



Internet2 End-to-End Performance Monitoring Initiative Update

Eric L. Boyd, Internet2

May 21, 2003



Agenda

- **Motivation**
- piPEs / AMI Overview
- piPEs / AMI Deployment Plan
- OWAMP Release
- Demo

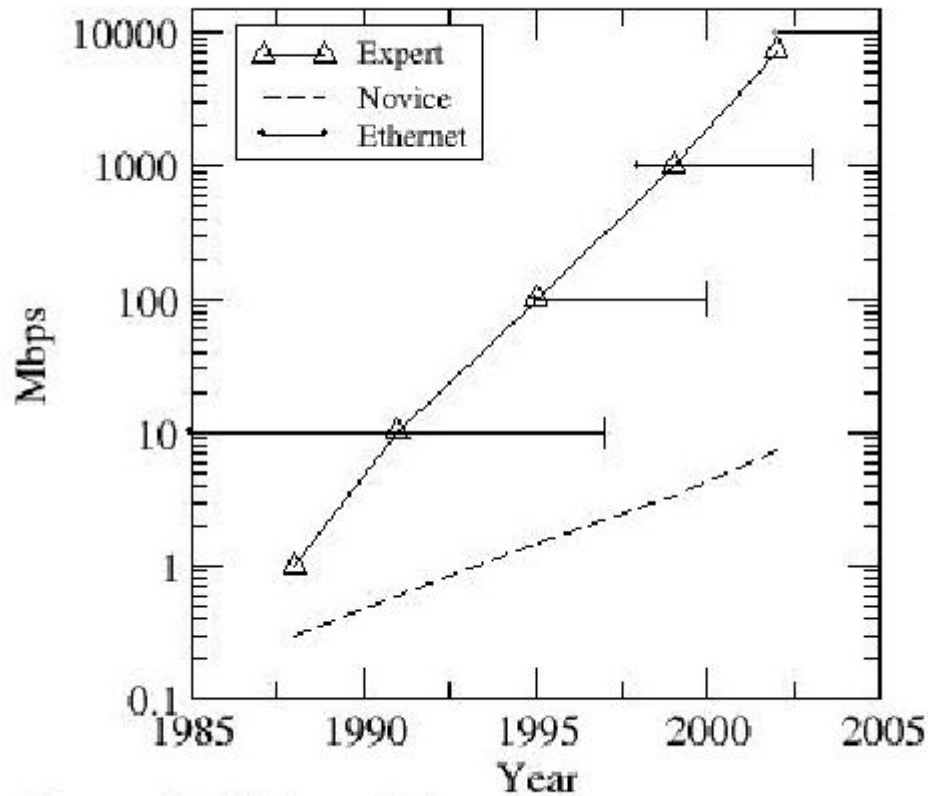


Context for E2E Performance

Even with high bandwidth network links, the Internet2 community often does not see expected performance.

