# FNAL E-Center Efforts

Phil DeMar, Maxim Grigoriev, David Eads
October 26, 2010

# Outline

- Project Overview
- Year 1 Progress
- Back-End Details
- Future

# PROJECT OVERVIEW

# Project Information:

E-Center Project Page(s), including documents & software repository:

http://code.google.com/p/ecenter/

Fermilab's project page: http://ecenter.fnal.gov

# Project Design Goals

- A central location to obtain detailed network path information of interest, and seek assistance with related concerns
- Coherent, user-understandable presentation of disjoint network path monitoring & measurement data
- ESnet backbone & site scope
  - But with a vision for deployment across other R&E domains

# FNAL E-Center Personnel & Expenditures

- Personnel:
  - PerfSONAR implementer & backend software developer: Maxim Grigoriev
  - Front-end software developer: David Eads
- Budget:
  - In FY10, spent at ~$260k of $350 year 1 budget
    - Almost all personnel effort costs
    - Very minimal travel (~$4k)
  - Includes $25k to co-fund DOE/NSF PerfSONAR Workshop
  - Expect to catch up to project spending profile in year 2

# E-Center: Year 1

Maxim Grigoriev

Fermilab

*25 October 2010*

# E-Center Capabilities

- Web portal for DOE Lab's users
- Web 2.0 technologies for closing on "wizard's" gap
- Novel approaches for network data visualization
- Utilization and aggregation of the decentralized distributed perfSONAR-PS monitoring data
- Forecasting and network monitoring data mining/analysis

# E-Center: Choices Made

- Choice of the **REST** Data services, see here
- Data cache **- MySQL** DB engine – ACID transactions, scalable, supports replication and clustering for the future extension
- **Modern Perl** for the backend and services – stable, 20+ years old language, great variety of APIs, mature web development frameworks, re-use of the perfSONAR-PS API
- **PHP, Drupal** for the front-end – great community, stable, well known and de-facto the most famous web-development language and most active Content Management System. Available out of the box support for most desired social networking features ( issue tracking, blogs, knowledge base, tags cloud)

# E-Center: Components

- Loosely coupled, flexible web framework
- Drupal CMS, PHP based Frontend
- REST based Data Webservices
- MySQL DB scalable  data fusion and cache
- Modern Perl Data API
- Documented client API for the data consumers – Forecasting, Data mining

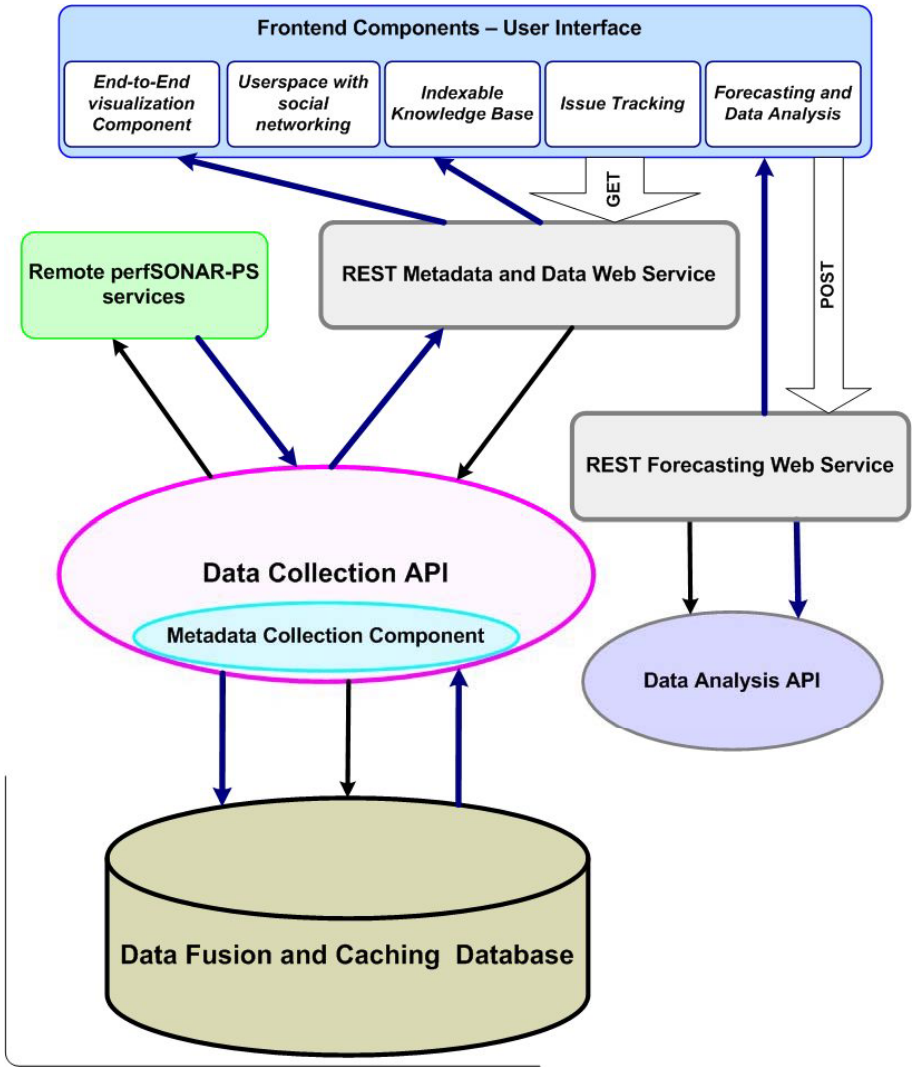

# E-Center: Data Collection

- Metadata collected and cached:
  - All ESnet hLSs
  - hLSs at FNAL, ORNL, LBL, BNL, ANL, SLAC
- Data collected on-fly:
  - OWAMP(delay), BWCTL(bw), PingER(RTT, packet loss) – from FNAL, ORNL, LBL, BNL, ANL, SLAC and ESnet
  - SNMP (utilization) – from ESnet centralized location
  - Topology - from ESnet centralized location
  - Traceroute – from ESnet  pS-PS boxes

# E-Center: Data Collection - continued

- Scheduled pull of all registered pS-PS metadata from ESnet & designated sites
- Scheduled pull of ESnet's Layer2 and Layer3 topology
- Permanent Caching of all metadata & ESnet topology elements
- Configurable temporary cache for pS-PS data > 1month
- On-fly asynchronous dispatching to remote pS-PS services via developed E-Center Data API

