


# ESnet Network Measurement Current Status

*Joe Metzger*

Energy Sciences Network  
Lawrence Berkeley National Laboratory

Jan 24th 2008  
ESCC meeting



*Networking for the Future of Science*



# Bandwidth Test Platforms

---

- Deployment
  - MANs: BNL, ANL, FNAL, LBL
  - Hubs: WASH, Starlight, ALBU, SUNN, AOFA, CHIC, ATLA
  - 20 more will be ordered, and most will be installed in FY08
- Current Uses
  - Acceptance Testing
  - Debugging network problems
- Additional uses once we resolve some problems and get more systems deployed
  - Regular scheduled bandwidth testing
    - Internal to ESnet IP & SDN Network
    - External to important peers, and ESnet sites that are interested
  - Dynamic end-user testing via
    - perfSONAR
    - ESnet Performance Center

# Latency Test Platforms

---

- Deployment Status
  - 11 in field
  - ~20 more to be ordered & deployed
- Most are using CDMA clocks
  - Signal not available at all locations
  - Signal degrades over time as POPs are populated
- IRIG
  - Original plan to get IRIG sources from Level 3 at several locations is not going smoothly.

# ESnet Statistics Collection Systems

---

- SNMP collection Systems
  - Raw Collectors
    - Used for monthly statistics reports
  - MRTG
    - Several different collections for:
      - ESnet Info
      - NetInfo
  - ESxSNMP & TSDB
    - ESxSNMP - SNMP collection system
    - TSDB Time Series Data Base
    - Developed in-house by Jon Dugan
    - Main focus is improved meta-data management
    - Open source ([code.google.com](http://code.google.com))
    - Should replace other ESnet SNMP stats collections systems in the next 6 months or so.
- NetFlow collection system

# perfSONAR Collaboration Status

---

- perfSONAR MDM 3.0
  - A bundle of all perfSONAR applications that complete the release management process
  - One goal is to include all of the parts necessary to support the GEANT MDM pilot
  - Substantial completion in February or March
- perfSONAR PS 0.06 Beta
  - A bundle of perfSONAR applications
  - Written in Perl, primarily in the US
  - Officially released Monday
    - But I have been running it for months.
- LHCOPN MDM
  - A proposal by Dante to deploy perfSONAR appliances at all of the LHC Tier1 sites.

# ESnet perfSONAR near-term goals

---

- NetInfo
  - Convert to use perfSONAR data sources
  - Dynamically generate all maps and navigation
  - Expand to utilize perfSONAR data from other domains
- Work with the LHCOPN community to deploy perfSONAR measurement services at the Tier 1 and 2 centers
- Work to get the perfSONAR collaboration focused on dynamic hybrid network issues
  - publication, visualization, and analysis of dynamic circuits

# What is perfSONAR

---

- A collaboration
  - Production network operators focused on designing and building tools that they will deploy and use on their networks to provide monitoring and diagnostic capabilities to themselves and their user communities.
- An architecture & a set of protocols
  - Web Services Architecture
  - Protocols based on the Open Grid Forum Network Measurement Working Group Schemas
- Several interoperable software implementations
  - Java & Perl
- A growing deployed measurement infrastructure

# perfSONAR Collaborators

---

- RNP
- ARNES
- BELNET
- CARNET
- CESNET
- CYNET
- DANTE
- DFN
- ESnet
- FCCN
- FERMI
- GARR
- GEANT
- GRNET
- HEAnet
- Internet2
- ISTF
- POZNAN
- UNINETT
- University of Delaware
- Renater
- RedIRIS
- SLAC
- SWITCH
- SURFnet

And anybody else I missed

# Widely Deployed perfSONAR Services

---

## Utilization MA:

- Internet2
- ISTF/ACAD
- CARNet
- Cesnet
- CYnet
- ESnet
- FERMI
- GARR
- GEANT2
- GRNET
- MREN
- PIONIER
- RedCLARA
- RENATER
- RNP
- SEEREN
- SWITCH
- Uninett

## Layer 2 Circuit Status for LHCOPN Circuits

- ASCC
- CERN
- ESnet
- GARR
- GEANT2
- NORDUnet
- Renater
- REDIris
- Sara/SURFnet
- SWITCH
- CANARIE
- USLHCnet
- UKERNA

# perfSONAR Architecture

---

- Interoperable network measurement middleware:
  - Modular
  - Web services-based
  - Decentralized
  - Locally controlled
- Integrates:
  - Network measurement tools
  - Network measurement archives
  - Discovery
  - Authentication and authorization
  - Data manipulation
  - Resource protection
  - Topology
- Based on:
  - Open Grid Forum Network Measurement Working Group schema.
- Each Domain
  - Collects measurement data internal to their domain
  - May collect cross-domain measurement data
  - **Controls publication of the information they collect and manage**

# perfSONAR codebases

---

- 2 Interoperable implementations
  - Java
  - Perl

# perfSONAR Services (1)

---

- Measurement Point Service
  - Enables the initiation of performance tests
  - Currently in use: bandwidth, latency, looking glass, circuit status
- Measurement Archive Service
  - Stores and publishes performance monitoring results
  - Currently in use: interface utilization and errors, bandwidth, latency, circuit status
- Lookup Service
  - Allows the client to discover the existing services and other LS services.
  - Dynamic: services registration themselves to the LS and mention their capabilities, they can also leave or be removed if a service goes down.
  - Currently available: service discovery in a single domain. Global integration getting close.
- Authentication Service
  - Based on Existing efforts: Internet2 MAT, GN2-JRA5
  - Authentication & Authorization functionality for the framework
  - Users can have several roles, the authorisation is done based on the user role.
  - Trust relationship between networks
  - Currently status: Some services work with Eduroam identities

## perfSONAR Services (2)

---

- Transformation Service
  - Transform the data (aggregation, concatenation, correlation, translation, etc).
  - Current status: perfSONAR\_PS includes a transform service that determines circuit status based on info from several other services
- Topology Service
  - Publish the network topology information via the framework.
  - Currently status: perfSONAR\_PS includes a useful server. CNIS exports a perfSONAR topology interface
- Resource protector
  - Arbitrate the consumption of limited resources.
  - Current status: not in perfSONAR, but some capabilities in underlying tools.

# Weather Maps - Internet2

DCN Weathermap - ESnet (New York)

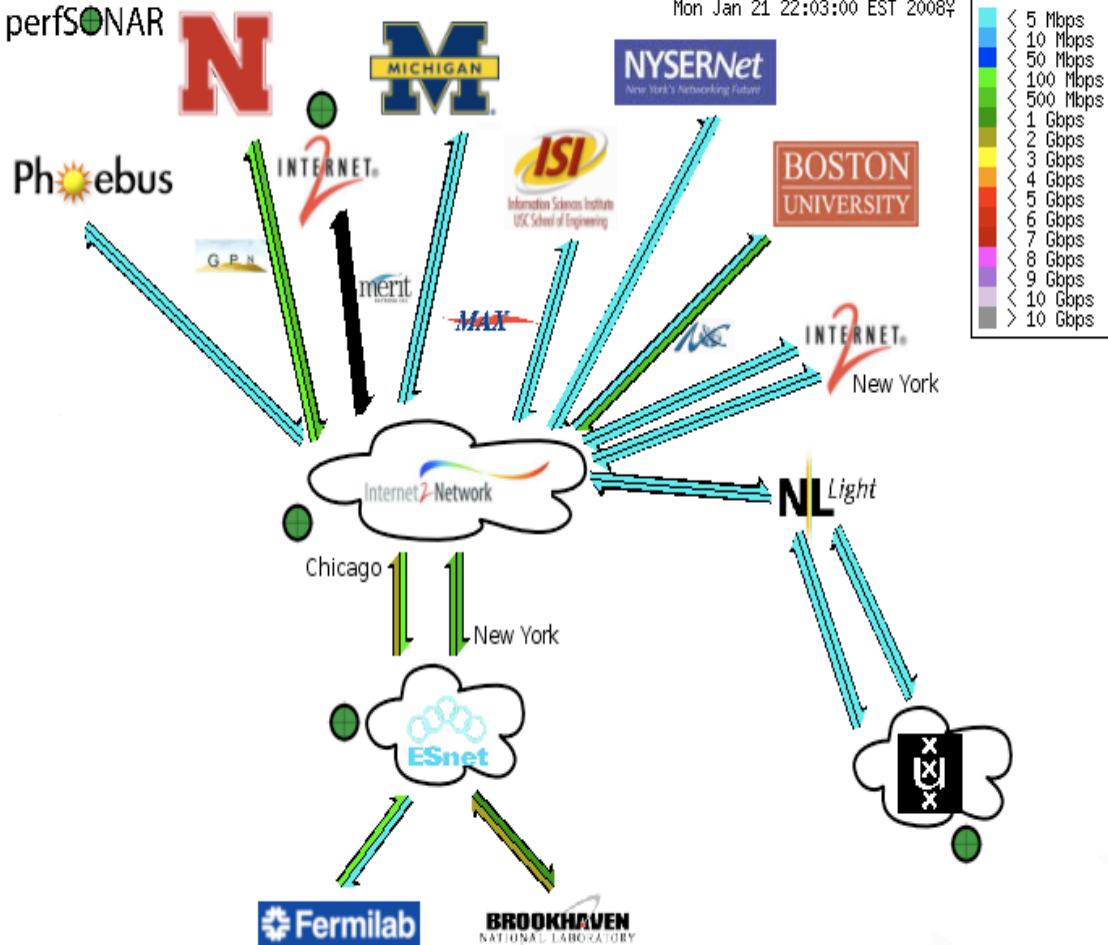
http://packrat.internet2.edu/~zurawski/TIP2008/ESnet\_New\_York.html

