



February 9<sup>th</sup> 2010, AIMS 2010

Matt Zekauskas, [matt@internet2.edu](mailto:matt@internet2.edu)

# perfSONAR Overview

# Outline

- Motivation
- What is *perfSONAR*?
- Who is involved in *perfSONAR*?
- Who is adopting *perfSONAR*?
- Why might *perfSONAR* be useful to you
- For more information...

# Motivation – A Typical Scenario

- User and resource are geographically separated
- Both have access to high speed communication network
  - LAN infrastructure - 1Gbps Ethernet
  - WAN infrastructure – 10Gbps Optical Backbone



# Motivation – A Typical Scenario

- User wants to access a file at the resource (e.g. ~600MB)
- Plans to use COTS tools (e.g. SCP, but could easily be something scientific like GridFTP or simple like a web browser)
- What are the expectations?
  - 1Gbps network (e.g. *bottleneck* speed on the LAN)
  - $600\text{MB} * 8 = 4,800 \text{ Mb}$  file
  - User expects *line rate*, e.g.  $4,800 \text{ Mb} / 1000 \text{ Mbps} = 4.8 \text{ Seconds}$
- What are the realities?
  - Congestion and other Network performance factors
  - Host performance
  - Protocol Performance
  - Application performance


# Motivation – A Typical Scenario

- Real Example (New York USA to Los Angeles USA):

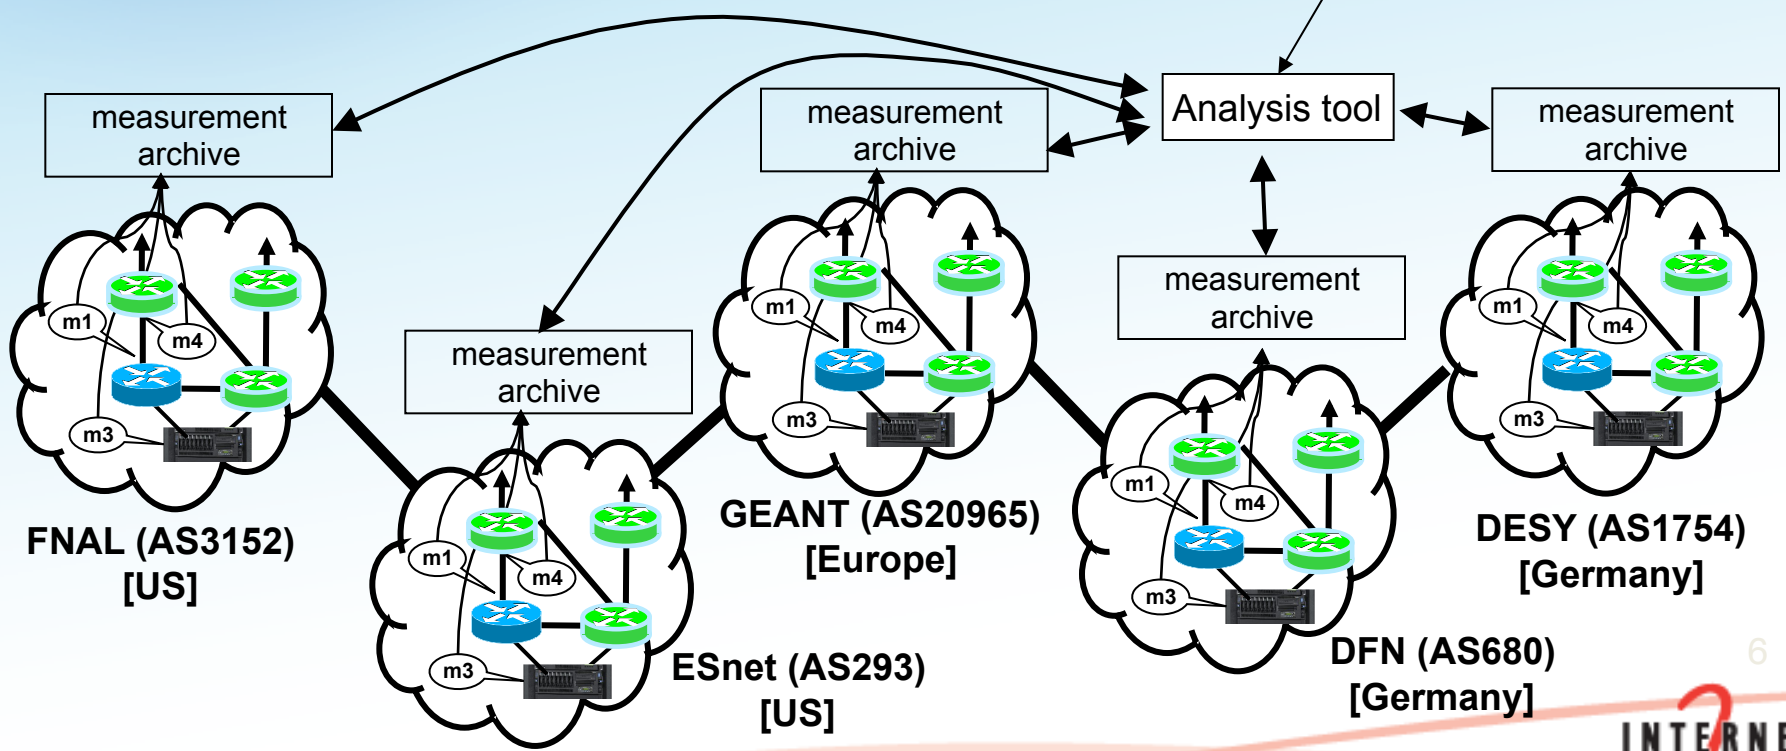

```
[zurawski@nms-rthr2 ~]$ scp zurawski@bwctl1.losa.net.internet2.edu:pS-Performance_Toolkit-3.1.1.iso .  
pS-Performance_Toolkit-3.1.1.iso      2%  17MB  1.0MB/s  10:05 ETA_
```