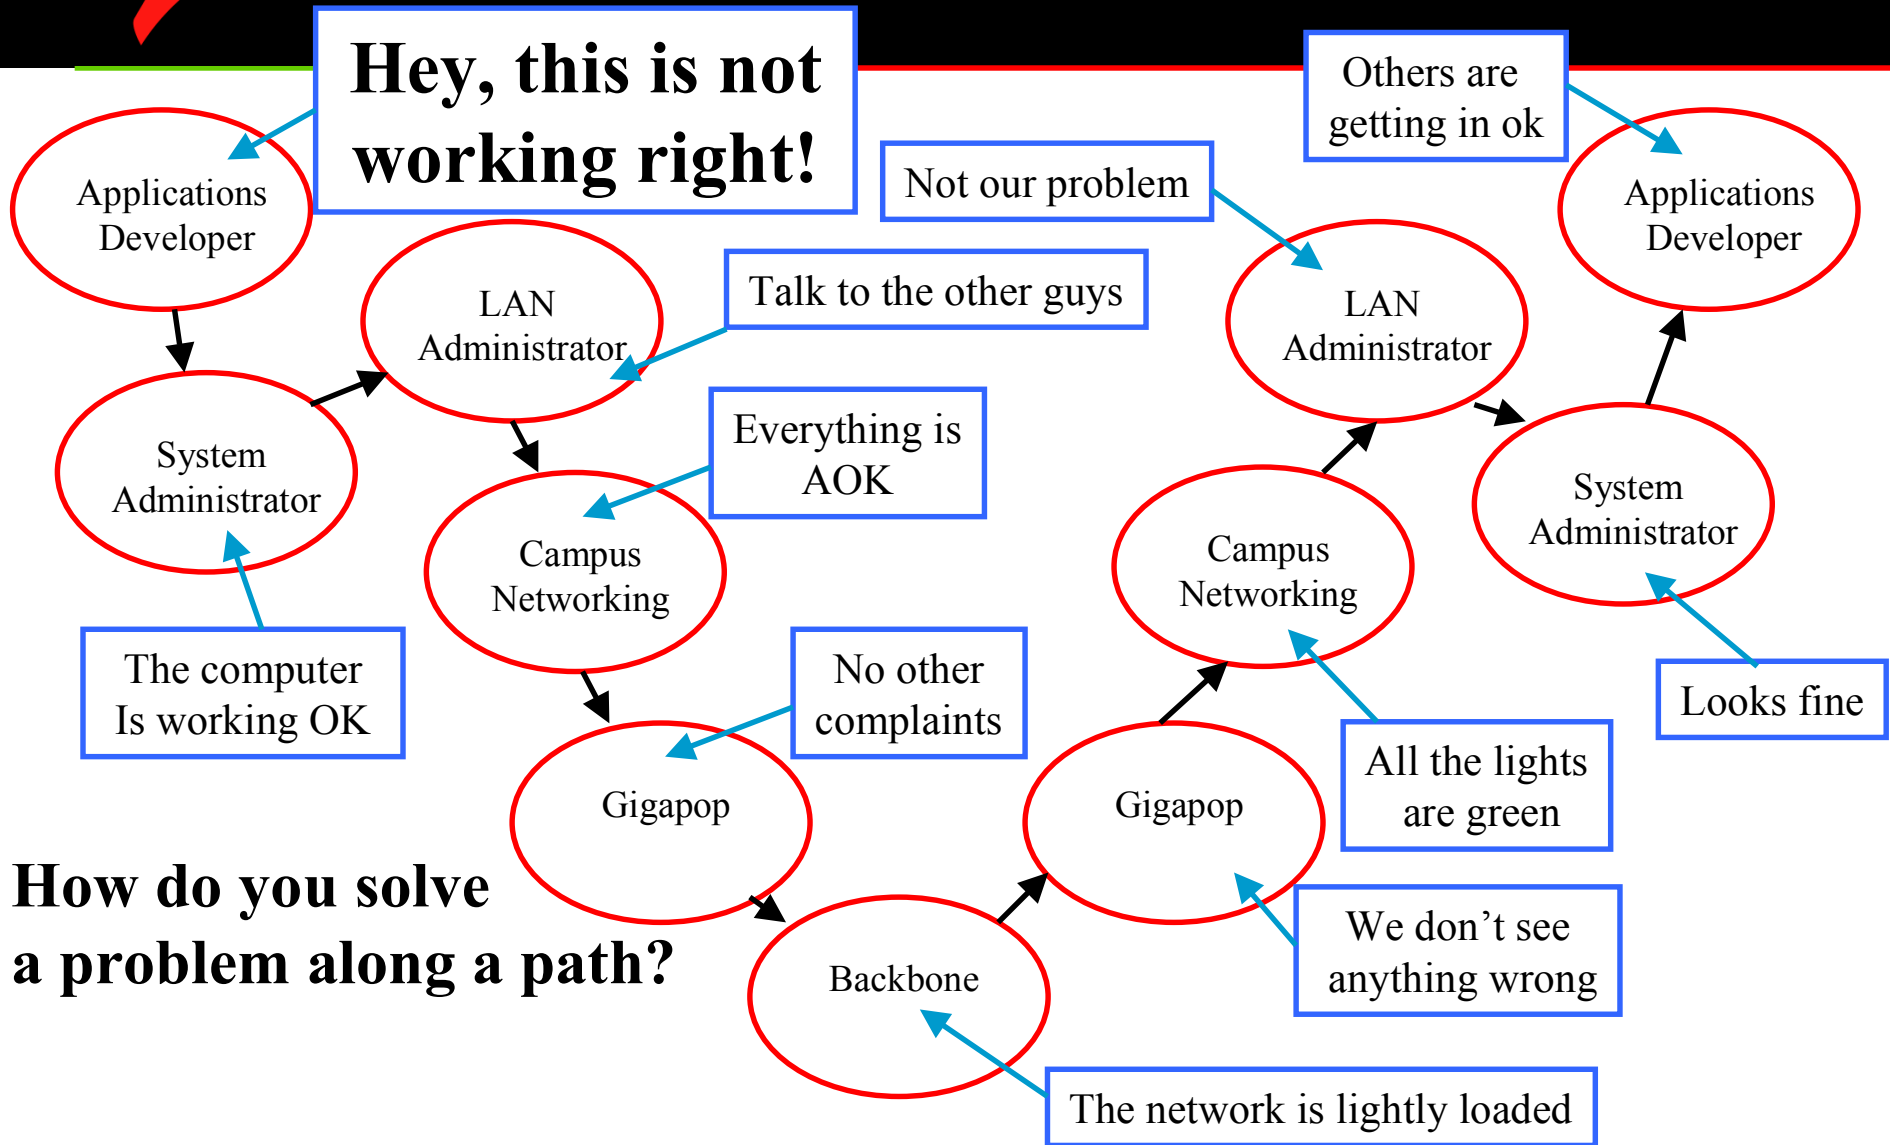


The Wizard Gap



Courtesy Matt Mathis, mathis@psc.edu

A Problem



How do you solve a problem along a path?



Good Work Already Going On

Many projects and research is underway by others investigating performance on portions of the end-to-end path.

E2Epi seeks to bring these efforts together to look at the end-to-end path as a complete system.



Can You Go Direct to the Problem?

How can you tell where is the problem?

Need a tool to tell you:

- Where the problem is.
- The type of problem
- Who to contact to get it fixed

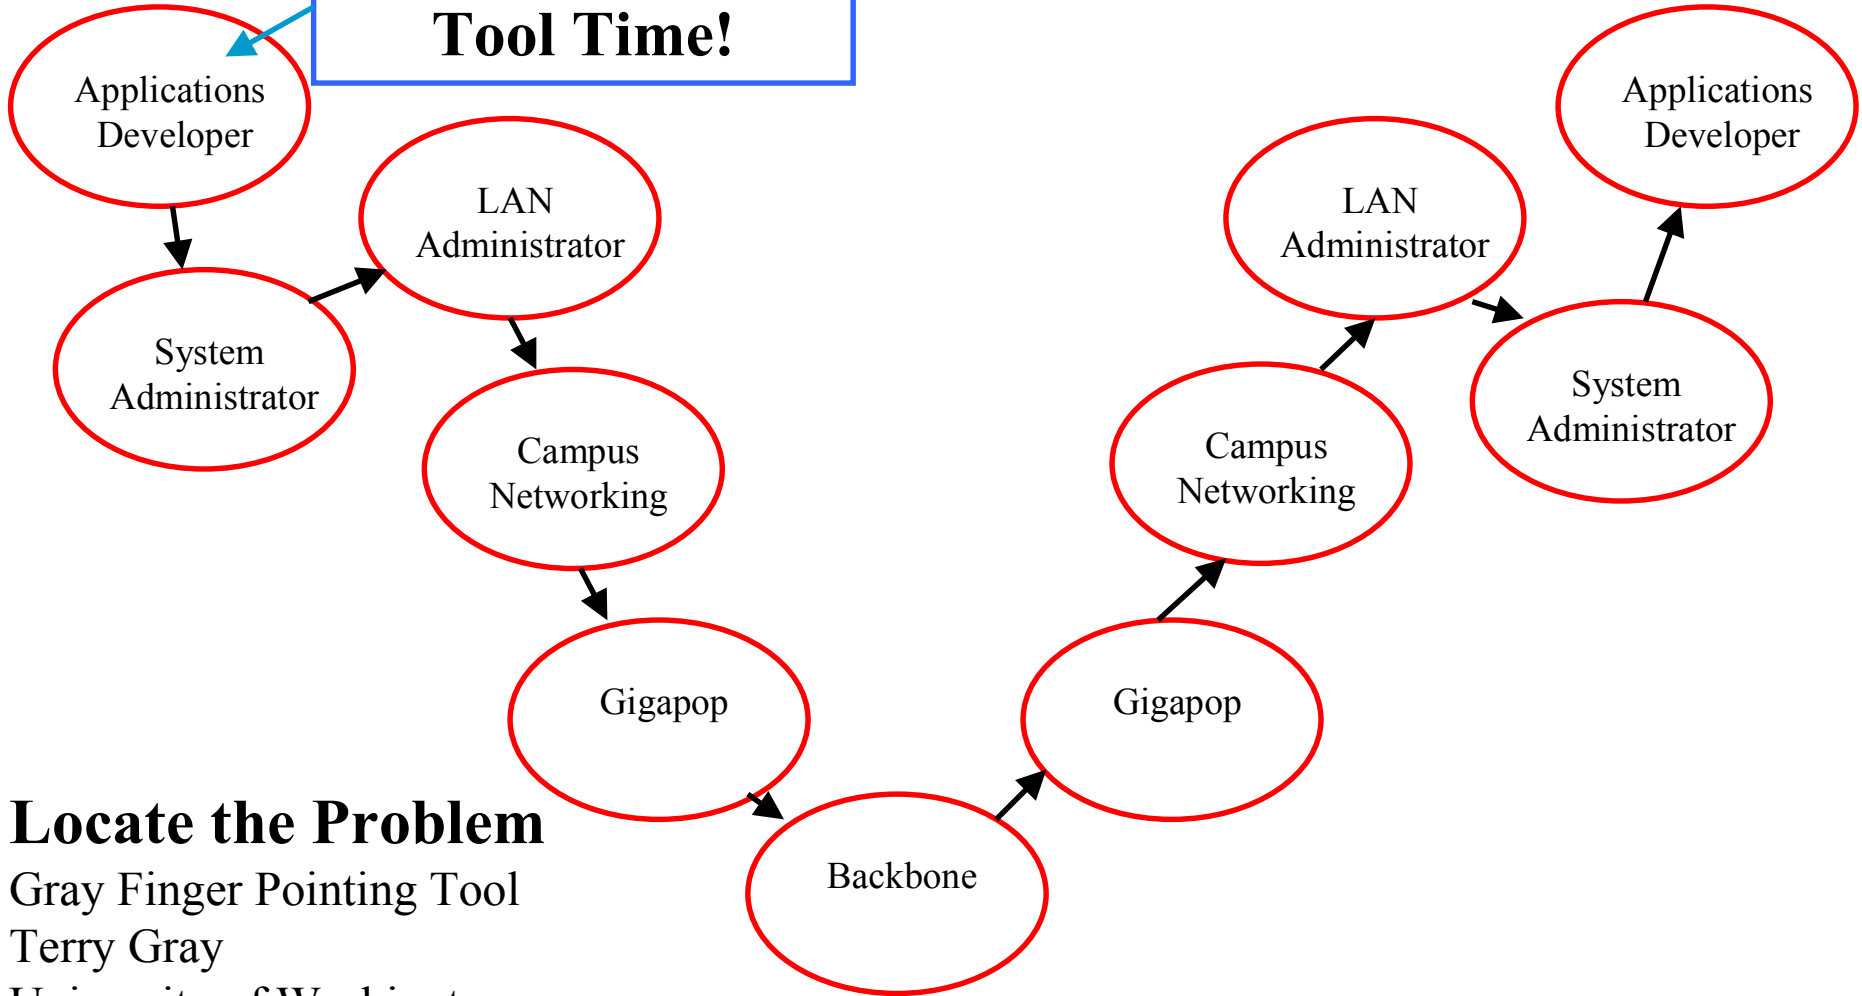
Terry Gray, University of Washington

“We Need a Finger Pointing Tool”