Getting Started Latest Headlines NetInfo-Internal Home - Internet2 Wiki perfSONARUI perfSONAR Wiki PerfSONAR Domain ... USLHCNet DOE Revie... Program for Winter 2...

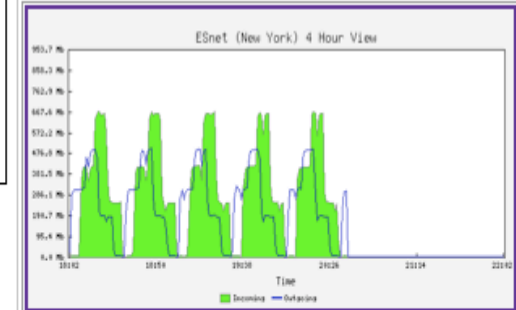
## ESnet (New York) Performance

Mon Jan 21 22:03:00 EST 2008

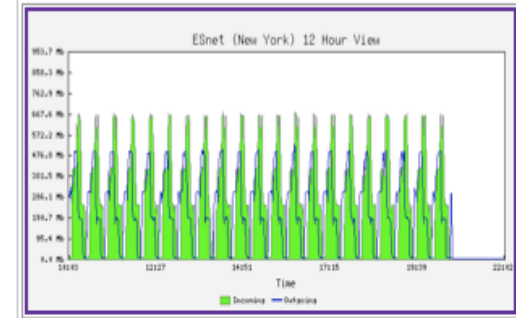
Powered By perfSONAR



### 4 Hour Graph

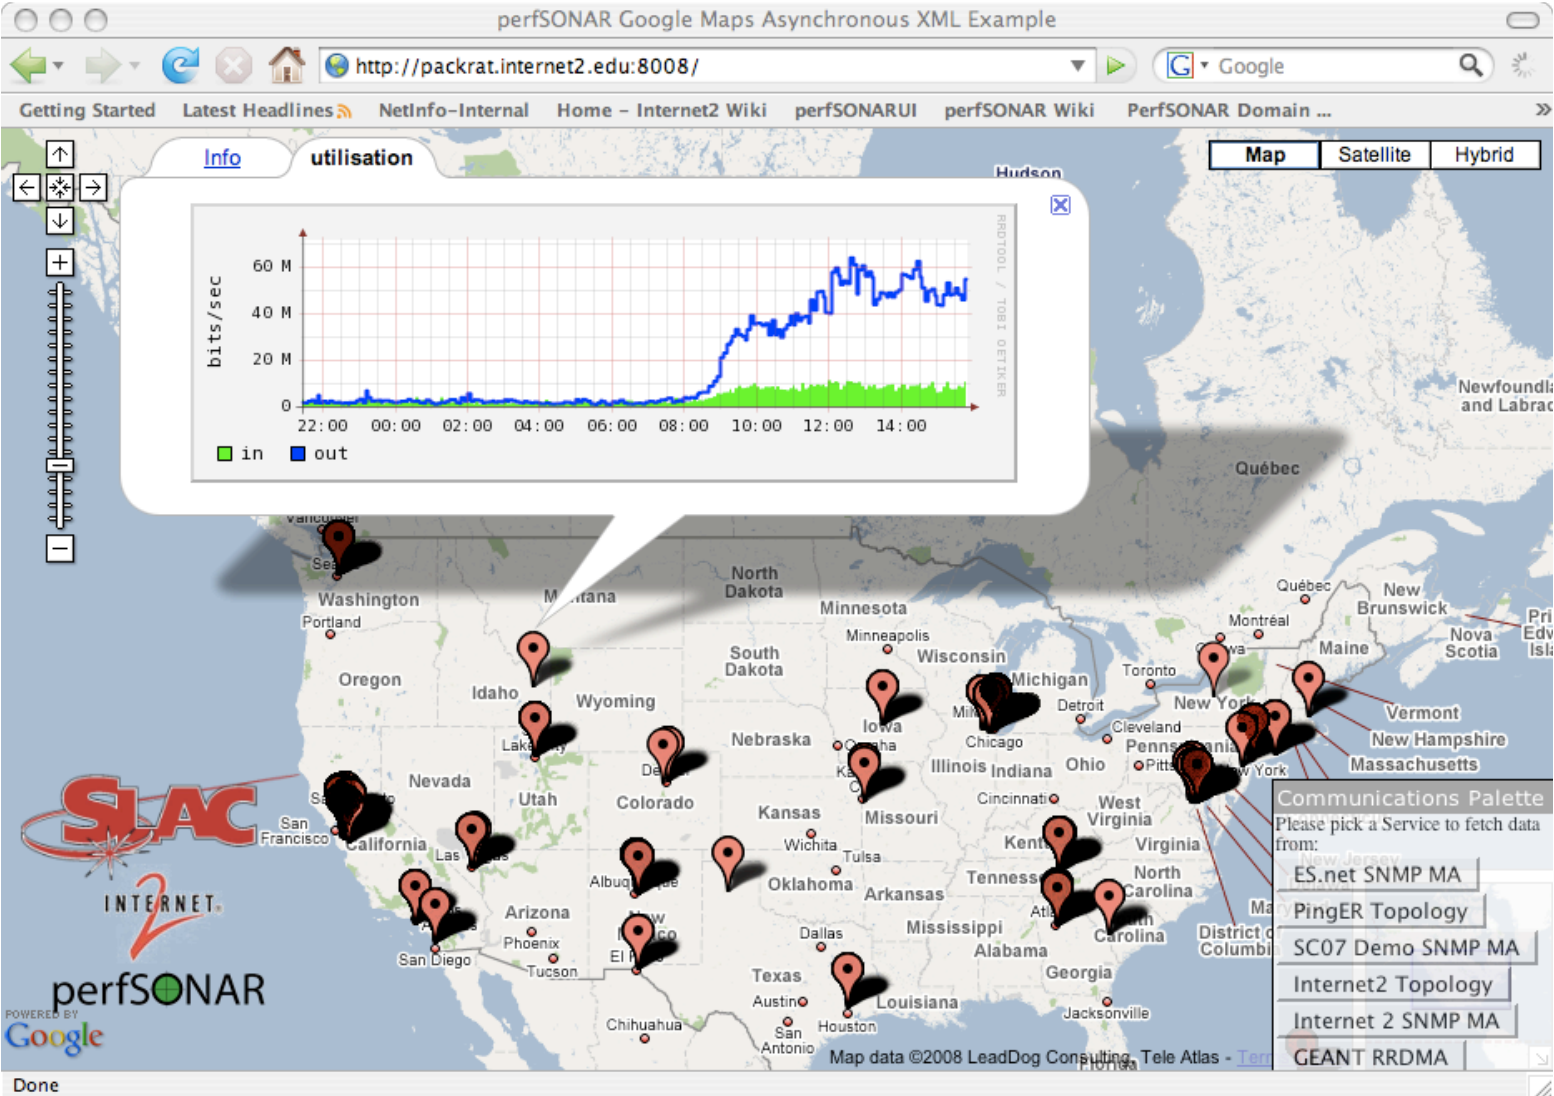


### 12 Hour Graph



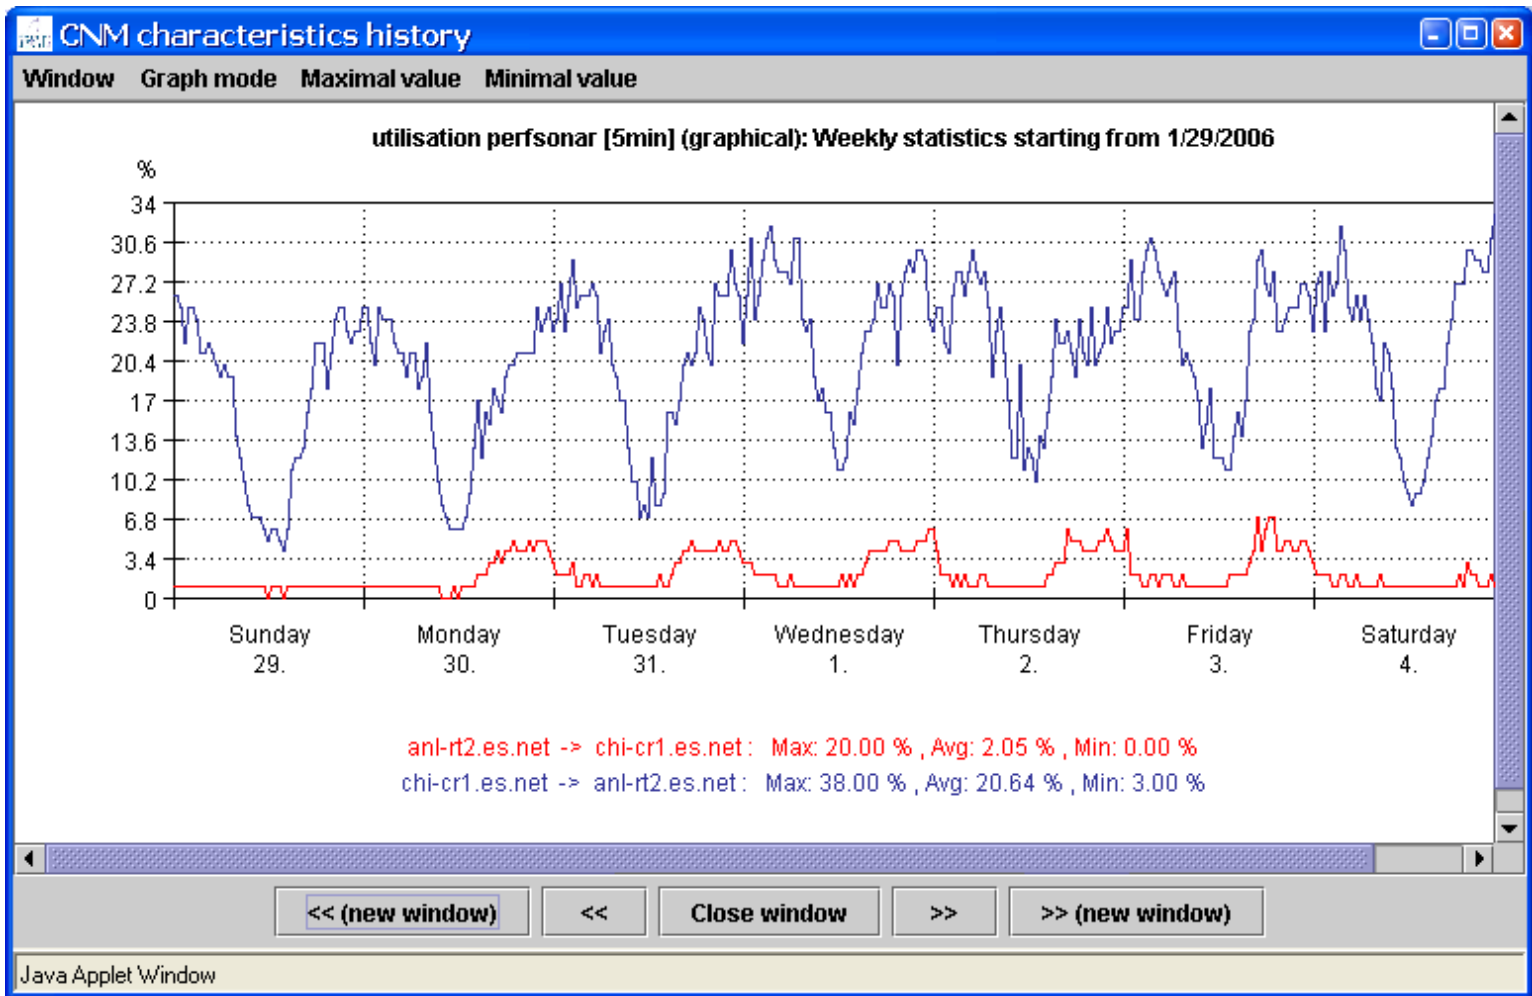
Done

# Gmaps from SLAC





# CNM from DFN



# perfSONARUI from acad.bg

PerfsonarUI-v0.11a

File Interfaces Circuits BWCTL Playground Help

Interfaces Circuits BWCTL Playground

perfsONAR

**Query & search options**

Options  
Service addresses

**Execute query**

Populate  
Clear  
Retrieve all  
Refresh

**Interface details**

GEANT2  
 Hostname rt1.cop.dk.ge...  
 IP address 62.40.124.45  
 Interface name so-4/0/0  
