

Collaborative DOE Enterprise Network Monitoring

Yee-Ting Li
SLAC National Accelerator Center

High Level

- Promote the use of network monitoring to aid troubleshooting of network and data issues
- perfSONAR deployed in WAN, now also need information about the LAN for complete end-to-end solution
- Make available such data in an effective, secure and applicable means to the DOE community

Participants

- BNL
- ORNL
- SLAC
- Fermilab
- ESnet

Goals

- Deployment of perfSONAR at the perimeter at End Sites
- Exposure of Site's Internal Network Performance
- Deployment of E-Center to provide AA for measurement data and archiving of network performance data
- Accessible presentation tools for inter-domain network performance
- Application / User level monitoring

Deploy perfSONAR at Site Perimeter

- Aims to achieve
 - promote and advance the evolution of perfSONAR use and field testing
 - publicise existence of easy to use and appropriate tools for network monitoring for network engineers and scientists alike
 - set expectations for users of (network) performance achievable end-to-end - for project planning and scoping
 - provide infrastructure to help identify, diagnose and report performance issues

Site Perimeter

- Required information:
 - Site contact for issues related to bad/anomalous behaviour
- Use perfSONAR to leverage existing work on distributed network monitoring
 - conform to common deployment path
- Site-to-Site
 - Full tests between SLAC-BNL-ORNL
 - Full tests between site perimeter to adjacent (ESnet) PoP

Perimeter Data Collection

- Provide passive and active tests between participants
 - One-way delay (OWAMP), Two way delay (Ping), achievable throughput (iperf)
 - Border router utilisations (snmp)
 - Traceroutes (Traceroute MA)
 - Establish mesh of both active and passive tests
- Provide well published endpoints for on-demand tests (as web portal etc)
 - Reverse traceroute
 - Ping
 - NDT

Site Perimeter Challenges

- Work with site's cyber-security to enable exposure of
 - network tests between participants
 - allow exposure of web services to consumer consumption
 - classification of network performance data for different types of users
 - approval of servers and services (port openings, kernel versions, logging and auditing of servers etc).
- Provide documentation to ease and expand perfSONAR deployment at DOE sites

