

perfSONAR Update

Jason Zurawski

October, 2007

OGF 21 – Seattle WA



Introduction

- Overview
- Motivation
- Architecture
- perfSONAR-PS
- Services
- Defining Standards
- Summary

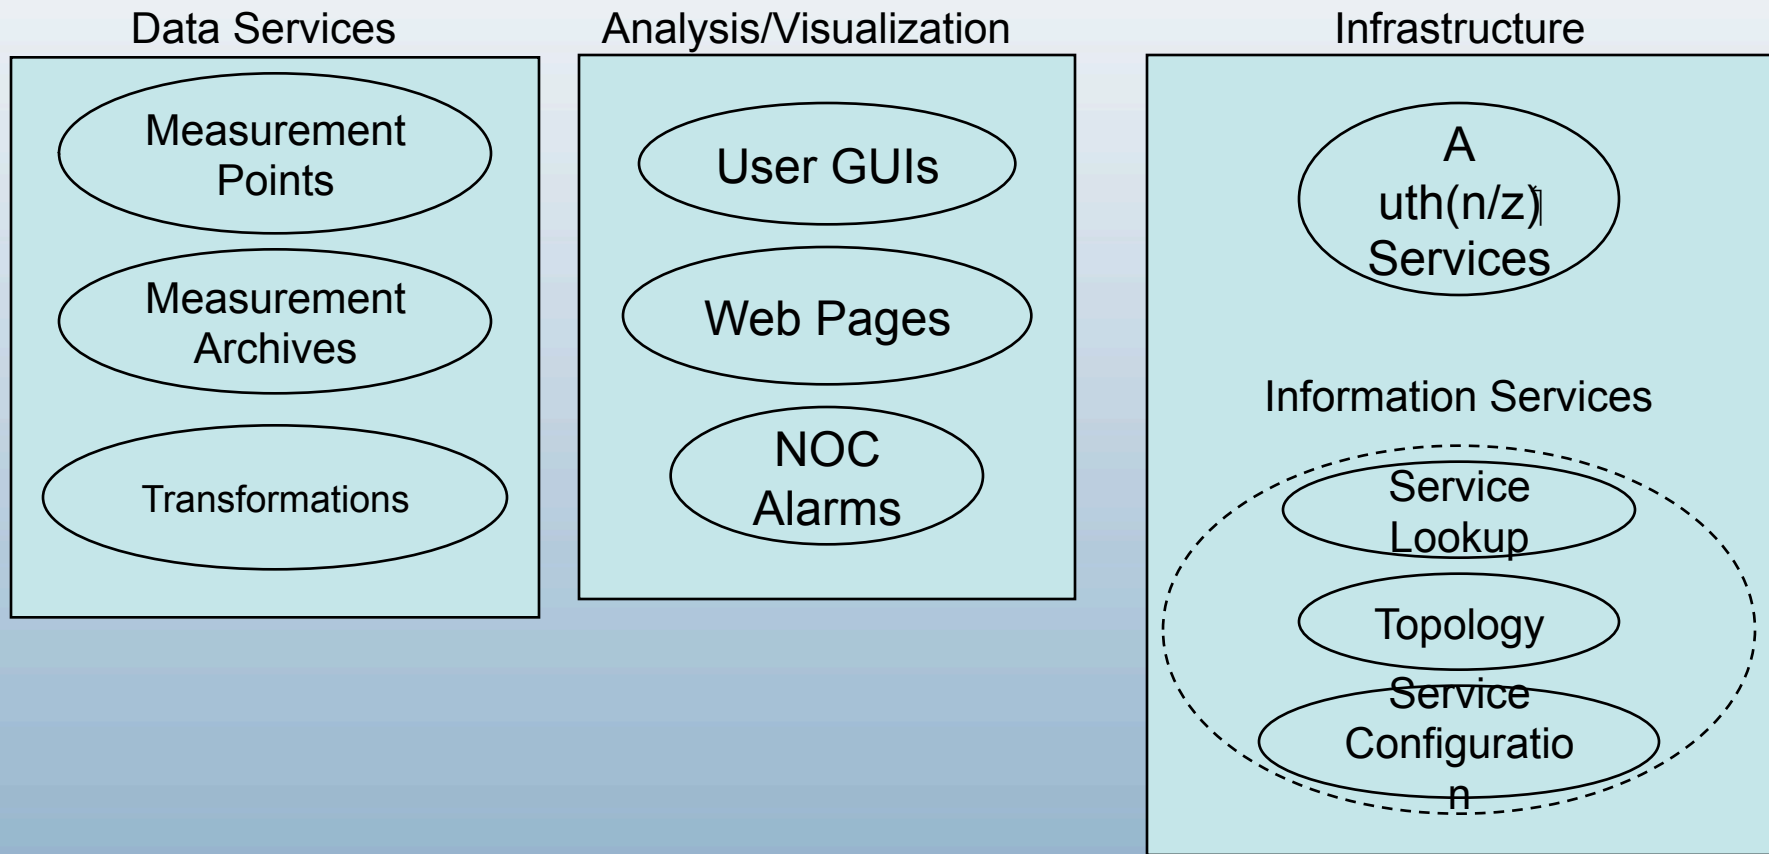
perfSONAR Overview

- Most organizations perform monitoring and diagnostics of their own network
- Networking is increasingly a cross-domain effort
- Monitoring and diagnostics must also become a cross-domain effort

perfSONAR Motivation

- A set of protocols and schemas for implementing a SOA for sharing and controlling network performance tools
 - Implementing NM-WG recommendations
 - Sometimes driving NM-WG direction
- A community of users and developers
- Multiple sets of interoperable software

perfSONAR Architecture



perfSONAR Parts and Partners

- Performance Monitoring Infrastructure
 - perfSONAR (ESnet, GEANT2 JRA1, Internet2 and Members, RNP)
- Distributed System Infrastructure
 - Topology Service (Internet2, University of Delaware)
 - Distributed Lookup Service

perfSONAR-PS

- Create “separate but equal” implementation of perfSONAR standard
 - Use same protocol/standards
 - Encourage new development
 - Proof of interoperability (strengthens the standard)
- Targeted for University NOC deployments
 - Lightweight
 - Easy to deploy/manage
 - Primary users couldn't

perfSONAR-PS Services

- Focus on development of major perfSONAR components
 - LS
 - Topology
 - Link/Circuit Status
 - SNMP Based MA
 - PingER (SLAC)
 - Visualization
 - New additions
 - OWAMP/BWCTL
 - Traceroute
- 8
- Contact if you are interested in testing pre-releases of anything



perfSONAR-PS Services

- Downloads
 - Currently available as individual 'Micro-releases' on the perfSONAR Wiki
 - wiki.perfsonar.net ('Downloads' Section)
 - After interoperability testing will be available as a 'Bundle' via the perfSONAR web site

- Installation

-

- Cu

9

Currently scripts to analyse dependency software

- Moving to CPAN based installation



Lookup Service

- 'Oracle' of perfSONAR deployments
 - Accept service registrations
 - Handle complex data/service queries
 - Manage the lifetimes of data and services to keep framework relevant
