



July 11<sup>th</sup> 2011 – Summer Joint Techs

Jeff Boote, Internet2 Assistant Director R&D

Aaron Brown, Internet2 Network Software Engineer

## **Dynamic Network Services**

# Internet2's New Strategic Focus

Advanced network and network services leadership

Services at scale:  
Services “above the network”

U.S. UCAN

National/Regional collaboration

Research community development and engagement

Industry partnership development and engagement

Global reach and leadership



INTERNET



# Dynamic Network Vision

- Want ability to partition network in ways to...
  - Enable isolation of large flows and small flows
  - Enable predictable performance
  - Enable virtual privacy
    - Connector defined semantics (ex. L2 vs L3)
    - Protocol experimentation
  - Enable incremental scalability of underlying resources

# Dynamic Network Motivation

- Data movement to support science:
  - Increasing in size (100s of TBs in the LHC World)
  - Time sensitivity (data may grow “stale” if not processed immediately)
- Flexible network engineering to support connectors (VLAN service)
- Flexible provisioning of backbone resources
  - Provision backbone resources similar to the way cloud CPU resources are scheduled
- Ability to access the above functionality across multiple collaborating domains

# DYNES and DYGIR Overview

- Planned deployments within the next year will extend service to around 50 domains
  - DYNES BOF will go into detail
    - Group A deployments should begin within the next month
  - DyGIR supported deployments will be happening as new hardware is deployed, likely before the end of the year

# Dynamic Services Potential

- Future Internet2 dynamic services are likely to be provided using planned NDDI substrate
- OS<sup>3</sup>E planned service
- Historical service offering ION could be provided over NDDI substrate
  - Well known defined service offering
  - Mature definition of inter-domain service requirements
  - Reasonable interim step while determining potential of OpenFlow inter-domain potential
- DyGIR extends these services globally
- DYNES extends these services within participating regionals and campuses
- Cloud model for computing and services likely to extend to network resources

# Dynamic Services Current Status


- New services to support E2E performance and data intensive science
- IDC protocol largely unchanged for 3 years
- ION service ~2 years old
- OSCARS 0.6 release (soon to be available for testing)

# OSCARS 0.6

- Simplified installation/configuration
- Wraps up foreseeable functional requirements from planned service offerings
  - Modest future development expected in support of additional switch/router hardware
- Additional monitoring functionality available
  - Second half of talk outlines this important functionality



# Previous View Of Circuits




- If a failure occurs, what can a user do?

