



# Internet2 E2E piPEs

End-to-End Performance Initiative  
Performance Improvement System

Eric L. Boyd

Internet2

13 April 2004

- What is piPEs?
- Goals
- E2E piPEs Measurement Infrastructure
- Abilene Measurement Domain
- Data Analysis Status
- LA <-> CERN Demo



# Internet2 E2E piPEs

- Project: End-to-End Performance Initiative Performance Environment System (E2E piPEs)
- Approach: Collaborative project combining the best work of many organizations, including DANTE/GEANT, Daresbury, EGEE, Georgia Tech, GGF NMWG, ITECs, MonALISA, NLANR/DAST, UCL, etc.
- Reaching out to many other groups including AMP, CENIC, ESNET, LBL, SLAC, PSC, UDel, etc.

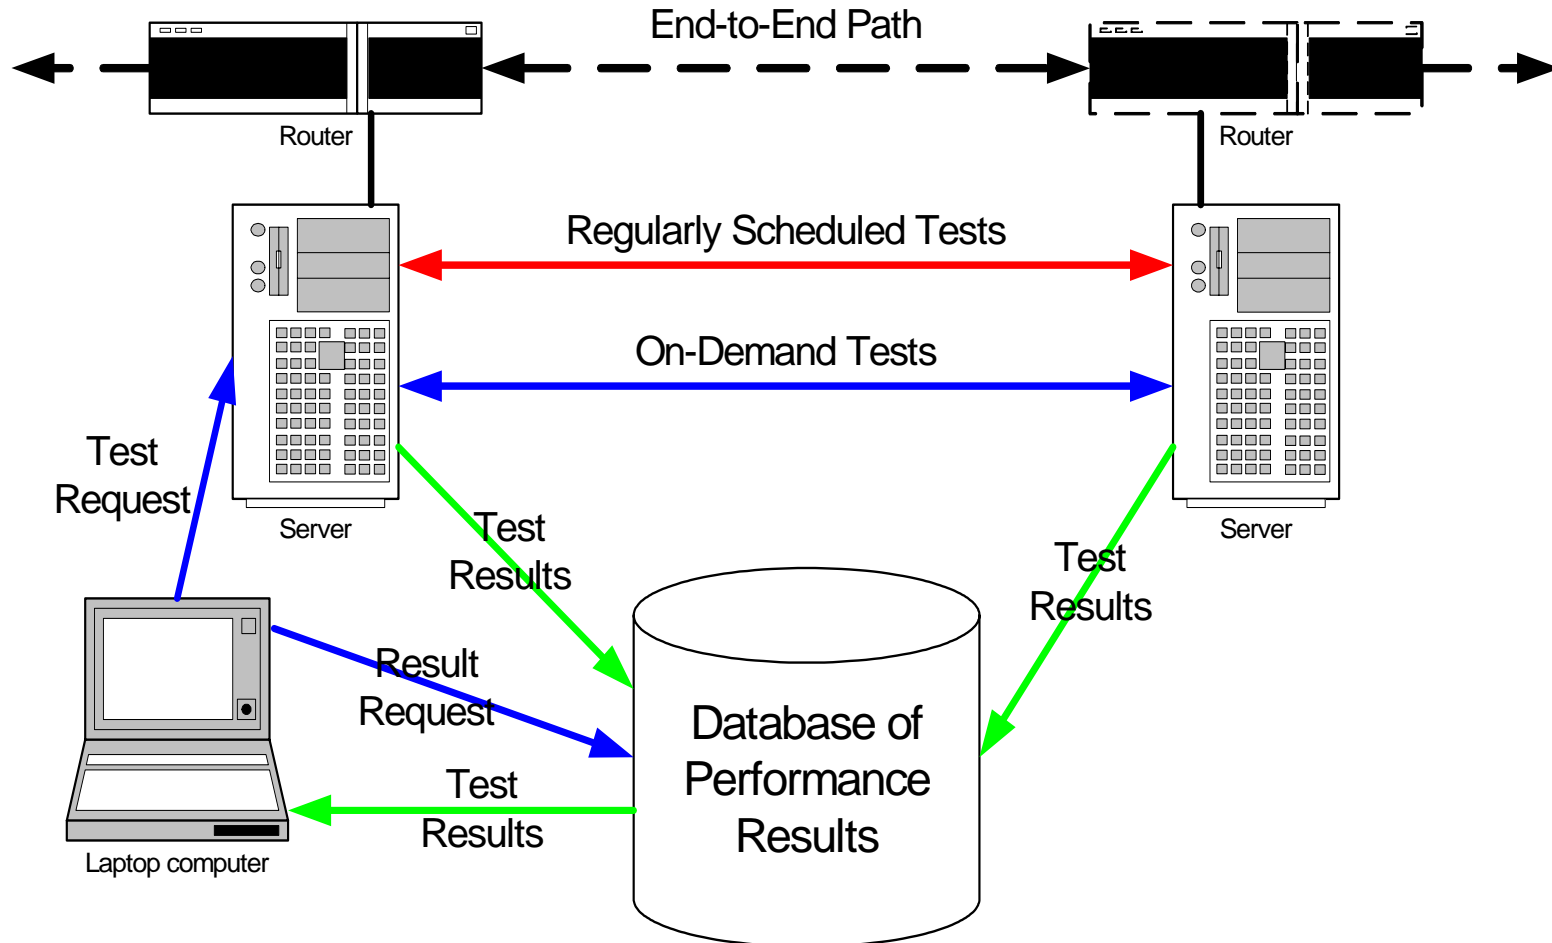


# Internet2 E2E piPEs Goals

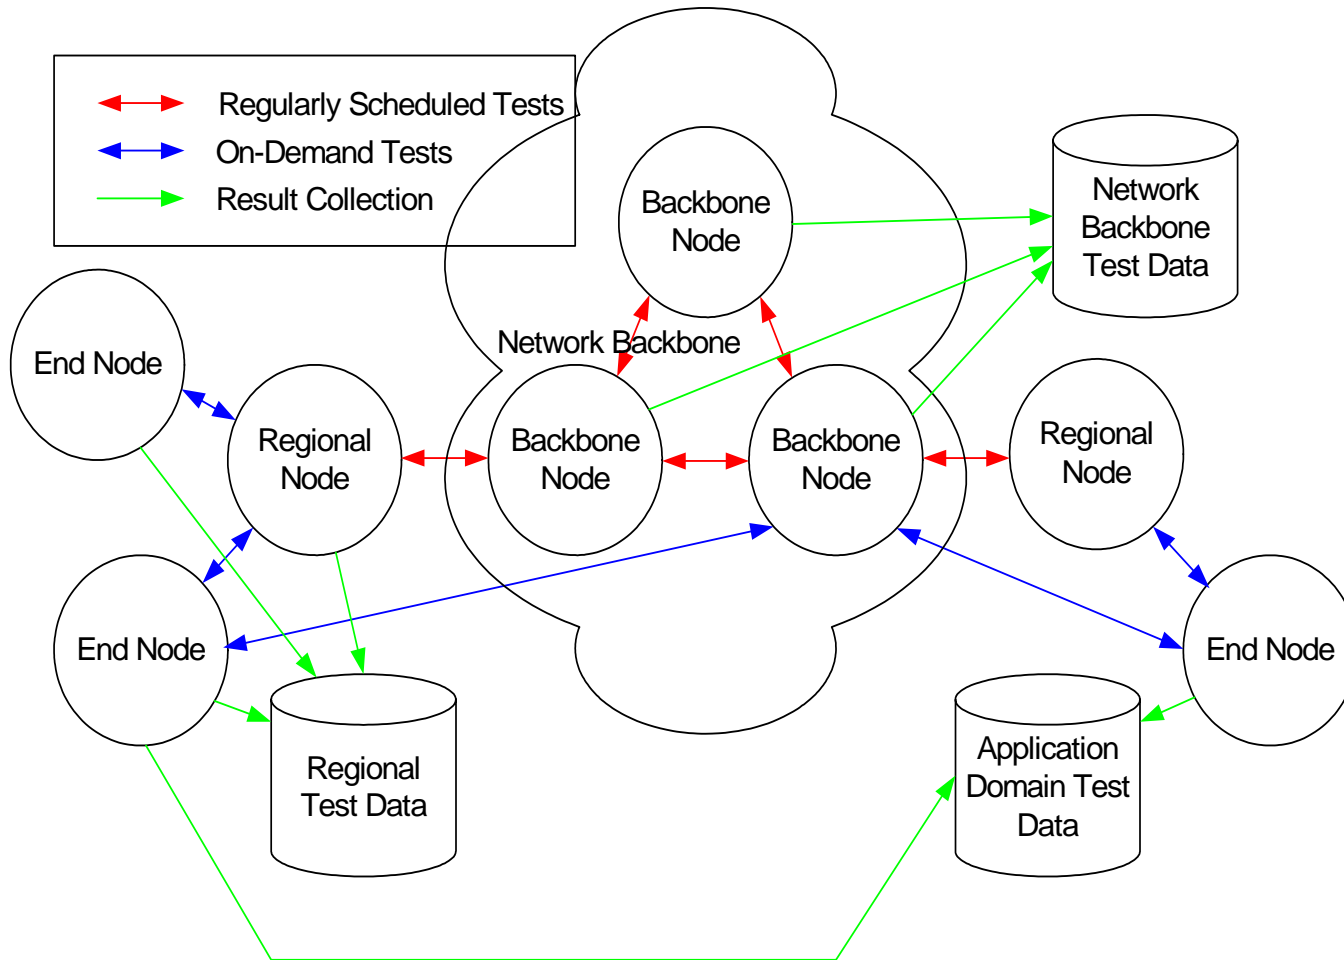
- Enable end-users & network operators to:
  - determine E2E performance capabilities
  - locate E2E problems
  - contact the right person to get an E2E problem resolved.
- Enable remote initiation of partial path performance tests
- Make partial path performance data publicly available
- Interoperable with other performance measurement frameworks



