

*Supporting Advanced Scientific Computing  
Research • Basic Energy Sciences • Biological  
and Environmental Research • Fusion Energy  
Sciences • High Energy Physics • Nuclear Physics*

# Network Monitoring and Visualization at ESnet

**Jon Dugan, Network Engineer  
ESnet Network Engineering Group**


**February 3, 2010**

**Winter Joint Techs, Salt Lake City, UT**



# Overview


Data Collection  
(ESxSNMP)




Data Visualization  
(Graphite)



Event/Metadata  
Log  
(Net Almanac)







# ESxSNMP: Goals

---

- Automate everything possible
- Provide summaries but don't lose raw data
  - Disk is cheap
  - It can be useful to take a hard look at the past
- Flexibility and scalability
- Minimize up front assumptions
- Protect data collection from DoS by users
- Make data easy to access and manipulate





# ESxSNMP: Polling

- Interface metadata
  - Automatically detects new interfaces
  - Automatically detects interface changes
  - Historical log of interface info
- Automatic addition of new devices
  - Detects new entries in our RANCID database
- Allow arbitrary transformations at poll time
  - Stored by ifDescr rather than ifIndex
    - ifHCInOctets.fxp0 vs ifHCInOctets.1
    - Sidesteps problem of ifIndex renumbering
  - Store firewall counters by name
  - Custom transformations via simple Python class
- High Performance
  - 7000 interfaces every 30 seconds
  - Storing the metrics is limiting factor






# ESxSNMP: Metrics Storage

---

- TSDB
  - RRD summarizes data
  - Optimized for retrieval by timestamp
  - Allows for multilevel storage
  - Similar interface to RRD, but fewer surprises
  - Distinct library
- Can be distributed
  - Disk I/O can be an issue
    - SSD
    - RAM disk
  - Allow many requests to be serviced
  - The design accounts for this, current deployment does not






# ESxSNMP: Metrics Retrieval

---

- Allow easy, consistent access to data
  - Data will be used in unanticipated ways
  - Language neutral
- Technical details
  - RESTful interface
  - URL hierarchy: eg, `core-rtr-1/interface/xe-0_0_0/in`
  - HTTP transport using HTTP semantics
  - Data returned in JSON format






# ESxSNMP: Tested platforms

---

- Standard MIB polling
  - Juniper
  - Cisco
  - Foundry
  - Force10
- Custom MIB polling
  - Juniper: firewall and class of service
  - Cisco: CPU utilization