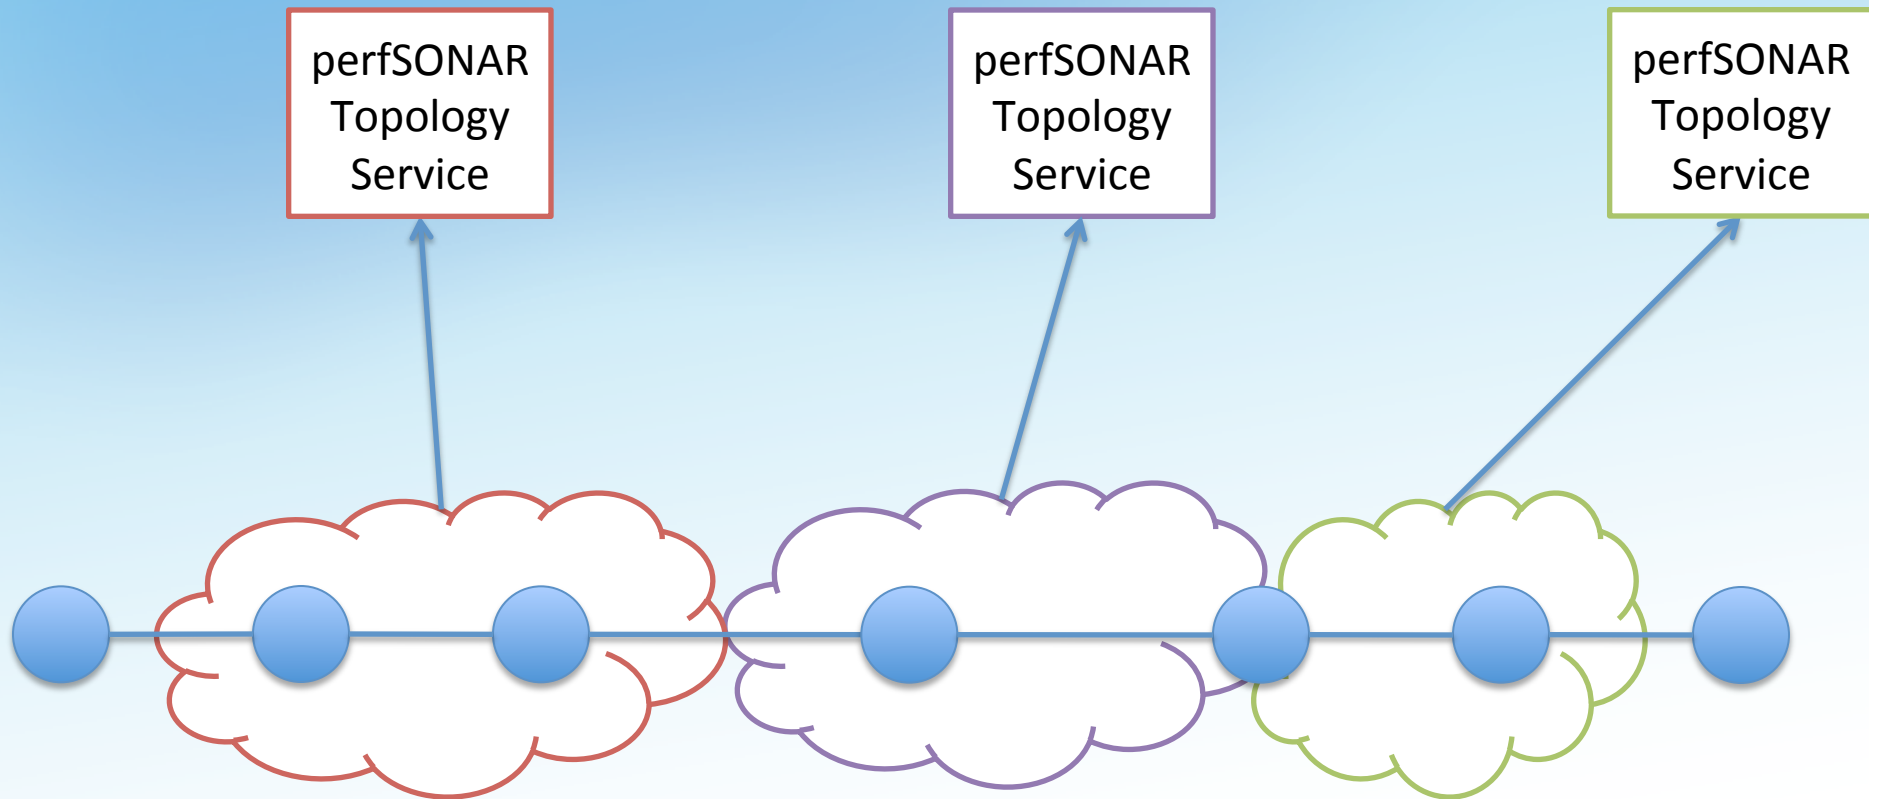
# Enabling Circuit Monitoring

- Goal: to enable users to get measurements in their circuits while allowing domains to provide as much or as little information to the user as the domain wants
- Develop a solution in collaboration with other groups and organizations including DANTE, ESnet, the Network Markup Language Working Group and the Network Measurement Control Working Group
  - Broad agreement ensures that users can monitor their circuits, no matter what domains they traverse
- Multi-faceted approach
  - Enable domains to export monitoring data about circuits
  - Enable users to discover the domains that make up their circuit, and the monitoring data those domains contain about the circuit
- Leverage the standard perfSONAR infrastructure when available