Gray Finger Pointing Tool

**Finger Pointing
Tool Time!**



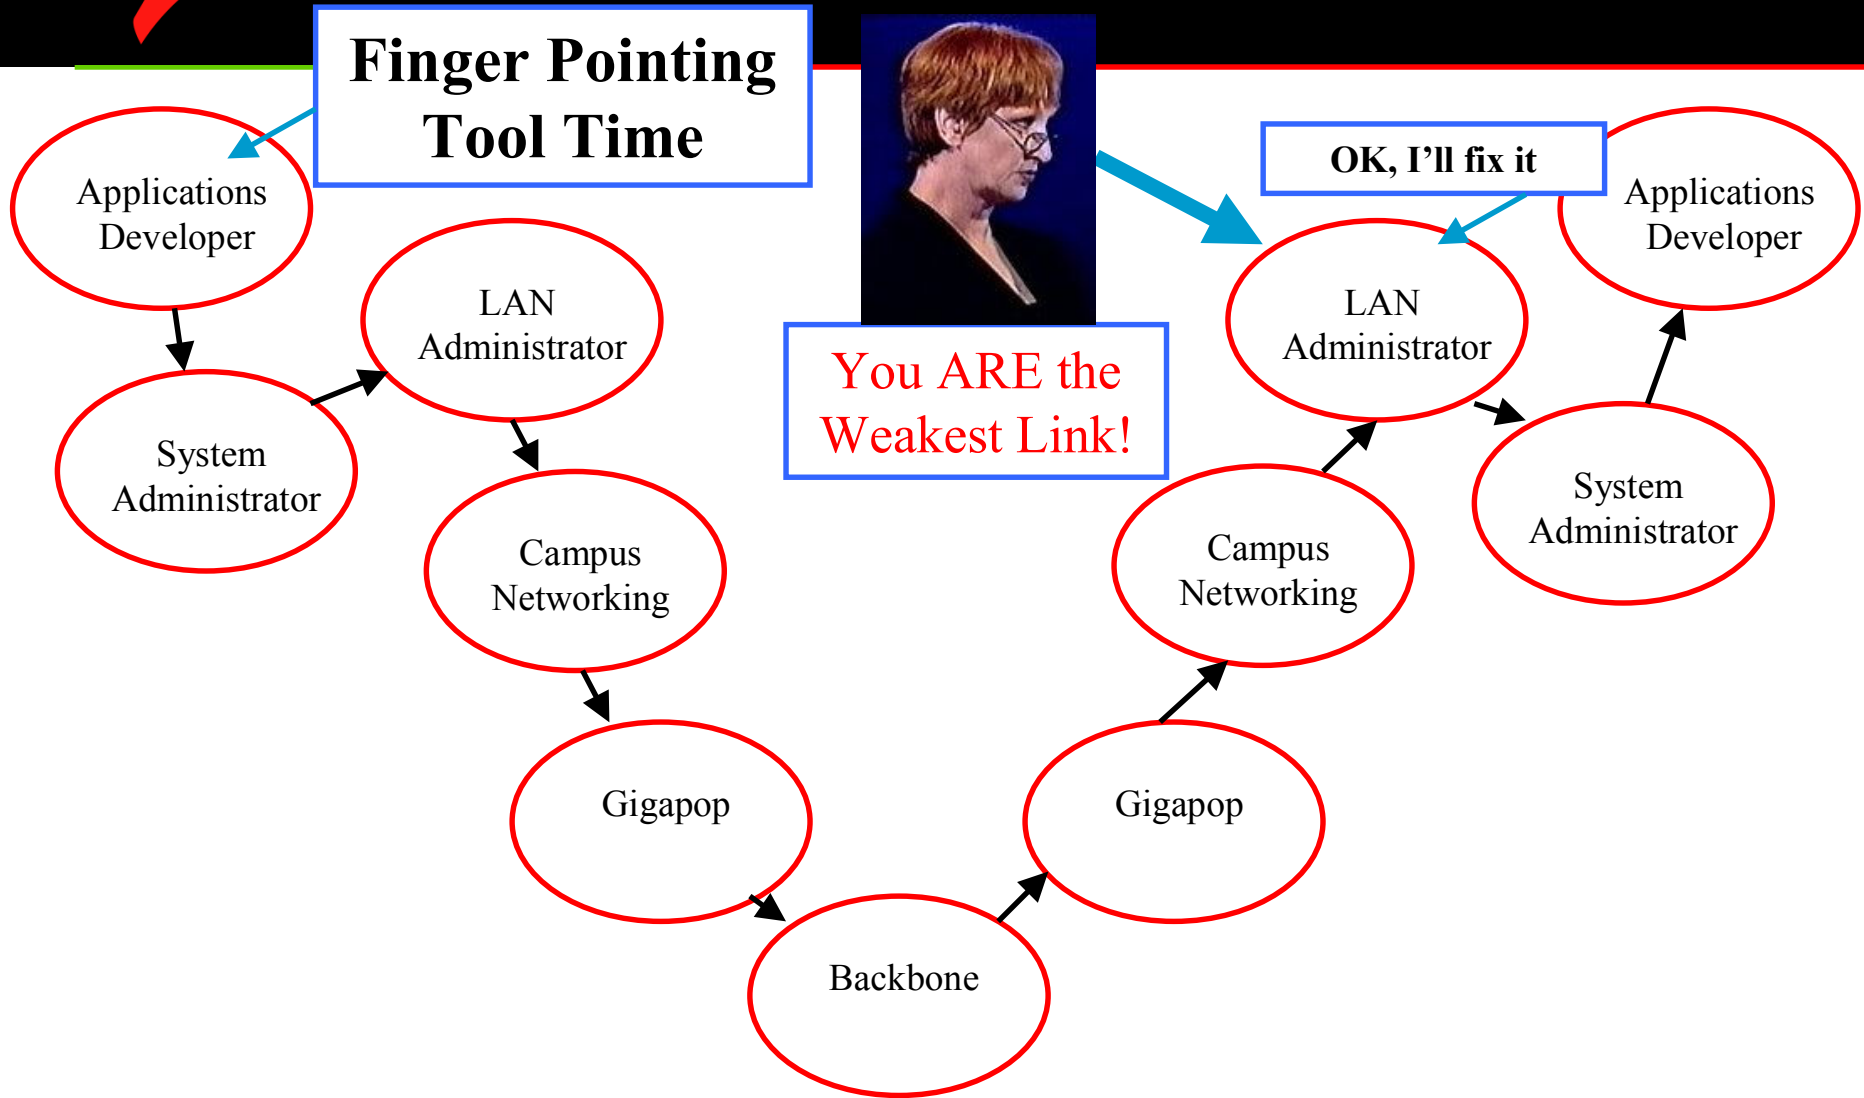
Locate the Problem

Gray Finger Pointing Tool

Terry Gray

University of Washington

Gray Finger Pointing Tool





Agenda

- Motivation
- piPEs / AMI Overview
- piPEs / AMI Deployment Plan
- OWAMP Release
- Demo



BNI AMI

- Project: Backbone Network Infrastructure Abilene Measurement Infrastructure
- Goal: Instrument Next Generation Abilene network with extensive performance measurement capabilities
- Approach: Collaborative project among BNI, Engineering, and E2E.



E2E piPES


- Project: End-to-End Performance Initiative Performance Environment System (E2E piPES)
- Goal: To allow end-users and network operators to determine performance capabilities, locate problems, and contact the right person to get a problem resolved.
- Approach: Collaborative project combining the best work of many organizations.