# Graphite: Data visualization

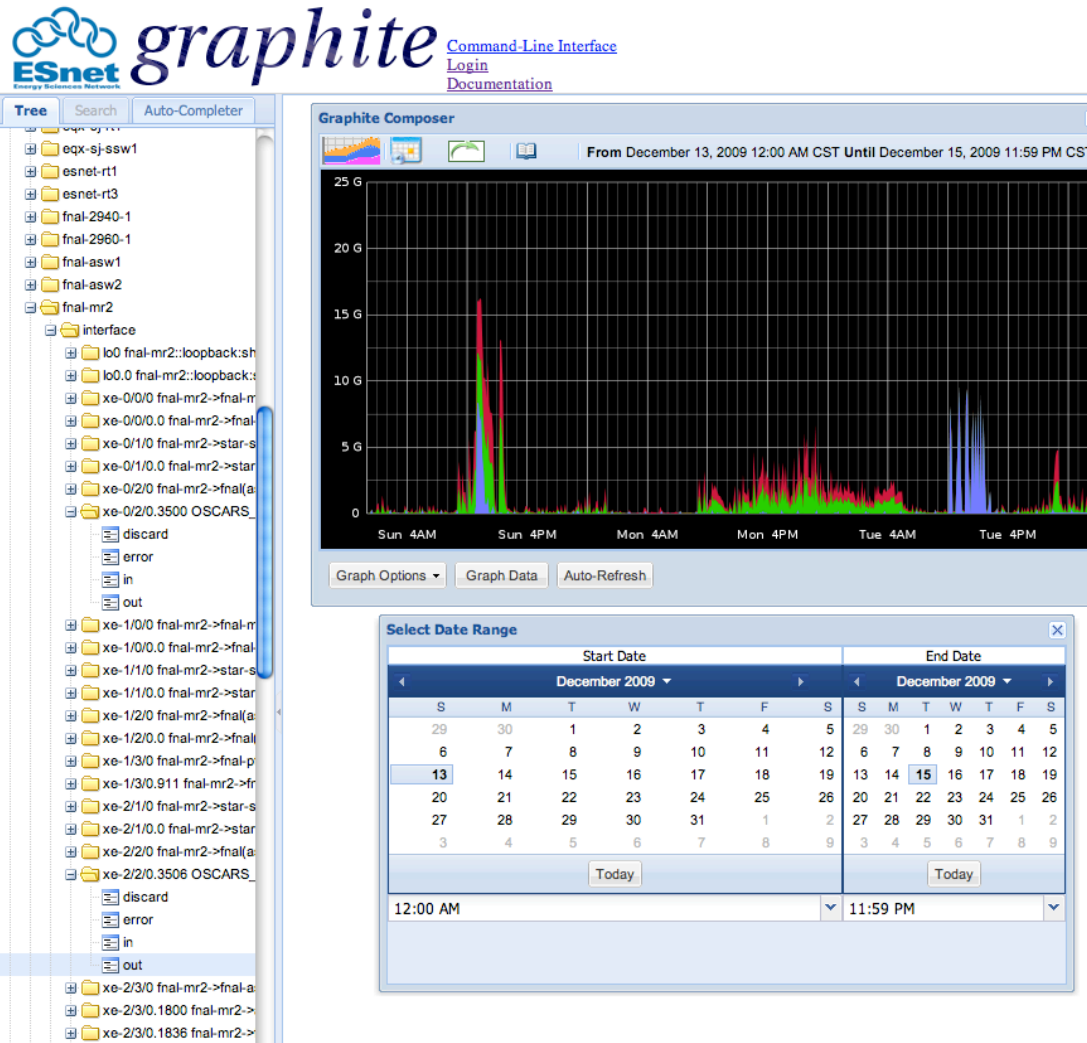
---

- Developed by Orbitz to visualize internal performance data
- Clean design allowed easy integration
- Flexible
- Bookmarkable
- Fast





# Graphite: Screenshot



**graphite** [Command-Line Interface](#)  
[Login](#)  
[Documentation](#)

**Tree** Search Auto-Completer

- exq-sj-ssw1
- esnet-rt1
- esnet-rt3
- fnal-2940-1
- fnal-2960-1
- fnal-asw1
- fnal-asw2
- fnal-mr2
  - interface
    - lo0 fnal-mr2::loopback:sh
    - lo0.0 fnal-mr2::loopback:sh
    - xe-0/0/0 fnal-mr2->fnal-mr2
    - xe-0/0/0.0 fnal-mr2->fnal-mr2
    - xe-0/1/0 fnal-mr2->star-s
    - xe-0/1/0.0 fnal-mr2->star-s
    - xe-0/2/0 fnal-mr2->fnal(a
    - xe-0/2/0.3500 OSCARS
    - discard
    - error
    - in
    - out
  - xe-1/0/0 fnal-mr2->fnal-m
  - xe-1/0/0.0 fnal-mr2->fnal
  - xe-1/1/0 fnal-mr2->star-s
  - xe-1/1/0.0 fnal-mr2->star
  - xe-1/2/0 fnal-mr2->fnal(a
  - xe-1/2/0.0 fnal-mr2->fnal
  - xe-1/3/0 fnal-mr2->fnal-p
  - xe-1/3/0.911 fnal-mr2->fr
  - xe-2/1/0 fnal-mr2->star-s
  - xe-2/1/0.0 fnal-mr2->star
  - xe-2/2/0 fnal-mr2->fnal(a
  - xe-2/2/0.3506 OSCARS
  - discard
  - error
  - in
  - out
  - xe-2/3/0 fnal-mr2->fnal-a
  - xe-2/3/0.1800 fnal-mr2->
  - xe-2/3/0.1836 fnal-mr2->