- Example:

- 10 minutes seems unreasonable given the investment in technology
  - Backbone network
  - High speed LAN
  - Capable hosts
- Performance realities as network speed decreases:
  - 100 Mbps Speed – 48 Seconds
  - 10 Mbps Speed – 8 Minutes
  - 1 Mbps Speed – 80 Minutes
- How could this happen?
- More importantly, why are there not more complaints?



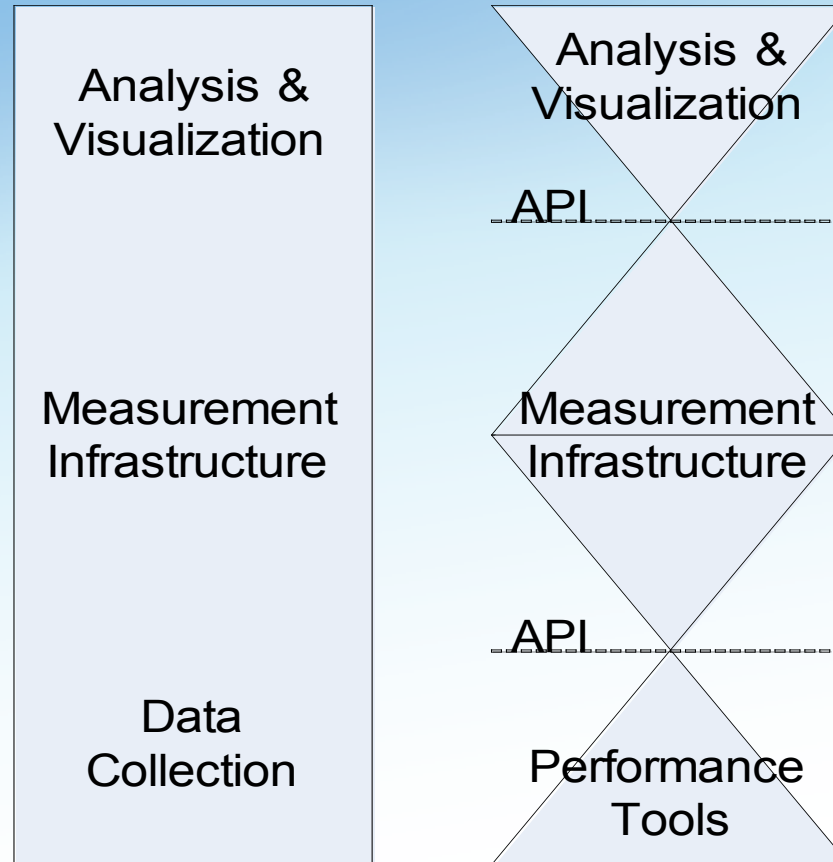
Many collaborations are inherently multi-domain, so for an end-to-end monitoring tool to work everyone must participate in the monitoring infrastructure