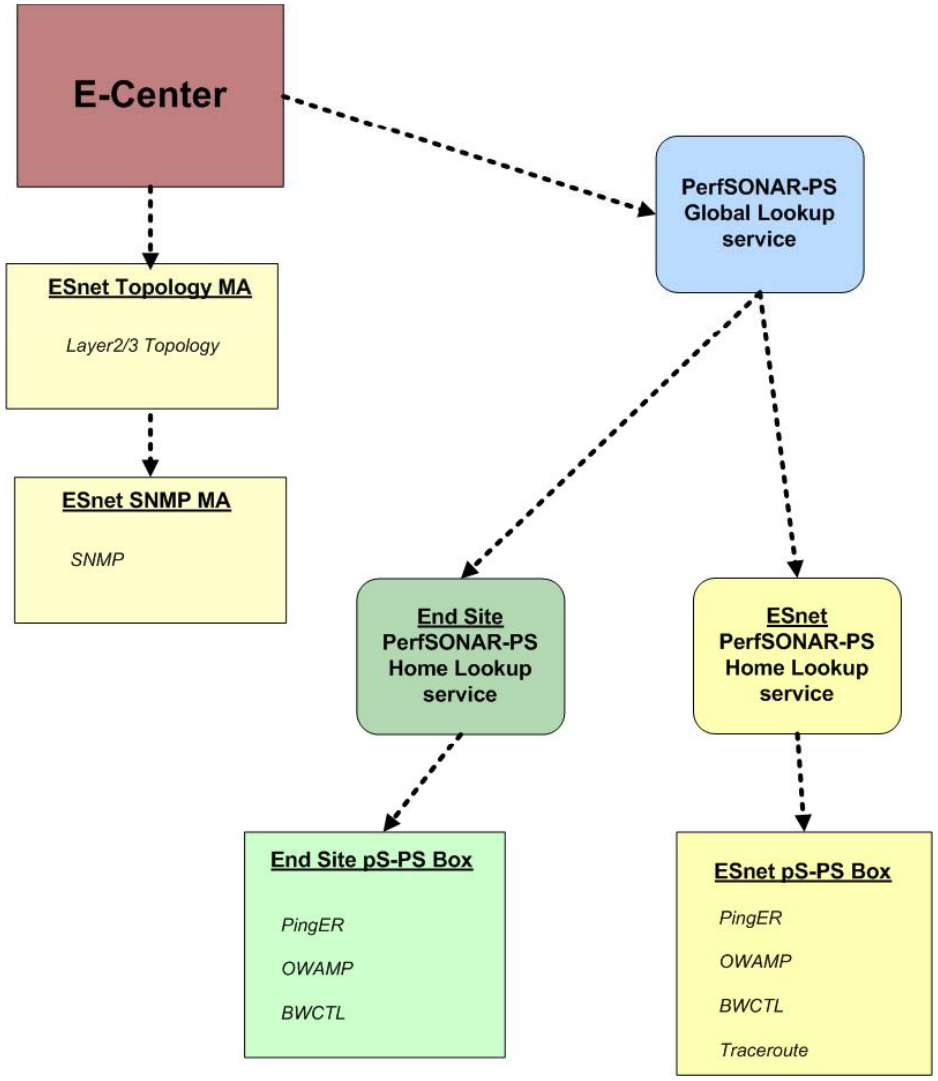

# E-Center: pS-PS Data API

- Re-usable modern PERL based OO API
- Support for any perfSONAR-PS data type
- Scheduled metadata cache updates based on perfSONAR-PS Global Lookup service Discovery
- Database Abstraction layer
- Any new perfSONAR-PS data type could be immediately plugged-in and supported

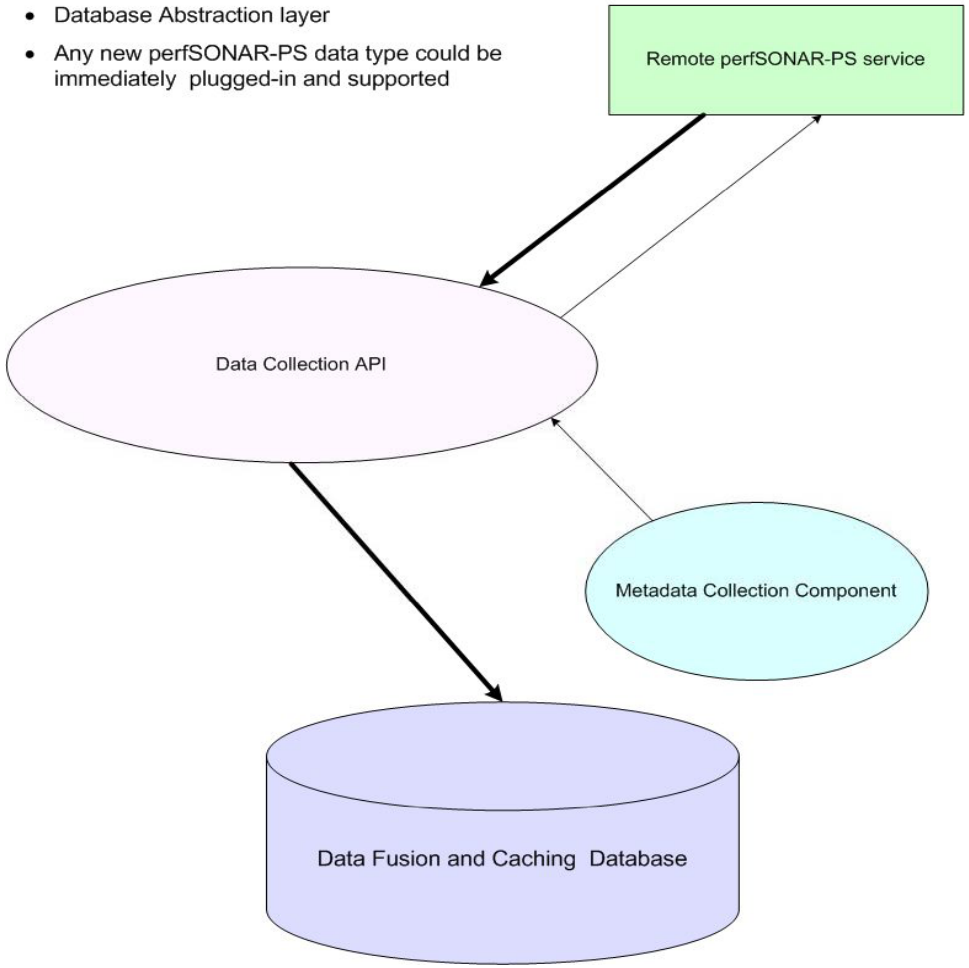

# E-Center: What is the HUB ?

- Logical aggregation of the End-Site pS-PS monitoring endpoints with ESnet border monitoring endpoint
- Follows ESnet established topology – HUB is the domain part of the node name
- More like a physical POP
- All monitored data is registered within unique HUB at the time of metadata pull
- Allows to run query for any data available for the HUB, for example if the End-Site's endpoint is having problems providing metadata then data request will be dispatched to the ESnet monitoring box within the same HUB

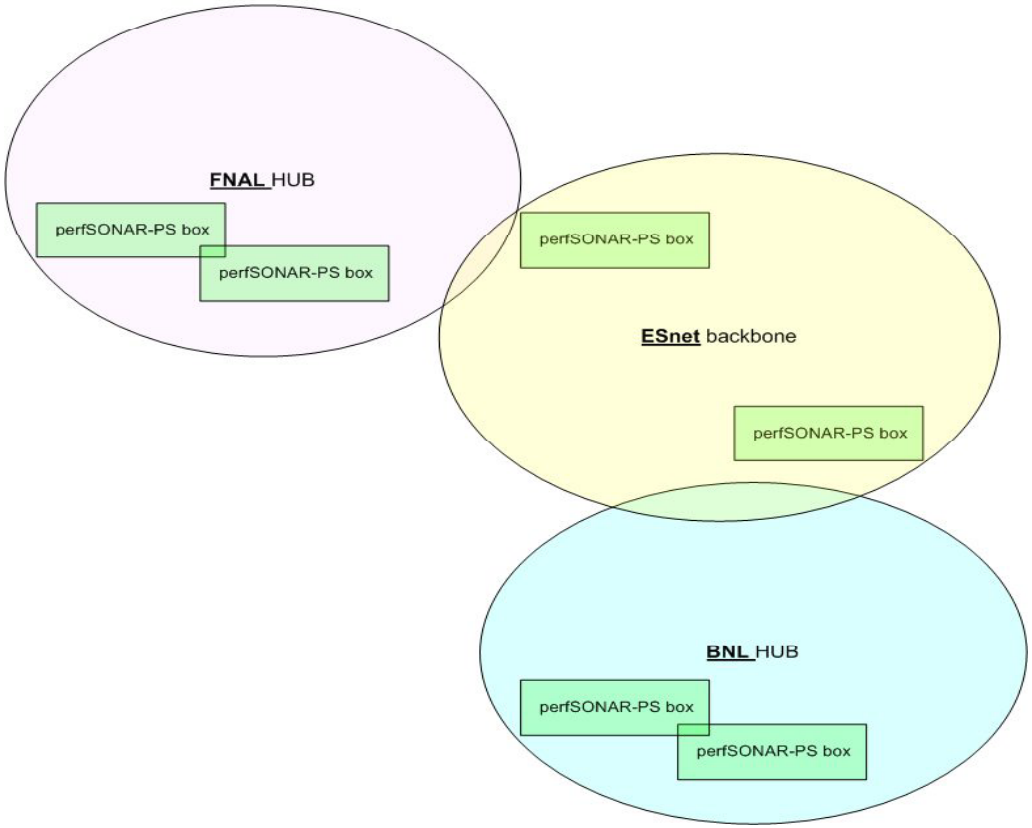

# E-Center: REST Data Webservice

- Re-usable modern PERL based OO API
- Follows REST standards
- Supports JSON, XML, YAML formats for the structured data response encoding
- Interface and Protocol are documented
- Supports End-to-end data queries for source/destination IP/hostname or E-Center HUB
- Provides a universal gateway for the E-Center data consumers (Analysis, Forecasting)