The Measurement System

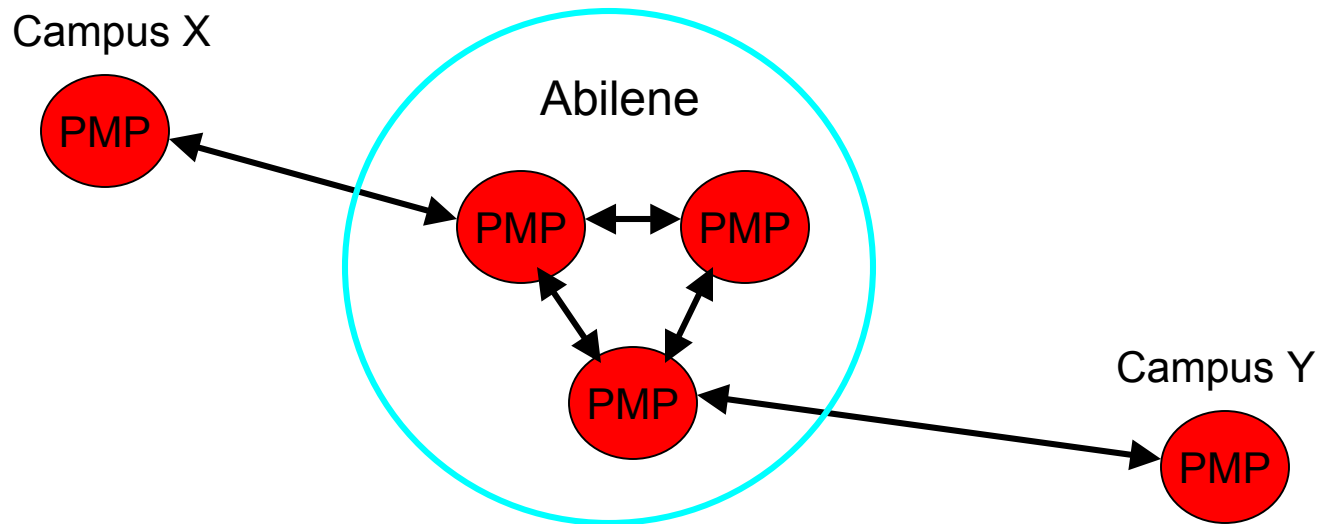
The New Abilene will have measurement devices as part of its structure

Abilene



PMP = Performance Measurement Point (at each Abilene Node)

Extend the System to Campuses




PMP at each **Campus Border**




Performance Reporting

Perform regular tests to provide web based Performance Reports across Abilene and for Campus-Abilene paths




Provide a Testing Tool for the End-Users Computer to test across campus and beyond



The tool reports performance of the path to a specified destination by using active tests (like to the campus border) or using stored test results from regularly scheduled tests. It also performs tests on the host system to check host tuning parameters

Add PMP along paths to provide additional detail

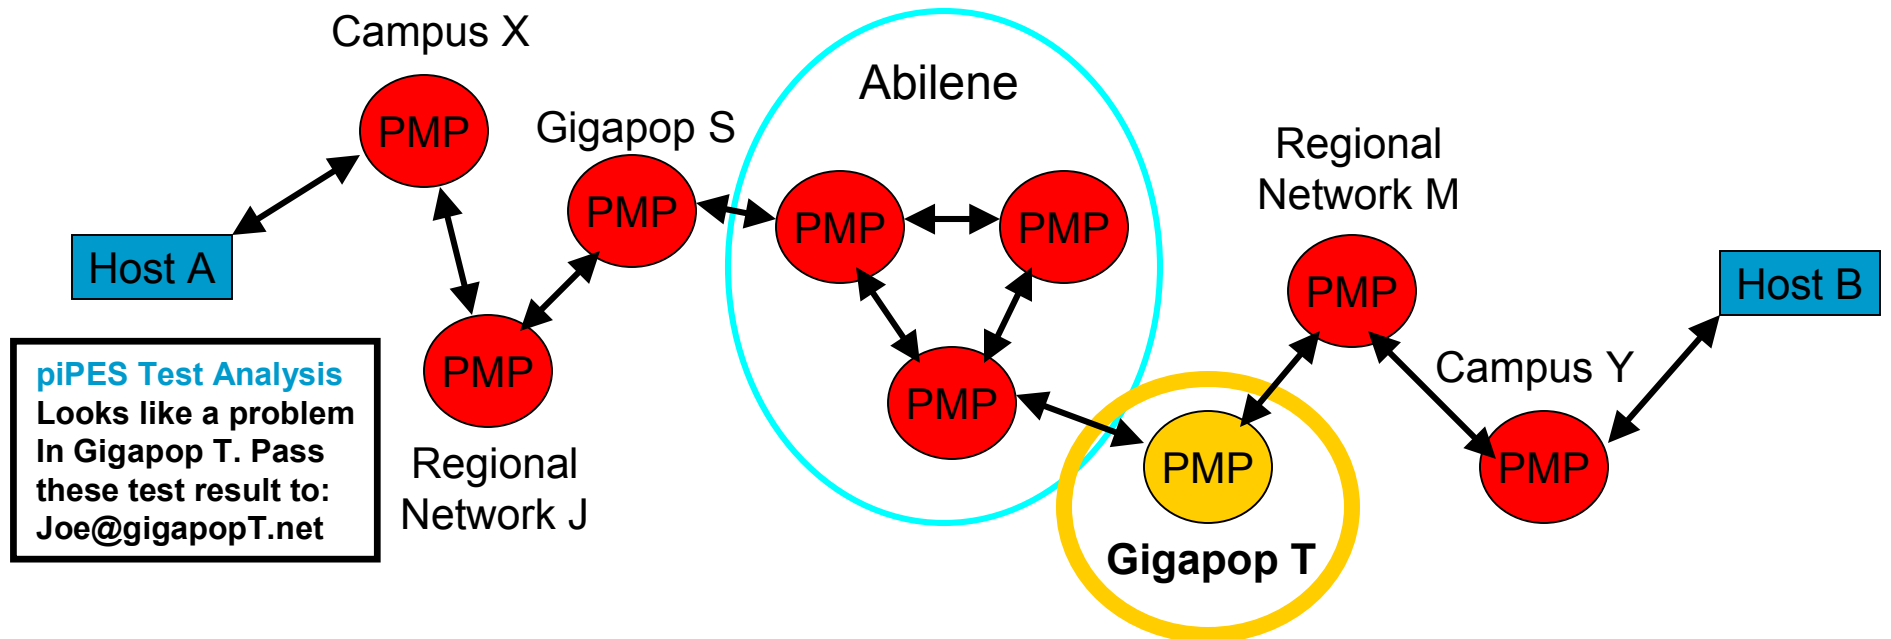


Regular tests are performed between adjacent PMP and stored in a Data Repository. This information is used by the End-User tool to reduce the number or active tests needed.



Who You Gonna Call?