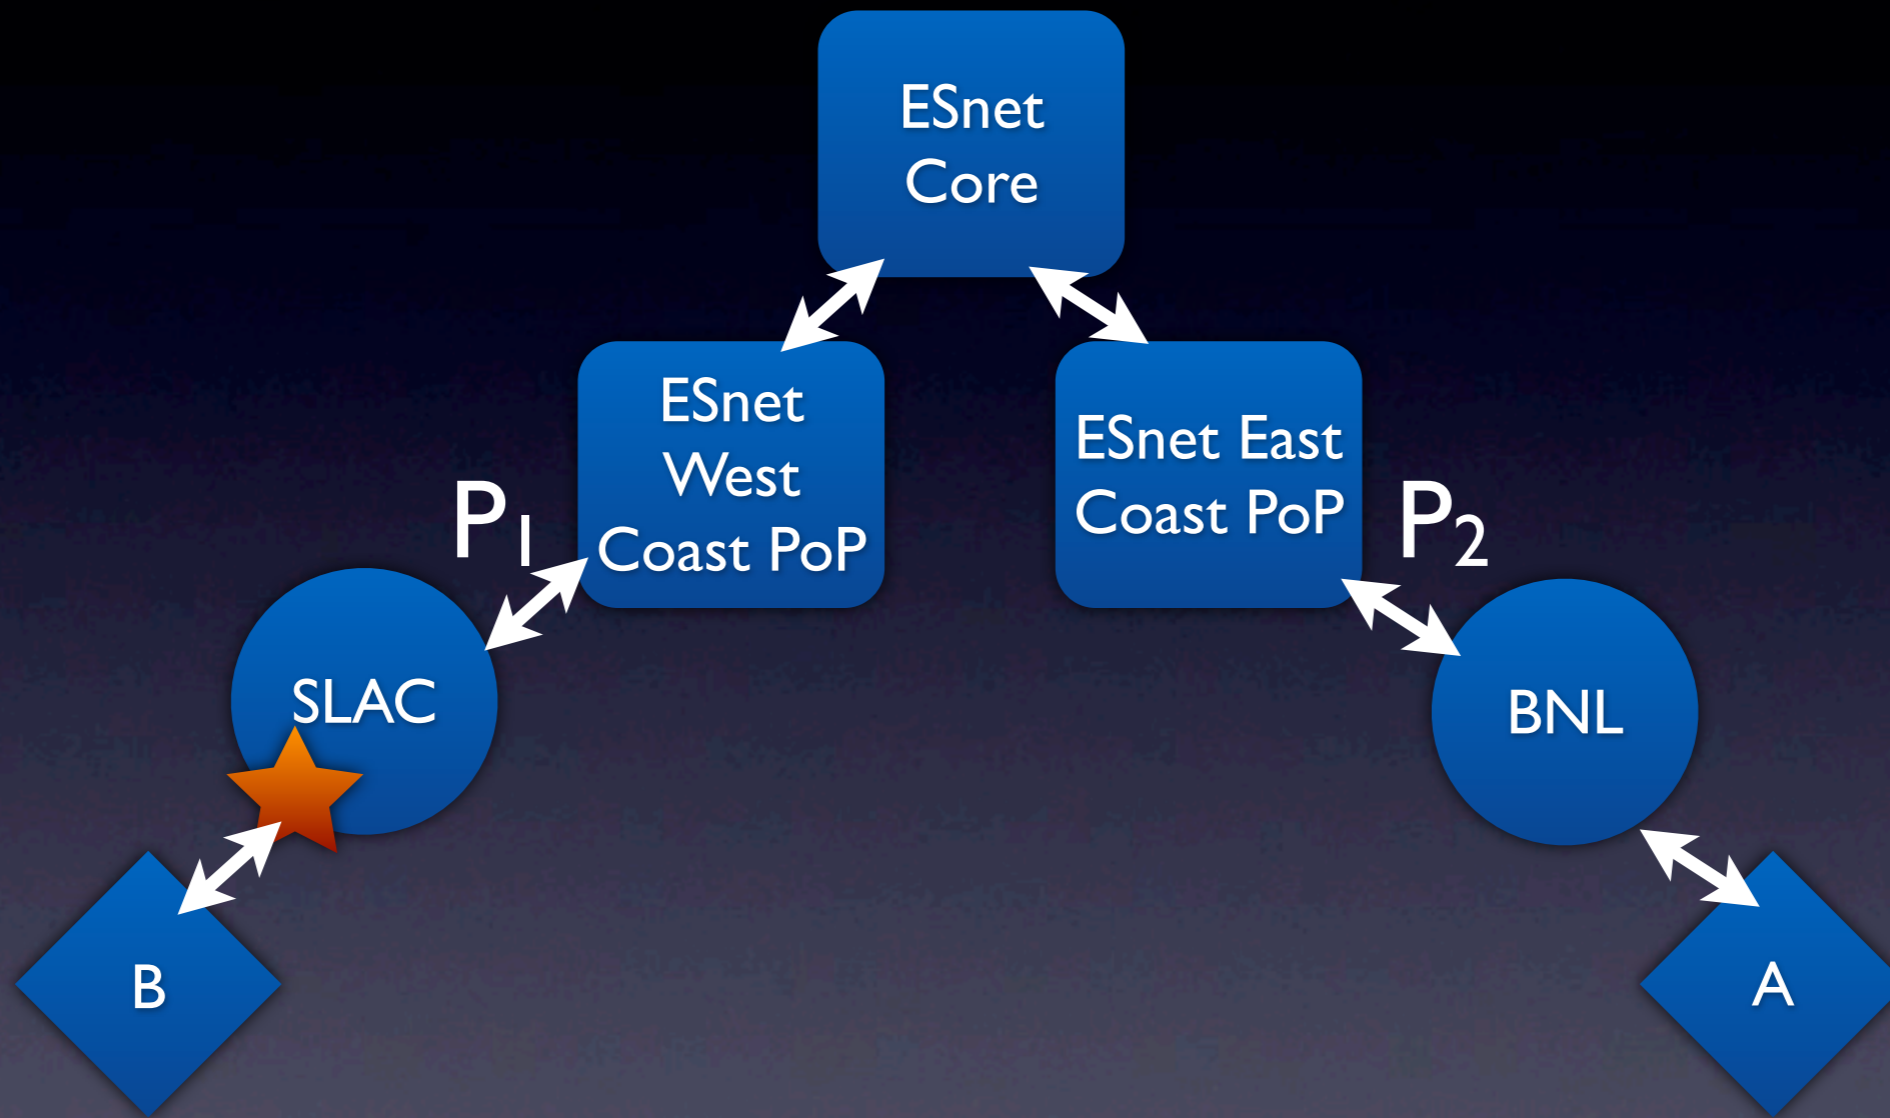
Exposure of Site's Internal LAN

- Aims to achieve
 - establish and provide true end-to-end monitoring and measurement capability
 - high level topology between routers, firewalls, switches and hosts
 - integrate with existing monitoring tools at End Sites to provide this information
 - provide and deploy consistent interface to performance data so developers can focus on data collection and analysis

perfSONAR at the WAN



End-to-End



Given the performance parameters of both bandwidth, packet lost rate, and RTT of the links BP_1 , P_1P_2 , and P_2A , what is the performance of AB ?

Use Cases

- Scientists and network users think of Host-to-Host, not Site-to-Site
 - just because the perimeter performance is okay, does not mean performance end-to-end is okay
- Possible issues
 - Slow router/firewall causes many packet drops
 - Duplex mismatches on old hosts/network devices
 - Large layer 2 spanning trees causes network inefficiencies
 - High CPU or low memory on hosts
 - High IO waits caused by slow disks etc.

Internal Data Collection

- Provide perfSONAR interface for data already being collected by existing tools at Sites
- Router utilisations (SNMP-MA)
- Topology information about interconnections of devices
 - cyber-security issues?
 - security through obscurity
 - sanitisation of topology information
 - how to deal with dynamic changes of host moves etc.?
- Host issues? Disk / CPU?
 - integrate with ganglia, nagios etc?

Internal Data Collection

- Balance between visibility and cost
- Provide read-only access to data already collected
 - switch and router utilisations (SNMP)
 - network topologies (NeDI, CDP/LLDP etc)
 - higher network layers? (vlan groups, subnets etc)
- (Partial) Mesh of Active tests
 - Place perfSONAR MP's or OWAMP/BWCTL responders at strategic locations in LAN
- Integrate existing monitoring, eg Nagios, Smokeping etc.

Enterprise Monitoring

- Integrated, consistent and easy to use UI
 - Dashboards
 - Graphs
- Reports
- (Automated) Alerts and Notifications
 - Trends (deviation)
 - Emails, paging etc
- Typically focused on application/service performance
- With perfSONAR, this becomes a presentation and transformation layer service
 - Data Access Layer becomes perfSONAR client API
 - GUI(s), alerts (pubsub)

Enterprise Monitoring

- Network Centric
 - MRTG
 - Cacti
 - NeDi
 - Netflow
 - collectd
- Host Centric
 - Nagios
 - Groundworks
 - OpenView
 - Spectrum
 - Tivoli



Internal Deployment Challenges

- (Cyber) security
 - identification of classes of data and who should be allowed to view it
 - different sites use different internal tools to collect performance and topology information
 - integrate with site's AA policies and procedures
- Cost of active tests may be prohibitive
 - both for servers, maintenance, and traffic loads
 - need to demonstrate usefulness with specific deployment cases
 - periodic tests and archives required (perfSONAR)

E-Center

- perfSONAR data
 - each domain's data is disjoint from the data being collected in other domains
 - difficulty in finding the data required; access issues to data
- E-Center will
 - provide building blocks for federated network services to enhance distributed computing capabilities of large-scale, global collaborations
 - advance use of perfSONAR
 - provide users with services to enable better understanding of wide-area network performance wrt to distributed computing efforts
 - present in coherent and understandable manner
 - web portals, blogs, wiki's etc.


E-Center

- Aim to:
 - Ensure continual access and appropriate responsiveness of local e-Center gateway services to central e-Center facility
 - Adapt and deploy e-center at each site
 - provide installation and deployment feedback
 - usability and customisations at each site
 - Quantify, resolve, and document site security issues involving site perfSONAR & E-Center
 - Contribute site perspective into the strategic planning of E-Center capabilities & services
 - Integrate local AA support with E-Center data collection & other WAN perfSONAR efforts
 - Use Grid User Management System (GUMS)
 - DoE Users
 - Site and ESnet network administrators
 - Site information (IT) officers
 - DoE program officers

Internal Deployment

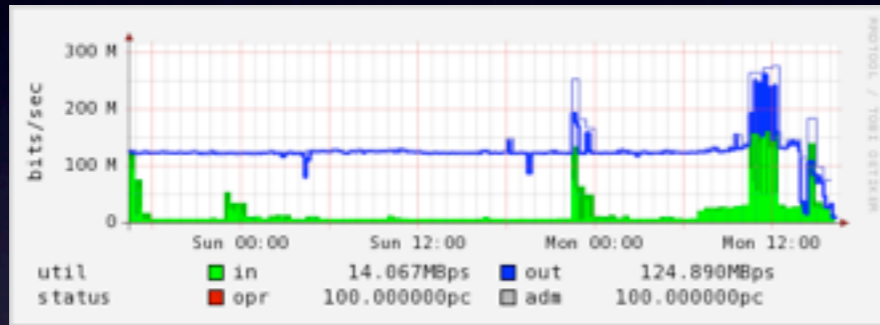
- Quantify, resolve, and document site security issues involving site perfSONAR
- Train local staff to use and support perfSONAR resources, including user education
- Provide comprehensive level of perfSONAR support at sites
 - creation of generate template for end site perfSONAR deployment
 - true end-to-end monitoring

Tying it all together... (I)




ATLAS scientist transfers file from BNL to SLAC - notices that recent transfers are taking longer than they did last week

Tying it all together... (2)