- Example of the data query: give me back all available monitoring data between E-Center HUB at LBL and HUB at BNL for some hour on 09/25/2010
- Where HUB is defined as collection of the perfSONAR-PS monitoring boxes at the End-Site and ESnet WAN border

***GET http://xenmon.fnal.gov:8055/data.json??src_hub=BNL&dst_hub=LBL&start=2010-09-25 17:50:00&end=2010-09-25 18:50:00&resolution=20***

- *Check local cache if all data ( within 30 minutes time period) is available*
- *Return available results ( if complete )*
- *If not complete then send query to the remote pS-PS service(s) and store response in the local cache*
- *Merge all received data*
- *Return all gathered results in the structured way*

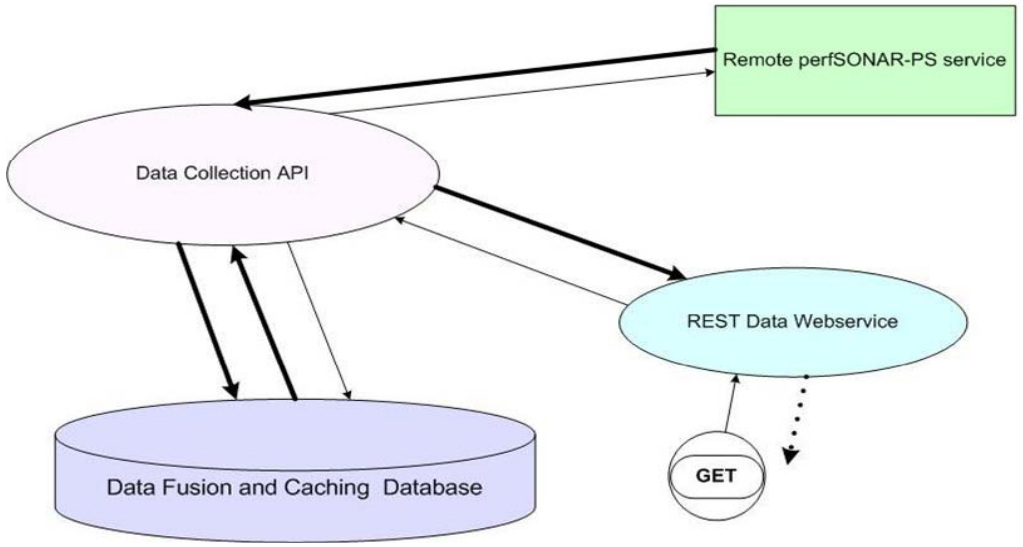

# E-Center: Proposed Integration of the Data Analysis and Forecasting

- User will be provided with special web-interface to select time period for the forecast and allowed to select time period for the data sample
- User will see results of the data analysis/forecast on the same time series graph
- When User clicks on **Forecast** button, the E-Center will send the same data structure it received from the Data Service with extra parameters to the Forecasting Service
- The response from the Forecasting Service will be displayed along the original time series
- Internally the Forecasting module will be implemented as standalone web service
- The Forecasting Service will be independent from the rest of the codebase and could be coded in any language with utilization of any external API
- The data protocol for the Forecasting Service will be documented