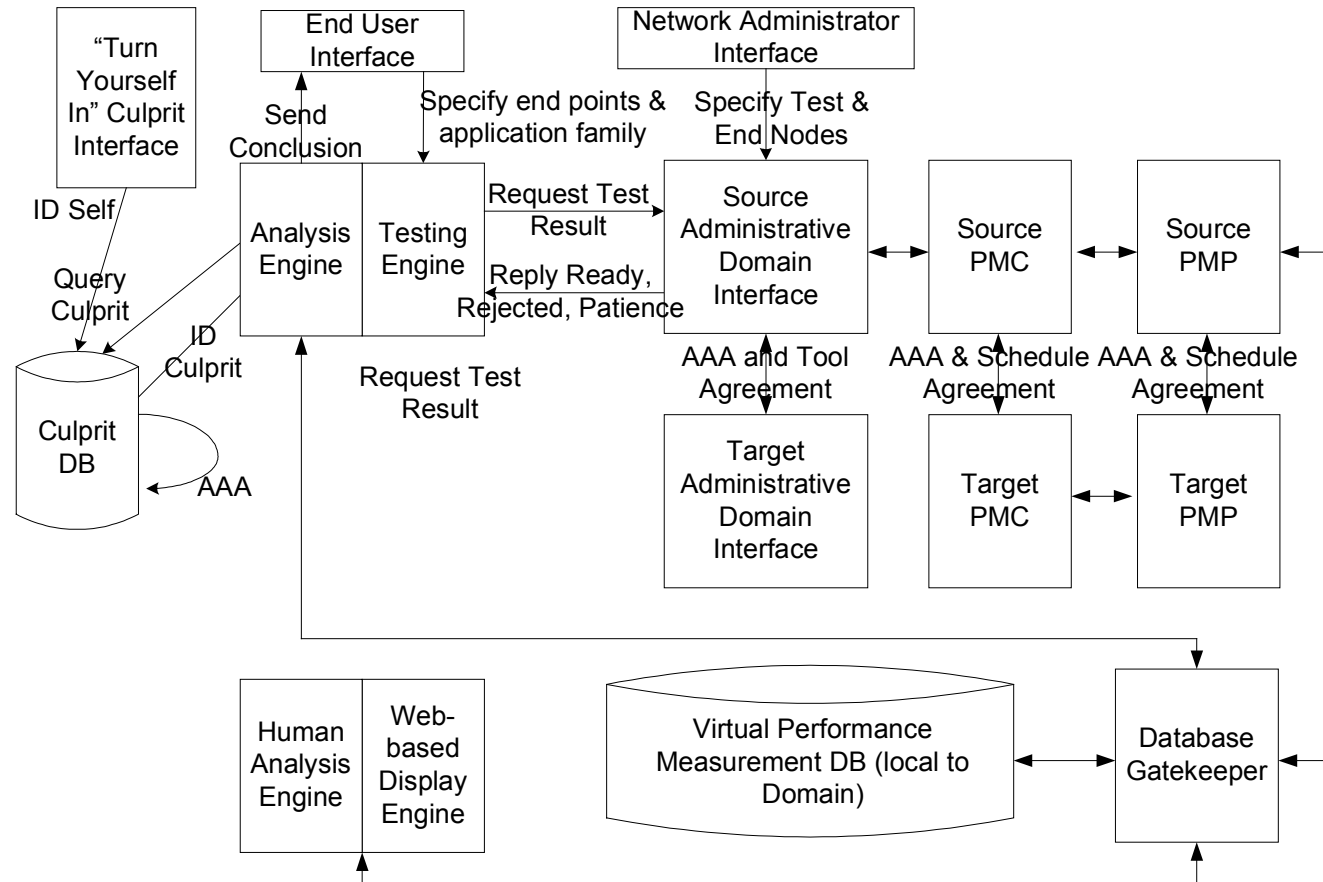
E2E piPES will provide contact information when a segment in the path appears to have a problem



The test results obtained by the End-User will be passed on to the contact person



E2E piPES Architecture





Testing / Analysis Engine

Problem: Encode Matt Z.'s Brain

Open Questions:

- What are measurement results are “acceptable” for a given application (family)?
- What tools generate those results?
- How do you handle incomplete data?
- How do you rank multiple result-generators?
- What is the iterative decision tree to understand E2E problem?



Measurement Schema

Measurement Types?

e.g. BW, Latency, Loss, Jitter

Measurement Units?

e.g. Seconds vs. Microseconds

Map Tools to Measurements?

e.g. OWAMP -> 1-way Latency

Measurement Metadata?

Database Table Design?



Access, Authentication, Authorization

Roles for Access

- Standard End-User?
- Near Neighbor (Test Buddy)?
- NOC Staff / Network Engineer?

Shibboleth for Implementation

- Each campus decides who in each role
- Others trust campus designation



Other R&D Issues

Design system to scale (eliminate centralization)?

How to “discover” PMPs?