**Graphite Composer** From December 13, 2009 12:00 AM CST Until December 15, 2009 11:59 PM CST

Graph Options Graph Data Auto-Refresh

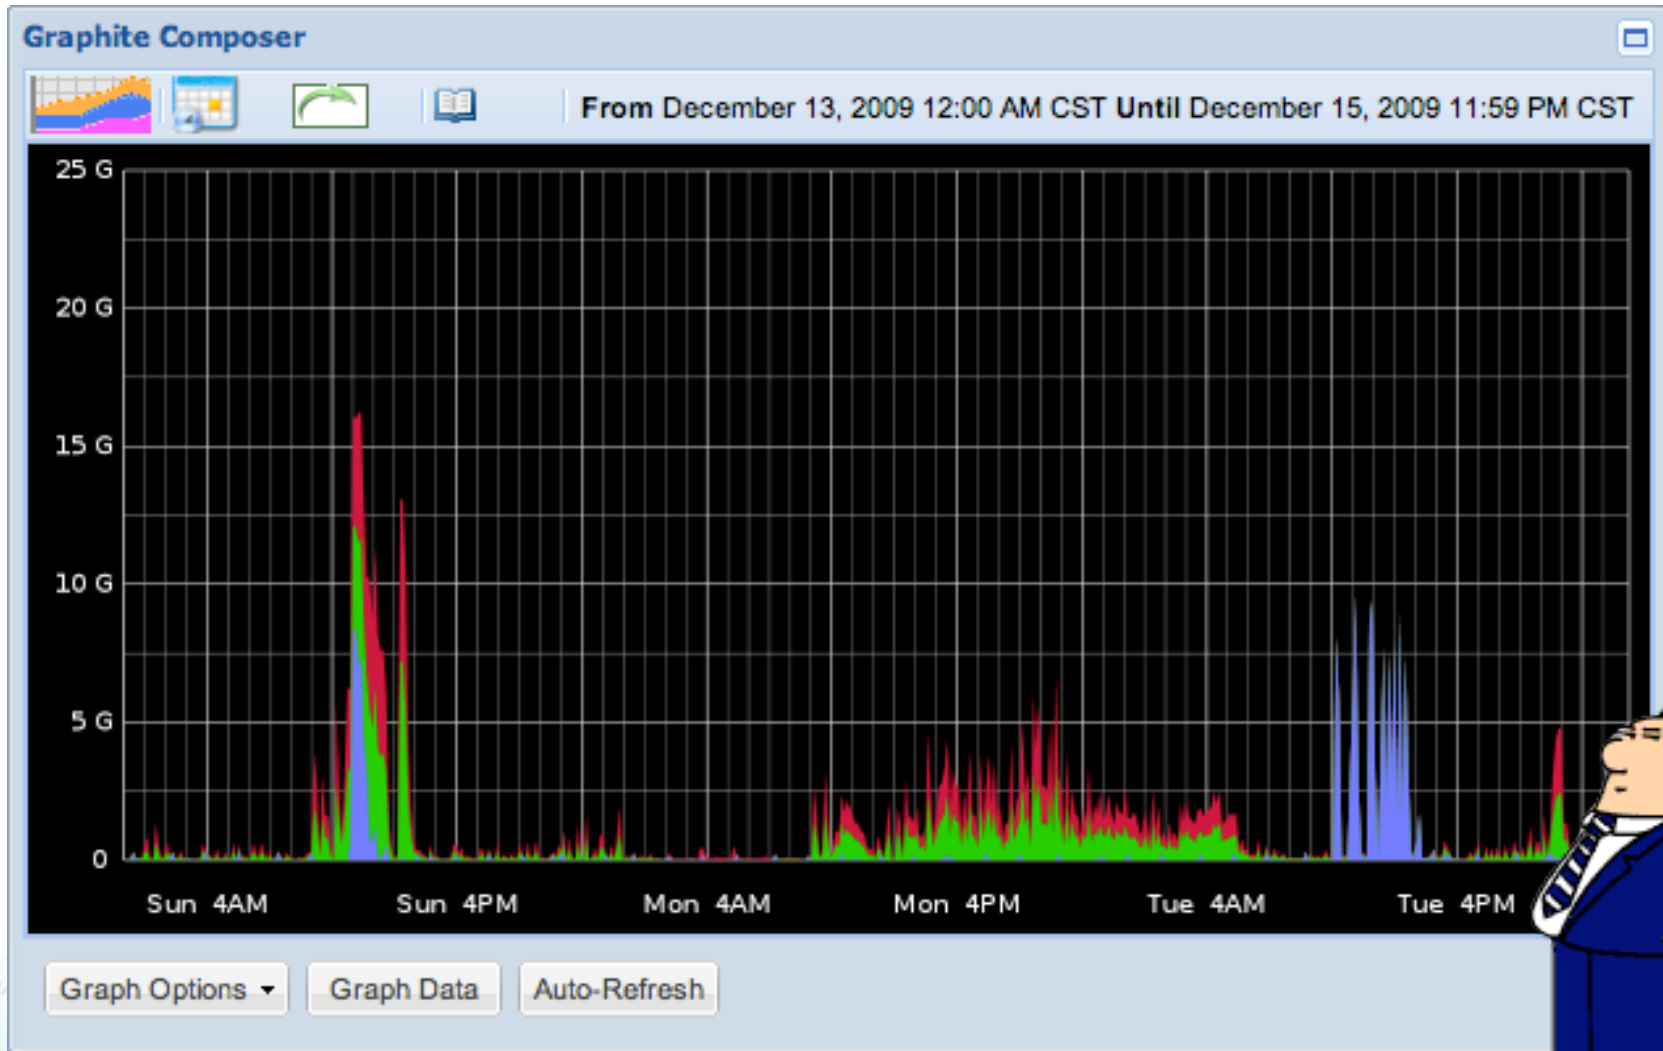
**Select Date Range**


Start Date							End Date						
December 2009													
S	M	T	W	T	F	S	S	M	T	W	T	F	S
29	30	1	2	3	4	5	29	30	1	2	3	4	5
6	7	8	9	10	11	12	6	7	8	9	10	11	12
13	14	15	16	17	18	19	13	14	15	16	17	18	19
20	21	22	23	24	25	26	20	21	22	23	24	25	26
27	28	29	30	31	1	2	27	28	29	30	31	1	2
3	4	5	6	7	8	9	3	4	5	6	7	8	9

Today Today

12:00 AM 11:59 PM

# What's that, right there?







# Net Almanac

---

“Why is there a traffic spike on this graph?”