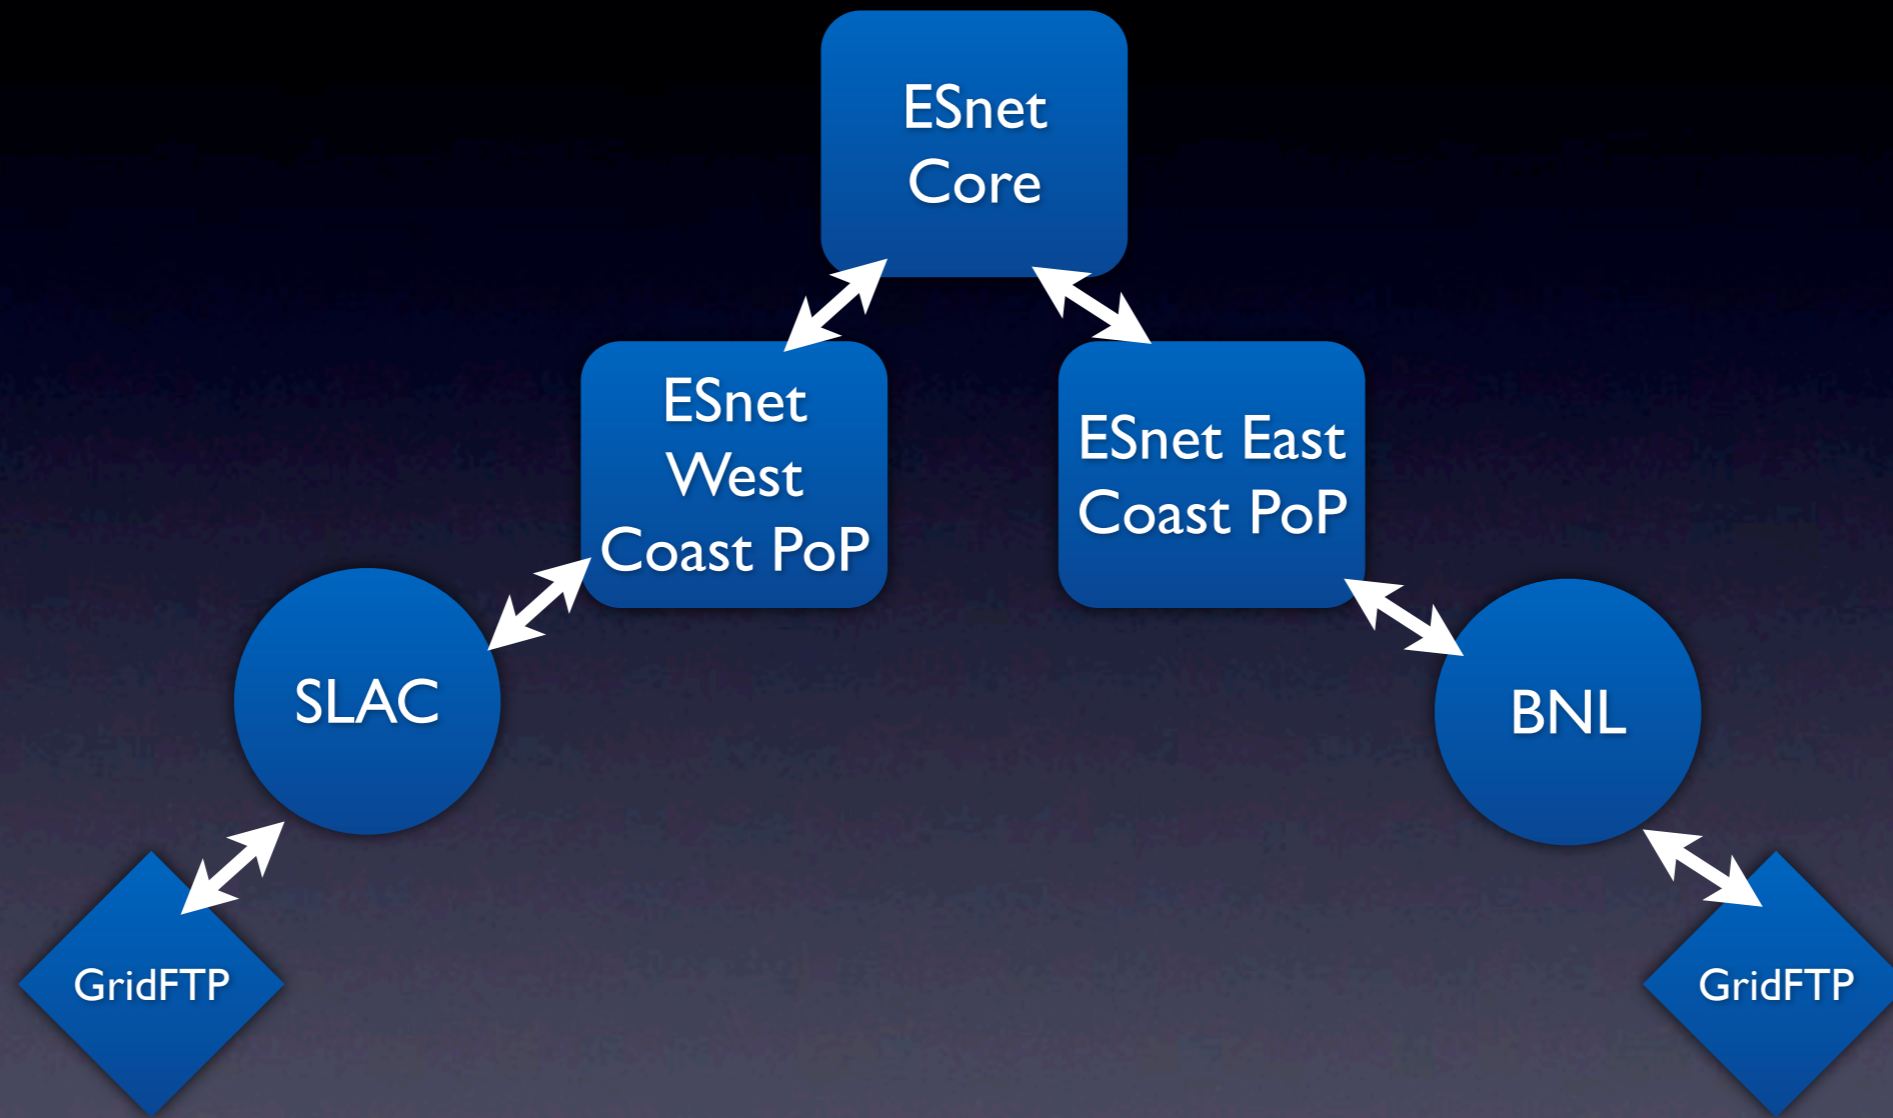
SLAC



BNL

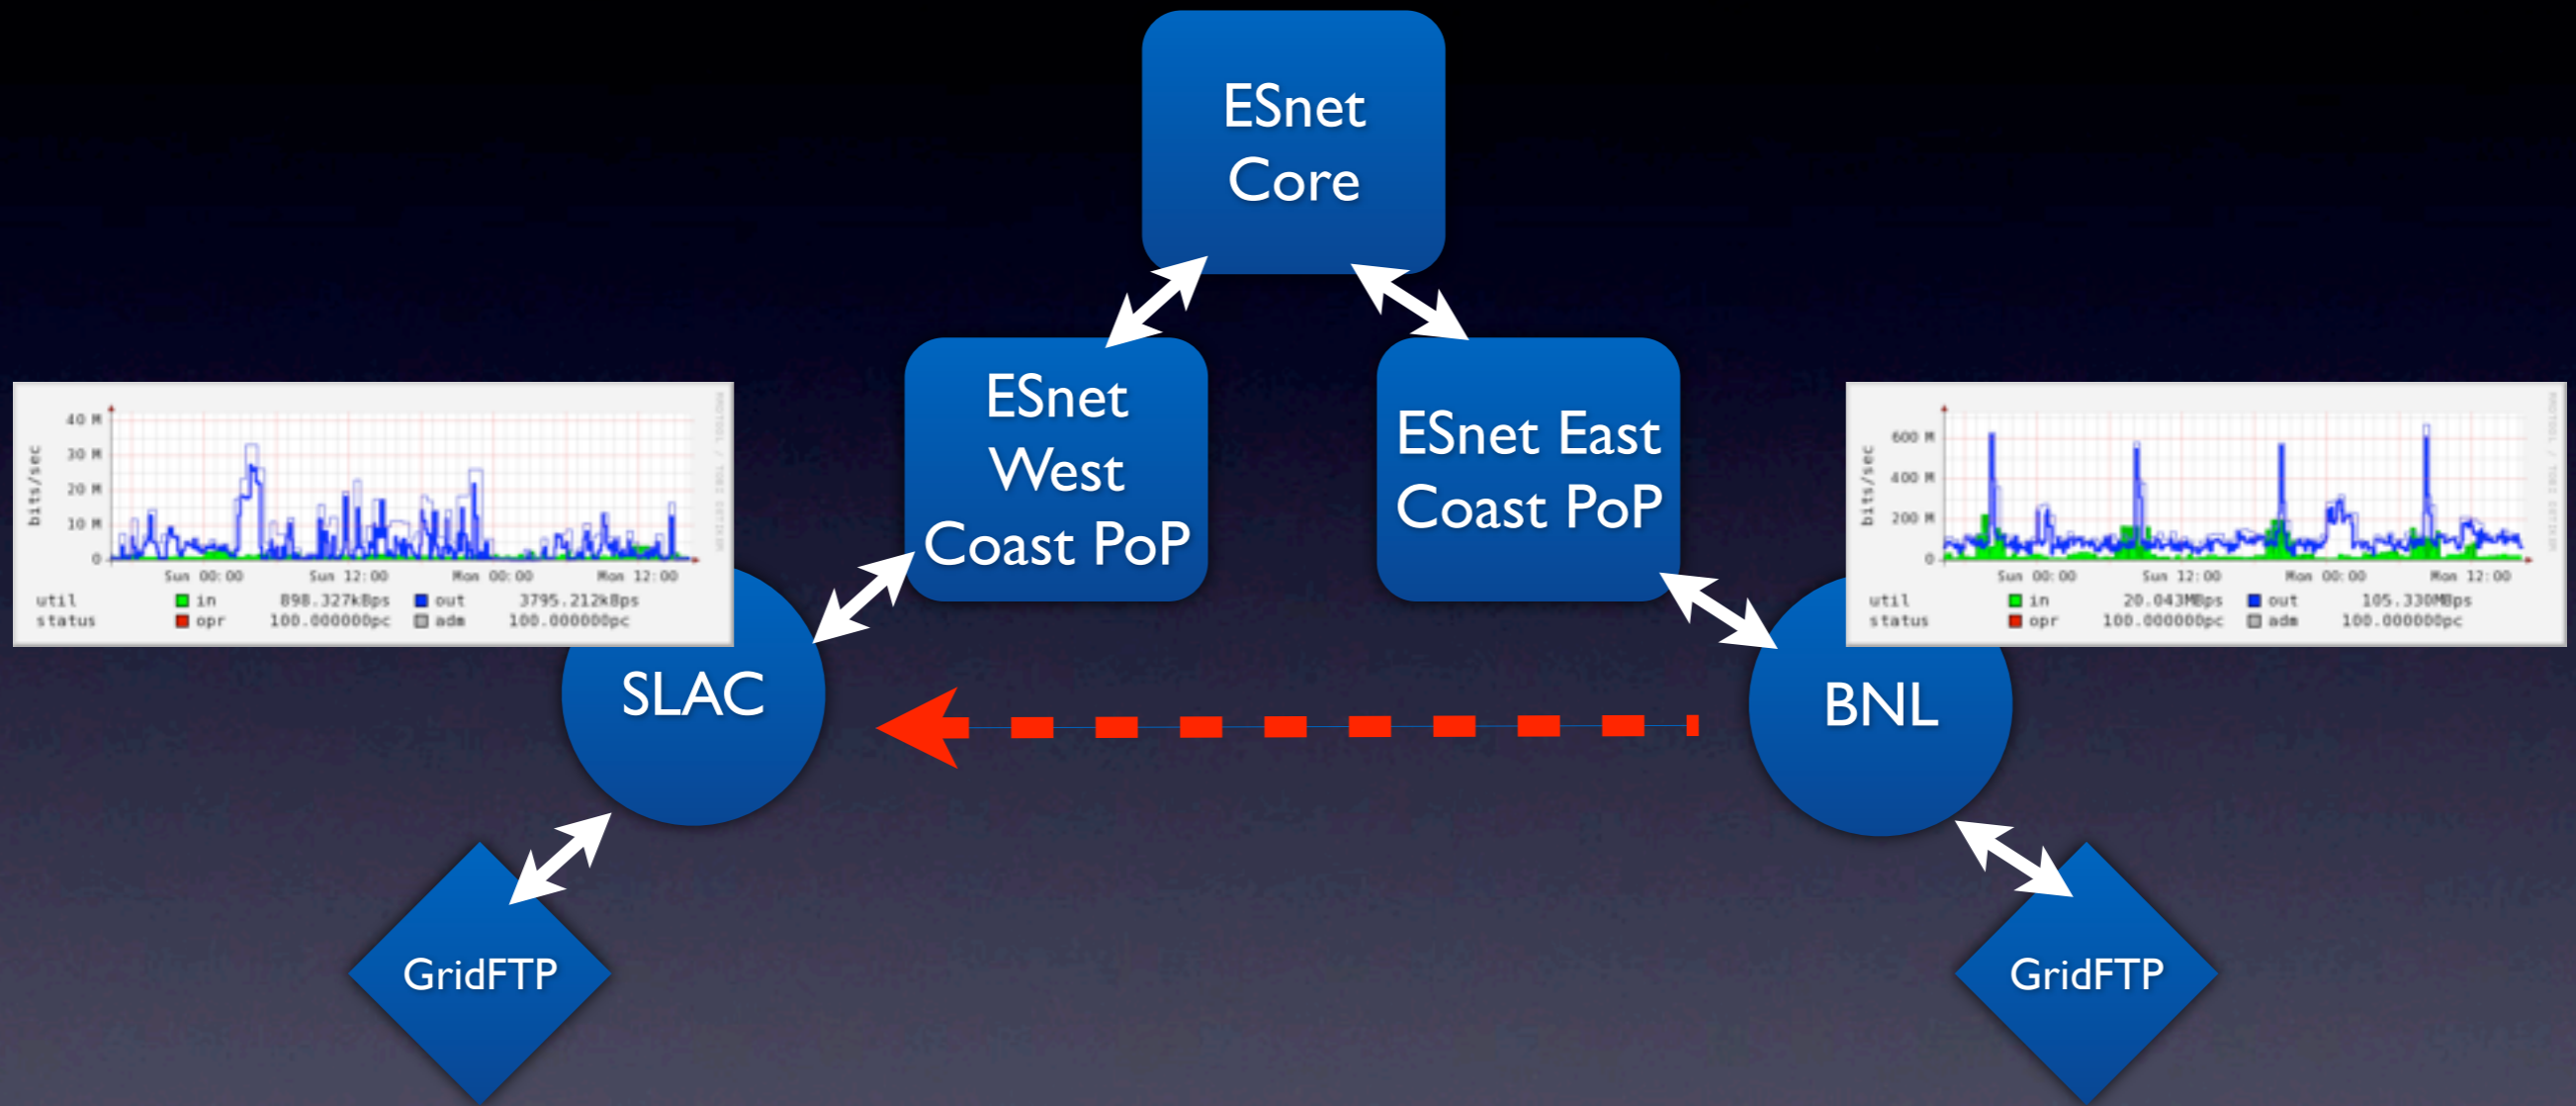
Use **GridFTP MA**, see all transfer from BNL to SLAC are slow

