighlighter}
```

Tables

First Header	Second Header
Edited cell	Content Cell
Content Cell	Content Cell

Classes

.content-align-center: Align center

Embedded media

<http://www.youtube.com/watch?v=aigN65VmHxI>

Path: h2

[Disable rich-text](#)

▸ **Input format**

▼ Revision information

Revision message:

Provide a short description of the changes made to this .

▼ Groups

Audience:

Fermi National Accelerator Lab (FNAL)

Show this post in this group.

Public

Show this post to everyone, or only to members of the groups checked above. Posts without any groups are always *public*.



Front-end: Measurement mesh health

- Shows which perfSONAR data is being collected by DRS
- Helps identify/troubleshoot problems in measurement infrastructure

Hub status

Hub	snmp	bwctl	traceroute	owamp	pinger
SLAC	■	■	■	■	■
FNAL	■	■	■	■	■
JGI	■	■	■	■	■
BNL	■	■	■	■	■
NERSC	■	■	■	■	■
SNLL	■	■	■	■	■
LBL	■	■	■	■	■
LLNL	■	■	■	■	■
PNWG	■	■	■	■	■
ANL	■	■	■	■	■
PPPL	■	■	■	■	■
ORNL	■	■	■	■	■

Metadata 87
Cached data 9584



Front-end: End-to-end UI & data visualization

- **Streamlined UI:**

Reduce clicking and "knob turning" for common cases (inspired by Google instant).

- **Provide many views of data:**

Subway map, geographical map, and end-to-end tables provide several perspectives on path data. Site-centric and measurement health status provide logical views of site status. Chart data magnifying glass feature.

- **Self-revealing:**

Tool-tips integrated throughout UI provide users with context for interpreting results.

- **Highlight critical information:**

Chart highlighting technology attempts to make finding and isolating abnormal results easy.



Dynamic Circuits Monitoring

- **Real-Time ESnet OSCAR circuits data:**
 - .Circuits Utilization is gathered in real-time, without caching
- **10 min delay for circuits pulling:**
 - .Circuits metadata pulled from ESnet topology service every 10 minutes to provide timely updates on active circuits



Devel module ▾ Groups My Unread My account My bookmarks Create content ▾ Administer ▾ E-Center Tests ▾ Log out

Source: NERSC ▾
 Destination: ANL ▾ ×
 Paste traceroute
 Show recent results:
 Last 3 hours Last 6 hours
 Last 12 hours Last day
 or
 Specify day:
 Format: 2012-01-11

NERSC to ANL | Circuit: **es.net-2930**

Jan 11, 2012 09:00 to Jan 11, 2012 12:00

<http://localhost/ecenter/lxba> **Link**

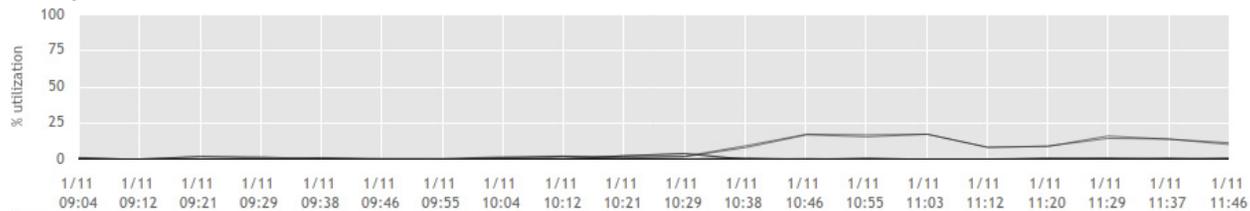
Timezone: ▾

[« Back to IP results](#)

Traceroute (logical) ■ Forward ■ Reverse 🔍 *Click nodes to see detail*



Per hop utilization



[Show data tables](#)



User's content and user's engagement

Devel module ▾ My account Create content ▾ Administer ▾ Help E-Center Tests ▾ Log out

Create Issue

Subject: *

Body: *



A few months ago, Deloitte, a large professional consultancy, invited me to speak at a conference held for their tech consultants in the Midwest region. I delivered this address on September 21st, 2007:

Good morning... I'm honored to be here. I'm Dave Eads. I am an activist, a programmer, and a designer. I'll tell you more about my story in a bit. First, a question...

In Steven King's excellent book "On Writing: A Memoir of the Craft", he concludes with his ultimate description of the writer's vocation:

"Writing," he says, "isn't about making money, getting famous, getting dates, ... or making friends. In the end, it's about enriching the lives of those who will read your work, and enriching your own life, as well. It's about getting up, getting well, and getting over. Getting happy, okay? Getting happy. ... Writing is magic, as much the water of life as any other creative art."

Can we say the same of being a geek, or a hacker in the broad sense, or even an IT consultant? What is being a geek about? Is it about making money or being famous? Can we say it's about getting up, getting well, or getting over? Can we say it's about getting happy? What can we say about the satisfactions of our work?

And if the technologies we are expert in are, for most people, indistinguishable from magic, have we become the magicians of our hyperconnected world? As experts in information technology, what is our responsibility to our fellow citizens, our society, and our world?

Today, I'll be talking about several projects that I have been involved with over the years -- as an independent consultant currently working for the [Chicago Technology Cooperative](#), as a founder and staff member at the [FreeGeek Chicago](#) computer recycling, training, and advocacy project, and as the accidental web guru for [The View From The Ground](#), a web publication in the tradition of classic human rights reporting, focused on the plight of Chicago's public housing residents during a time of historic and often tragic change. As we go, I'll be focusing on the ways we strategically applied free and open source software in innovative ways, and asking some hopefully poignant questions about the nature of information technology work, and what we can hope from our profession. The photographs you'll see were taken by [Patricia Evans](#), [Jason Reblando](#), the [Basel Action Network](#), and myself.

In the spring of 2000, six months after I moved to Chicago, I heard about a crazy guy named Jamie Kalven

▾ **Groups**

Audience:

- Brookhaven National Lab (BNL)
- E-Center
- E-Center Administrators
- Energy Science Net (ESNet)
- Fermi National Accelerator Lab (FNAL)
- IPv6
- Joint Genome Institute (JGI)
- Lawrence Berkley National Lab (LBL)
- Lawrence Livermore National Lab (LLNL)
- National Energy Research Scientific Computir

Show this post in these groups.

Public

Show this post to everyone, or only to members of the groups checked above. Posts without any groups are always *public*.



Year 3 Plans

- Hardening E-Center data retrieval system
- Circuit-based network path monitoring – *in progress*
- On-demand measurement test capabilities:
 - Traceroute
 - BWCTL
- Alerts and updates notifications framework - *completed*
- Mobile device support: Smartphone (Android, iPhone) and tablet (iPad) interface
- Work with DOE labs and roll out wide-scale deployment of the perfSONAR-PS services under the E-Center management.



Future Ideas

- Standalone DRS with full AAA stack
 - Deployed at any end-site, will allow data protection and anonymization of the site's internal Layer2/Layer3 data
 - Higher level integration between remote DRSs to provide authorized users with public and private network monitoring data for end-host to end-host path
- Integration or/and collaboration with myESnet portal
 - Leveraging developed components and technologies
 - Efficient utilization of resources
 - Single Point of Contact for Esnet users



Project Information

- E-Center Project Site, including documents & software repository:

<https://cdcvs.fnal.gov/redmine/projects/ecenter>

- Live Demo: <https://ecenter.fnal.gov>



Demo