 Interface desc... Link to NORD...  
 Capacity 9953280000  
 Resolution N/A  
 Key localhost.local...

Search by interface IP

File  
Paste  
Search  
Traceroute & search

Summary

Select interface

No.	hostName	ifAddress	ifName	ifDescription	Cap...	Inbo...	Out...	Inbo...	Out...	Inpu...	Out...	Dom...	Status
32	rt1.bud.hu.geant2.net	62.40.124.105	so-0/2/0	Link to RoEduNet 1 (Pr:T-Systems...	62...	94...	N/A	58...	N/A	N/A	N/A	GE...	OK
31	rt1.bud.hu.geant2.net	62.40.124.245	so-0/3/0	Link to RoEduNet 2 (Pr:T-Systems...	62...	94...	38...	58...	24...	N/A	N/A	GE...	OK
34	rt1.ath.gr.geant2.net	62.40.124.89	ge-7/0/0	Link to GRnet (Pr:GRnet Id:in-house)	10...	33...	12...	3,3...	1...	N/A	N/A	GE...	OK
16	rt1.lic.nt.geant2.net	62.40.124.185	so-7/0/0	Link to FCCN (Pr:FCCN Id:local Fib...	2,4...	22...	21...	55...	53...	N/A	N/A	GE...	OK

Summary for all interfaces ( Last 15 min )

