



ESnet Network Measurements ESCC Feb 15 2007

Joe Metzger
metzger@es.net



- Users dependence on the network is increasing
 - Distributed Applications
 - Moving Larger Data Sets
 - The network is becoming a critical part of large science experiments
- The network is growing more complex
 - 6 core devices in 05', 25+ in 08'
 - 6 core links in 05', 40+ in 08', 80+ by 2010?
- Users continue to report performance problems
 - 'wizards gap' issues
- The community needs to better understand the network
 - We need to be able to demonstrate that the network is good.
 - We need to be able to detect and fix subtle network problems.

- perfSONAR is a global collaboration to design, implement and deploy a network measurement framework.
 - Web Services based Framework
 - Measurement Archives (MA)
 - Measurement Points (MP)
 - Lookup Service (LS)
 - Topology Service (TS)
 - Authentication Service (AS)
 - Some of the currently Deployed Services
 - Utilization MA
 - Circuit Status MA & MP
 - Latency MA & MP
 - Bandwidth MA & MP
 - Looking Glass MP
 - Topology MA
 - This is an **Active** Collaboration
 - The basic framework is complete
 - Protocols are being documented
 - New Services are being developed and deployed.

perfSONAR Collaborators

- ARNES
- Belnet
- CARnet
- CESnet
- Dante
- University of Delaware
- DFN
- ESnet
- FCCN
- FNAL
- GARR
- GEANT2
- Georga Tech
- GRNET
- Internet2
- IST
- POZNAN Supercomputing Center
- Red IRIS
- Renater
- RNP
- SLAC
- SURFnet
- SWITCH
- Uninett

* Plus others who are contributing, but haven't added their names to the list on the WIKI.

perfSONAR Deployments

16+ different networks have deployed at least 1 perfSONAR service (Jan 2007)



- ESnet Deployed Services
 - Link Utilization Measurement Archive
 - Virtual Circuit Status
- In Development
 - Active Latency and Bandwidth Tests
 - Topology Service
 - Additional Visualization capabilities
- perfSONAR visualization tools showing ESnet data
 - Link Utilization
 - perfSONARUI
 - <http://perfsonar.acad.bg/>
 - VisualPerfSONAR
 - https://noc-mon.srce.hr/visual_perf
 - Traceroute Visualizer
 - <https://performance.es.net/cgi-bin/level0/perfsonar-trace.cgi>
 - Virtual Circuit Status
 - E2EMon (for LHCOPN Circuits)
 - http://cnmdev.lrz-muenchen.de/e2e/lhc/G2_E2E_index.html

- Central Monitoring Software
 - Uses perfSONAR protocols to retrieve current circuit status every minute or so from MAs and MPs in all the different domains supporting the circuits.
 - Provides a web site showing current end-to-end circuit status
 - Generates SNMP traps that can be sent to other management systems when circuits go down
- MA & MP Software
 - Manages the perfSONAR communications with the central monitoring software
 - Requires an XML file describing current circuit status as input.
- Domain Specific Component
 - Generates the XML input file for the MA or MP
 - Multiple development efforts in progress, but no universal solutions
 - CERN developed one that interfaces to their abstraction of the Spectrum NMS DB
 - DANTE developed one that interfaces with the Acatel NMS
 - ESnet developed one that uses SNMP to directly poll router interfaces
 - FERMI developed one that uses SNMP to directly poll router interfaces
 - Others under development

E2Emon Central Monitoring Software

http://cnmdev.lrz-muenchen.de/e2e/lhc/G2_E2E_index.html

E2E Link Monitoring System - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

[Start page](#)

E2ECU view

[All E2E Links](#)
[Problem Links](#)

Domain view

[CERN](#)
[DFN](#)
[ESNET](#)
[FERMI](#)
[GARR](#)
[GEANT2](#)
[IN2P3](#)
[SURFNET](#)
[SWITCH](#)
[USLHCNET](#)

Project view

[IGTMD](#)
[LHCOPN](#)

Availability Statistics

[Current Month](#)
[Last Month](#)

E2E Links for Project LHCOPN

E2E Link ID	State Oper	State Admin	Additional Info
CERN-GRIDKA-LHCOPN-001	Up	Normal Oper.	Error: E2E Link is not contiguous (End Point missing or gap found)
CERN-FERMI-LHCOPN-002	Up	Normal Oper.	Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
CERN-CNAF-LHCOPN-001	Down	Maintenance	
CERN-RAL-LHCOPN-001	Up	Normal Oper.	Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
CERN-FERMI-LHCOPN-001	Up	Normal Oper.	
CERN-BNL-LHCOPN-002	Up	Normal Oper.	Error: E2E Link is not contiguous (End Point missing or gap found) Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
GRIDKA-SARA-LHCOPN-001	Up	Normal Oper.	Error: E2E Link is not contiguous (End Point missing or gap found) Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
CERN-NDGF-LHCOPN-001	Up	Normal Oper.	Error: E2E Link is not contiguous (End Point missing or gap found) Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
CERN-IN2P3-LHCOPN-001	Up	Normal Oper.	Error: E2E Link is not contiguous (End Point missing or gap found) Warning: Operational state is known not for all involved links Warning: Administrative state is known not for all involved links
CERN-BNL-LHCOPN-001	Up	Normal Oper.	
CNAF-GRIDKA-LHCOPN-001	Up	Normal Oper.	
CERN-SARA-LHCOPN-001	Up	Normal Oper.	