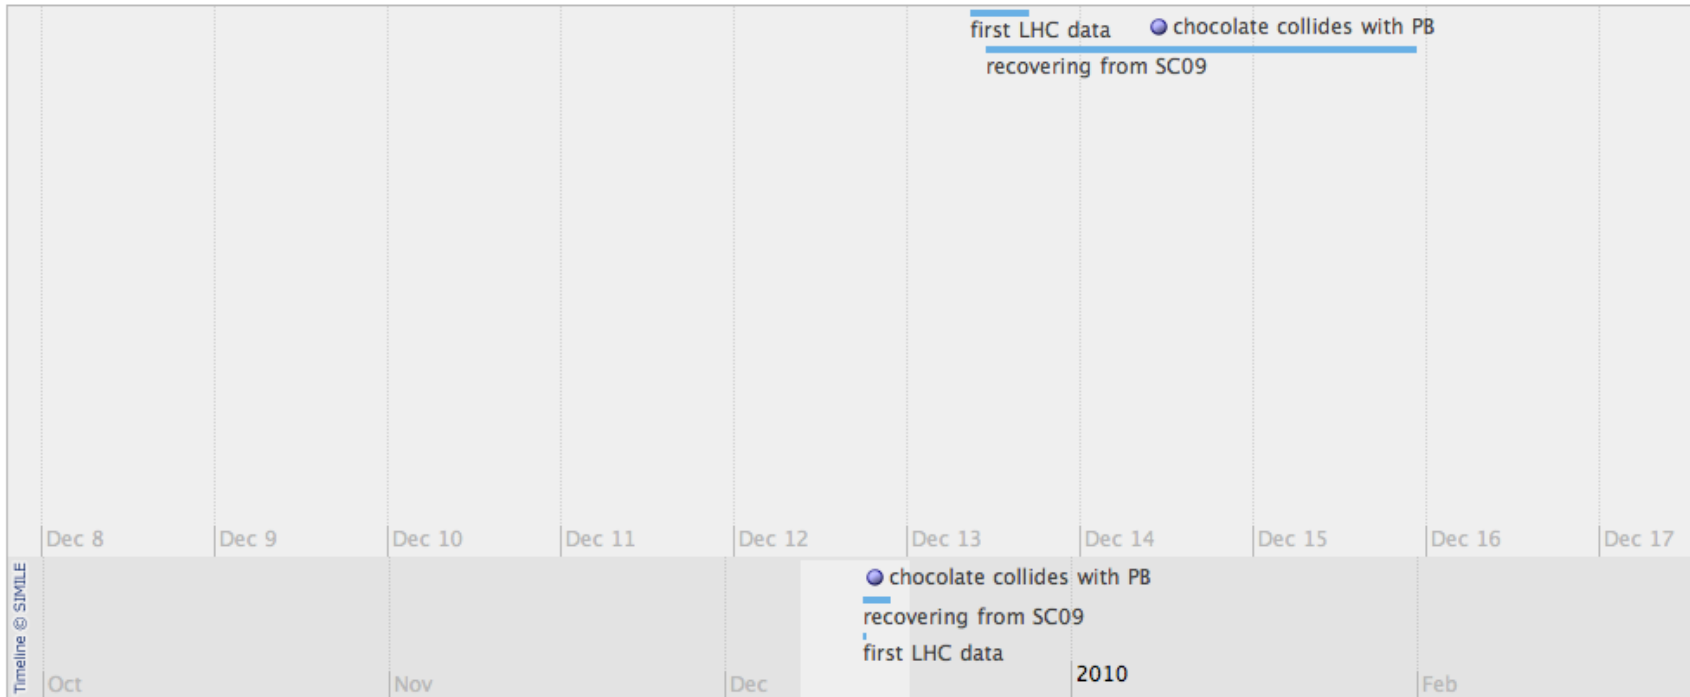
- Long term memory for events
  - conferences, data trials, etc
  - outages, maintenance
  - interface up/down
  - configuration changes
- Human interface
- Machine interfaces: REST/JSON