Utilisation, Mb/s

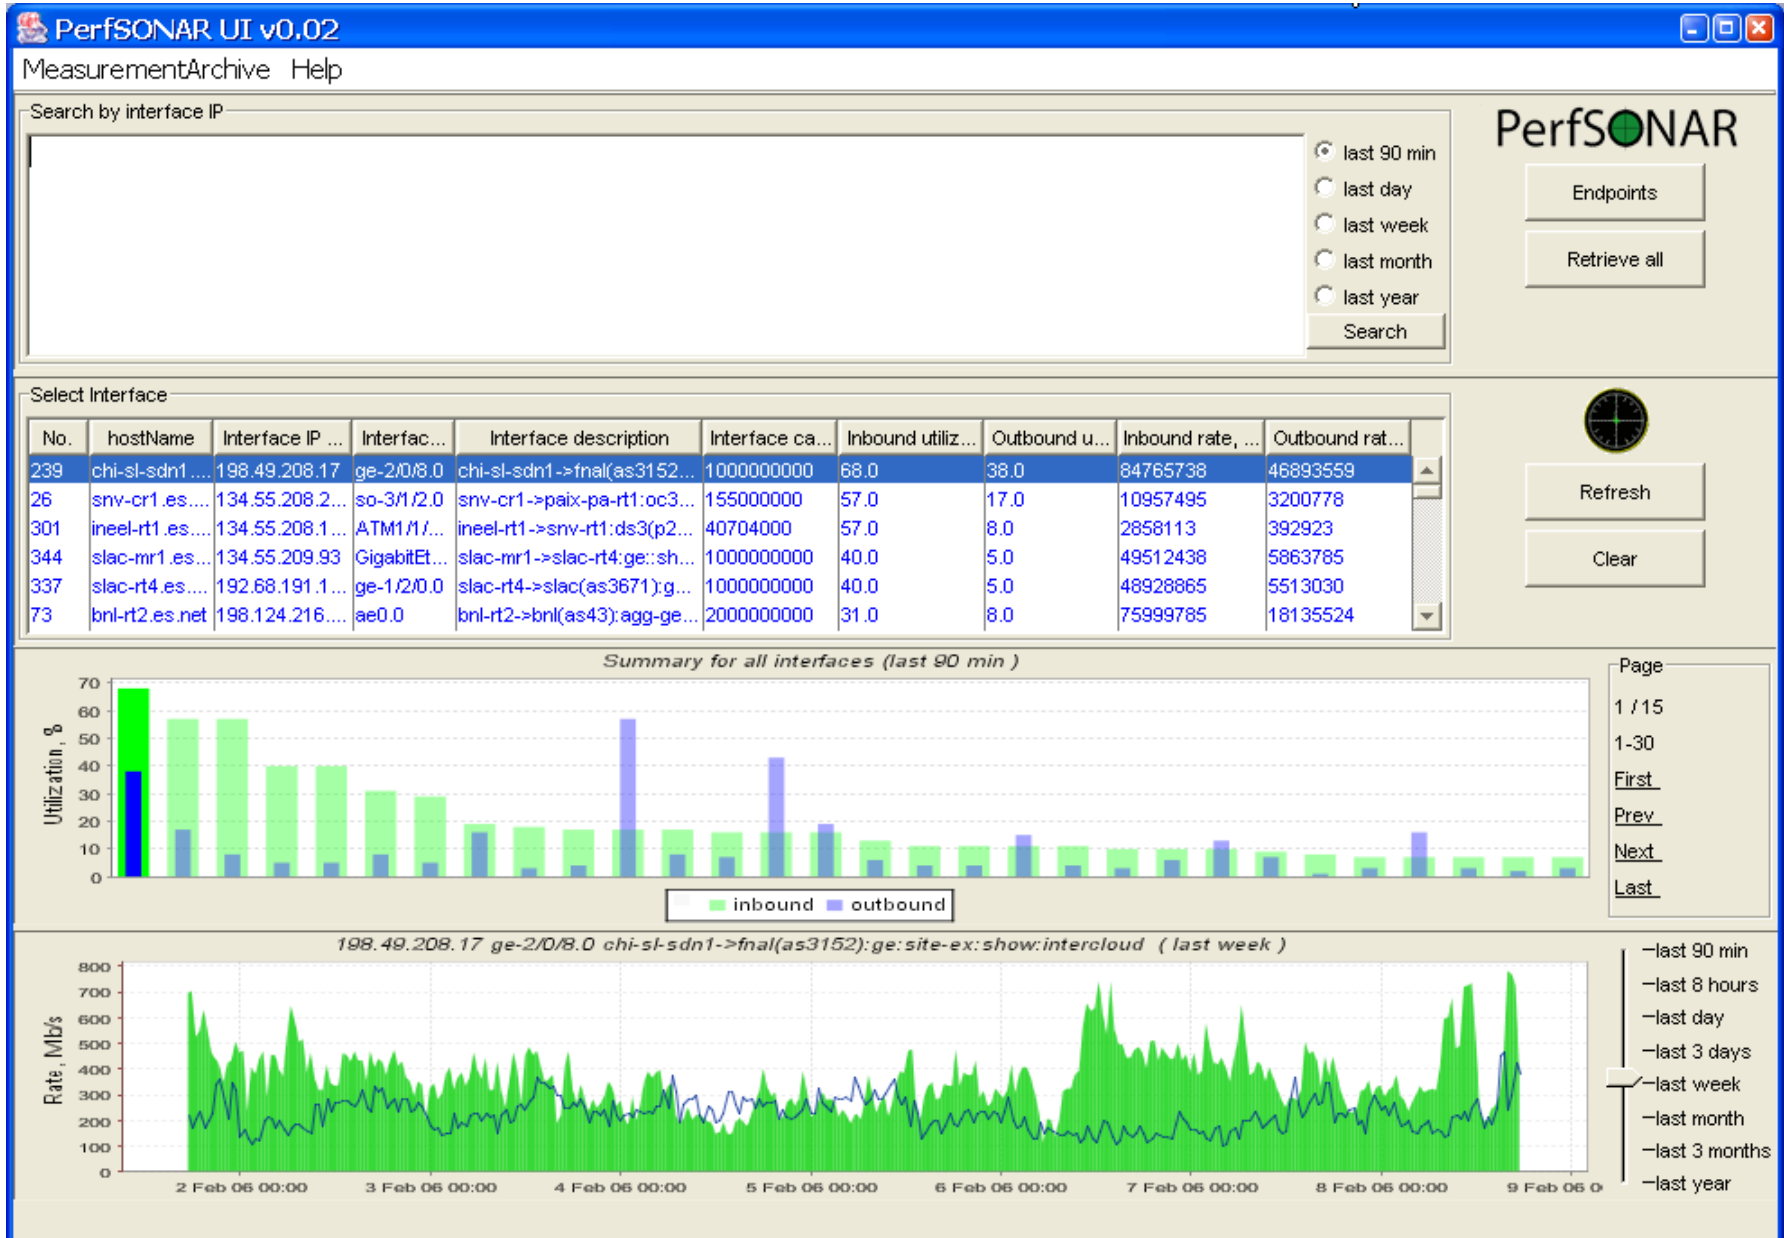
Input errors/Output drops

Page 1 / 2  
1-30  
First  
Prev  
Next  
Last

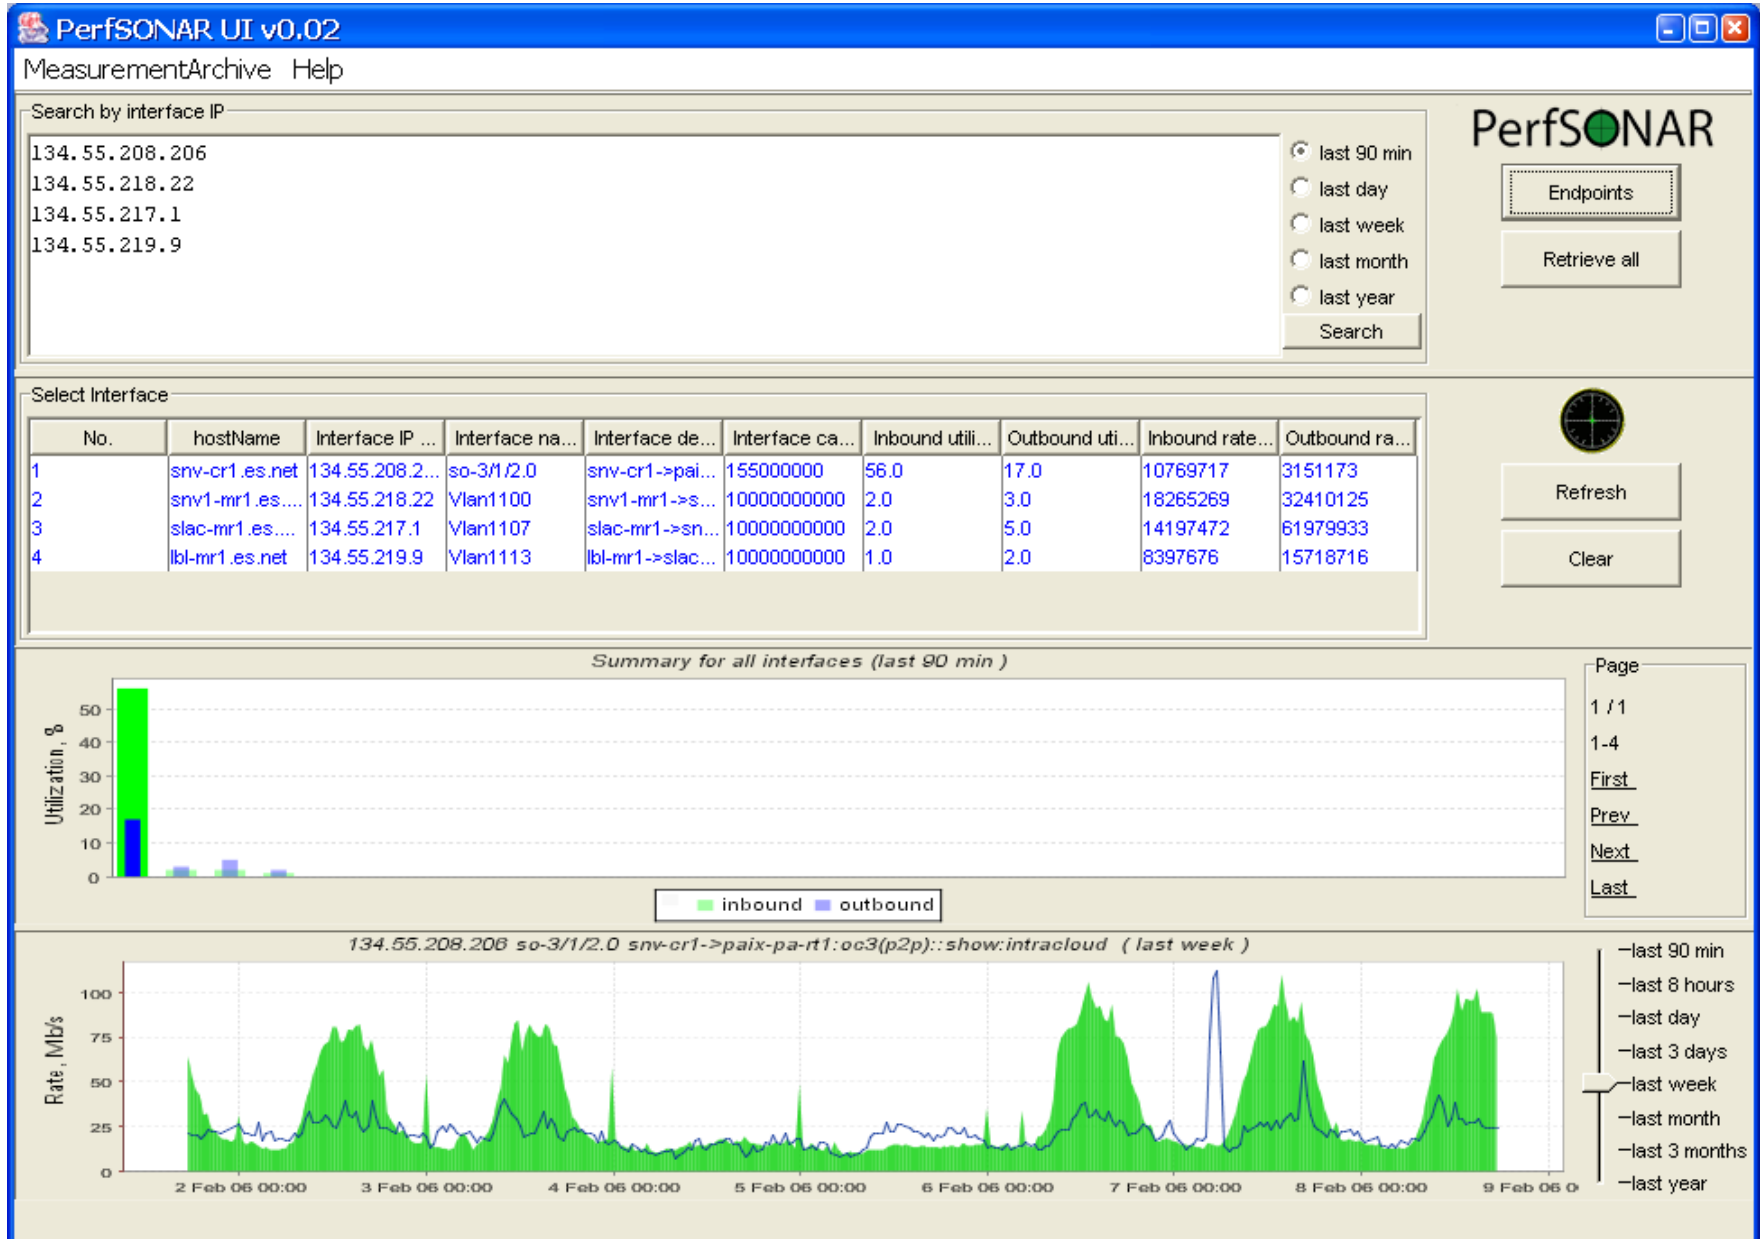
Time  
Last 15 min  
Last 90 min  
Last 8 hours  
Last 24 hours  
Last 3 days  
Last week  
Last month

Completed.

# PerfsonarUI 1



# PerfsonarUI 2



# PerfsonarUI 3

PerfSONAR UI v0.02

MeasurementArchive Help

Search by interface IP

134.55.208.206  
134.55.218.22  
134.55.217.1  
134.55.219.9

last 90 min  
 last day  
 last week  
 last month  
 last year

PerfSONAR

Select Interface

No.	hostName	Interface IP ...	Interface na...	Interface de...	Interface ca...	Inbound utili...	Outbound uti...	Inbound rate...	Outbound ra...
1	snv-cr1.es.net	134.55.208.2...	so-3/1/2.0	snv-cr1->pai...	1550000000	56.0	17.0	10769717	3151173
2	snv1-mr1.es....	134.55.218.22	Vlan1100	snv1-mr1->s...	100000000000	2.0	3.0	18265269	32410125
3	slac-mr1.es....	134.55.217.1	Vlan1107	slac-mr1->sn...	100000000000	2.0	5.0	14197472	61979933
4	lbl-mr1.es.net	134.55.219.9	Vlan1113	lbl-mr1->slac...	100000000000	1.0	2.0	8397676	15718716

Summary for all interfaces (last 90 min)