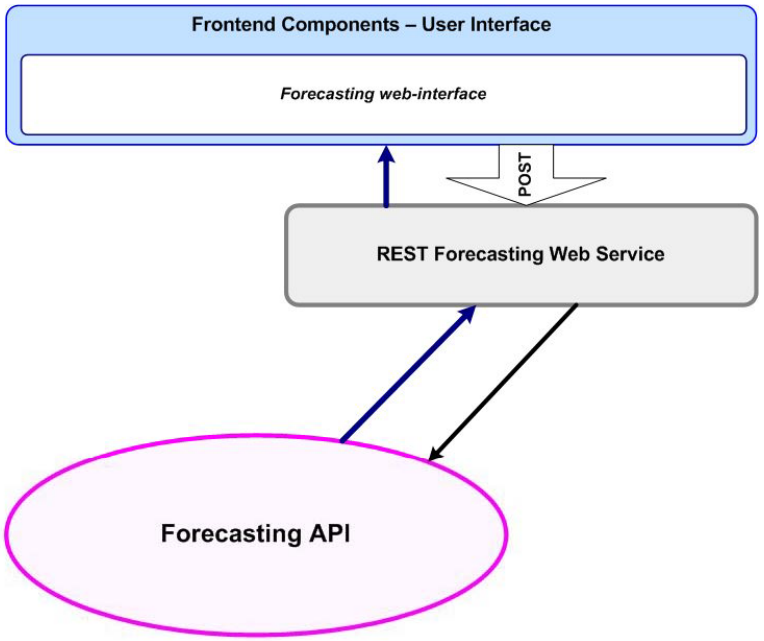

# E-Center: Forecasting web-interface mockup

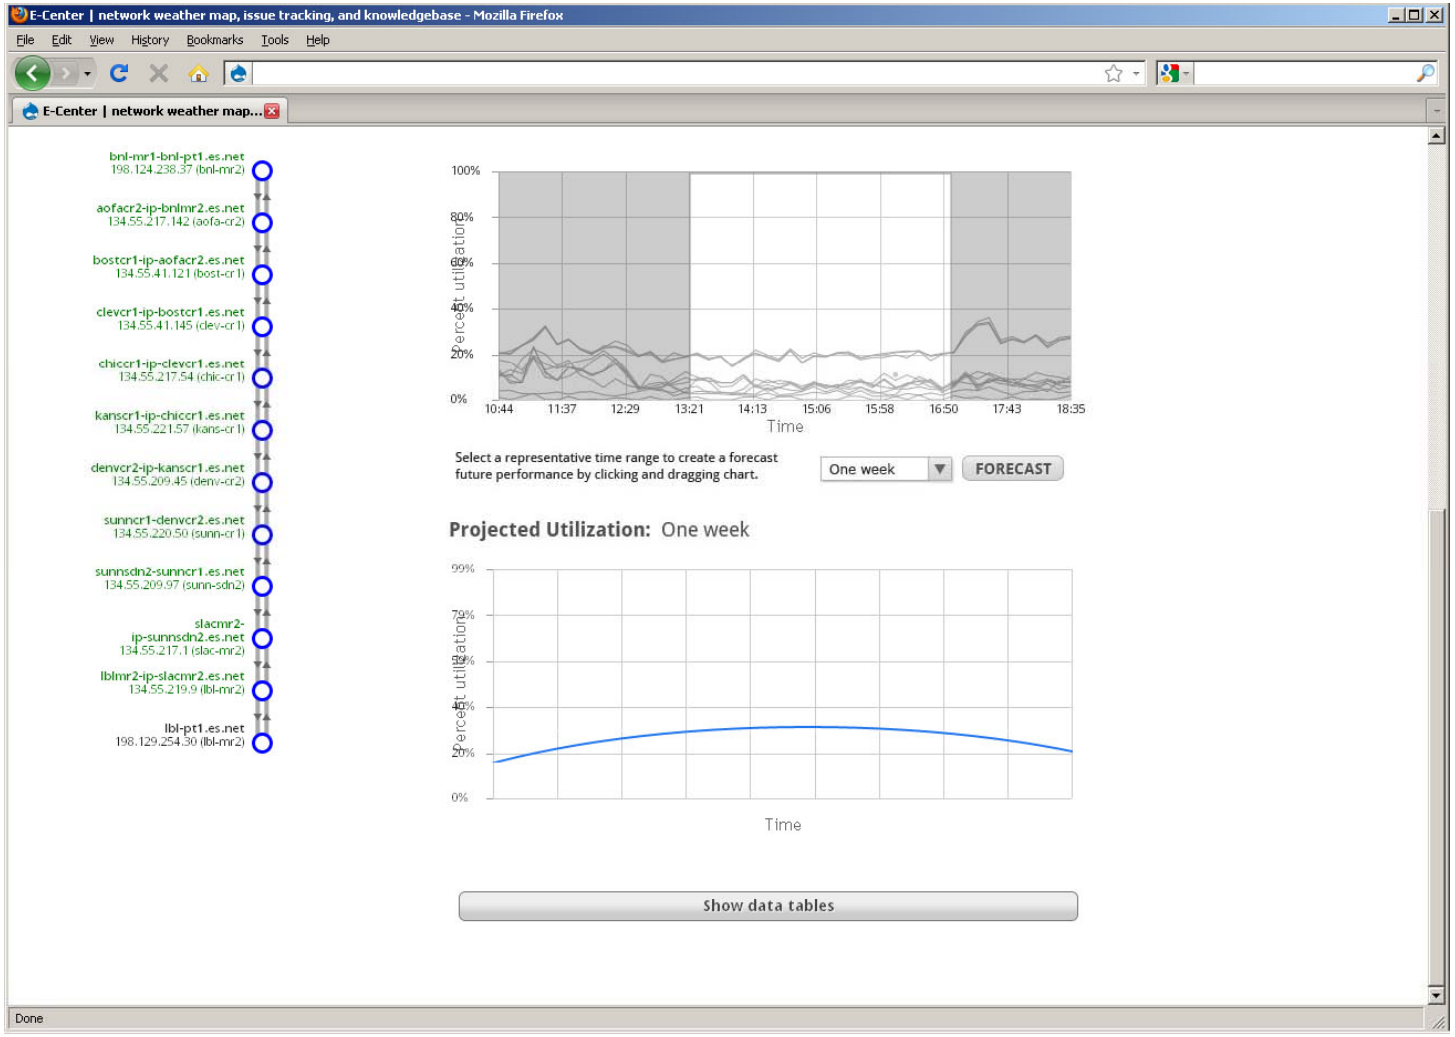

# E-Center: Prototype

- Data Collection and Data API is done
- REST Data service is working
- Prototype of the Front-end visualization done
- Prototype Issue tracking is implemented

# E-Center: Issues

- Stability of the remote pS-PS boxes, “health” monitoring
- Active monitoring box might not be “visible” on the pS-PS Lookup Service(LS), need for “inventory” service
- Failover assurance for the ESnet based centralized SNMP and Topology services *(informational bottleneck)*
- Scalability of the Data Service cache in case of full mesh monitoring effort, stress testing E-Center’s frontend and backend services
- Extending E-Center beyond DOE labs (possible collaboration with I2,OSG sites, US LHC Tier2 sites) ?

# FRONT-END DETAILS

David Eads

Fermilab

*25 October 2010*

# E-Center Frontend

## Platform choice: Why Drupal 6?

**Vs. Rapid Application Development frameworks (Django, Rails, CodeIgniter, etc):** Short term, a RAD framework might be more efficient. Long term, Drupal's popularity, strong community ecosystem and social tools offer more possibilities for a reasonable investment.

**Vs. Drupal 7 (to be released November, 2010):** The changes in the upcoming version of Drupal are appealing for this project. Currently, Drupal 7 and essential 3rd party modules are too unstable for production sites.

**Drupal 6 strengths:** Mature 3rd party module library, strong community, well-understood by many developers (including me), stable and secure.

**Drupal 6 weaknesses:** Fussy and weak AJAX library, many competing mapping solutions, reliance on a specific third party module (Content Construction Kit – CCK) makes complex data types (such as required by issue tracker) hard to implement and hard to debug.

# E-Center Front-end

## Key technical components

**OpenLayers:** A Javascript library for displaying map data, with good Drupal integration. OpenLayers is agnostic with respect to underling mapping solution – works with Google, Yahoo, Bing, MapBox, and other mapping solutions to provide tiles and base layers. E-Center has aggressively extended the Drupal OpenLayers project.

**Other essential contributed modules:** Views, CCK, AHAH Helper, Features.

**Custom modules**

- **Weathermap:** Provides primary weathermap user interface and interaction with backend.
- **Issues:** Provides issue tracking and query archiving.
- **UI modules:** Combobox module provides “searchable select box” based on jQuery UI library.
- **jqPlot:** Integrates jqPlot charting library with Drupal. Provides simple jQuery plugin that scrapes HTML tables and turns them into charts.

# E-Center Front-end

## Social functionality

**Models and inspirations:** StackOverflow and Stack Exchange network of sites, Afghanistan Election Data, Haiti Aid Map, ChartBeat, Google Analytics, NOAA Gulf Response map, ESNet Weathermap, Twitter, Reddit.

**Issue tracking as social networking:** Simple, fast issue tracking functionality allows for rapid, free-form conversations about various network issues, with original data always immediately accessible.

**Notifications:** Subscription-based model allows users to be notified of site activity via a wide variety of channels (email, SMS text messaging, RSS, Twitter, Facebook), and in some cases to interact with the site via social tools.

**Authentication support:** Users will be able to login and identify themselves using OAuth, OpenID, and Drupal's native authentication mechanisms. Authentication can be extended to support other sign-on technologies as needed.

**Planning and strategy needed:** What social media strategies will be useful and used by our audience?

# E-Center Front-end

**A user may get data for a path by clicking on the map, or selecting the time span and source and destination manually. Users may query based on site (such as BNL) or IP address.**