How to insert On-Demand tests into Regularly-Scheduled test set

Balance centralization and distributed database requirements

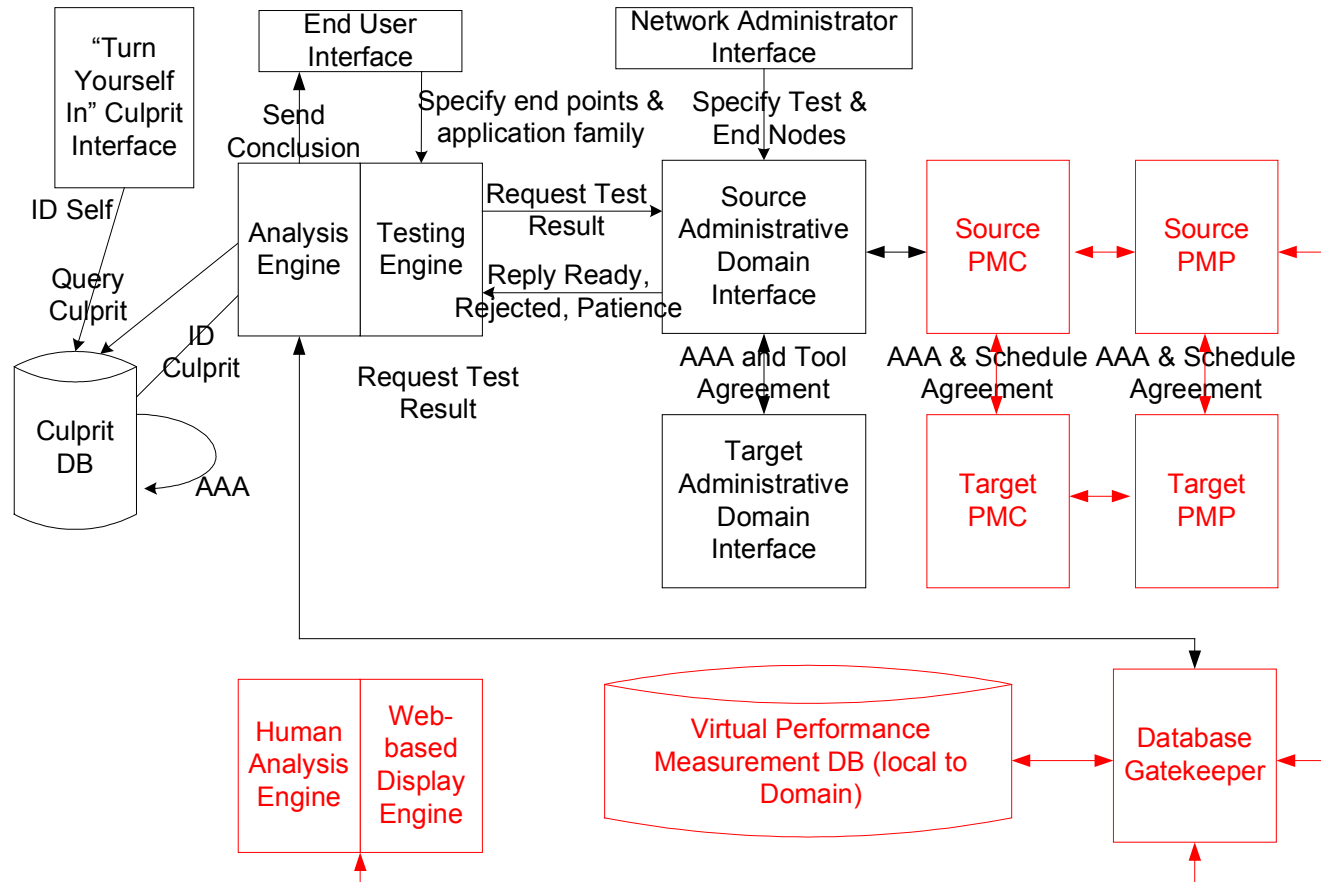


Agenda

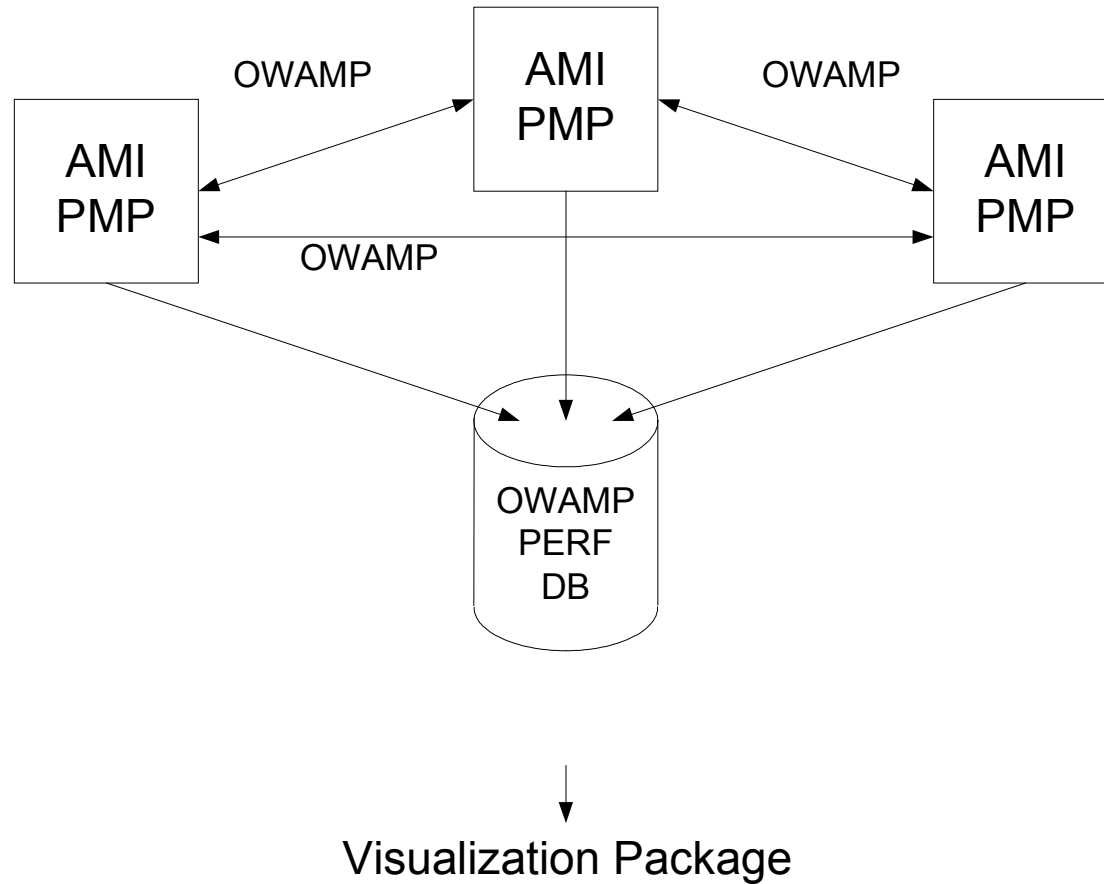
- Motivation
- piPEs / AMI Overview
- piPEs / AMI Deployment Plan
- OWAMP Release
- Demo