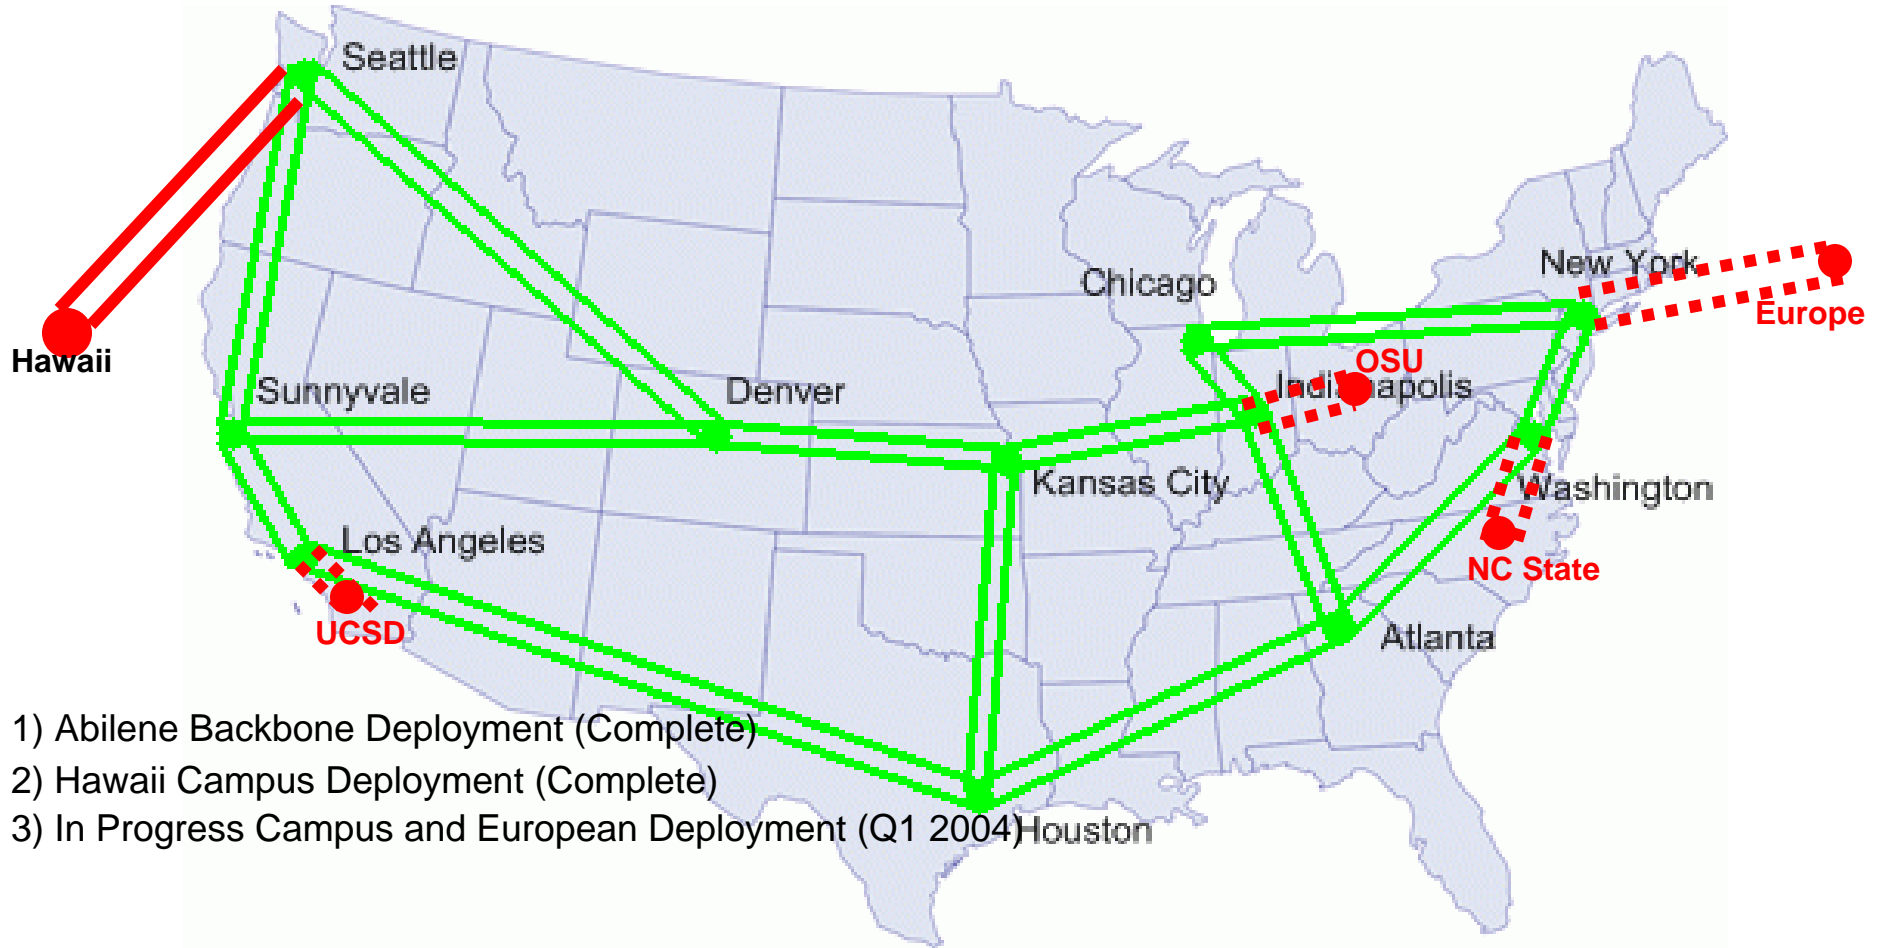
# Measurement Infrastructure Components



# Sample piPEs Deployment



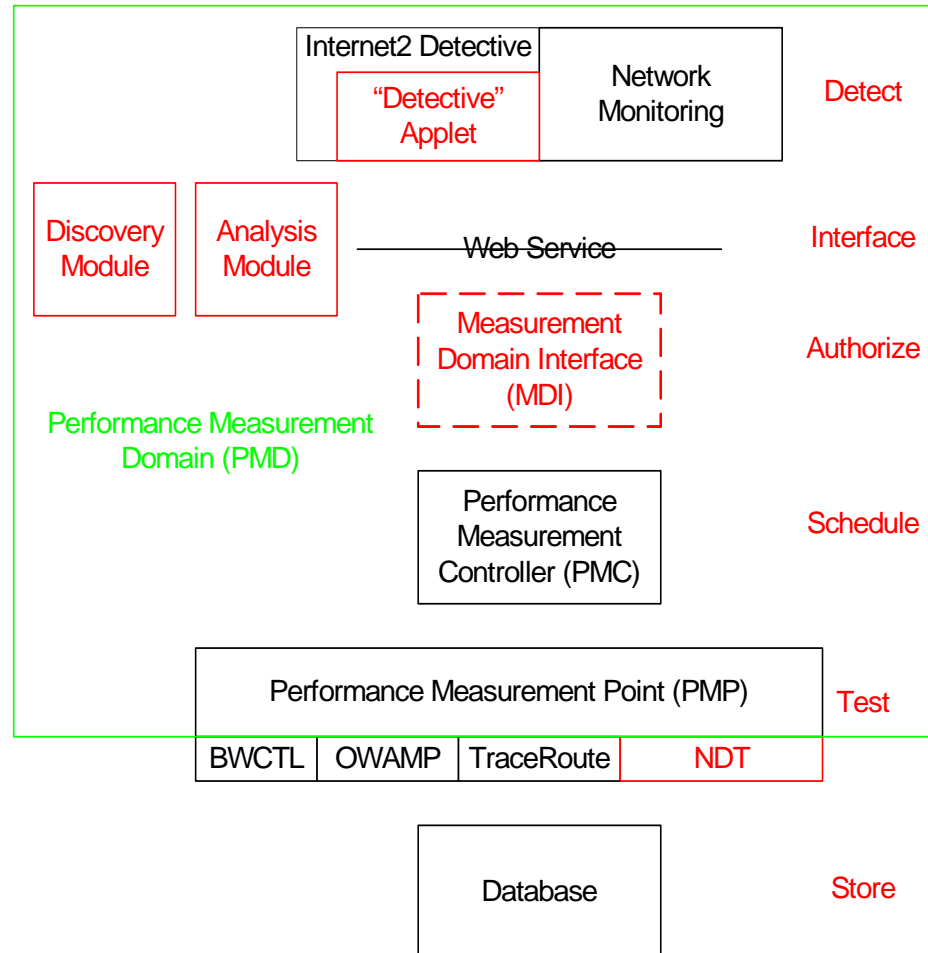
# piPEs Deployment



- 1) Abilene Backbone Deployment (Complete)
- 2) Hawaii Campus Deployment (Complete)
- 3) In Progress Campus and European Deployment (Q1 2004)



# Measurement Software Components







# Abilene Measurement Domain

- Part of the Abilene Observatory:  
<http://abilene.internet2.edu/observatory>
- Regularly scheduled OWAMP (1-way latency) and BWCTL (Iperf wrapper) Tests
- Web pages displaying:
  - Latest results [http://abilene.internet2.edu/ami/bwctl\\_status.cgi/TCP/now](http://abilene.internet2.edu/ami/bwctl_status.cgi/TCP/now)  
“Weathermap”  
[http://abilene.internet2.edu/ami/bwctl\\_status\\_map.cgi/TCP/now](http://abilene.internet2.edu/ami/bwctl_status_map.cgi/TCP/now)