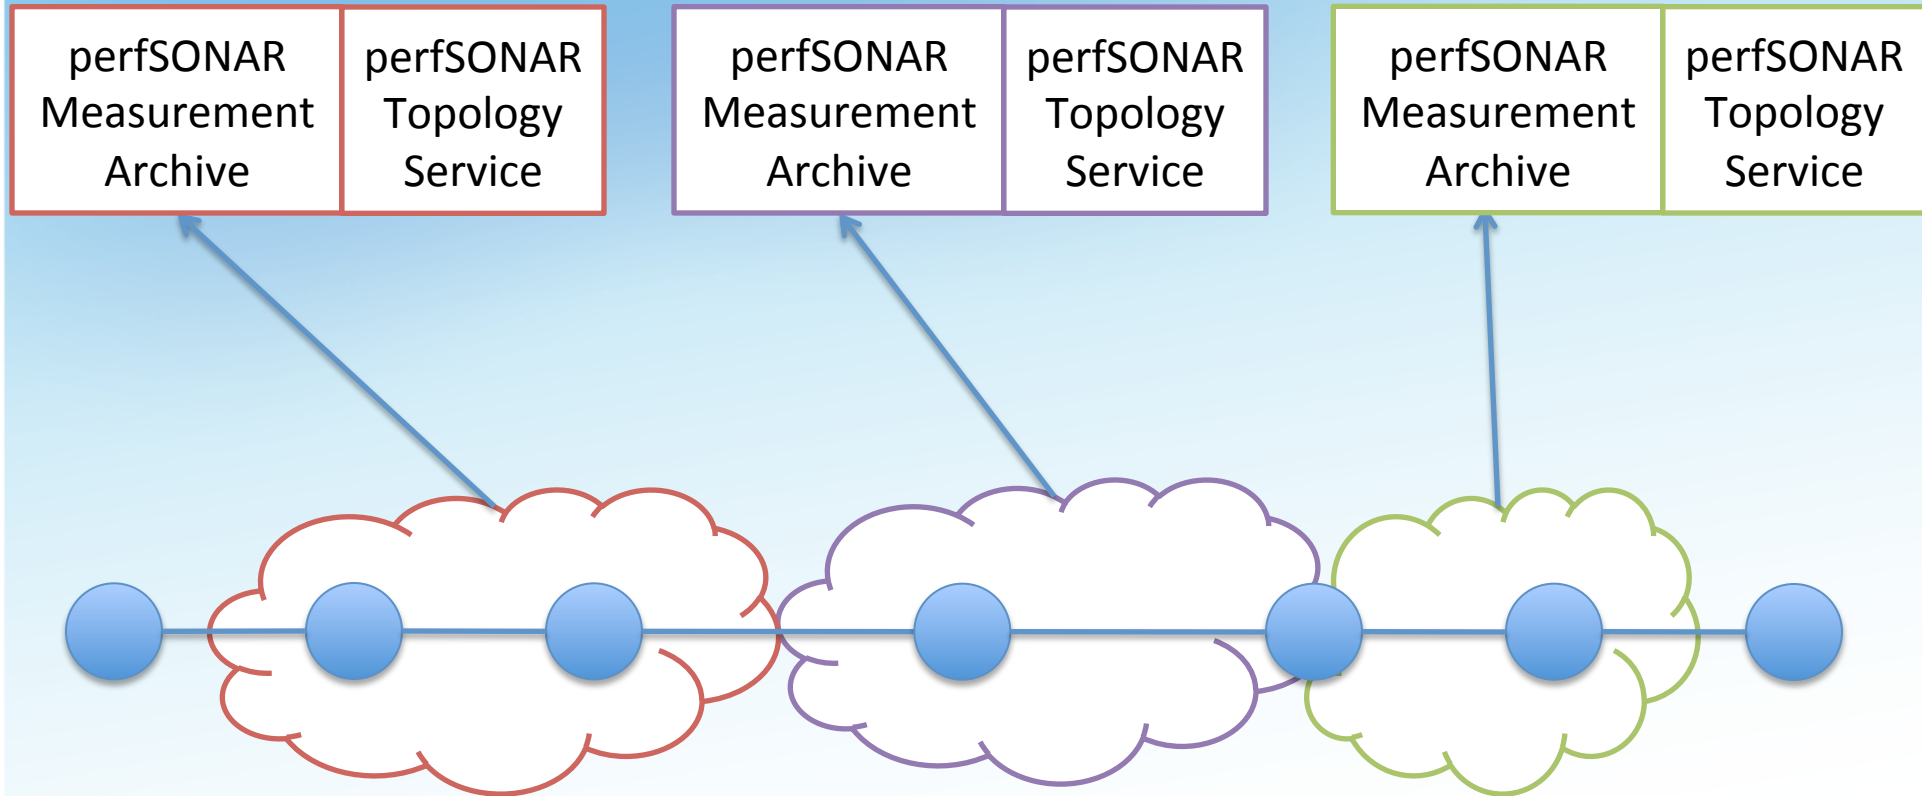
# Circuit Monitoring Architecture



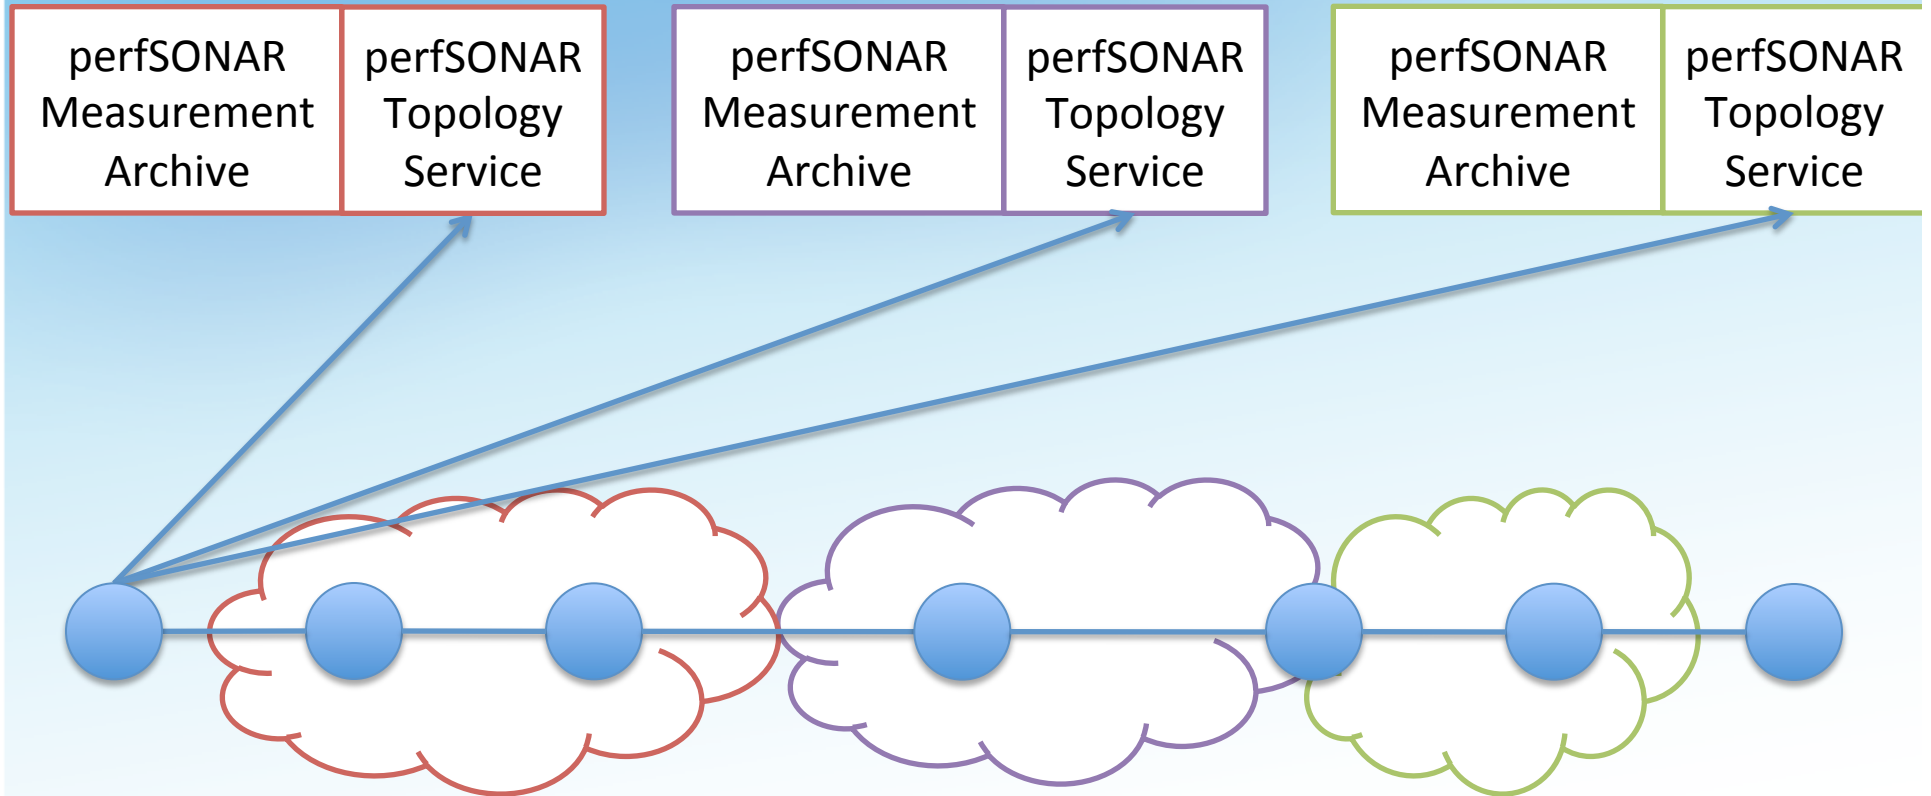
# Circuit Monitoring Architecture



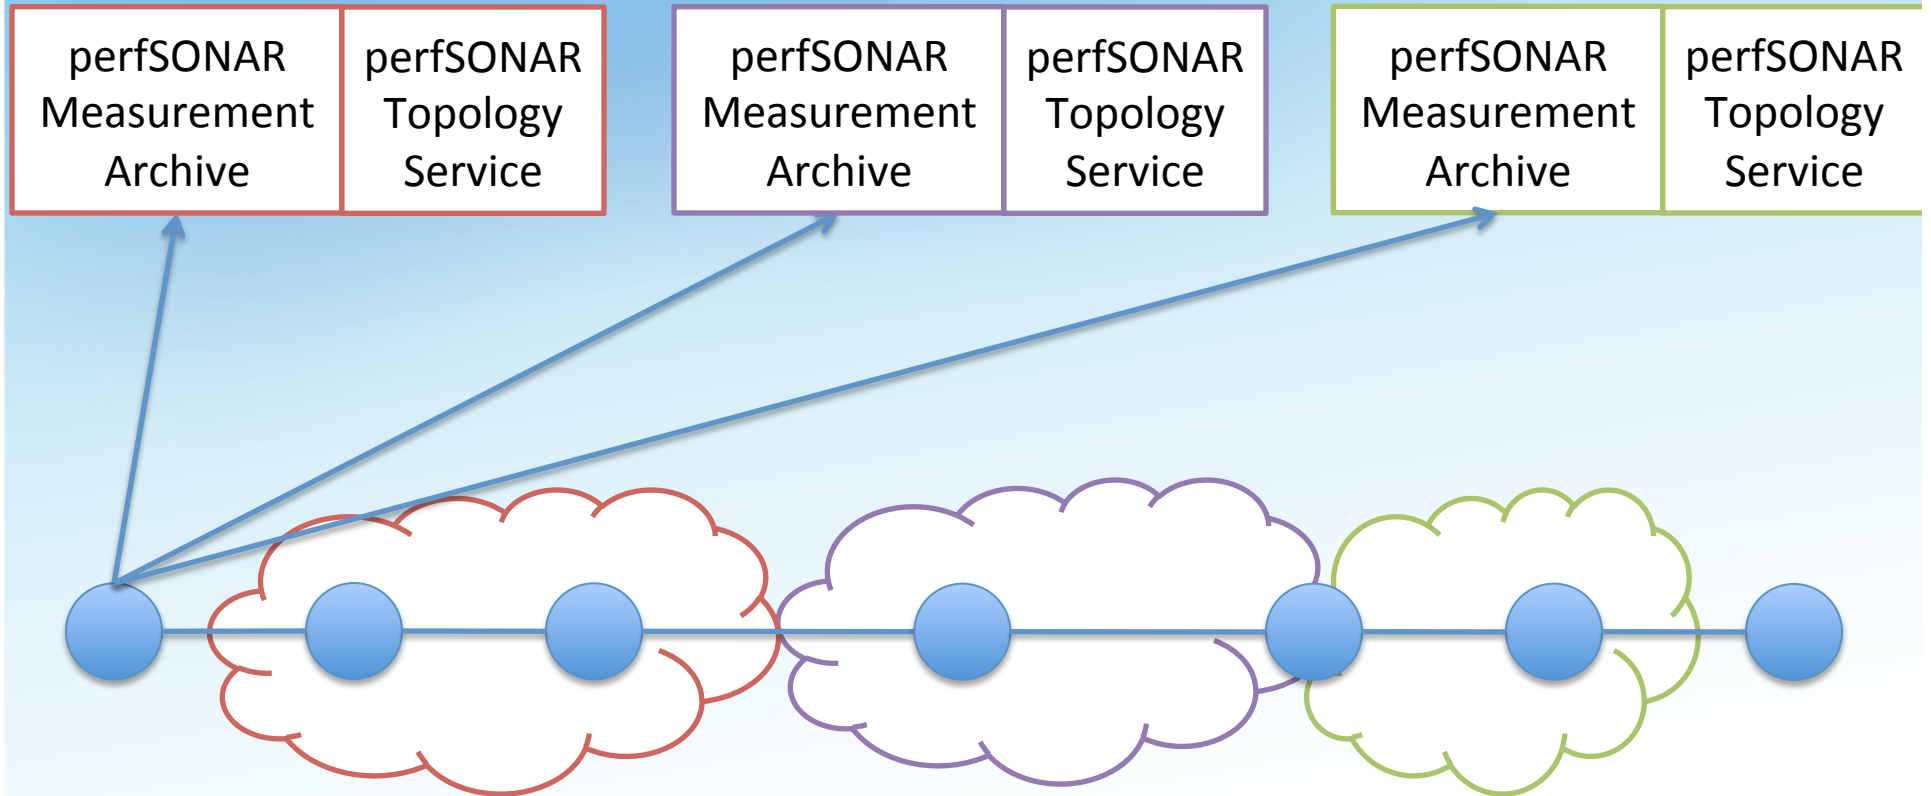
# Circuit Monitoring Architecture



# Circuit Monitoring Architecture



# Circuit Monitoring Architecture



# Implementation

- Two broad components
  - Domain Software Agent
    - Creates Domain's circuit descriptions
    - Registers these descriptions with a perfSONAR Topology Service
  - A Web UI client that displays statistics about a Domain's current and historical circuits



# Circuit Monitoring Agent

- This agent is the “glue” that connects together a Domain’s provisioning software (OSCARS) and monitoring infrastructure with the perfSONAR services so that users can find information about circuit statistics
- When new circuits are brought up, the agent looks at the intra-domain path for the circuit, and builds a description of that path.
  - This description is then registered into a perfSONAR Topology Service
- Needs to know how the domain monitors its devices to ensure an appropriate description of the circuit
  - If configured, the agent can use a user-defined script to start circuit monitoring