# Possible Solutions

- Desirable design features for any solution
  - Component Based
    - Functionality should be split into logical units
    - Each function (e.g. visualization) should function through well defined communication with other components (e.g. data storage)
  - Modular
    - Monolithic designs rarely work
    - Components allow choice of how to operate a customized end solution.
  - Accessible
    - Well defined interfaces (e.g. APIs)
- Initial design should facilitate future expansion

# Possible Solutions





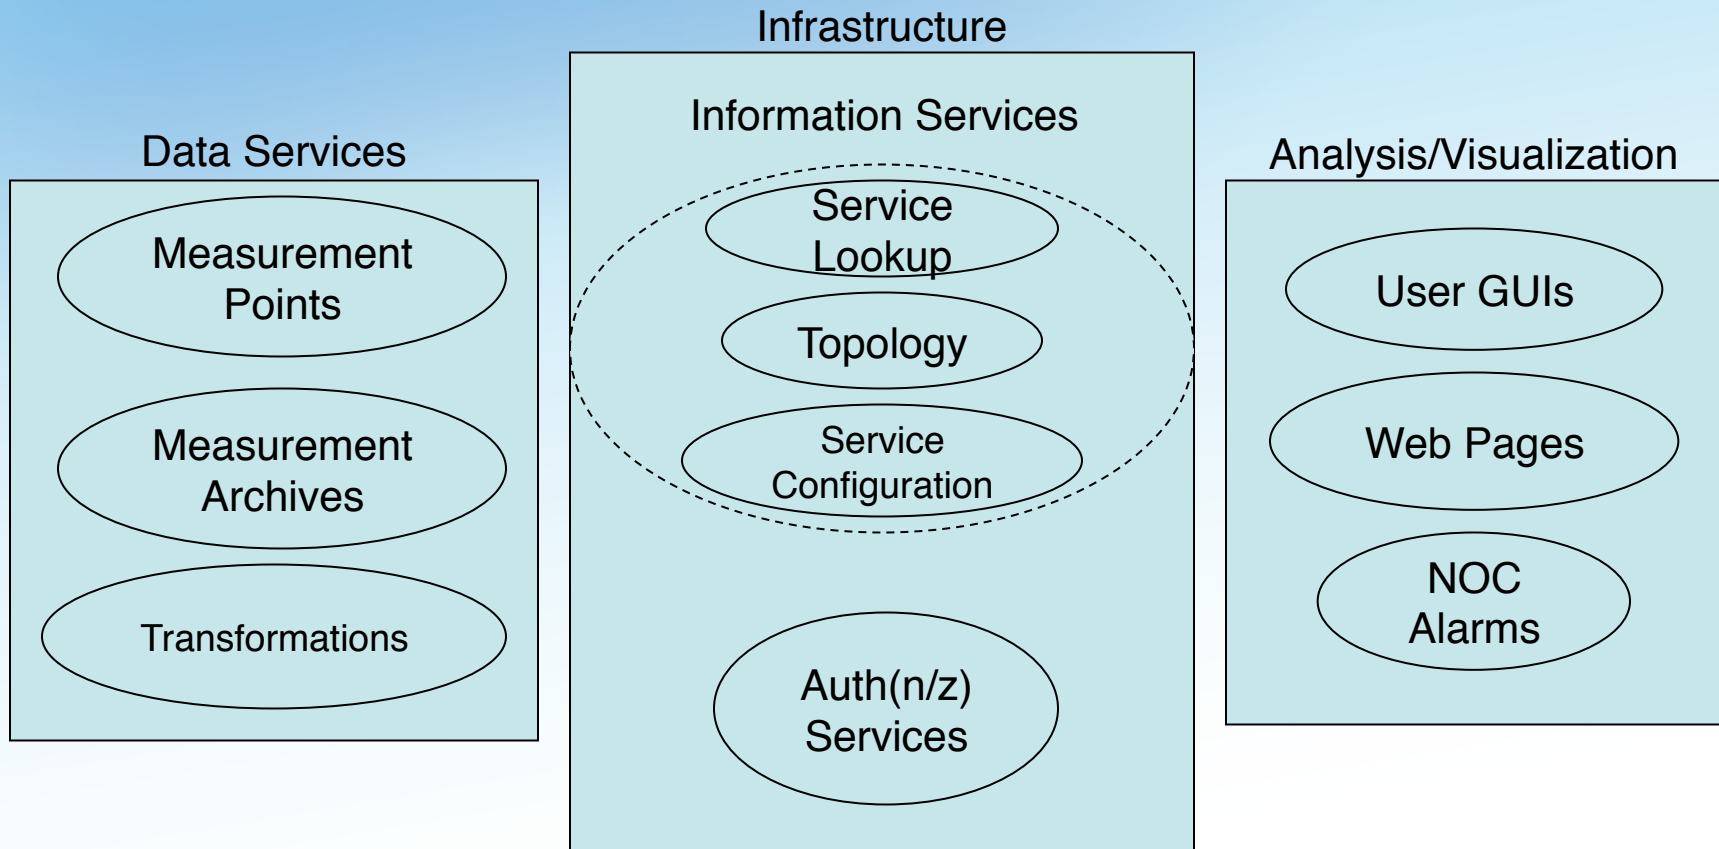
# perfSONAR Architecture Overview

- Interoperable network measurement middleware designed as a Service Oriented Architecture (SOA):
  - Each component is modular
  - All are Web Services (WS) based
  - The global *perfSONAR* framework as well as individual deployments are decentralized
  - All *perfSONAR* tools are Locally controlled
- *perfSONAR* Integrates:
  - Network measurement tools and archives (e.g. stored measurement results)
  - Data manipulation
  - Information Services
    - Discovery
    - Topology
  - Authentication and authorization

# perfSONAR Architecture Overview

- The key concept of *perfSONAR* is that each entity performs a service
  - Each service provides a limited set of services, e.g. collecting measurements between arbitrary points or managing the registration and location of distributed services
  - The service is a self contained entity and provides functionality on its own as well as when deployed with the remainder of the framework
- Services interact through protocol exchanges
  - Standardized message formats
  - Standardized exchange patterns
- A collection of *perfSONAR* services within a domain is a deployment
  - Deploying *perfSONAR* can be done *À la carte*, or through a complete solution
- Services federate with each other, locally and globally