  - Worst 10 Performing Links  
[http://abilene.internet2.edu/ami/bwctl\\_worst\\_case.cgi/TCP/now](http://abilene.internet2.edu/ami/bwctl_worst_case.cgi/TCP/now)
- Data available via web service:  
<http://abilene.internet2.edu/ami/webservices.html>

- E2E piPEs Overview
  - What is piPEs?
  - Goals
  - E2E piPEs Measurement Infrastructure
  - Abilene Measurement Domain
- **Data Analysis Status**
- Preliminary Data Discovery Approach
- AMI Web Service

# Data Collection / Correlation

- Collection Today:
  - Iperf (Throughput)
  - OWAMP (1-Way Latency, Loss)
  - SNMP Data
  - Anonymized Netflow Data
  - Per Sender, Per Receiver, Per Node Pair
  - IPv4 and IPv6
- Collection in the Future
  - NTP (Data)
  - Traceroute
  - BGP Data
  - First Mile Analysis
- Correlation Today:
  - “Worst 10” Throughputs
  - “Worst 10” Latencies
- Correlation in the Future:
  - 99<sup>th</sup> Percentile Throughput over Time
  - Throughput/Loss for all E2E paths using a specific link
  - Commonalities among first mile analyzers
  - Sum of Partial Paths vs. Whole Path

# Data Analysis

- Analysis Today:
  - Throughput over Time
  - Latency over Time
  - Loss over Time
  - Worrisome Tests? (Any bad apples in “Worst Ten”?)
  - “Not the Network” (If “Worst Ten” is good enough)
- Analysis in the Future:
  - Latency vs. Loss
  - How good is the network?
  - Do common first mile problems exist?
  - Does a link have problems that only manifest in the long-haul?
  - Is the network delivering the performance required by a funded project?