Page generated at 2007-02-09, 17:25:06 MET

[Start page](#)

E2ECU view

[All E2E Links](#)
[Problem Links](#)

Domain view

[CERN](#)
[DFN](#)
[ESNET](#)
[FERMI](#)
[GARR](#)
[GEANT2](#)
[IN2P3](#)
[SURFNET](#)
[SWITCH](#)
[USLHCNET](#)

Project view

Status of E2E Link CERN-FERMI-LHCOPN-001

Oper. State: **Up**
Admin. State: **Normal Oper.**

Domain	CERN				USLHCNET				ESNET				FERMI			
Link Structure	EP	←.....→→	DP	←.....→	DP	←.....→→	DP	←.....→	DP	←.....→→	DP	←.....→	EP
Type	EndPoint	ID Part.Info	ID Part.Info	Demarc	Domain Link	Demarc	ID Part.Info	ID Part.Info	Demarc	Domain Link	Demarc	ID Part.Info	ID Part.Info	Demarc	Domain Link	EndPoint
Local Name	CERN-TO	S513-C-BE1	CERN-FERMI-LHCOPN-001-GVA-CERN	USLHCNET-GEN	CERN-FERMI-LHCOPN-001-GVA-CHI	USLHCNET-CHI	CERN-FERMI-LHCOPN-001-CHIESNET	CERN-FERMI-LHCOPN-001-STARLIGHT-Tail	ESNET-STARLIGHT	CERN-FERMI-LHCOPN-001-FERMI-STARLIGHT	ESNET-FERMI	CERN-FERMI-LHCOPN-001-Site-Tail	md8	FERMI-ESNET	md2	FERMI-T1
State Oper.	-	Up	Up	-	Up	-	Up	Up	-	Up	-	Up	Up	-	Up	-
State Admin.	-	Normal Oper.	Normal Oper.	-	Normal Oper.	-	Normal Oper.	Normal Oper.	-	Normal Oper.	-	Normal Oper.	Normal Oper.	-	Normal Oper.	-
Timestamp	-	2007-01-26 T13:15:22+01:00	2007-02-06 T17:31:23+01:00	-	2007-02-09 T17:20:33+01:00	-	2007-01-26 T13:15:19+01:00	2007-02-09 T16:20:36.0	-	2007-02-09T16:20:36.0	-	2007-02-09 T16:20:36.0	2007-02-09 T16:20:00.0-6:00	-	2007-02-09 T16:20:00.0-6:00	-

- Latency

- 1U Server with one of:

- EndRun Praecis CT CDMA Clock
 - Meinberg TCR167PCI IRIG Clock
 - Symmetricom bc637PCI-U IRIG Clock

- Bandwidth

- 4U dual Opteron server with one of:

- Myricom 10GE NIC
 - 9.9 Gbps UDP streams
 - ~6 Gbps TCP streams
 - Consumes 100% of 1 CPU
 - Chelsio S320 10GE NIC
 - Should do 10G TCP & UDP with low CPU Utilization
 - Has interesting shaping possibilities
 - Still under testing...

- **SNMP Interface Utilization**
 - Collected every minute
 - For MRTG & Monthly Reporting
- **Circuit Availability**
 - Currently based on SNMP Interface up/down status
 - Limited to LHCOPN and Service Trial circuits for now
- **NetFlow Data**
 - Sampled on our boundaries
- **Latency**
 - OWAMP

- Web Interface to run Network Measurements
- Available to ESnet sites
- Supported Tests
 - Ping
 - Traceroute
 - IPERF
 - Pathload, Pathrate, Pipechar
 - (Only on GE systems)
- Test Hardware
 - GE testers in Qwest hubs
 - TCP iperf tests max at ~600 Mbps.
 - 10GE testers are being deployed in ESnet4 hubs
 - Deployed in locations where we have Cisco 6509 10GE Interfaces
 - Available via Performance Center when not being used for other tests
 - TCP iperf tests max at 6 Gbps.

ESnet Measurement Summary

- Standards / Collaborations
 - PerfSONAR
- LHCOPN
 - Circuit Status Monitoring
- Monitoring Hardware in ESnet 4 Hubs
 - Bandwidth
 - Latency
- Measurements
 - SNMP Interface Counters
 - Circuit Availability
 - Flow Data
 - One Way Delay
 - Achievable Bandwidth
- Visualizations
 - PerfSONARUI
 - VisualPerfSONAR
 - NetInfo