Tying it all together... (3)



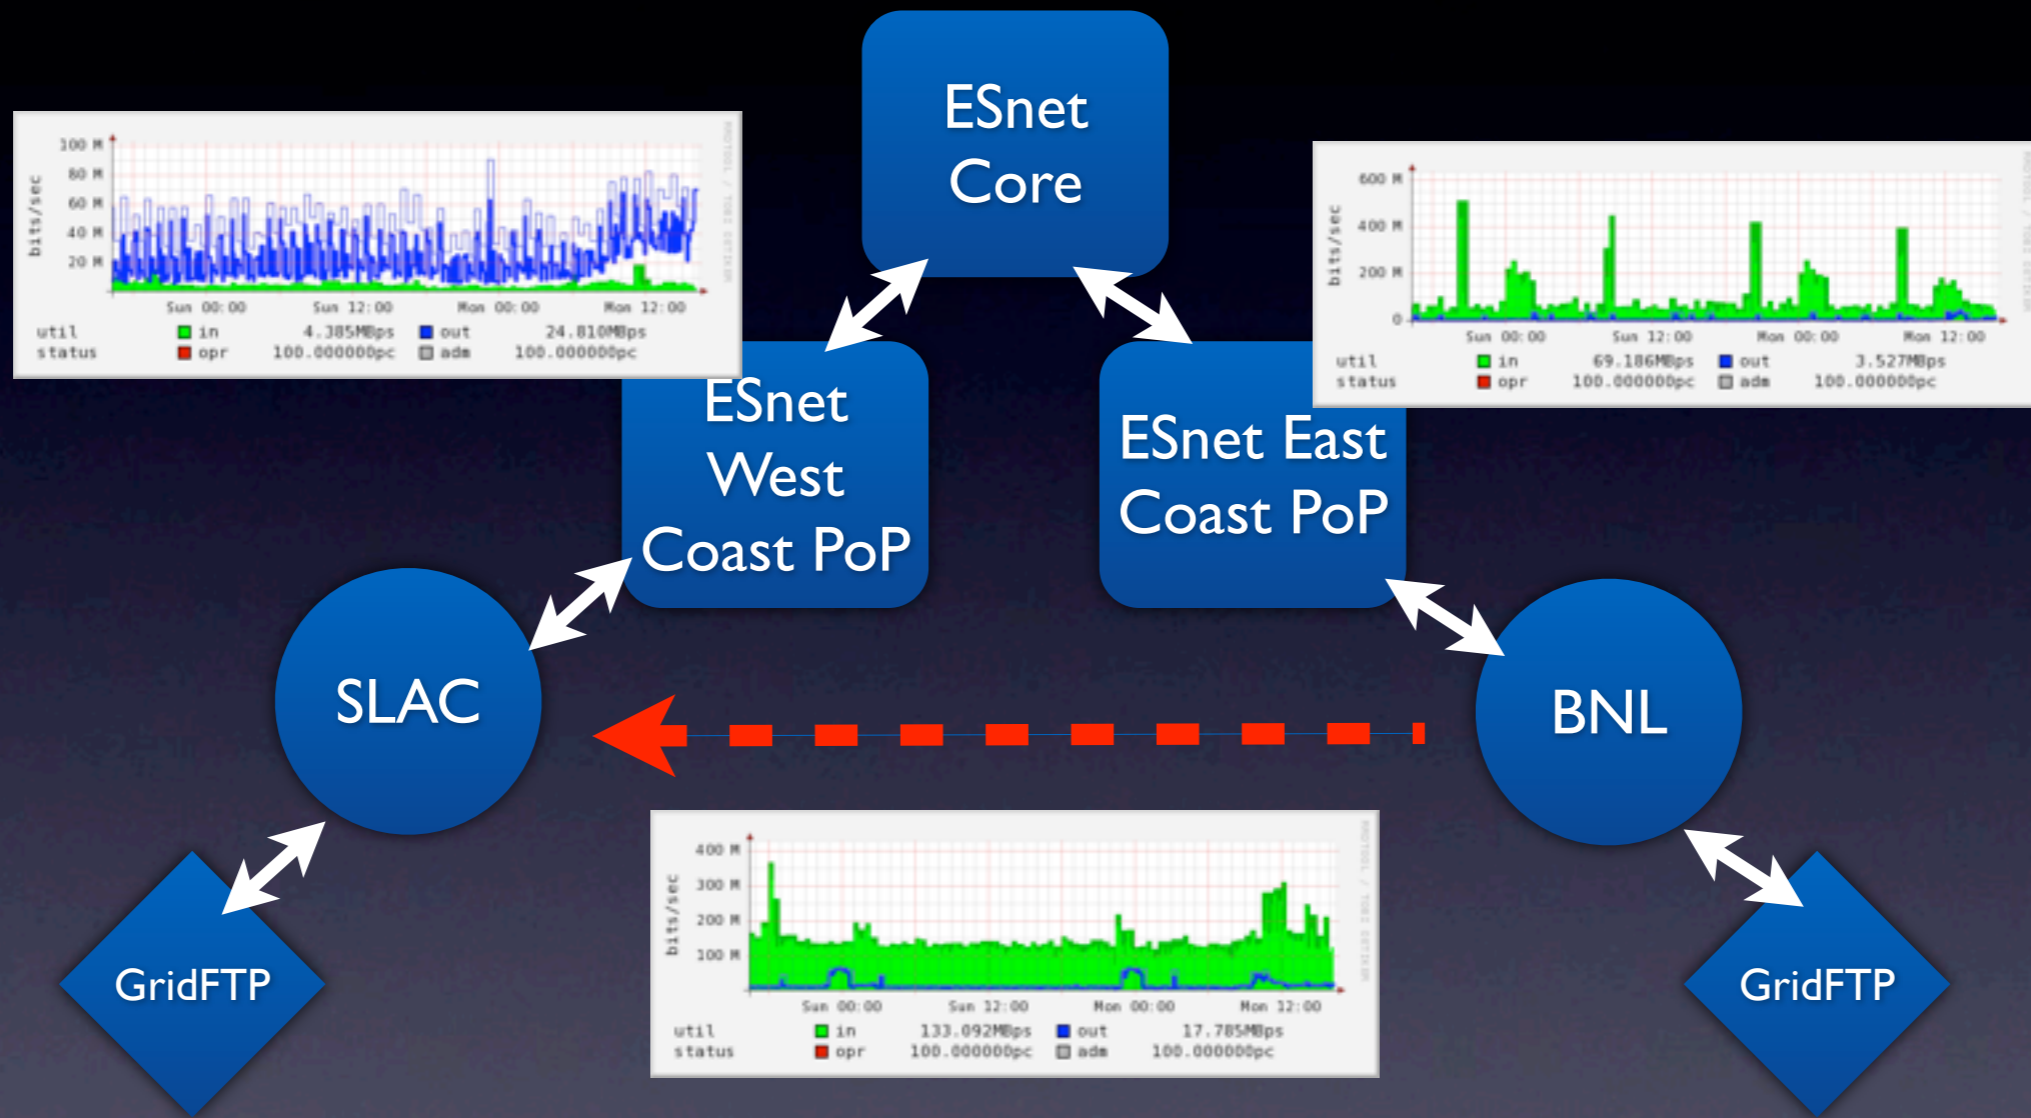
Use **Traceroute MA** service to determine end-to-end path