Utilization, %

Rate, Mb/s

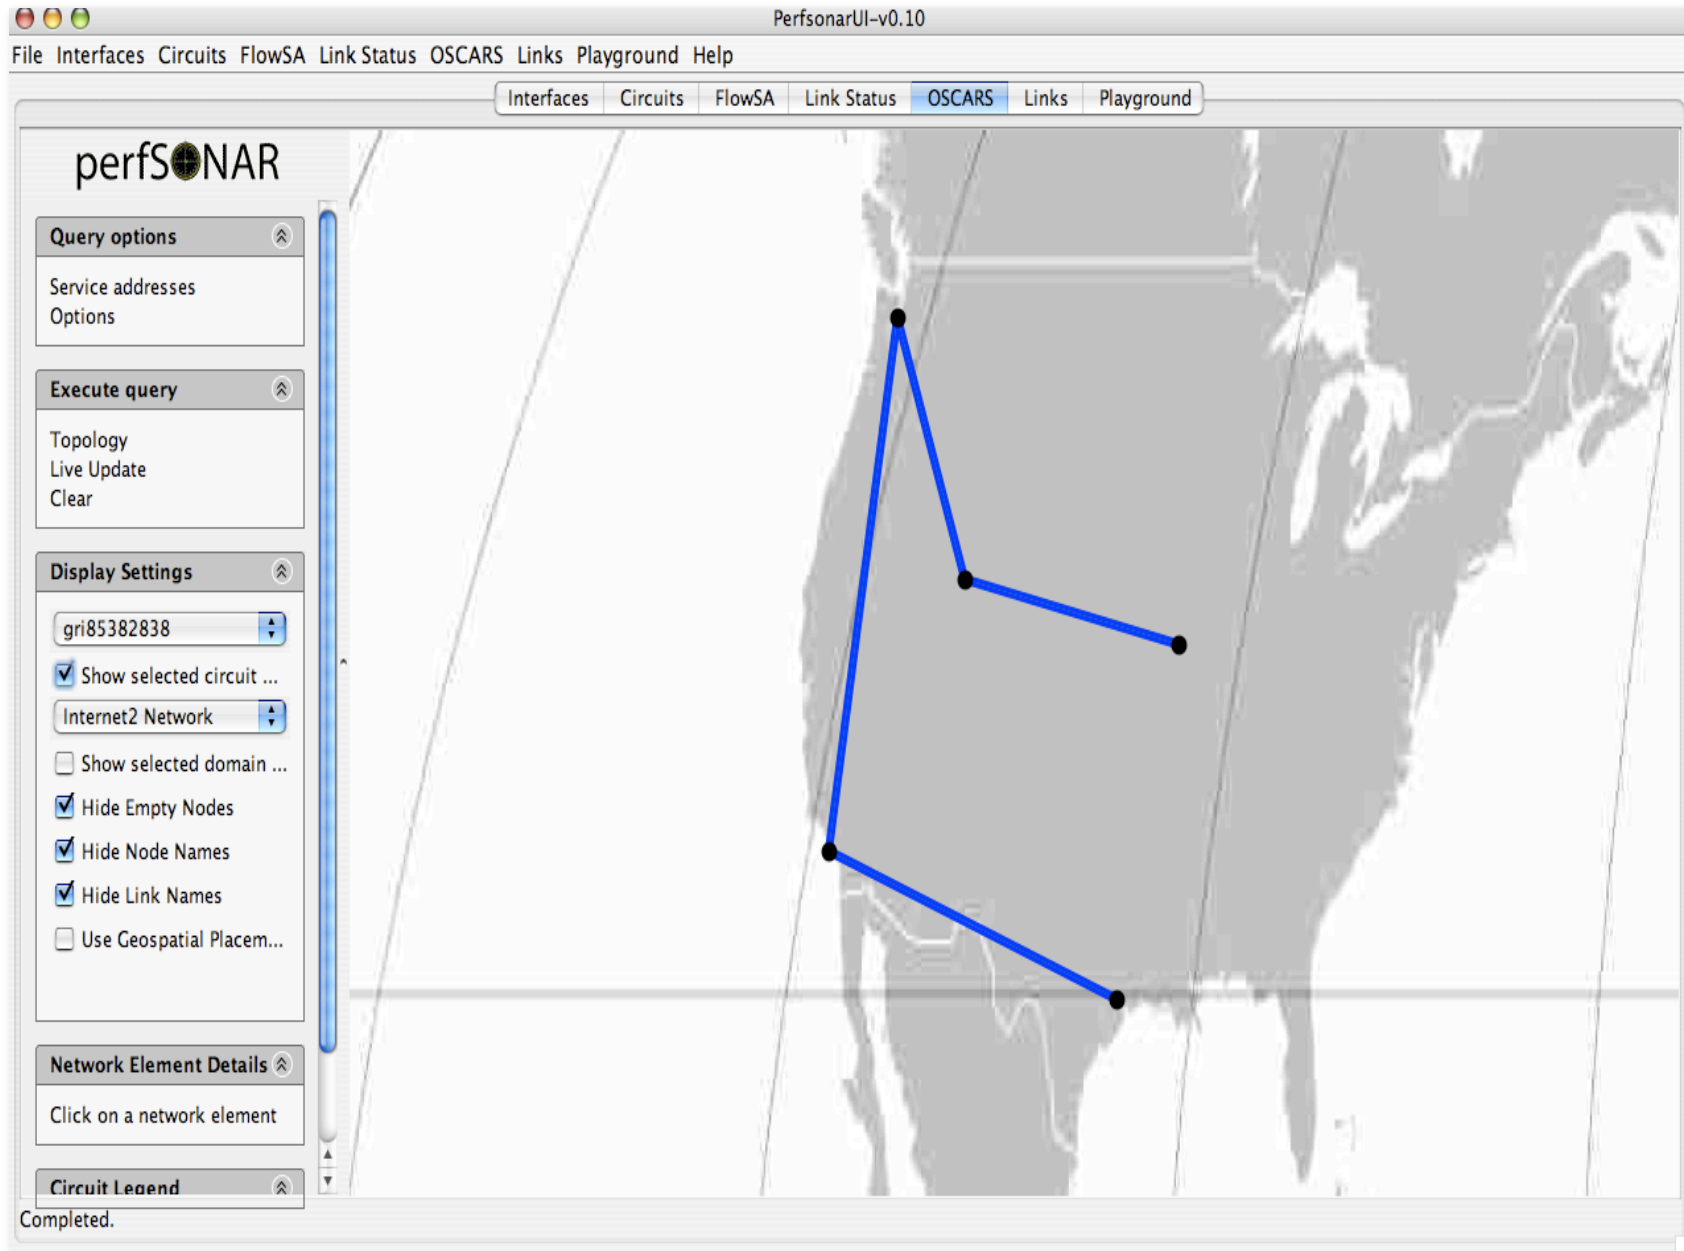
134.55.208.206 so-3/1/2.0 snv-cr1->paix-pa-rt1:oo3

Measurement archives

Select MA(s)

Use	Measurement Archive	URL
<input type="checkbox"/>	GEANT	<a href="http://mu.dante.org.uk:8090/axis/s...">http://mu.dante.org.uk:8090/axis/s...</a>
<input type="checkbox"/>	ISTF	<a href="http://selena.acad.bg:8090">http://selena.acad.bg:8090</a>
<input type="checkbox"/>	SWITCH	<a href="http://archive.sonar.net.switch.ch...">http://archive.sonar.net.switch.ch...</a>
<input type="checkbox"/>	UNINETT	<a href="http://mi6.uninett.no:8090">http://mi6.uninett.no:8090</a>
<input checked="" type="checkbox"/>	ESNET	<a href="http://mea1.es.net:8080/axis/servi...">http://mea1.es.net:8080/axis/servi...</a>
<input type="checkbox"/>	Internet2	<a href="http://thunderbird.internet2.edu.80...">http://thunderbird.internet2.edu.80...</a>

# Oscars Circuit plugin - Internet2



PerfsonarUI-v0.10

File Interfaces Circuits FlowSA Link Status OSCARS Links Playground Help

Interfaces Circuits FlowSA Link Status OSCARS Links Playground

perfSONAR

Query options

Service addresses  
Options

Execute query

Topology  
Live Update  
Clear

Display Settings

gri85382838

Show selected circuit ...

Internet2 Network

Show selected domain ...

Hide Empty Nodes

Hide Node Names

Hide Link Names

Use Geospatial Placem...