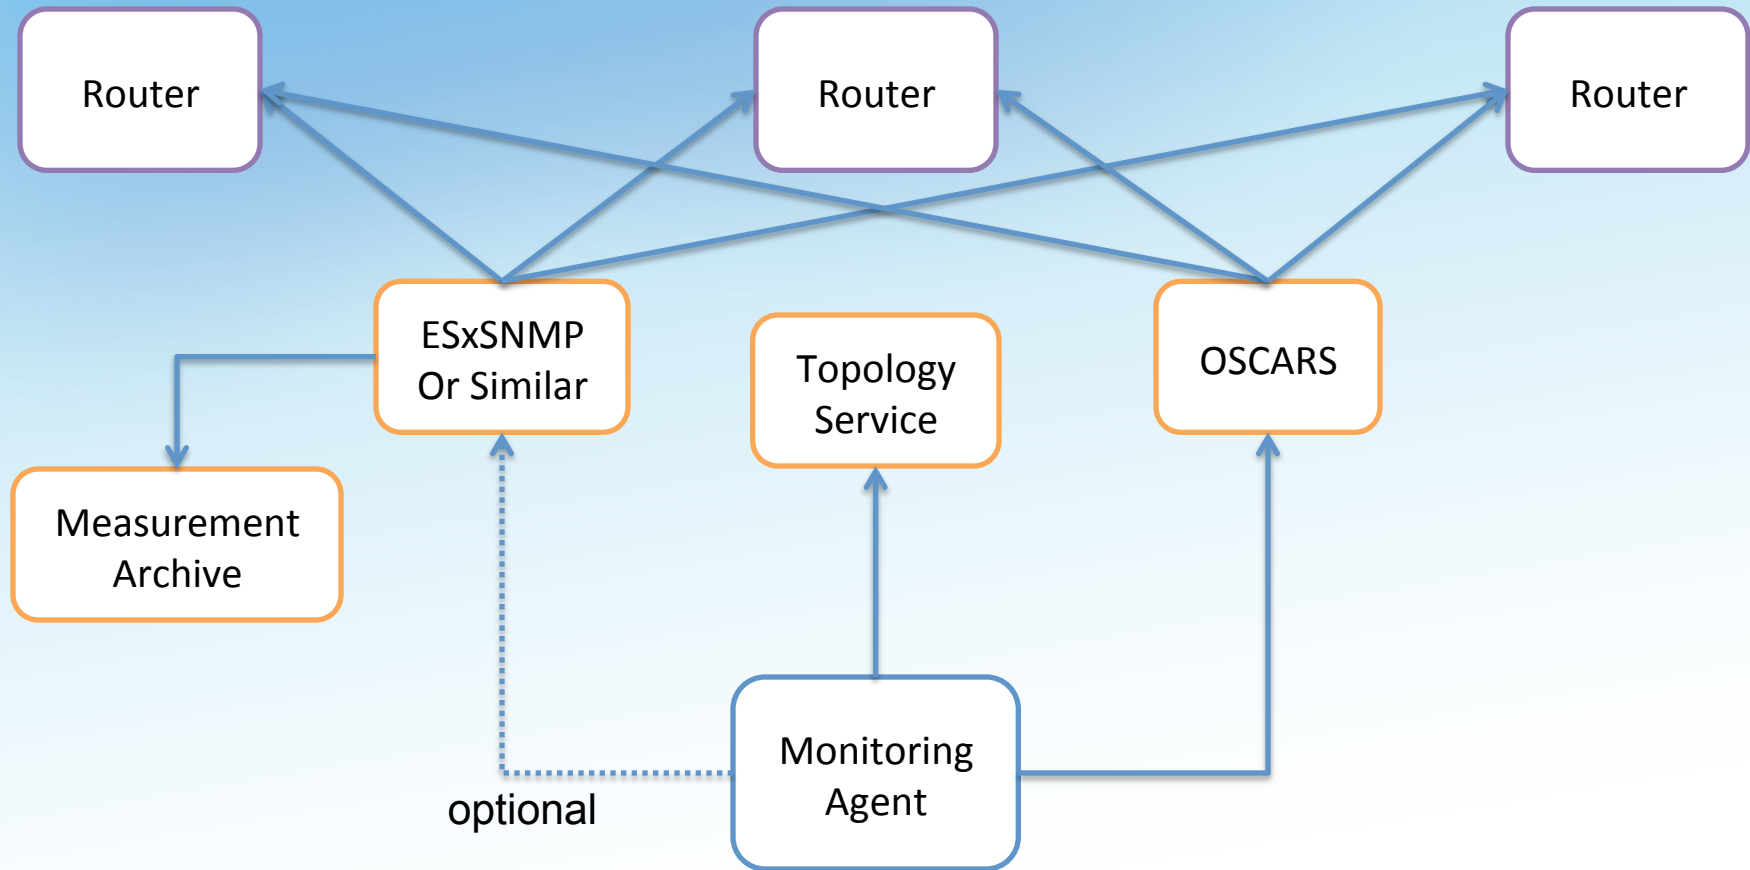
# Router/Switch Monitoring Component

- Everyone has their own method of monitoring their hardware
- Define the needed functionality instead of requiring a specific solution
  - Offer a specific solution to users who want to use it
- Requirements:
  - Software that can measure the operational status and utilization of the elements making up the circuit
  - These measurements are made available using standard perfSONAR protocols
- As long as the monitoring meets the above requirements, it can be made to work in the Circuit Monitoring infrastructure

# Router/Switch Monitoring Component

- Specific Solution: ESxSNMP
  - Developed by Jon Dugan at ESnet
  - Uses SNMP to monitor operational status and utilization statistics for all equipment elements, including physical interfaces, VLAN interfaces and LSPs
  - These interface statistics are then made available using the perfSONAR-PS SNMP MA
  - This software will be packaged for easy installation