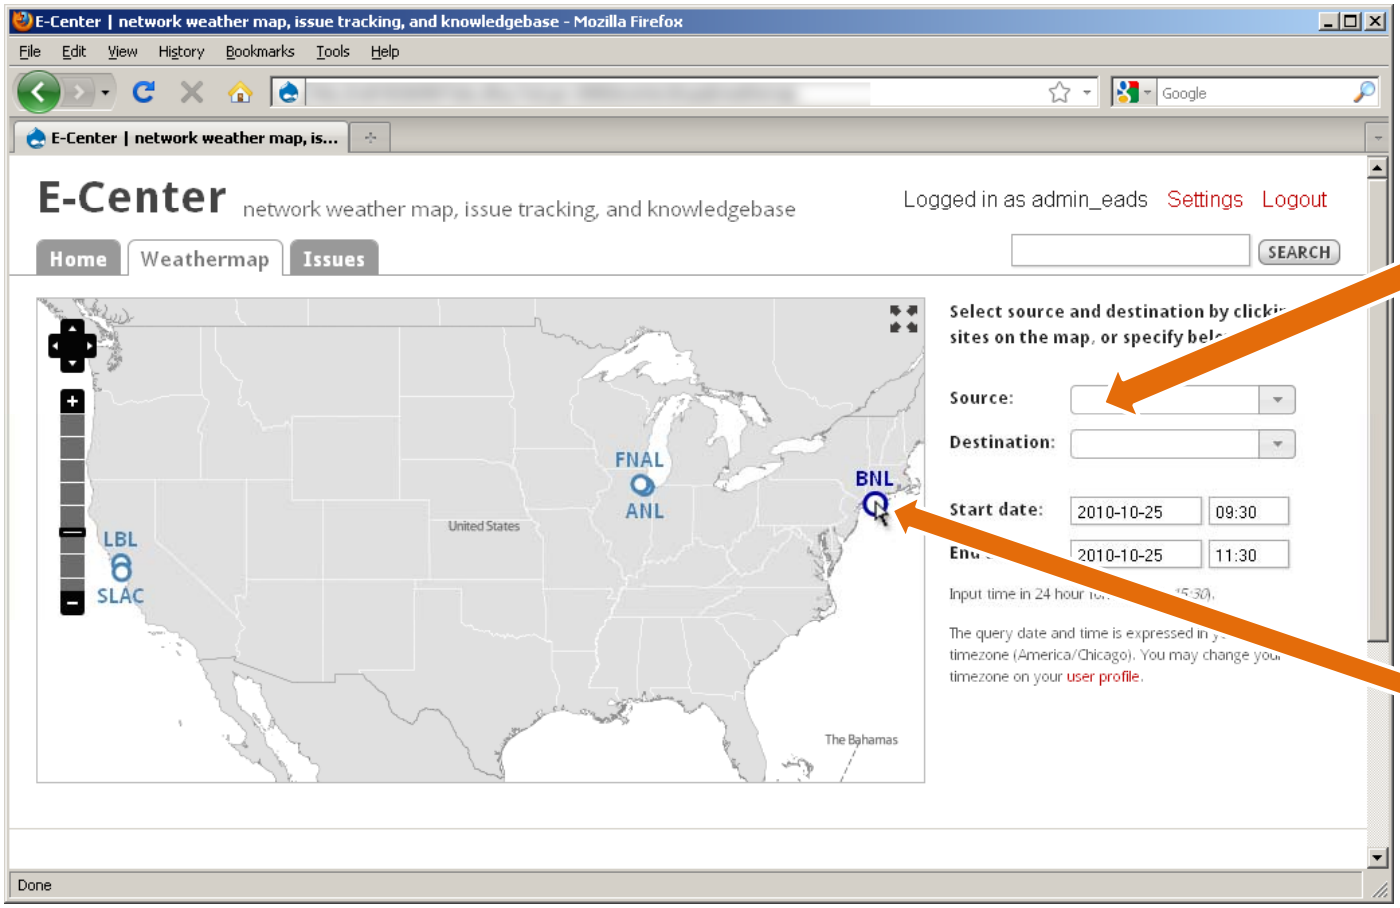


Select source and destination via by typing first few letters of site name or selecting from list.

Select source and destination by clicking map

# E-Center Front-end

**When source, destination, and a valid time range are provided, the frontend queries the data retrieval service for information about the provided path, and processes and caches the result for display or to be stored with an issue.**

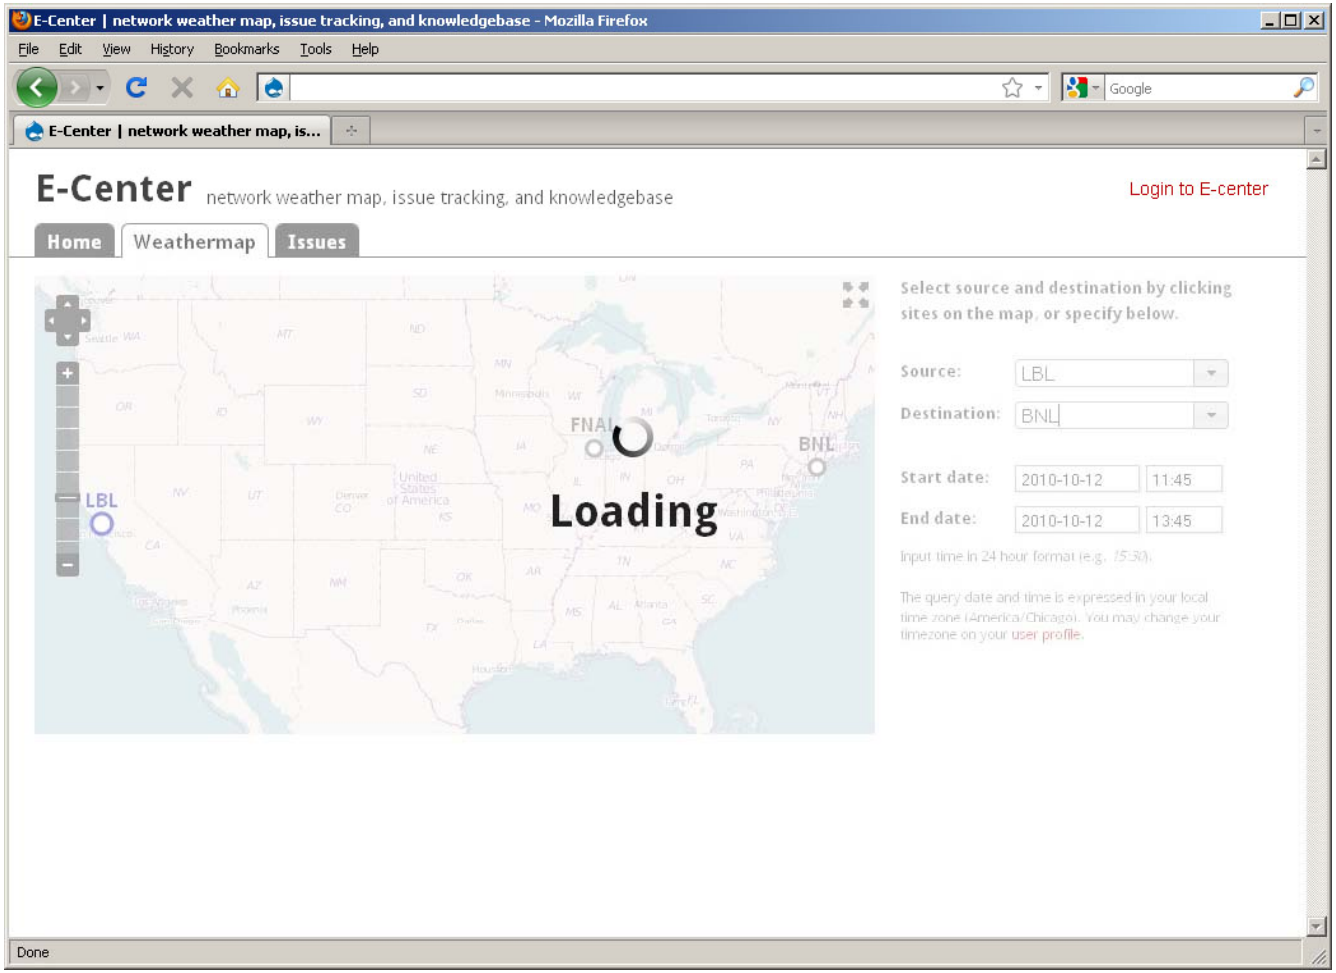

# E-Center Front-end

**We provide a variety of ways of viewing and interacting with network data.**

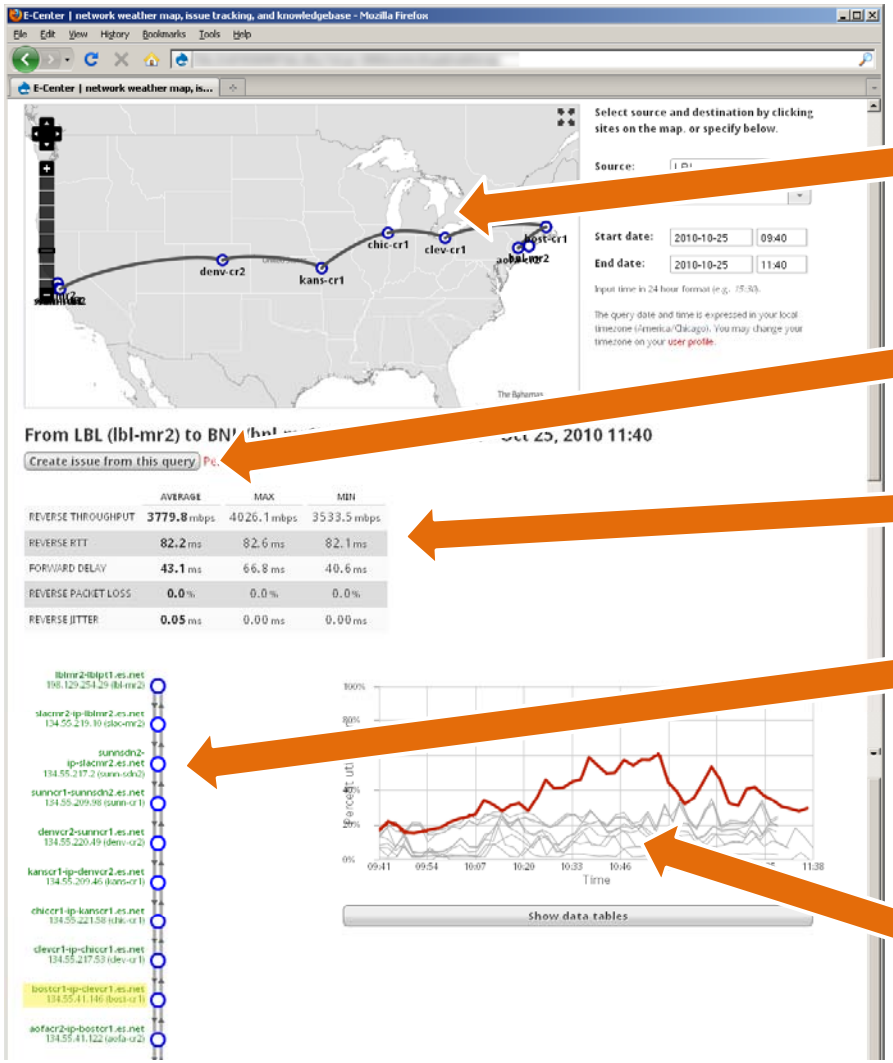


Traceroute plotted on map based on geographical location of hops.