# Net Almanac: Example

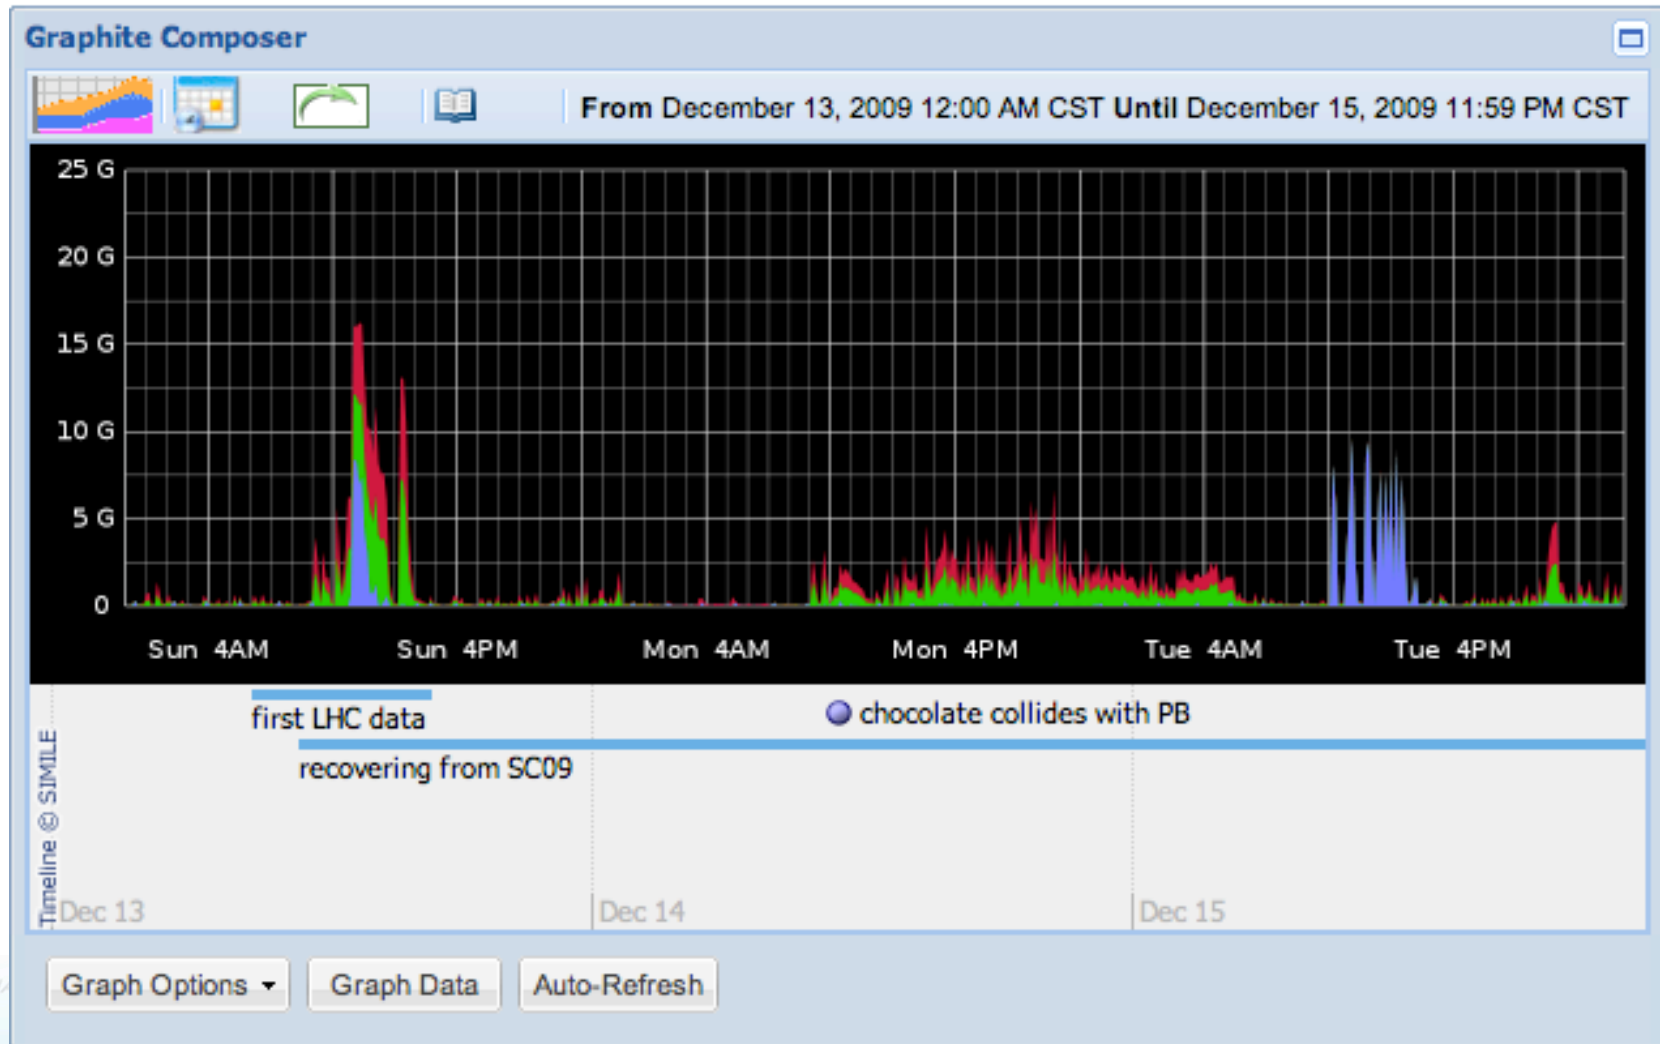
Timeline of Events between 2009-12-13 and 2009-12-16; 3 events total