# Data Discovery / Interoperability

## ■ Discovery in the Future:

- Where are the measurement nodes corresponding to a specific node?
- Where are the test results for a specific partial path?

## ■ Interoperability in the Future:

- Can I have a test within or to another measurement framework?
- Can I have a measurement result from within or to another measurement framework?



# American/European Collaboration Goals

- Awareness of ongoing Measurement Framework Efforts / Sharing of Ideas (Good / Not Sufficient)
- Interoperable Measurement Frameworks (Minimum)
  - Common means of data extraction
  - Partial path analysis possible along transatlantic paths
- Open Source Shared Development (Possibility, In Whole or In Part)
- End-to-end partial path analysis for transatlantic research communities
  - VLBI: Onsala, Sweden  $\leftrightarrow$  Haystack, Mass.
  - HENP: CERN, Switzerland  $\leftrightarrow$  Caltech, Calif.



# American/European Demonstration Goals

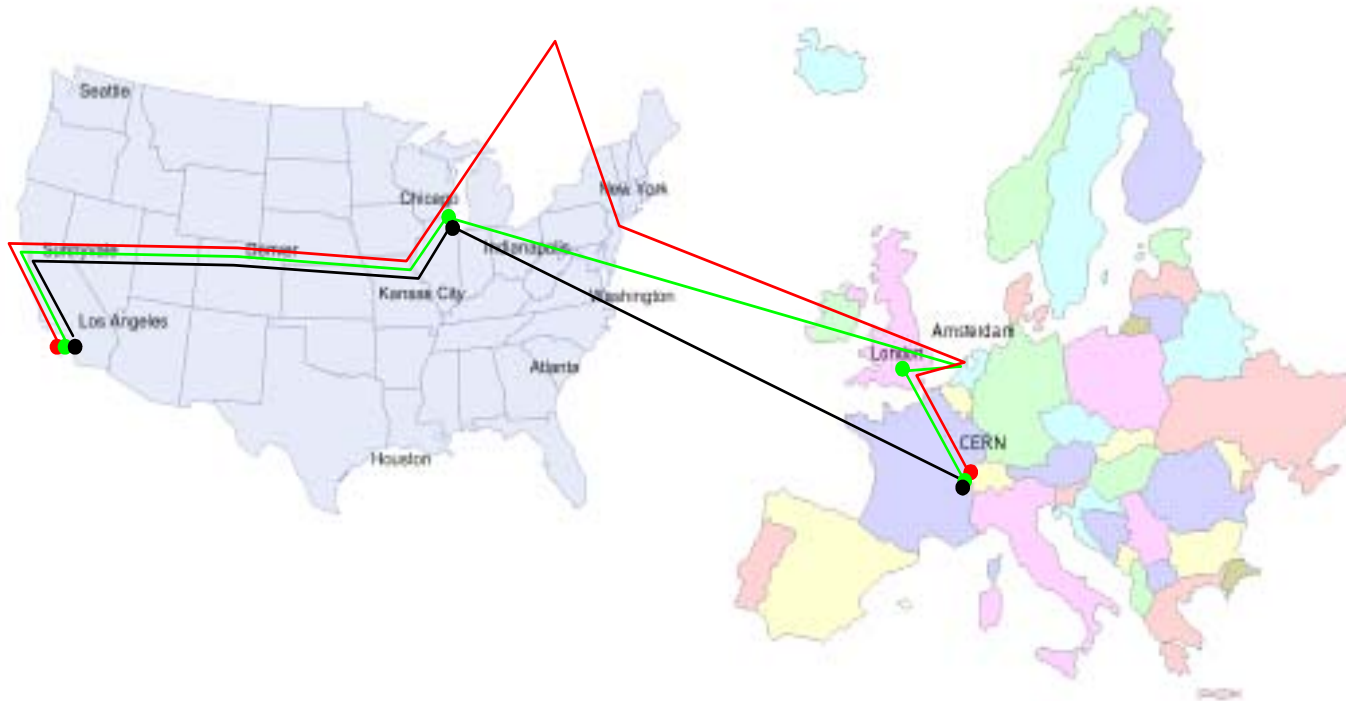
- Demonstrate ability to do partial path analysis between “Caltech” (Los Angeles Abilene router) and CERN.
- Demonstrate ability to do partial path analysis involving nodes in the GEANT network.
- Compare and contrast measurement of a “lightpath” versus a normal IP path.
- Demonstrate interoperability of piPEs and analysis tools such as Advisor and MonALISA