Create a new issue based on query.

End-to-end data provides overview of path performance.

Traceroute “subway map” provides logical view of path, and shows path asymmetry.

SNMP utilization chart shows utilization for each hop over selected time period.

# E-Center Front-end

## User interface and visualization strategy

**Built for speed/fluidity:** Users should not have to make many decisions, click many buttons – queries at auto-submitted if a source and destination provided, with a sane default time window.

**Provide many views of data:** Subway map and geographical map show two different ways of thinking about same data.

**Self-revealing:** Proposed tool-tips in end-to-end table can help educate end-users about networking concepts and how to read the data. Similar strategies can be used throughout design.

**Highlight critical information:** Evolving color-coding strategy will allow users and network experts to quickly find network problems without digging.

**Data exploration:** Once areas of interest are identified, provide quick, fluid ways to get granular data, and make useful comparison.

# E-Center Front-end

## Interaction with network data

Traceroute “subway map” provides combined logical view of forward and reverse paths

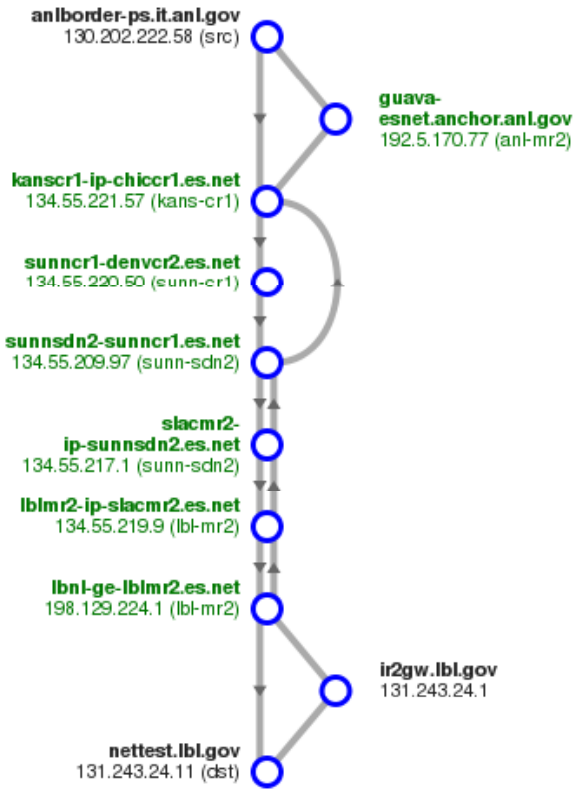


Easily zoom in/out on any portion of charts, series highlighting tied to subway map and geographical map

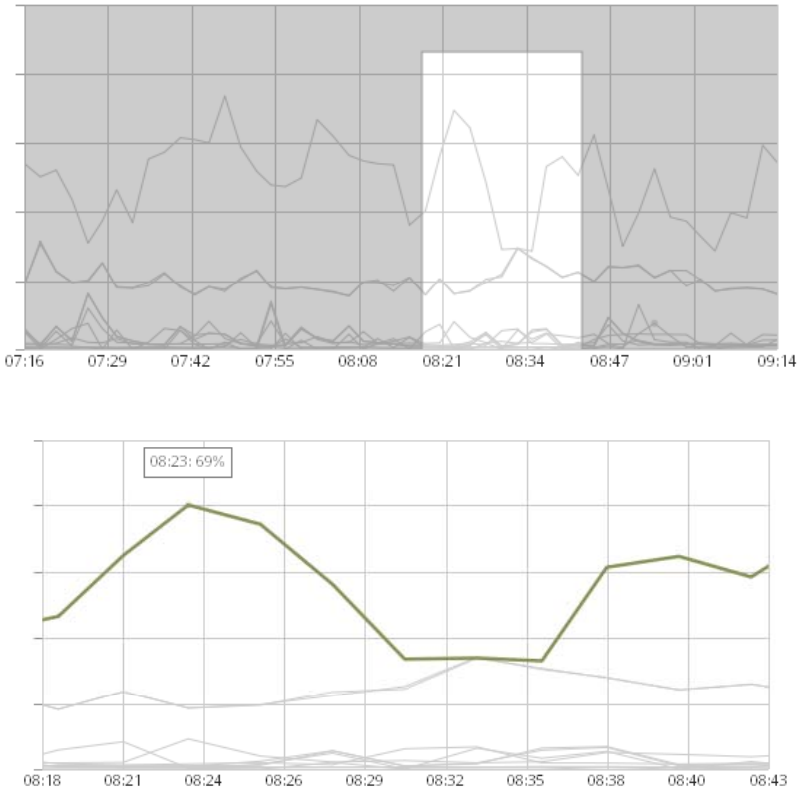

# E-Center Front-end

## Issue tracker

**Short, effective discussions about network issues:**
Combines blog, micro-blog (aka Twitter) and message forum functionality.

**Archive query data:** The issue tracker archives query results, enabling advanced querying and analysis (e.g. “when did most network problems occur?”)

**Subscription based, variety of notification channels:**
Users will be able to subscribe to all issues, specific issues, tags, or search terms. Notifications can be sent via email or alternative channels (social media and mobile integration possible).

**Knowledge base:** Knowledge base functionality currently envisioned as collections of issue threads based on tags and administrator flagging.