- Web Service interface to XML Database
 - Berkeley DB XML
 - Service Info/Data kept in native formats
- Draw away the complex query tasks from otherwise 'busy' services

Lookup Service

- XML based configuration/protocol
 - Native storage/query mechanisms [XPath/XQuery]
 - Message format to exchange the data
- Targeted at single domain deployment
 - Single instance to manage multiple services
 - Multi-domain 'distributed' LS (dLS) in development
- Client applications aware of LS to find services
 - perfSONAR-UI
 - perfAdmin

Lookup Service

- Compatible with Java based XML LS
- Current Deployment:
 - Internet2 (Ann Arbor)
 - University of Delaware
 - TransPAC2
- Planned Deployment:
 - IU for Internet2 network and regional connectors
 - Planetlab (dLS Research w/ University of Delaware)

Topology Service

- Provides a repository for obtaining topology information about a domain
 - Currently obtain the entire network
 - XQuery interface allows the construction of arbitrarily complex queries about the network
- Topology is specified according to the schema in development for NMWG/
NML/DICE

Topology Service

- Current Deployments
 - Internet2
- Planned Deployments
 - Internet2 DCN (Dynamic Circuits)
 - SLAC (PingER Topology Information)

Link Status Measurement Archive

- Provide access to up/down status information about layer2 links
 - Data stored in a SQL database
 - Database schema allows for storing time ranges during which a link had a certain status
 - Minimizes storage costs for rarely changing links
- Communication/Configuration via XML
- Target audience is network operators and users interested in obtaining the status of the links over which their data flows

Link Status Measurement Archive

- Collector
 - Allows for the periodic collection of the status of one or more links
 - Can use SNMP, Scripts or simply Constants
 - Can store results directly into a database or into a remote Measurement Archive

Link Status Measurement Archive

- Visualization
 - A perfSONAR-UI Plugin is available that
c
an display a network and the status of its links
- Current Deployment
 - Internet2 Network
 - HOPI (in2p3 circuit)
- Planned Deployment
 - SLAC
 - Fermilab