Click and drag with your mouse to navigate the timeline. You may click on an event for more detail. All times are in Pacific Standard Time.



# You got your chocolate in my PB!






# Free your data

---

- The web has made a lot of data human accessible
- Needs to be more machine accessible without sacrificing usability
- Stop reinventing the wheel
- HTTP and JSON/XML ubiquitous
- RESTful Services





# RESTful Integration Successes

---

- Graphite
  - Consumes data from ESxSNMP
  - Consumes data from Net Almanac
  - Provides data as plots, CSV, or JSON
- Net Almanac
  - Consumes data from syslog, outage calendar
  - Provides data as JSON
- ESnet Weathermap
  - Consumes data from ESxSNMP
  - Java / Python living together
  - <http://weathermap.es.net/>
- Traceroute visualizer
  - Consumes Graphite plots
  - Consumes perfSONAR topology information






# perfSONAR Integration

---

- ESxSNMP and Graphite used at SC09
  - Primary SNMP polling for SCinet
  - Used to judge Bandwidth Challenge
- Implement a bridge between ESxSNMP and perfSONAR in about 45 minutes
  - Perl and Python living together
- Native perfSONAR interface on the way
  - Python perfSONAR library in development









# RESTful Services: Examples

---

- Possible future services
  - Outage Notifications
  - Contact Information (NOCs, etc)
  - Read twitter feeds
  - perfSONAR?
  - OSCARS
  - Access to other report data (monthly stats)







# Lessons Learned

---

- Don't reinvent the wheel
- Sometimes you need a different kind of wheel
- Simplicity requires effort
- Everything is a struggle
- Programmers are optimists (sort of)
- Simple, language neutral APIs easily accommodate unexpected use cases






# Links and whatnot

---

- Services
  - <http://stats1.es.net/graphite/>
  - <http://weathermap.es.net/>
- Code
  - <http://code.google.com/p/esxsnmp/>
  - <http://code.google.com/p/tsdb/>
  - <http://code.google.com/p/net-almanac/>
  - <http://code.google.com/p/esnet-weathermap/>
  - <http://graphite.wikidot.com/>
- REST
  - <http://www.infoq.com/articles/rest-introduction>
  - <http://tomayko.com/writings/rest-to-my-wife>
- Me
  - Jon Dugan <[jdugan@es.net](mailto:jdugan@es.net)>



# Extra Slides





# RESTful Services

---

- Representational State Transfer
  - Fielding's PhD Thesis
  - Provides an “architectural style”
- Common Usage
  - Exposed resources
  - Multiple representations
    - Human: HTML/CSS/PNG, etc
    - JSON