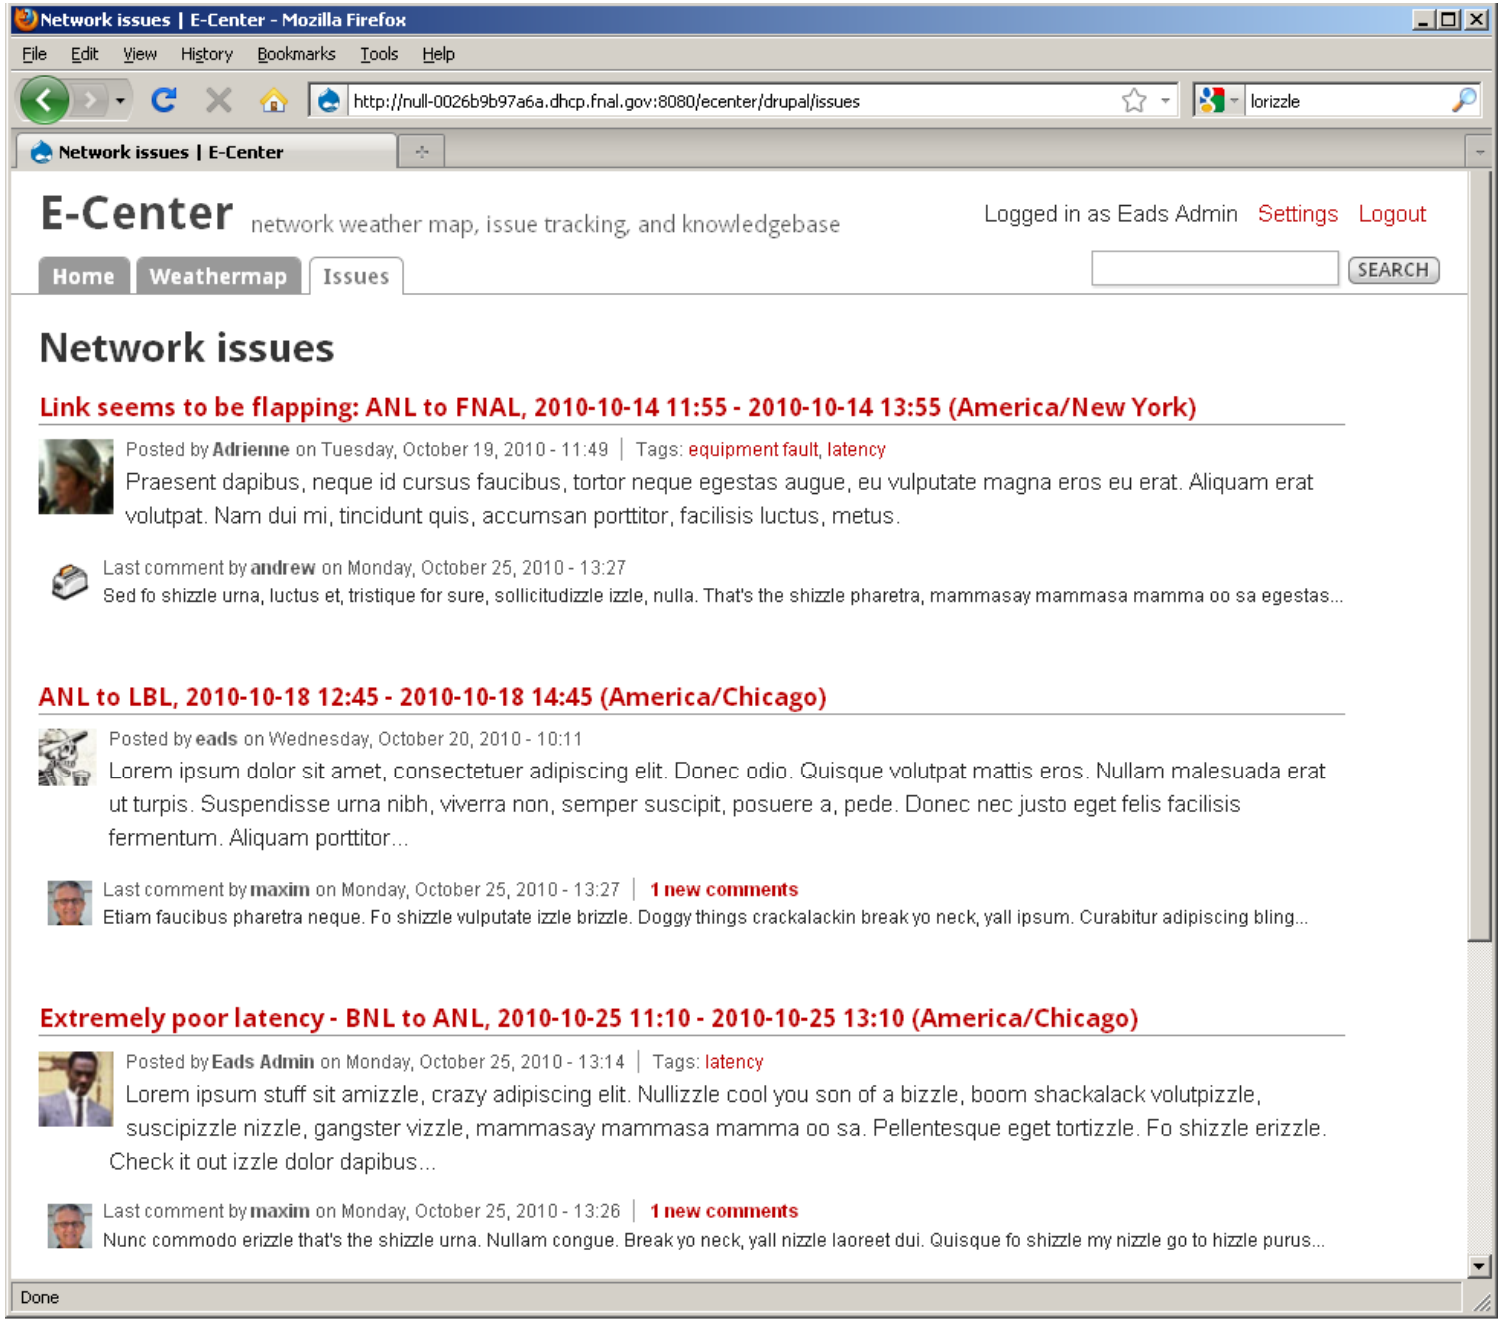
Network issues | E-Center - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://null-0026b9b97a6a.dhcp.fnal.gov:8080/ecenter/drupal/issues
lorizzle
Network issues | E-Center
E-Center
network weather map, issue tracking, and knowledgebase
Logged in as Eads Admin Settings Logout
Home Weathermap Issues
SEARCH
Network issues
Link seems to be flapping: ANL to FNAL, 2010-10-14 11:55 - 2010-10-14 13:55 (America/New York)
Posted by Adrienne on Tuesday, October 19, 2010 - 11:49 | Tags: equipment fault, latency
Praesent dapibus, neque id cursus faucibus, tortor neque egestas augue, eu vulputate magna eros eu erat. Aliquam erat volutpat. Nam dui mi, tincidunt quis, accumsan porttitor, facilisis luctus, metus.
Last comment by andrew on Monday, October 25, 2010 - 13:27
Sed fo shizzle urna, luctus et, tristique for sure, sollicitudizzle izzle, nulla. That's the shizzle pharetra, mammasay mammasa mamma oo sa egestas...
ANL to LBL, 2010-10-18 12:45 - 2010-10-18 14:45 (America/Chicago)
Posted by eads on Wednesday, October 20, 2010 - 10:11
Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Donec odio. Quisque volutpat mattis eros. Nullam malesuada erat ut turpis. Suspendisse urna nibh, viverra non, semper suscipit, posuere a, pede. Donec nec justo eget felis facilisis fermentum. Aliquam porttitor...
Last comment by maxim on Monday, October 25, 2010 - 13:27 | 1 new comments
Etiam faucibus pharetra neque. Fo shizzle vulputate izzle brizzle. Doggy things crackalackin break yo neck, yall ipsum. Curabitur adipiscing bling...
Extremely poor latency - BNL to ANL, 2010-10-25 11:10 - 2010-10-25 13:10 (America/Chicago)
Posted by Eads Admin on Monday, October 25, 2010 - 13:14 | Tags: latency
Lorem ipsum stuff sit amizzle, crazy adipiscing elit. Nullizzle cool you son of a bizzle, boom shackalack volutpizzle, suscipizzle nizzle, gangster vizzle, mammasay mammasa mamma oo sa. Pellentesque eget tortizzle. Fo shizzle erizzle. Check it out izzle dolor dapibus...
Last comment by maxim on Monday, October 25, 2010 - 13:26 | 1 new comments
Nunc commodo erizzle that's the shizzle urna. Nullam congue. Break yo neck, yall nizzle laoreet dui. Quisque fo shizzle my nizzle go to hizzle purus...
Done