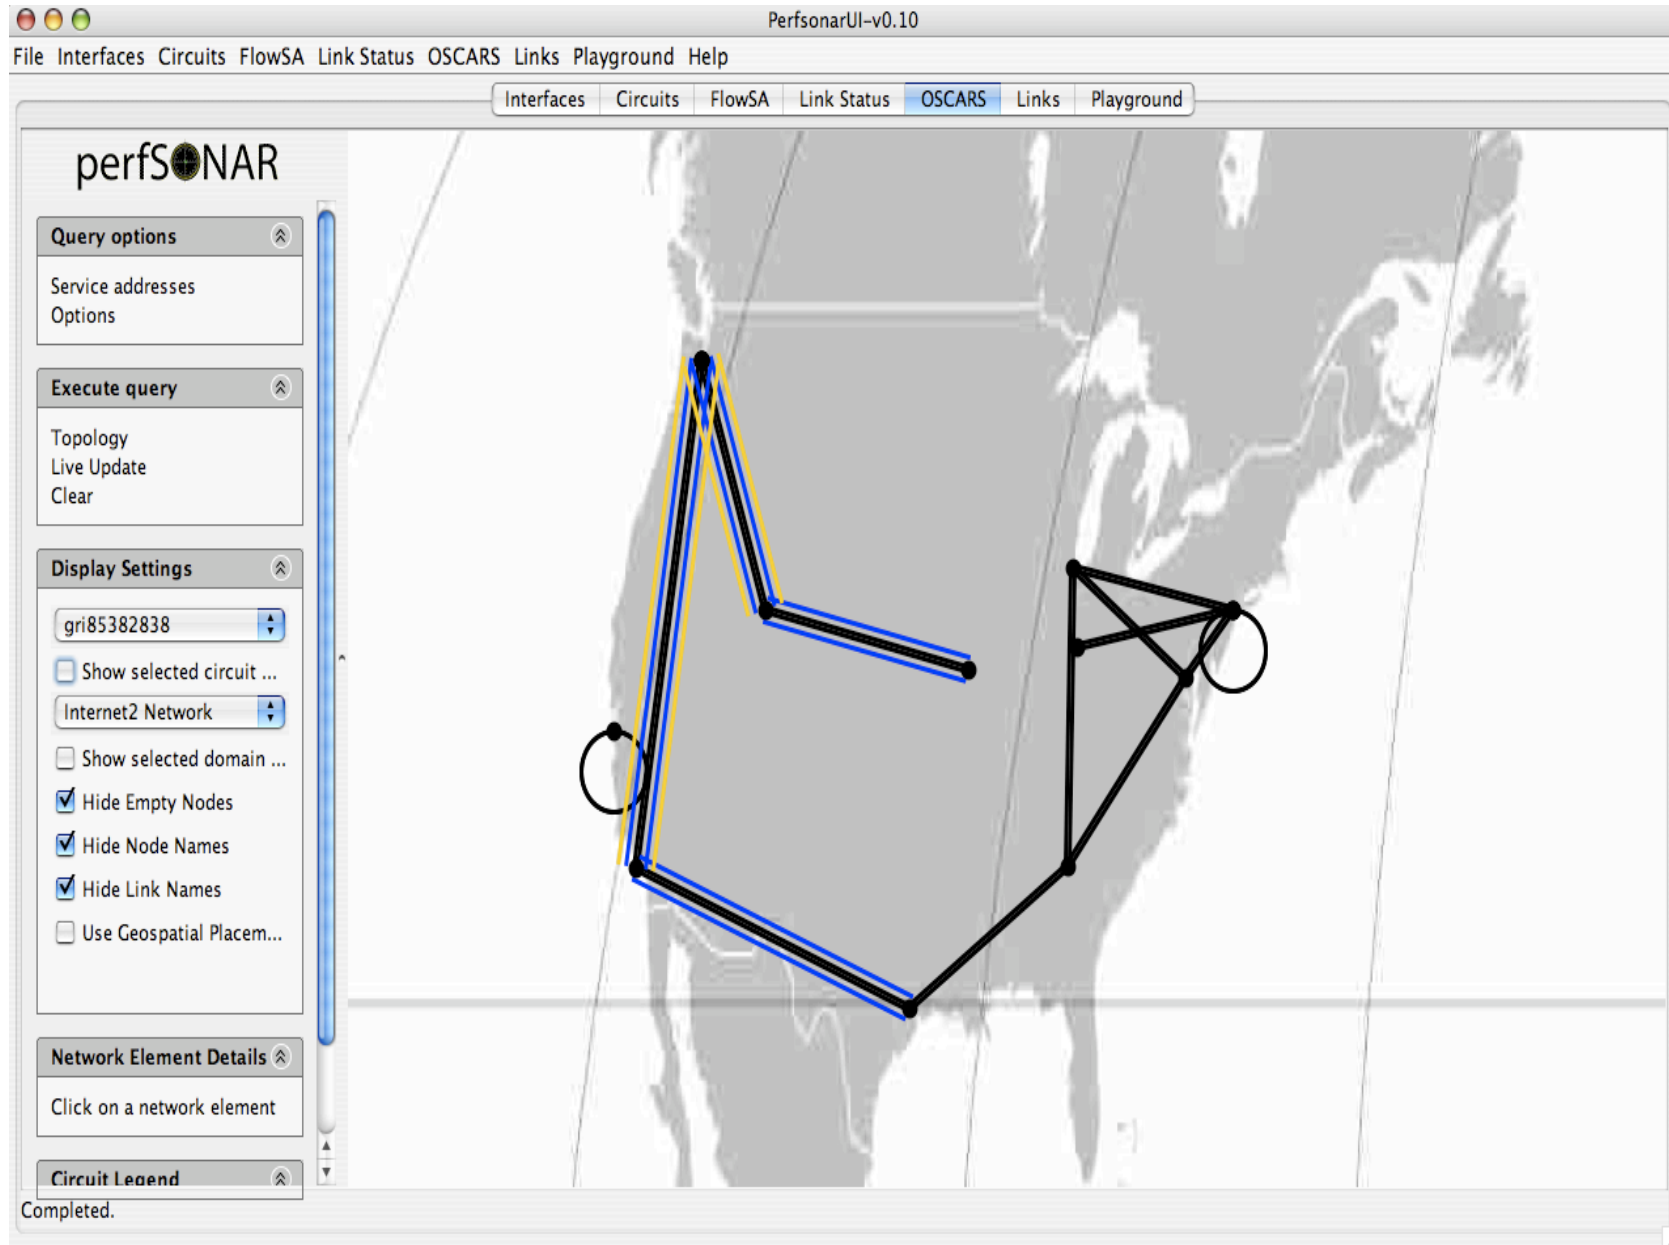
Network Element Details

Click on a network element

Circuit Legend

Completed.

# Oscars circuit plugin



PerfsonarUI-v0.10

File Interfaces Circuits FlowSA Link Status OSCARS Links Playground Help

Interfaces Circuits FlowSA Link Status **OSCARS** Links Playground

**perfSONAR**

**Query options** ⤴

Service addresses  
Options

**Execute query** ⤴

Topology  
Live Update  
Clear

**Display Settings** ⤴

gri85382838

Show selected circuit ...

Internet2 Network

Show selected domain ...

Hide Empty Nodes

Hide Node Names

Hide Link Names

Use Geospatial Placem...

**Network Element Details** ⤴

Click on a network element

**Circuit Legend** ⤴

Completed.

# E2Emon - Monitoring Circuits

E2E Link Monitoring System

http://stats.geant2.net/e2emon/mon/G2\_E2E\_index\_PROD.html

Getting Started Latest Headlines NetInfo-Internal Home - Internet2 Wiki perfSONARUI perfSONAR Wiki PerfSONAR Domain ... USLHCNet DOE Revie... Program for Winter 2...

Prod. Non Prod. All Links

**E2E Links for Project LHCOPN (Prod.)**

E2E Link ID	State Oper	State Admin	Additional Info
<a href="#">CERN-ASGC-LHCOPN-001</a>	Up	Normal Oper.	<b>Error:</b> E2E Link is not contiguous (End Point missing or gap found) <b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-ASGC-LHCOPN-002</a>	Up	Normal Oper.	<b>Error:</b> E2E Link is not contiguous (End Point missing or gap found) <b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-BNL-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-BNL-LHCOPN-002</a>	Up	Normal Oper.	<b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-CNAF-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-FERMI-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-FERMI-LHCOPN-002</a>	Up	Normal Oper.	<b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-GRIDKA-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-IN2P3-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-NDGF-LHCOPN-001</a>	Up	Normal Oper.	<b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-PIC-LHCOPN-001</a>	Up	Normal Oper.	<b>Error:</b> E2E Link is not contiguous (End Point missing or gap found) <b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-RAL-LHCOPN-001</a>	Up	Normal Oper.	<b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links
<a href="#">CERN-SARA-LHCOPN-001</a>	Up	Normal Oper.	
<a href="#">CERN-TRIUMF-LHCOPN-001</a>	Up	Normal Oper.	<b>Error:</b> E2E Link is not contiguous (End Point missing or gap found) <b>Warning:</b> Operational state is not known for all involved links <b>Warning:</b> Administrative state is not known for all involved links

Start page

**E2ECU view**

[E2E Links](#)  
[Mon. Links](#)  
[Problem E2E Links](#)  
[Problem Mon. Links](#)

**Domain view**

[CANARIE](#)  
[CERN](#)  
[CESNET \(?\)](#)  
[DFN](#)  
[ESNET](#)  
[FERMI](#)  
[GARR](#)  
[GEANT2](#)  
[HOPI](#)  
[IN2P3](#)  
[INTERNET2](#)  
[NETHERLIGHT](#)  
[PSNC](#)  
[REDIRIS](#)  
[RENATER](#)  
[SWITCH](#)  
[USLHCNET](#)