  - Services are designed to automatically discover the presence of other *perfSONAR* components
  - Clients are designed with this distributed paradigm in mind

# perfSONAR Architecture Overview



# perfSONAR Architecture Overview

- A *perfSONAR* deployment can be any combination of services
  - An instance of the *Lookup Service* is required to share information
  - Any combination of data services and analysis and visualization tools is possible
- *perfSONAR* services automatically federate globally
  - The *Lookup Service* communicates with a confederated group of directory services (e.g. the *Global Lookup Service*)
  - Global discovery is possible through APIs
- *perfSONAR* is most effective **when all paths are monitored**
  - Debugging network performance must be done *end-to-end*
  - Lack of information for specific domains can delay or hinder the debug process

# Who is perfSONAR?

- The *perfSONAR* Consortium is a joint collaboration between
  - ESnet
  - Géant
  - Internet2
  - Rede Nacional de Ensino e Pesquisa (RNP)
- Decisions regarding protocol development, software branding, and interoperability are handled at this organization level
- There are two independent efforts to develop software that is compatible with *perfSONAR*
  - perfSONAR-MDM
  - perfSONAR-PS
- Each project works on an individual development roadmap and works with the consortium to further protocol development and insure compatibility

# Who is perfSONAR-MDM?

- [perfSONAR-MDM](#) is made up of participants in the Géant project:
  - Arnes
  - Belnet
  - Carnet
  - Cesnet
  - CYNet
  - DANTE
  - DFN
  - FCCN
  - GRNet
  - GARR
  - ISTF
  - PSNC
  - Nordunet (Uninett)
  - Renater
  - RedIRIS
  - Surfnet
  - SWITCH
- perfSONAR-MDM is written in Java primarily and was designed to serve as the monitoring solution for the Large Hadron Collider (LHC) project.
- perfSONAR-MDM is available as Debian or RPM packages.