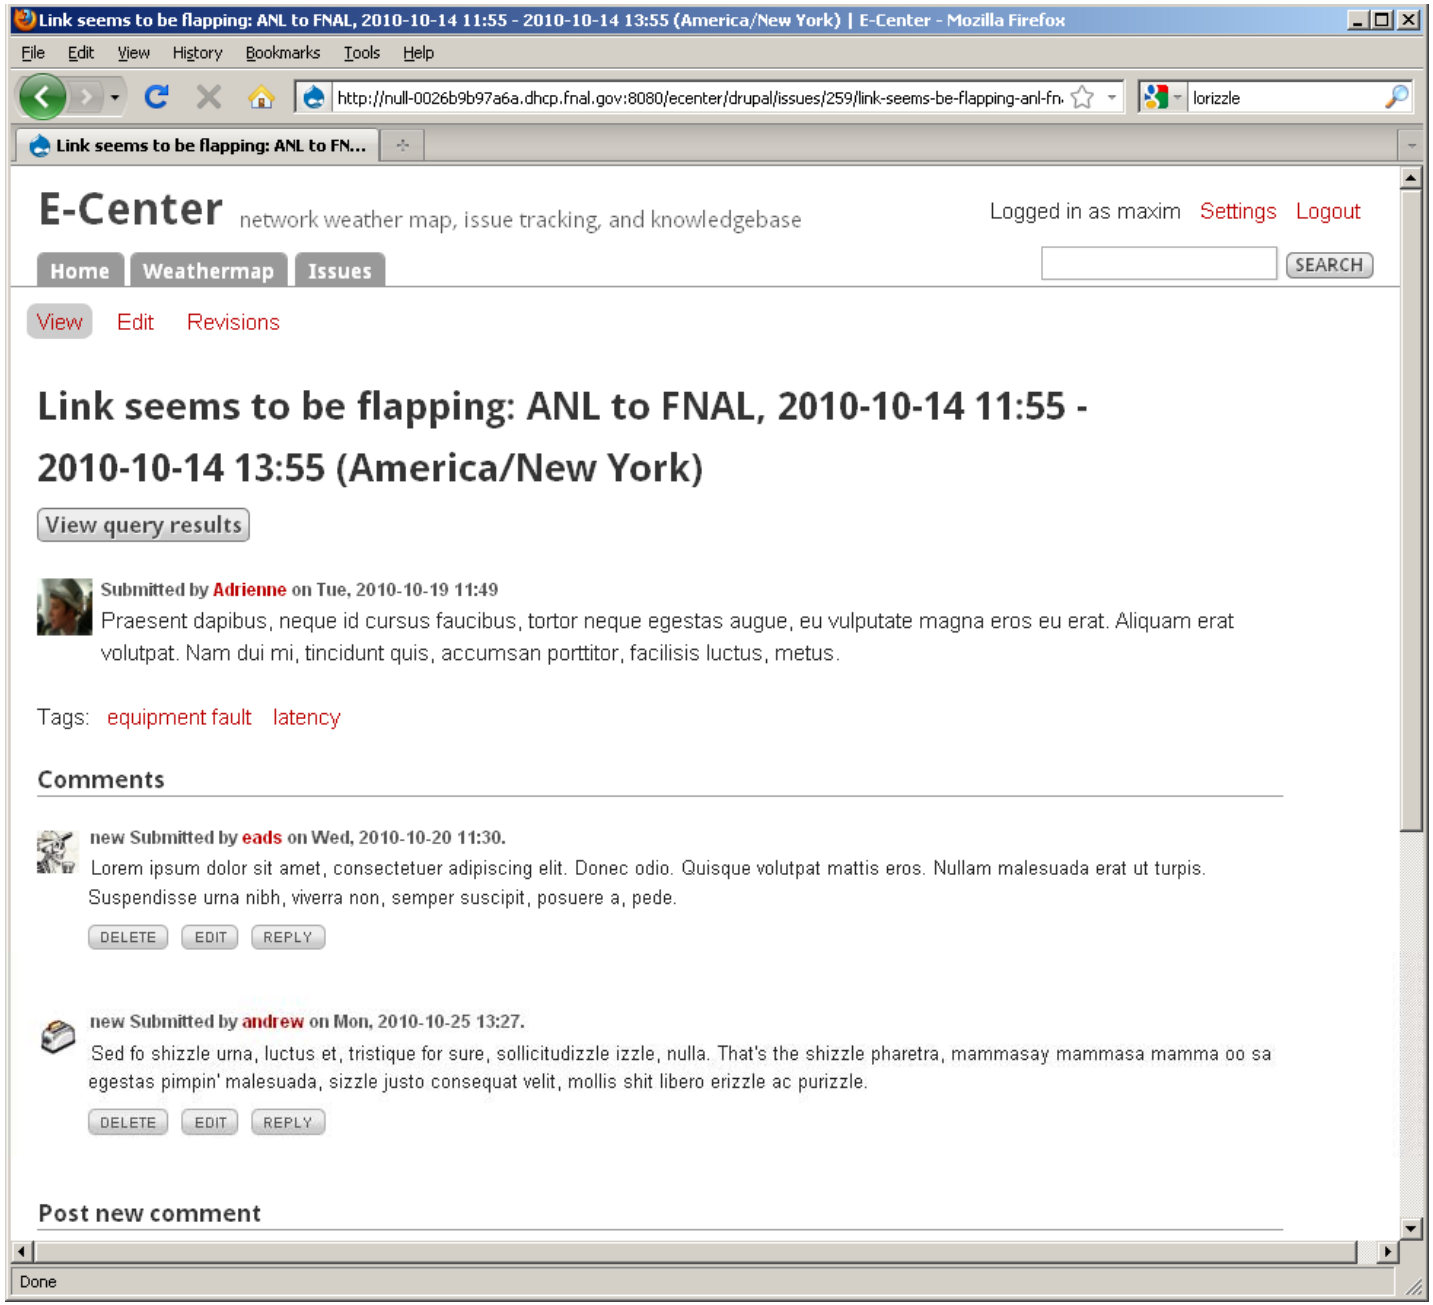
Link seems to be flapping: ANL to FNAL, 2010-10-14 11:55 - 2010-10-14 13:55 (America/New York) | E-Center - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://null-0026b9b97a6a.dhcp.fnal.gov:8080/ecenter/drupal/issues/259/link-seems-be-flapping-anl-fn
lorizzle
Link seems to be flapping: ANL to FN...
E-Center
network weather map, issue tracking, and knowledgebase
Logged in as maxim Settings Logout
Home Weathermap Issues
SEARCH
View Edit Revisions
Link seems to be flapping: ANL to FNAL, 2010-10-14 11:55 - 2010-10-14 13:55 (America/New York)
View query results
Submitted by Adrienne on Tue, 2010-10-19 11:49
Praesent dapibus, neque id cursus faucibus, tortor neque egestas augue, eu vulputate magna eros eu erat. Aliquam erat volutpat. Nam dui mi, tincidunt quis, accumsan porttitor, facilisis luctus, metus.
Tags: equipment fault latency
Comments
new Submitted by eads on Wed, 2010-10-20 11:30.
Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Donec odio. Quisque volutpat mattis eros. Nullam malesuada erat ut turpis. Suspendisse urna nibh, viverra non, semper suscipit, posuere a, pede.
DELETE EDIT REPLY
new Submitted by andrew on Mon, 2010-10-25 13:27.
Sed fo shizzle urna, luctus et, tristique for sure, sollicitudizzle izzle, nulla. That's the shizzle pharetra, mammasay mammasa mamma oo sa egestas pimpin' malesuada, sizzle justo consequat velit, mollis shit libero erizzle ac purizzle.
DELETE EDIT REPLY
Post new comment
Done

# E-Center Front-end

## Future Directions

1. **Subway map:** New visualization strategies for unstable traceroute paths
2. **Auto-updating:** “Real time” weathermap with periodic data updates
3. **Mobile device integration:** Smartphone (Android, iPhone) and tablet (iPad) interface
4. **Animation:** Animate historical performance on geographic map and subway map
5. **Charts:** Work with network experts to refine and improve charts
6. **Drupal distribution:** For replication. Harder than it should be with current Drupal technology.
7. **Map:** Enhanced map features (magnifying glass tool, “unclustering”)
8. **Enhanced backend services:** Integration with alternative data backends, multiple data backends (including prediction services)
9. **Enhanced social features:** True wiki/knowledgebase, sharing functionality, user customization options

# FUTURE CAPABILITIES

# E-Center: Development Roadmap

- Production system procurement and deployment
  - Dual x 12 core AMD™ Opteron 6172, 32GB memory, 4TB disk array
- Stress testing production environment, failover procedures for the E-Center portal and Data Webservice
- Implementation and Integration of Forecasting & Data Analysis
- Web 2.0 components – issues tracking deployment, searchable knowledge base, user's space, support for the mobile
- Rolling out DOE testbed; extending scope to remaining DOE labs
  - Centralized Configuration Management, baseline config for all participated perfSONAR-PS hosts
  - Traceroute MA deployment at participating DOE labs
  - Topology (and SNMP) MA deployments at participating DOE labs

# Future Capabilities - Strategic

- Support for user-provided traceroute feeds
- Circuit-based network paths
- Support for other external services (anomaly, etc.)
- On-demand measurement test capabilities
  - Starting with TraceRoute MA
- User Management
- IPv6

Demo