Tying it all together... (4)



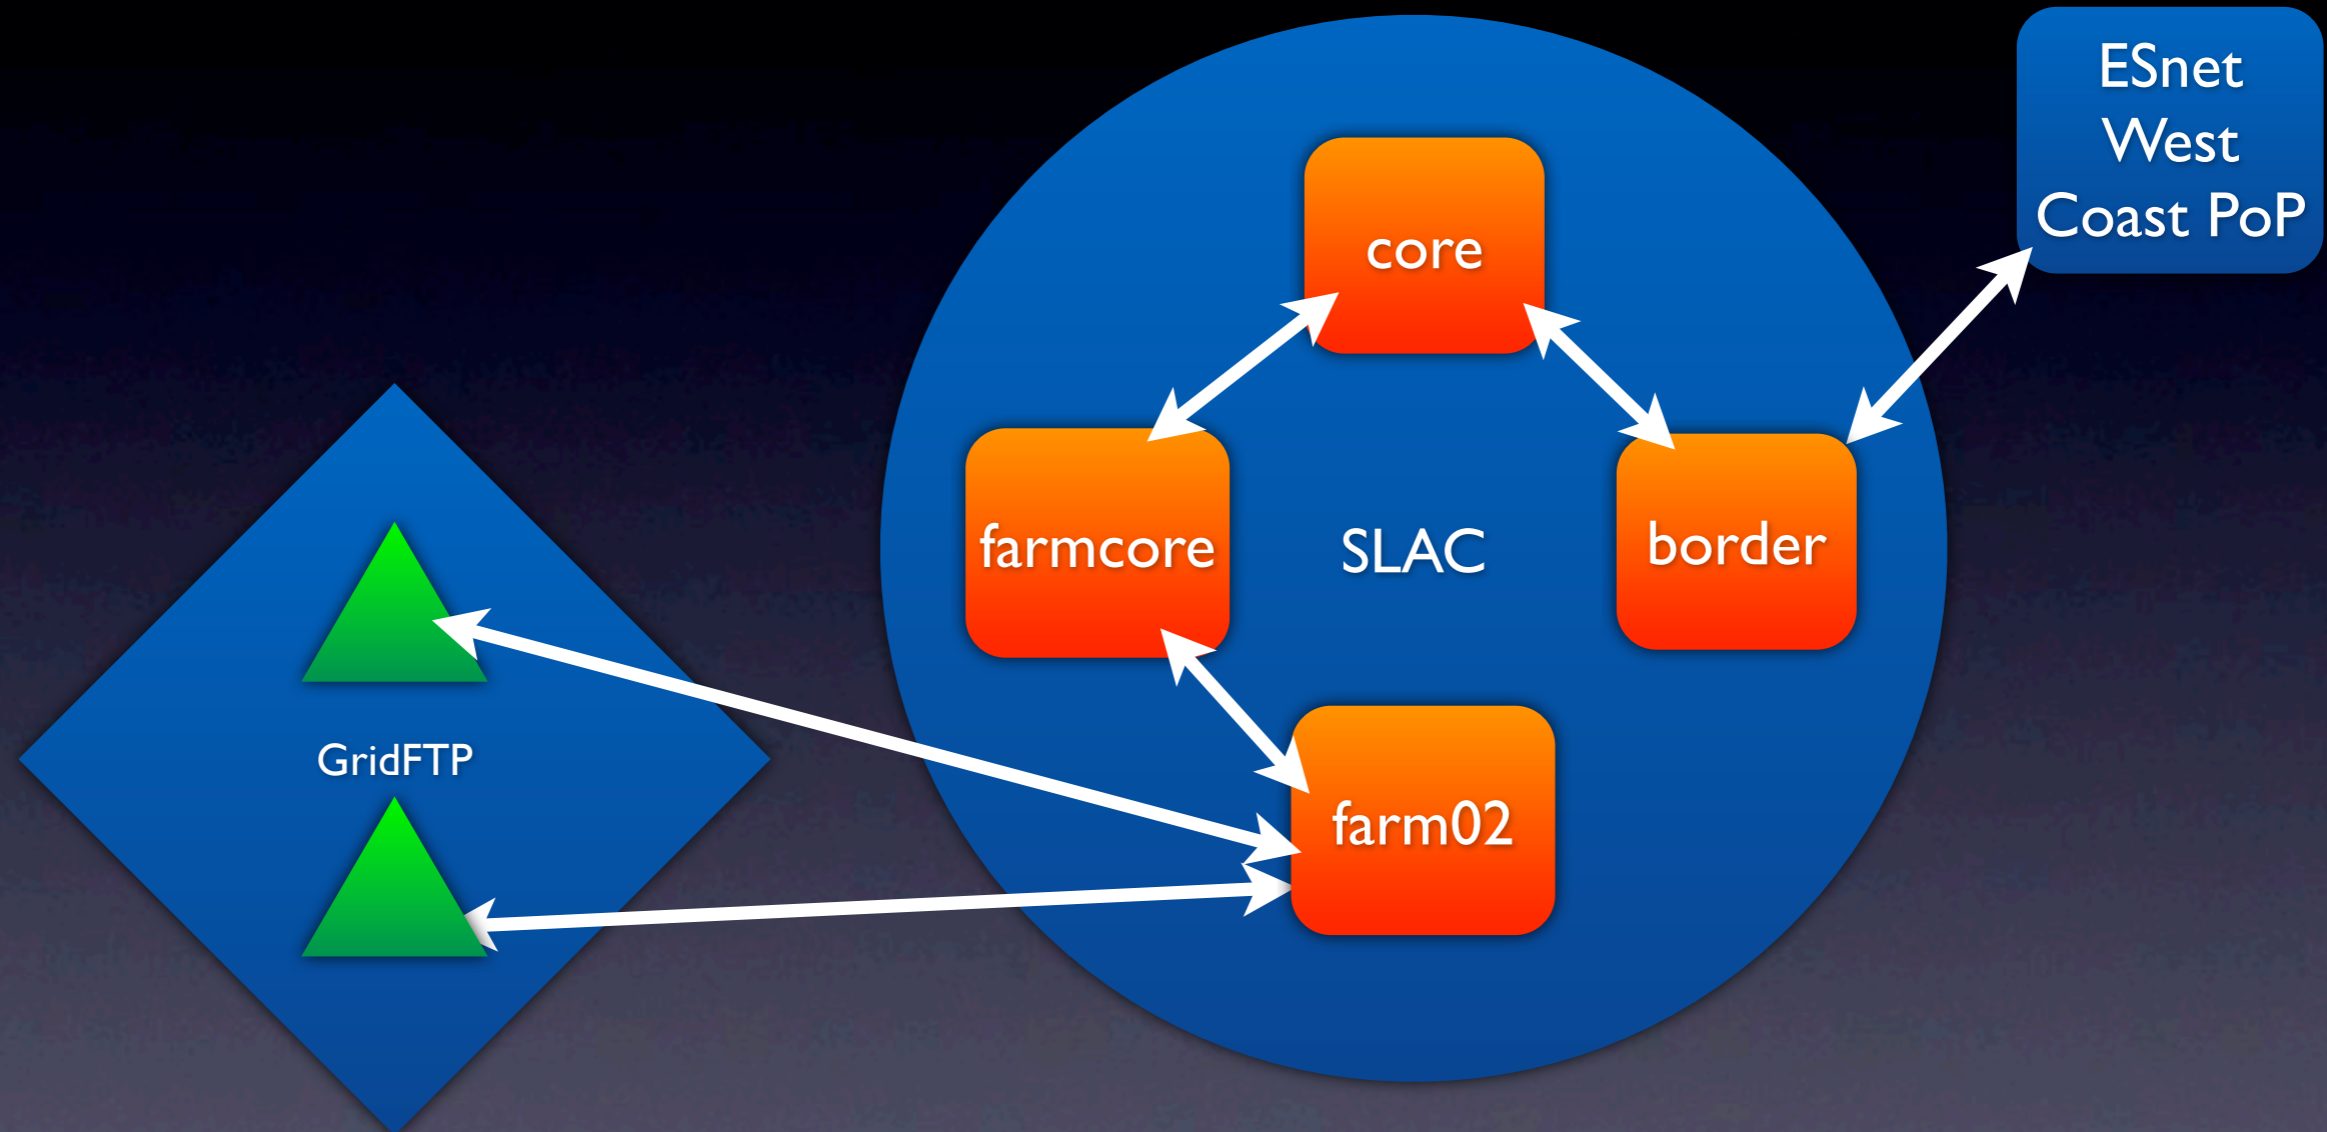
Use **SNMP-MA** to determine site perimeters looks okay