# Who is perfSONAR-PS?

- [perfSONAR](#)-PS is comprised of several members:
  - ESnet
  - Fermilab
  - Georgia Tech
  - Indiana University
  - Internet2
  - SLAC
  - The University of Delaware
- perfSONAR-PS products are written in the perl programming language and are available for installation via source or RPM packages
- perfSONAR-PS is also a major component of the Internet2 pS Performance Toolkit – A bootable Linux CD containing measurement tools.



# perfSONAR Adoption

- *perfSONAR* is gaining traction as an interoperable and extensible monitoring solution
- Adoption has progressed in the following areas:
  - R&E networks including backbone, regional, and exchange points
  - Universities on an international basis
  - Federal labs and agencies in the United States (e.g. *JET* nets)
  - Scientific Virtual Organizations, notably the LHC project
- Recent interest has also accrued from:
  - International R&E network partners and exchange points
  - Commercial Providers in the United States
  - Hardware manufactures



# perfSONAR Adoption

- Networks
  - APAN, CENIC, CSTNET, ESnet, Geant, Gloriad, GPN, Internet2, JGN2, LONI, MAX, NOX, NSERNET, RNP, Starlight, Transpac2, UEN
- Labs
  - ANL, BNL, FNAL \*\*, NERSC, PNNL, PSC, SLAC
- International Sites
  - Chinese University of Hong Kong, Chonnam National University (Korea), KISTI (Korea), Monash University (Melbourne, Victoria, Australia), MRREE (Lima, Peru), NCHC (Taiwan), NICT (Japan), Simon Frazier (Burnaby, BC, Canada), Thaisarn Nectec (Bangkok, Thailand), UNIFACS (Salvador, Bahia, Brazil)
- Other
  - Cobham, Northrop Gruman, Ocala Electric, Philadelphia Orchestra, REDDnet
- Current
  - <http://www.perfsonar.net/activeServices/IS/>
  - <http://stats1.es.net/cgi-bin/tree.cgi>
  - <http://stats1.es.net/cgi-bin/directory.cgi>
- Universities
  - Boston University \*
  - College of William and Mary
  - George Mason Univ
  - Georgia Tech University
  - Hope College
  - Indiana University \*
  - Leeward Community College
  - Louisiana State University
  - Michigan State University \*
  - Middle Tennessee State University
  - Northwestern \*\*
  - Oregon State
  - Penn State University
  - Southern Methodist University \*
  - Syracuse
  - Texas A&M University \*
  - Tufts \*
  - University of California Los Angeles
  - University of California San Diego \*\*
  - University of Chicago \*
  - University of Connecticut
  - University of Delaware
  - University of Hawaii
  - University of Michigan \*
  - University of Northern Iowa
  - University of Oklahoma \*
  - University of Texas \*
  - University of Utah
  - University of Wisconsin (Condor)
  - University of Wisconsin (Madison) \* \*\*
  - Vanderbilt \*\*
  - University of Florida \*\*

\* USATLAS

\*\* USCMS

# What is currently available

Services useful for solving end-to-end performance problems

- Link Utilization
- Traceroute beacons/"looking glass"es
- Active measurements and active measurement results
  - Latency (and loss): one-way and round-trip
  - Throughput

# How might perfSONAR be useful to you?

- A data source for your research
- A place to try out new algorithms and visualizations, and feed them back to production networks
- A place to install new data sources with the cooperation of production networks
- Plug: NSF CISE is sponsoring a workshop currently planned for this summer (thinking early or late July) to bring together researchers and R&E network operators to spread perfSONAR deployment and see how to create some lasting relationships between researchers and R&E operators in order to create conduits for the last two bullet points above. Contact me if you are interested.

# For more information

- General: [www.perfsonar.net](http://www.perfsonar.net)
- The perl implementation: <http://psps.perfsonar.net>
- perfSONAR tools and software: <http://software.internet2.edu>
- A hook to the global lookup service:  
<http://stats1.es.net/cgi-bin/tree.cgi>
- More human-readable list of services:  
<http://stats1.es.net/cgi-bin/directory.cgi>



## **perfSONAR Overview**

February 9<sup>th</sup> 2010, AIMS 2010

Matt Zekauskas, [matt@internet2.edu](mailto:matt@internet2.edu)

For more information, visit [psps.perfsonar.net](http://psps.perfsonar.net)