E2E piPES Architecture v1.0

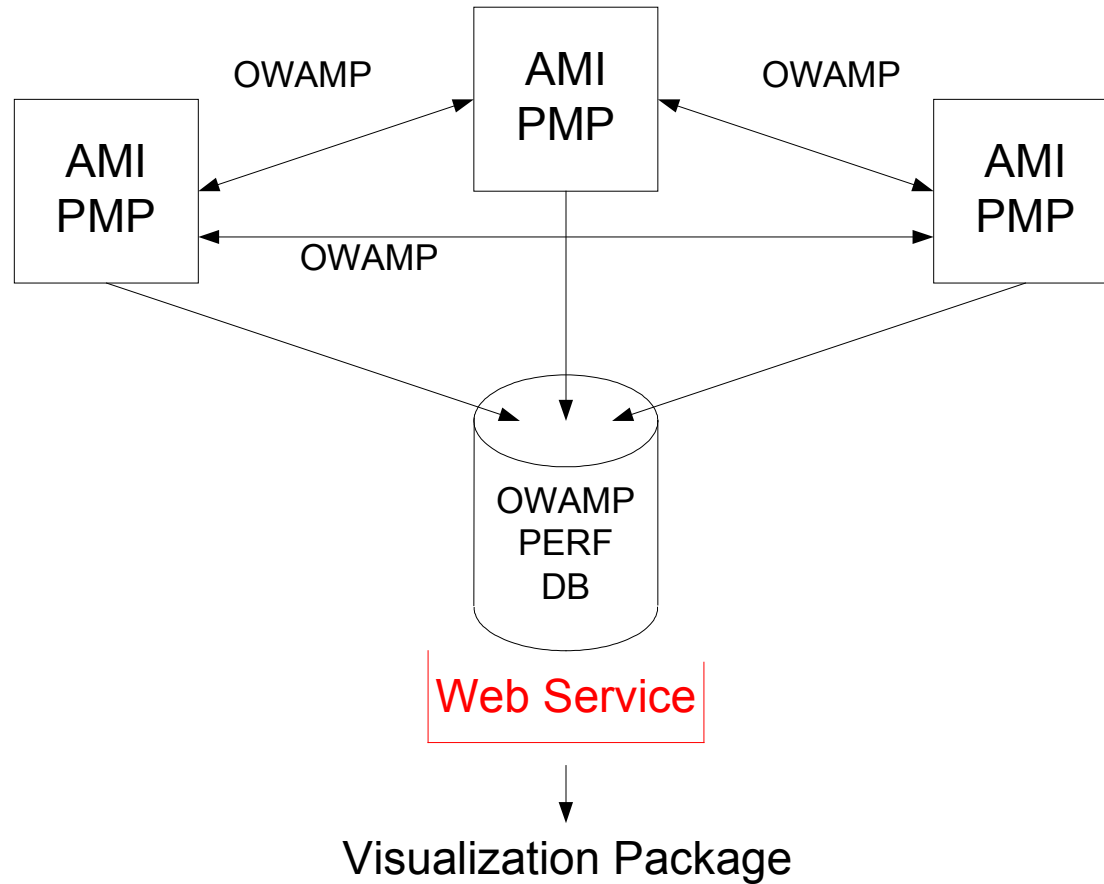


piPEs / AMI Rollout



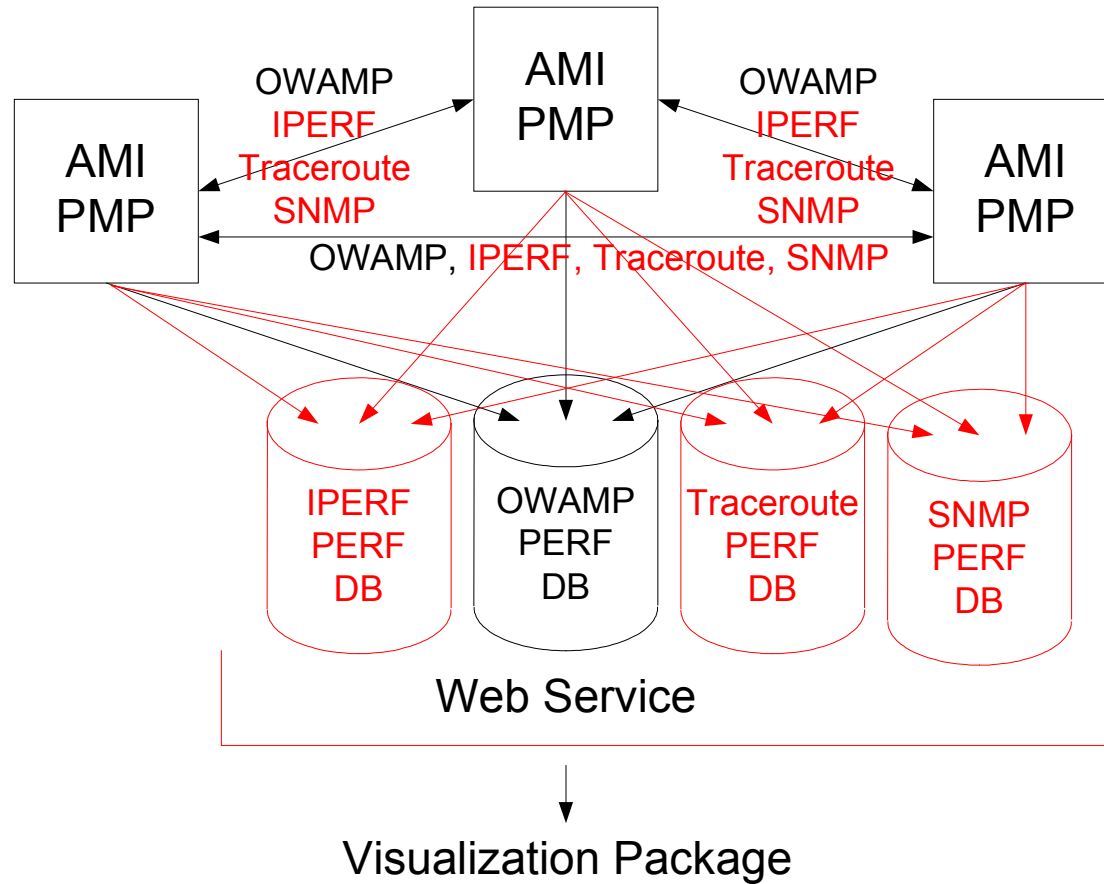


piPEs / AMI Rollout



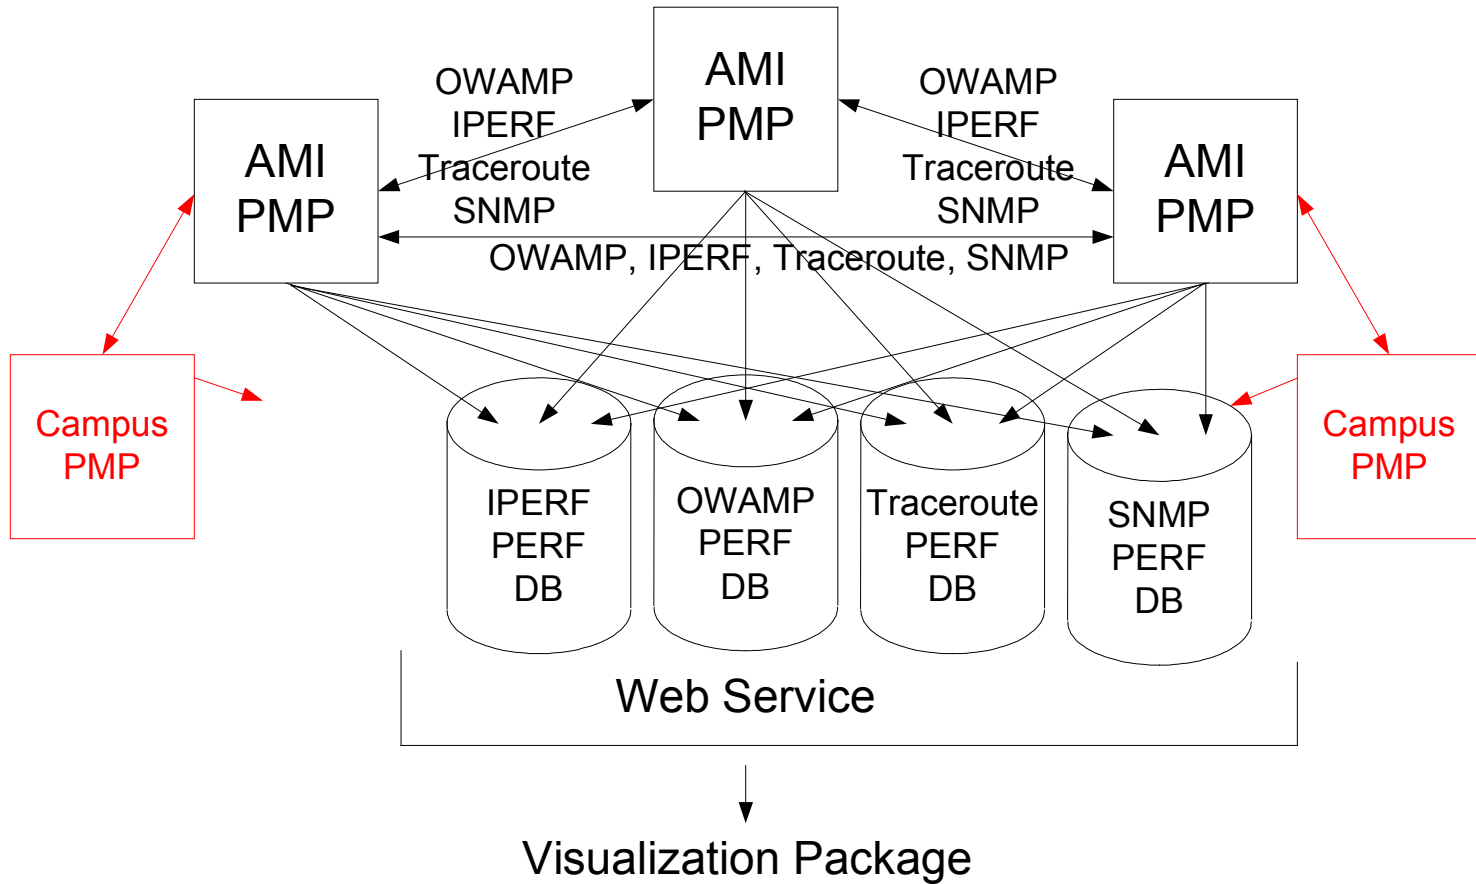


piPEs / AMI Rollout



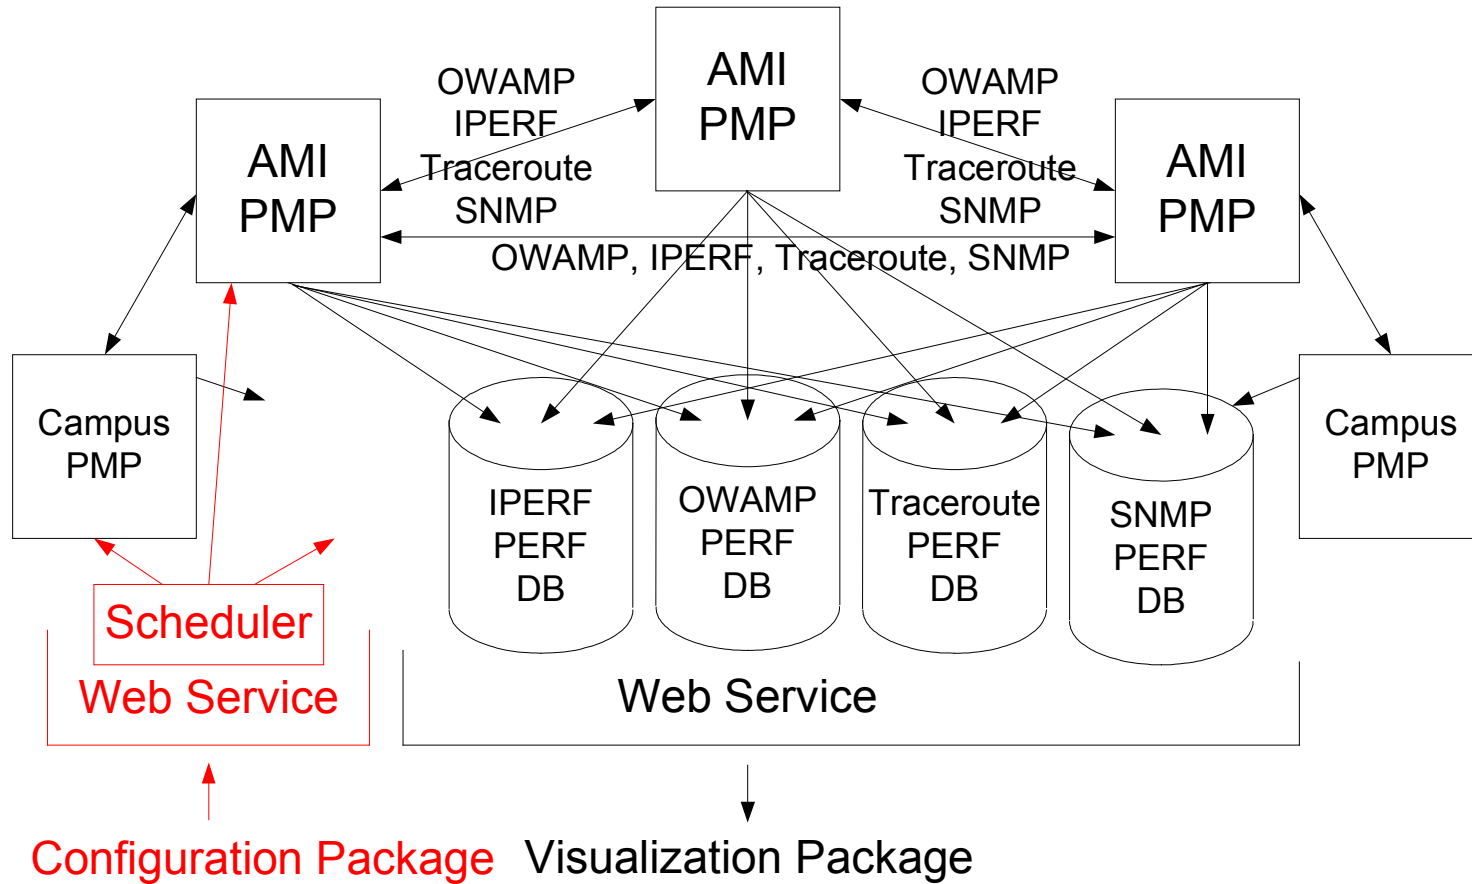


piPEs / AMI Rollout



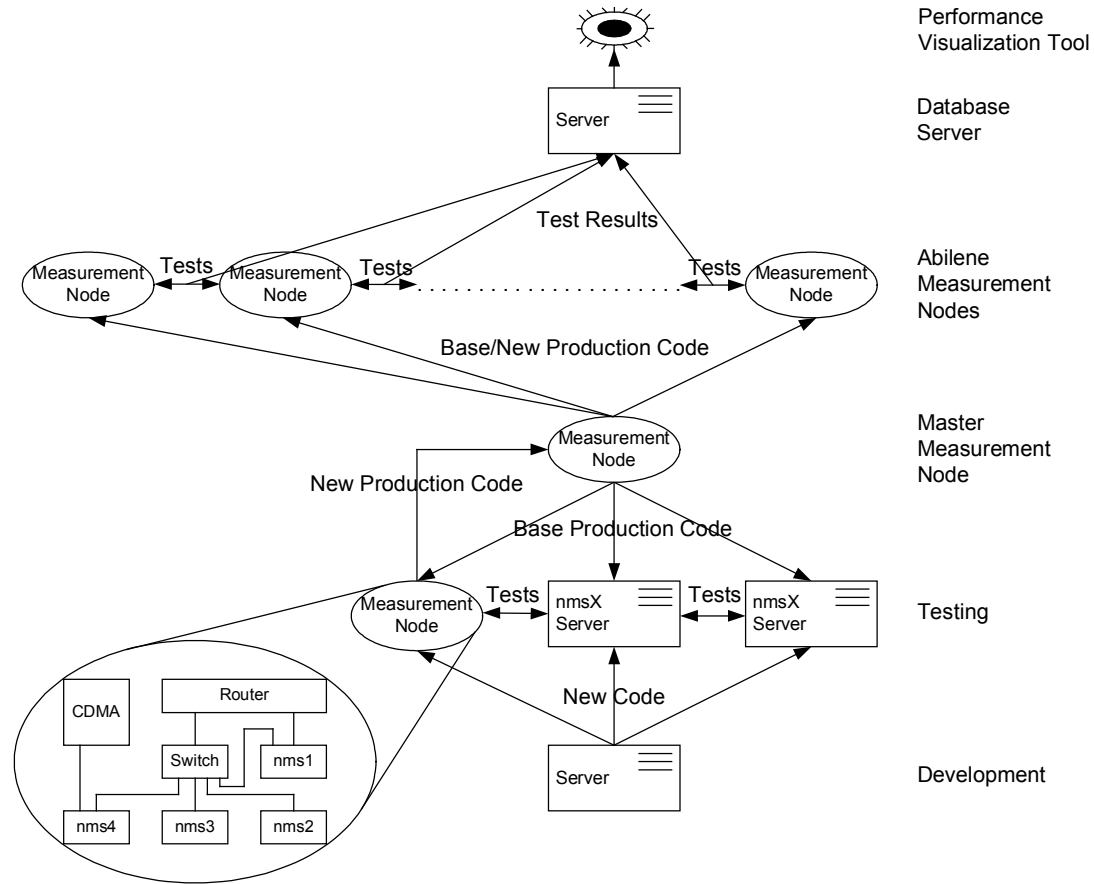


piPEs / AMI Rollout





piPEs / AMI Development & Testing Environment





piPEs / AMI Internal Resources

Eric Boyd – piPEs development, project coordination

Jeff Boote – OWAMP, piPEs development

Prasad Calyam – piPEs development

Chris Heermann – AMI

Matt Zekauskas - AMI

Susan Evett - Documentation

Russ Hobby – Campus Deployment

George Brett – Schema, “Grid Service”



piPEs External Resources

Yee-Ting Li – UCL (50%)

Paul Mealor – UCL (up to 100%)

PhD Student – UCL (contingent on grant)

Warren Matthews – SLAC

Jim Ferguson et al – NLANR / DAST

Nicolas Simar et al – DANTE

Many early adopter volunteers



Agenda

- Motivation
- piPEs / AMI Overview
- piPEs / AMI Deployment Plan
- **OWAMP Release**
- Demo



Latency Tests

Problem: Existing platforms are not interoperable (SURVEYOR,RIPE,...)

Solution: Standards

One-Way-Active-Measurement-Protocol