**Project view**

Done

# E2Emon: Status of E2E link CERN-LHCOPN-FNAL-001

Oper. State: **Up**

Admin. State: **Normal Oper.**

Domain	CERN			USLHCNET			
Link Structure	<b>EP</b>	←.....	.....→	<b>DP</b>	↔	<b>DP</b>	←.....
Type	EndPoint	ID Part.Info	ID Part.Info	Demarc	Domain Link	Demarc	ID Part.Info
Local Name	CERN-T0	S513-C-BE1	CERN-FERMI-LHCOPN-001-GVA-CERN	USLHCNET-GEN	CERN-FERMI-LHCOPN-001-GVA-CHI	USLHCNET-CHI	CERN-FERMI-LHCOPN-001-CHI-ESNET
State Oper.	-	<b>Up</b>	<b>Up</b>	-	<b>Up</b>	-	<b>Up</b>
State Admin.	-	<b>Normal Oper.</b>	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>
Timestamp	-	2007-04-08 T05:04:08+02:00	2007-04-08 T05:04:11+02:00	-	2007-04-08 T05:04:53+02:00	-	2007-04-08 T05:03:59+02:00

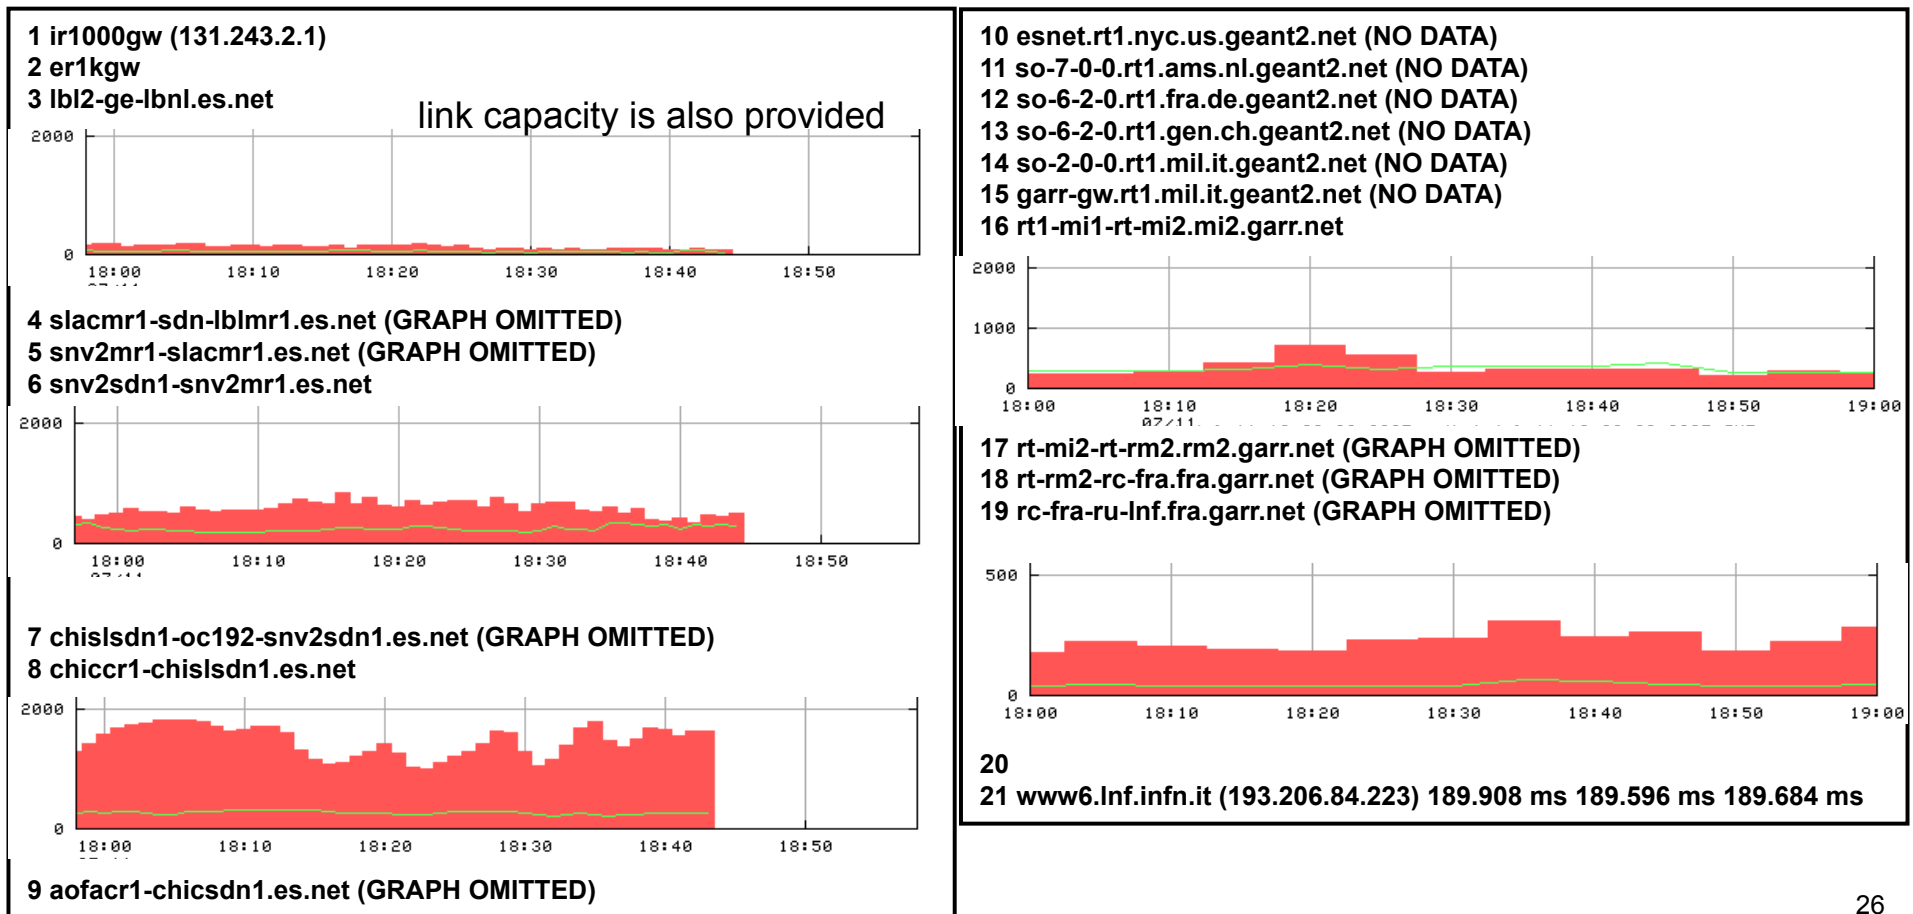
Page 1 of 1

ESNET				FERMI				
.....→	<b>DP</b>	↔	<b>DP</b>	←.....	.....→	<b>DP</b>	↔	<b>EP</b>
ID Part.Info	Demarc	Domain Link	Demarc	ID Part.Info	ID Part.Info	Demarc	Domain Link	EndPoint
CERN-FERMI-LHCOPN-001-STARLIGHT-Tail	<b>ESNET-STARLIGHT</b>	CERN-FERMI-LHCOPN-001-FERMI-STARLIGHT	<b>ESNET-FERMI</b>	CERN-FERMI-LHCOPN-001-Site-Tail	md8	<b>FERMI-ESNET</b>	md2	<b>FERMI-T1</b>
<b>Up</b>	-	<b>Up</b>	-	<b>Up</b>	<b>Up</b>	-	<b>Up</b>	-
<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	<b>Normal Oper.</b>	-	<b>Normal Oper.</b>	-
2007-04-08 T01:40:37.0	-	2007-04-08T01:40:37.0	-	2007-04-08 T01:40:37.0	2007-04-08 T01:40:01.0-6:00	-	2007-04-08 T01:40:01.0-6:00	-

E2Emon generated view of the data for one OPN link [E2EMON]

# Traceroute Visualizer

- Forward direction bandwidth utilization on application path from LBNL to INFN-Frascati (Italy)
  - traffic shown as bars on those network device interfaces that have an associated MP services (the first 4 graphs are normalized to 2000 Mb/s, the last to 500 Mb/s)



## **In Summary, perfSONAR is:**

---

- A global collaboration
- An architecture and a set of protocols
- Several interoperable software implementations
- A measurement infrastructure that is deployed in over a dozen different networks