# Demonstration Details

- Path 1: Default route between LA and CERN is across Abilene to Chicago, then across Datatag circuit to CERN
- Path 2: Announced addresses so that route between LA and CERN traverses GEANT via London node
- Path 3: “Lightpath” (discussed earlier by Rick Summerhill)
- Each measurement “node” consists of a BWCTL box and an OWAMP box “next to” the router.



# All Roads Lead to Geneva



Path 1 — DataTag — Default Route

Path 2 — Eurolink — "Cooked" Alternate Route

Path 3 — Lightpath — "Cooked" Alternate Route

Circles Correspond to OWAMP / BWCTL Measurement Node Pair

- **BWCTL:**

[http://abilene.internet2.edu/ami/bwctl\\_status\\_eu.cgi/BW/14123130651515289600\\_14124243902743445504](http://abilene.internet2.edu/ami/bwctl_status_eu.cgi/BW/14123130651515289600_14124243902743445504)

- **OWAMP:**

[http://abilene.internet2.edu/ami/owamp\\_status\\_eu.cgi/14123130651515289600\\_14124243902743445504](http://abilene.internet2.edu/ami/owamp_status_eu.cgi/14123130651515289600_14124243902743445504)

- **MONALISA**

- **NLANR Advisor**

The logo features a large, stylized red number '2' that is partially overlaid by the word 'INTERNET'. The '2' starts from the bottom left, curves upwards and to the right, then loops back down and to the left, ending under the 'E' in 'INTERNET'.

**INTERNET**®

[www.internet2.edu](http://www.internet2.edu)