Circuit Status Measurement Archive

- An e2emon-compatible service
 - Integrates with the Link Status MA to provide the information stored in MAs
 - Can work with local MAs directly or with remote MAs
 - Can use the Topology MA to obtain necessary information about nodes
 - Can use a Lookup Service to lookup the MA containing information on each link
- Target audience is administrators who want to publish circuit status information to e2emon clients

Circuit Status Measurement Archive

- Visualization
 - Tools that are compatible with e2emon should work with this service
- Current Deployment
 - Internet2 Network
 - HOPI (in2p3 circuit)
- Planned Deployment
 - SLAC
 - Fermilab

SNMP Measurement Archive

- Provide access to network performance data
 - Utilization
 - Errors
 - Discards
- Numerous tools already exist to collect passive measurements (via SNMP):
 - MRTG
 - Cacti
 - Cricket
- Expose archives from RRD files

SNMP Measurement Archive

- XML based configuration/protocol
 - Configuration file to expose the data
 - Message format to exchange the data
- Targeted at researchers, NOC staff debugging intra/inter domain network problems
- Many client applications available to graphically display the results:
 - perfSONAR-UI
 - perfAdmin
 - perfOMeter

SNMP Measurement Archive

- Compatible with Java RRD MA
- Current Deployment:
 - Internet2 Network
 - ESnet
 - Georgia Tech/SOX
 - Fermilab
 - TransPAC2
- Planned Deployment:
 - Regional Networks (MCNC, Merit, Magpi)
 - Campus Networks (Rutgers, UDel, UW)

PingER MP/MA

- Tool capable of collecting/storing ping measurements between sites
- Original Development by SLAC
 - pS integration by SLAC, Fermilab
- Target release by LHC deadlines

Visualization

- perfOMeter
 - Javascript/AJAX based utilization monitor
 - Polls (known w/ n) SNMP/RRD MA instances for utilization data
 - Integration w/ Topology/Lookup Services (Planned)
 - Special purpose for now, targeted more at Network novices than NOC staff

Visualization

- perfAdmin (Spyglass)
 - C
o
n
tacts known LS instances for service listing
 - Can query for service data
 - Beta 'utilization' graphing available
 - Can perform routine tasks on LS data
 - 'Keepalive' services
 - 'Deregister' services
 - Needs authentication (bad)

Visualization

- perfSONAR-UI Plugins
 - Developed by Internet2 intern Mark Gordon (UMich)
 - Utilization based on traceroute path
 - Status MA visualization
 - Able to interact with Status MA/Utilization MA/LS

Visualization

- Google Maps Mashup
 - Developed by Yee-Ting Lee at SLAC
 - Geographically place monitored instances on a Google map
 - Status MA visualization
 - Able to interact with Utilization MA/LS

Other Services (in development)

- **OWAMP MA**
 - Based on Internet2's One Way Active Measurement Protocol
 - Provided access to archived results
 - Schedule new tests
- **BWCTL MA**
 - Based on Internet2's BWCTL tool
 - Provided access to archived results
 - Schedule new tests
- **Traceroute MP/MA**
 - Joint effort between Internet2, UDel, and SLAC
 - Integrate with visual traceroute tools (SLAC, ESnet)
 - Currently in testing

perfSONAR-PS Releases

- Currently exploring testing and packaging strategies
 - Releases will ultimately be CPAN based
 - Standing issues include configuration/automation help.

- Versions

29

of SNMP MA and LS will be available this fall

INTERNET®

- Pre-CPAN versions will be available for LHC.com

Defining Standards

- Protocol Standards
 - Measurement Representation
 - Service communication (message formats)
 - NMWG/IETF
- Service Standards
 - E
 - X
 - pected behavior (implementation agnostic)
 - Current state + Future goals

Defining Standards

- The perfSONAR Topology schema is also used in the DCN control plane
 - We've spent quite a bit of effort harmonizing these
- The obvious win is that we have the measurement system have immediate access to dynamic circuits
- The broader impact is that we're approaching a unified network interaction model (UNIM)

Defining Standards

- Aiming to standardize pS specific schemata as “profiles” of current work on NM-WG Base schema work
- The topology efforts have spawned a new WG called the Network Markup Language WG (NML-WG)

Summary

- perfSONAR/perfSONAR-PS development is aiming for LHC targets
- Software development has driven schemata work up until now
 - Time to go the other direction for interoperability sake
- Infrastructure nearly complete, bring in the scientists
 - Visualization of data is important, we need help knowing what you want!

www.internet2.edu