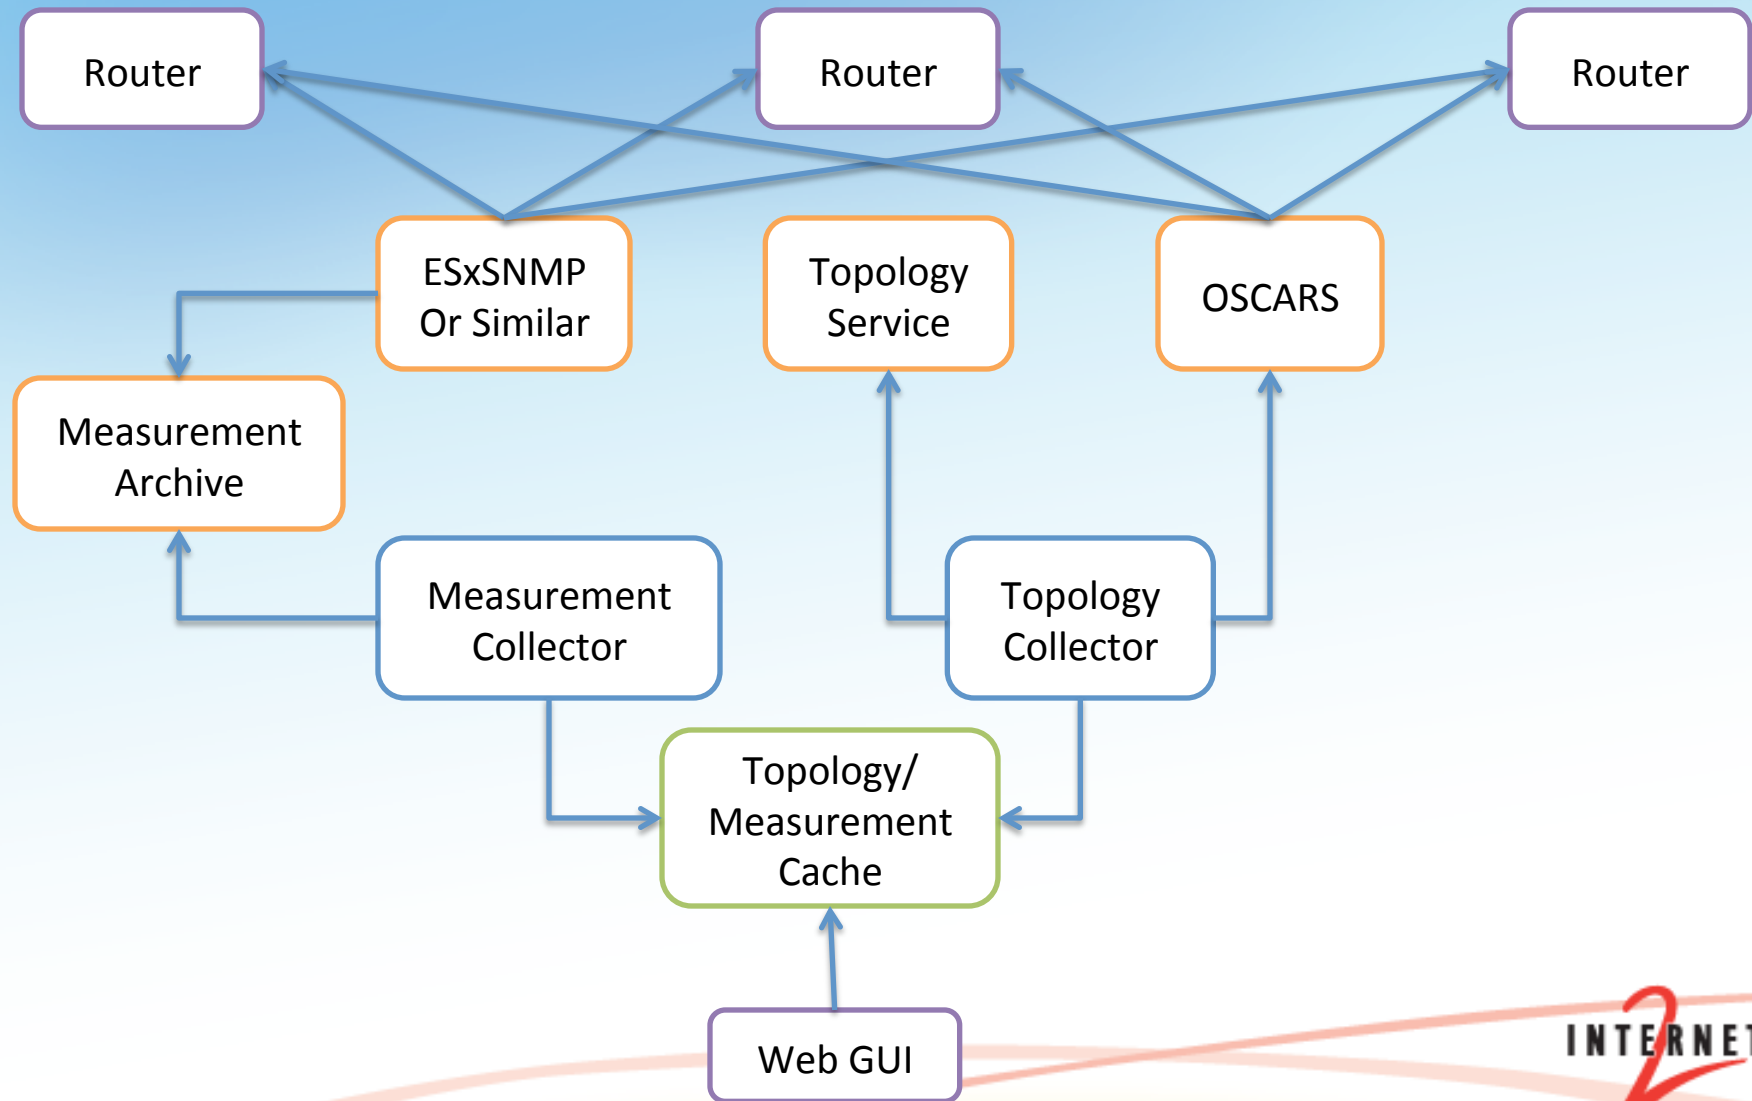
# Architecture - Monitoring Agent



# Web UI Client

- The distributed nature of the data doesn't mesh well with the near-instantaneous results expected by web users
- The Web UI includes collector daemons that lookup and retrieve topology and measurements about the end-to-end circuit from each domain, and store them in a local cache
- The Web UI then uses the cached information to display end-to-end circuit statistics to the end user

# Architecture – Web UI Client



# Web UI Client – Inter-Domain Path

## Circuit Monitoring

### Active Circuits

[caltech.edu-1001](#)  
[ion.internet2.edu-1048](#)

### Historical Circuits

[ion.internet2.edu-1047](#)

Circuit	
Name	caltech.edu-1001
Bandwidth	100
Status	ACTIVE
Start Time	Thu Jun 09 2011 14:18:00 GMT-0500 (CDT)
End Time	Thu Jun 09 2011 17:23:00 GMT-0500 (CDT)

Inter-Domain Path										
->	<a href="#">caltech.edu</a>	->	<a href="#">cenic.org</a>	->	<a href="#">ion.internet2.edu</a>	->	<a href="#">uslhcnet.org</a>	->	<a href="#">cern.ch</a>	->
<-	<a href="#">cern.ch</a>	<-	<a href="#">uslhcnet.org</a>	<-	<a href="#">ion.internet2.edu</a>	<-	<a href="#">cenic.org</a>	<-	<a href="#">caltech.edu</a>	<-

# Web UI Client – Intra-Domain Path

## Circuit Monitoring

### Active Circuits

[caltech.edu-1001](#)  
[ion.internet2.edu-1048](#)

### Historical Circuits

[ion.internet2.edu-1047](#)

Circuit	
Name	caltech.edu-1001
Bandwidth	100
Status	ACTIVE
Start Time	Thu Jun 09 2011 14:18:00 GMT-0500 (CDT)
End Time	Thu Jun 09 2011 17:23:00 GMT-0500 (CDT)

Inter-Domain Path										
->	<a href="#">caltech.edu</a>	->	<a href="#">cenic.org</a>	->	<a href="#">ion.internet2.edu</a>	->	<a href="#">uslhcnet.org</a>	->	<a href="#">cern.ch</a>	->