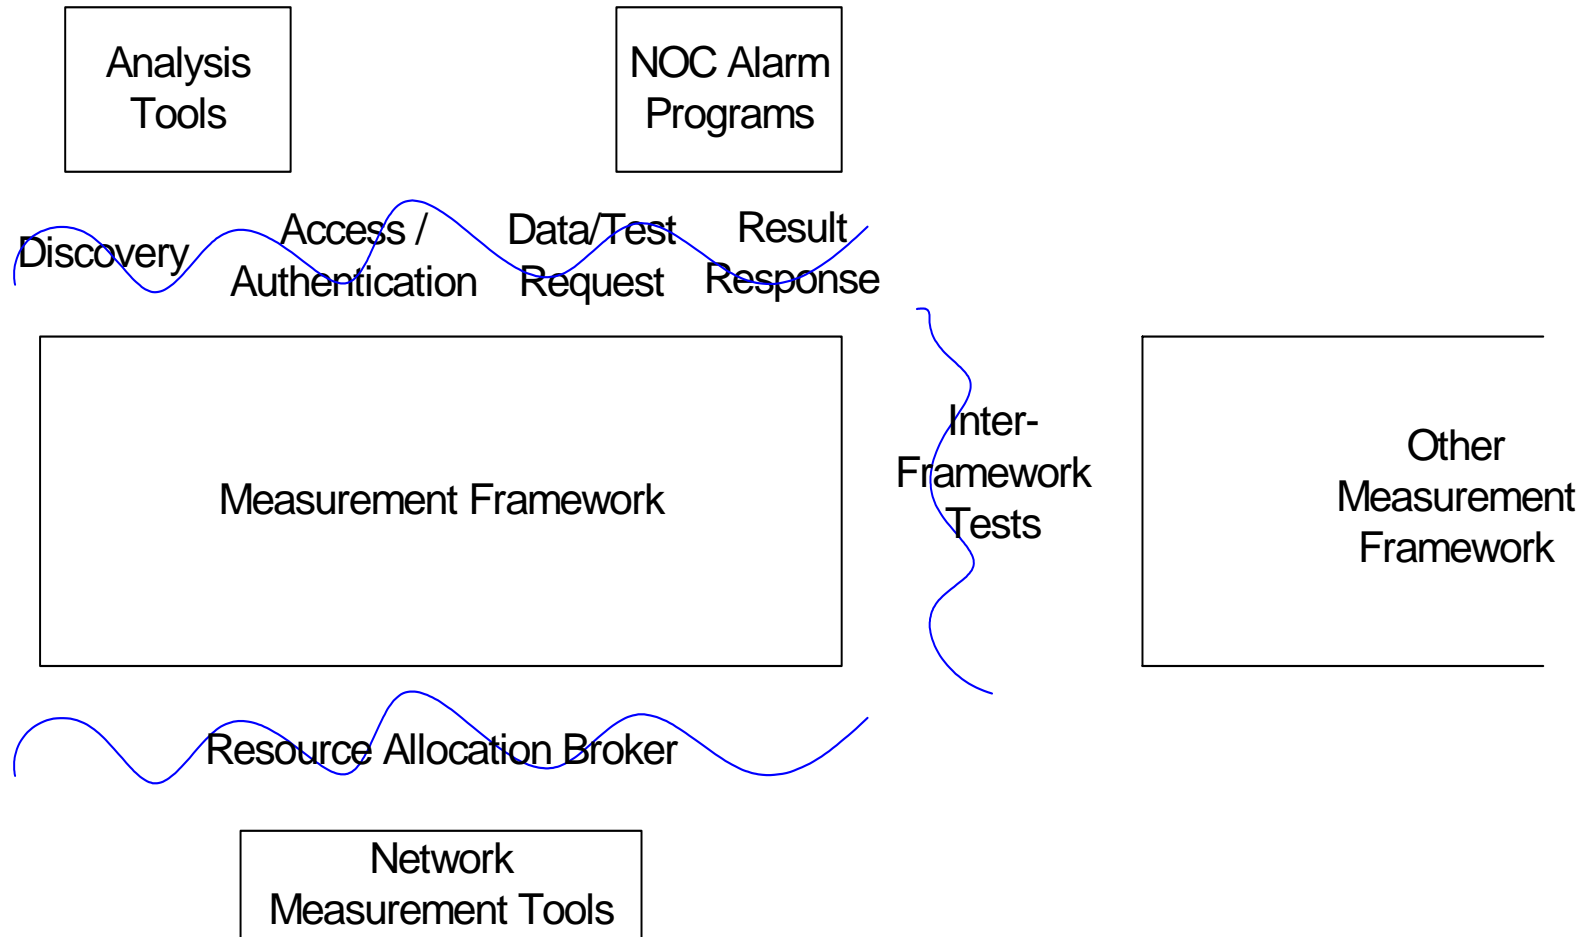


# Measurement Infrastructure Federation

- Why a Federation?
  - Multiple measurement frameworks in existence and under development (piPEs, NLANR Advisor, NLANR AMP, etc.).
  - No static “best practice” measurement framework is likely to emerge, given academics being academics.
  - Future measurement frameworks can build on shoulders of current efforts, not feet.
- Performance Measurement Architecture Workshop (NSF Grant # ANI-0314723)



# Measurement Infrastructure Federation Interfaces





# Measurement Infrastructure Federation Requirements

- Agreement on Characteristic Names
- Access and Authentication
- Discovery (Measurement Frameworks, Domains, Nodes, Databases)
- Test/Data Request Schema
- Result Report Schema
- Inter-Framework Tests
- Resource Allocation Broker for Tools
- Concatenation of Homogeneous Characteristics Results Gathered by Heterogeneous Tools



# GGF Network Measurement Working Group

- Hierarchy of Network Performance Characteristics
- Request Schema Requirements and Sample Implementation
- Report Schema Requirements and Sample Implementation





# Establishing a Performance Measurement Mesh

Issues include:

- Scheduling in the presence of scarce resources
- Making the tool bidirectional
- Adding security
- Ensuring correct source/target pairs
- Data collection / mining / analysis / display

Example:

- BWCTL for Iperf plus prototype PMD



# Open Research Issues

- Access and Authentication
- Discovery of Measurement Nodes (“Super-Traceroute”)
- Discovery of Measurement Databases
- Inter-framework Testing
- Compilation of results on partial paths
- Normalization of identical characteristics gathered by heterogenous tools