- Possible due to growing availability of good time sources
- Wide deployment of “open” servers would allow measurement of one-way delay to become as commonplace as measurement of RTT using ICMP tools such as ping.
- Current Draft: [draft-ietf-ippm-owdp-05.txt](#)
–Shalunov, Teitelbaum, Zekauskas



OWAMP Implementation

Basically:

- NTP system call interface

- Multiple processes for recv/send loops

- Written as an API to allow one-off implementations



OWAMP implementation status

Sample implementation

<http://owamp.internet2.edu/>

- Beta Release:
 - No “policy”
 - No authentication/encryption
 - FreeBSD only



Abilene OWAMP deployment

2 overlapping full meshes (IPv4 & IPv6)

- 11 measurement nodes = 220 ongoing tests
- UDP singletons
- Rate: 10 packets/second*
- Packet size: (32 byte payload)*
- Results are continuously streamed back to “Measurement Portal” for long-term archive and data dissemination (Near real-time)



Agenda

- Motivation
- piPEs / AMI Overview
- piPEs / AMI Deployment Plan
- OWAMP Release
- Demo



More Internet2 Information

On the Web

- www.internet2.edu
- e2epi.internet2.edu

Email

- Info-e2epi@internet2.edu



Questions?





INTERNET®

www.internet2.edu