<-	<a href="#">cern.ch</a>	<-	<a href="#">uslhcnet.org</a>	<-	<a href="#">ion.internet2.edu</a>	<-	<a href="#">cenic.org</a>	<-	<a href="#">caltech.edu</a>	<-

Intra-Domain Path for ion.internet2.edu								
	rtr.newy32aoa.net.internet2.edu		rtr.losa.net.internet2.edu					
->	<a href="#">xe-0/1/3</a>	->	<a href="#">xe-0/0/0</a>	->	<a href="#">xe-0/1/3</a>	->	<a href="#">xe-1/0/1</a>	->



# Web UI Client – Intra-Domain Path

**Circuit Monitoring**

Active Circuits  
[caltech.edu-1001](#)  
[ion.internet2.edu-1048](#)

Historical Circuits  
[ion.internet2.edu-1047](#)

Circuit	
Name	caltech.edu-1001
Bandwidth	100
Status	ACTIVE
Start Time	Thu Jun 09 2011 14:18:00 GMT-0500 (CDT)
End Time	Thu Jun 09 2011 17:23:00 GMT-0500 (CDT)

Inter-Domain Path

->	<a href="#">caltech.edu</a>	->	<a href="#">cenic.org</a>	->	<a href="#">ion.internet2.edu</a>	->	<a href="#">uslhcnet.org</a>	->	<a href="#">cern.ch</a>	->
<	<a href="#">cern.ch</a>	<	<a href="#">uslhcnet.org</a>	<	<a href="#">ion.internet2.edu</a>	<	<a href="#">cenic.org</a>	<	<a href="#">caltech.edu</a>	<

Intra-Domain Path for ion.internet2.edu

rtr.newy32aoa.net.internet2.edu	rtr.losa.net.internet2.edu							
->	<a href="#">xe-0/1/3</a>	->	<a href="#">xe-0/0/0</a>	->	<a href="#">xe-0/1/3</a>	->	<a href="#">xe-1/0/1</a>	->

Utilization at port xe-0/0/0 on host rtr.newy32aoa.net.internet2.edu

## Final thought

Dynamic Network Services are Maturing and the community's previous development and deployment experiences can be leveraged to create new services to support better E2E performance and data intensive science



## **Dynamic Network Services**

April 20<sup>th</sup> 2011, Internet2 Spring Member Meeting

Jeff Boote, Internet2 Assistant Director R&D

Aaron Brown, Internet2 Network Software Engineer

For more information, visit <http://www.internet2.edu/performance>