Tying it all together... (5)



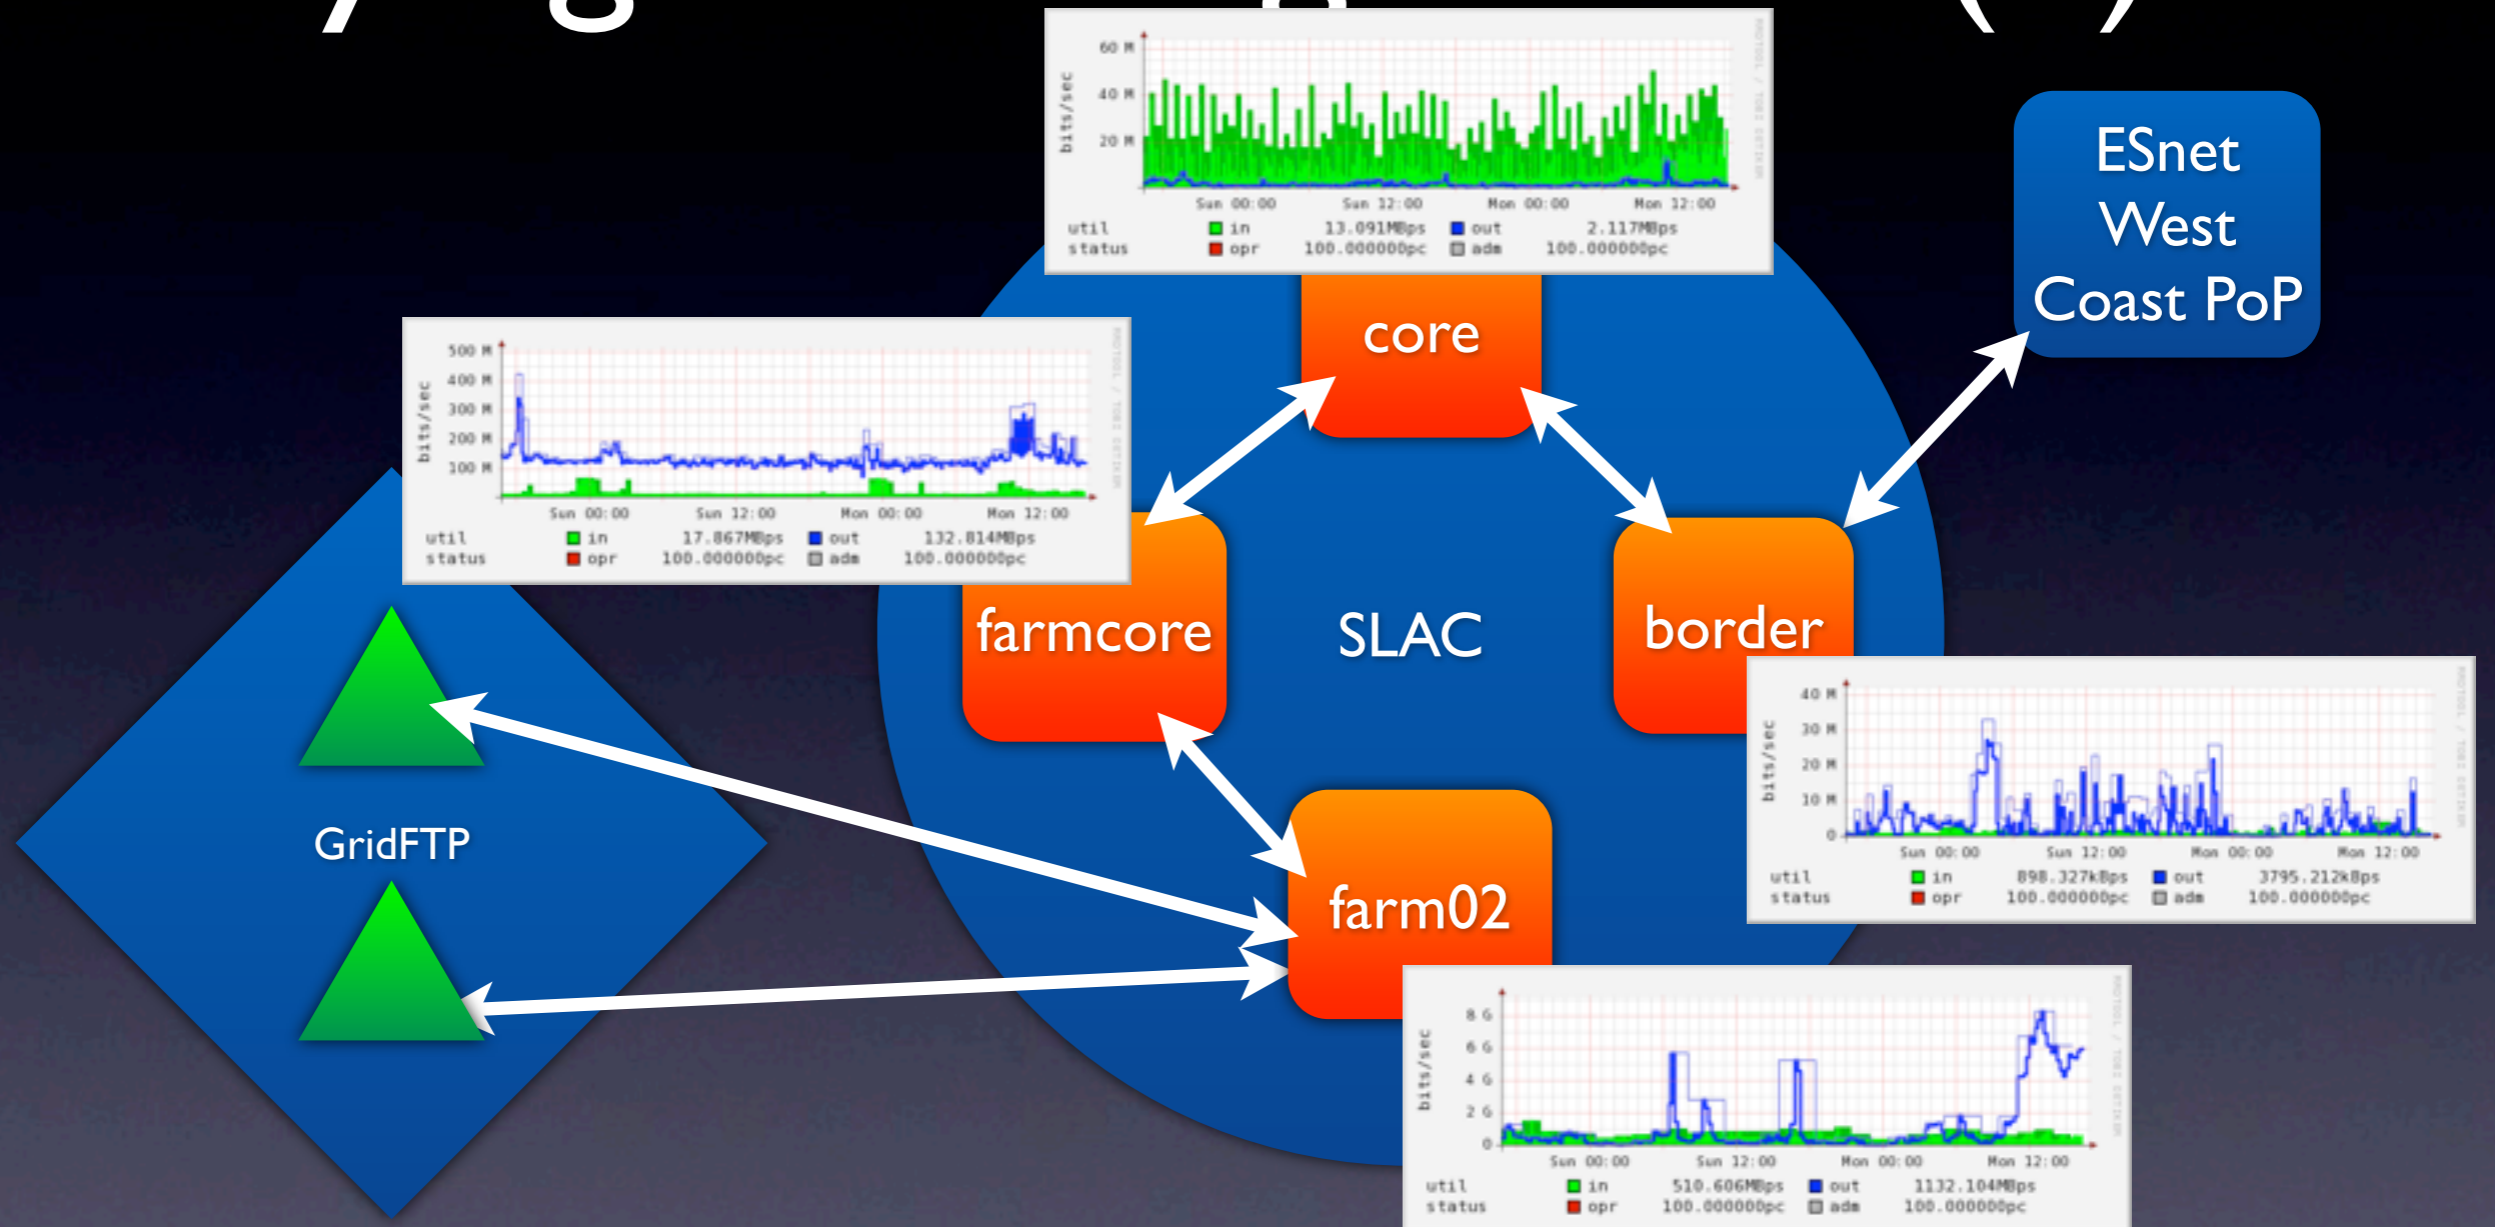
Use **SNMP-MA** to ensure ESnet POPs and cores looks okay

Tying it all together... (6)



Use SLAC's **Topology Service** to determine
Internal network path

Tying it all together... (7)



Use **SNMP-MA** to determine that SLAC *farm02* is overloaded

Description of Work

- Deploy and configure perfSONAR and related tools
- Adapt and deploy E-Center at each site
- Work with cybersecurity to classify network monitoring information and assign ACLs to each category
- Work with site users and admins to promote use of perfSONAR and e-Center

Deliverables

- Year 1:
 - full perimeter perfSONAR deployment at each site
 - passive: border snmp statistics
 - active: owamp, bwctl, ping, traceroute to ESnet PoP and other sites
 - support on-demand owamp and bwctl at perimeter
 - ACL based access to services and data
 - draft risk assessment for services
 - draft training document for local staff
 - draft site support document for AA for e-Center
 - coordinate data collection for e-Center

Deliverables

- Year 2:
 - leverage and adapt other monitoring tools
 - configure internal topology service
 - provide snmp data for internal network devices
 - implement persistent full mesh active measurements between sites
 - implement persistent active data to ESnet PoPs
 - support on-demand internal active measurements to off-site authorised users
 - document site security policies for topology and data
 - draft user support document for local admins
 - coordinate data collection of internal data into e-Center
 - help develop and feedback into e-Center tools and presentation
 - implement site local AA with e-Center

(Potential) Deliverables

- Year 3 (unfunded):
 - GridFTP MA
 - build interactive and integrated web GUI(s) to allow users to navigate, retrieve, visualise, manage and control data presentation
 - visualisation tools to run large amounts of monitoring data into easy-to-understand forms for human visual analytics
 - turn data into knowledge and intelligence
 - conduct performance and behaviour analysis of network traffic and build profiles of reach application and site performance
 - identify bottlenecks in site-to-site performance that may lead to system-wide failures
 - traffic analysis to identify highly-utilised network resource and hub, provide recommendations for improvement

Questions?

Other things...

Alarms and Network Alerts

- Publication of service tickets
 - correlate-able database of known performance and or service outages
 - eg: one ESnet path is down, traffic may go through a much more congested path
 - eg: end site taking down servers which will lead to downtime of services
- how to determine dependencies on each host /service?
 - does end-sites even have such information available?
 - eg: afs, dns, nis etc?
 - potential to use nagios or similar as data gather, present in perfSONAR based schemata?

Application / User Level Monitoring

- GridFTP
 - Working with LBNL
 - Already have alpha build with schemata etc.
 - Looking to try to deploy at a couple of sites
 - Face many challenges with respect to the types of data exposed