



Performance (Debugging) Measurement Architecture Workshop

Matt Zekauskas, matt@internet2.edu

JET Roadmap Workshop

13 April 2004

Background

- Working with DANTE, TF-NGN, UCL, GGF, NLANR (DAST on Advisor; MNA on AMP/PMA), and knowing a bunch of more basic research projects... widen the conversation on how we can work more cooperatively?
- Is there an overall plan/architecture?
- What's common, what's missing?

- Performance (Debugging) Measurement Architecture Workshop 2003
- 1.75 day workshop last December
- Bring together a bunch of (mostly) NSF-funded projects
- Including a related one: end-to-end middleware diagnostics (start with authorization and authentication, but can expand)
- And some network-specific [BB, Campus]

Not a new idea

- “Challenge: evolve the field of Internet measurement from a set of disjoint, independent activities of network researchers, operators, and users toward an integrated service of the network that can support diagnosis of problems in real-time.”
 - kc claffy, ed. “Measurement Roadmap” from the NREN/NASA workshop in 2000 on gigabit networking

Our leverage

- Projects specifically about the integration
- R&E backbone networks are interested in supporting
- A few big demanding users (e.g. HENP) with diverse endpoints

Existing Architectures

- Grid Measurement
- piPEs
- TF-NGN Performance Monitoring Group
- Internet2 Middleware End-to-End Diagnostics



Internet2 Middleware E2E

- Federated Security – Shibboleth; Authenticate locally, Authorize globally
- Based on trust of institutional authentication
- Pass minimum amount of information necessary
- If failure
 - What piece of the distributed infrastructure is to blame?
 - Want to debug without releasing private information
- Diagnostic backplane with common event record

Current Projects

- **CaINGI / NPACI**
 - Accessible to researchers
- **NIMI**
 - Authentication focus
 - Experiments
- **SLAC**
 - pinger, iepm-bw
 - netflow {local}
- **NLANR DAST: Advisor**
- **NLANR MNA**
 - AMP, PMA {oc192 soon}
- **piPEs, as implemented**
- **TeraGrid (SAMI)**
- **Esnet Perf. Monitoring**
- **LBL: NTAF, SCNM**
- **IU: tools for campuses**
 - Flowscan, gcube, L2TP
- **Mathis: NPAD**

My view of results

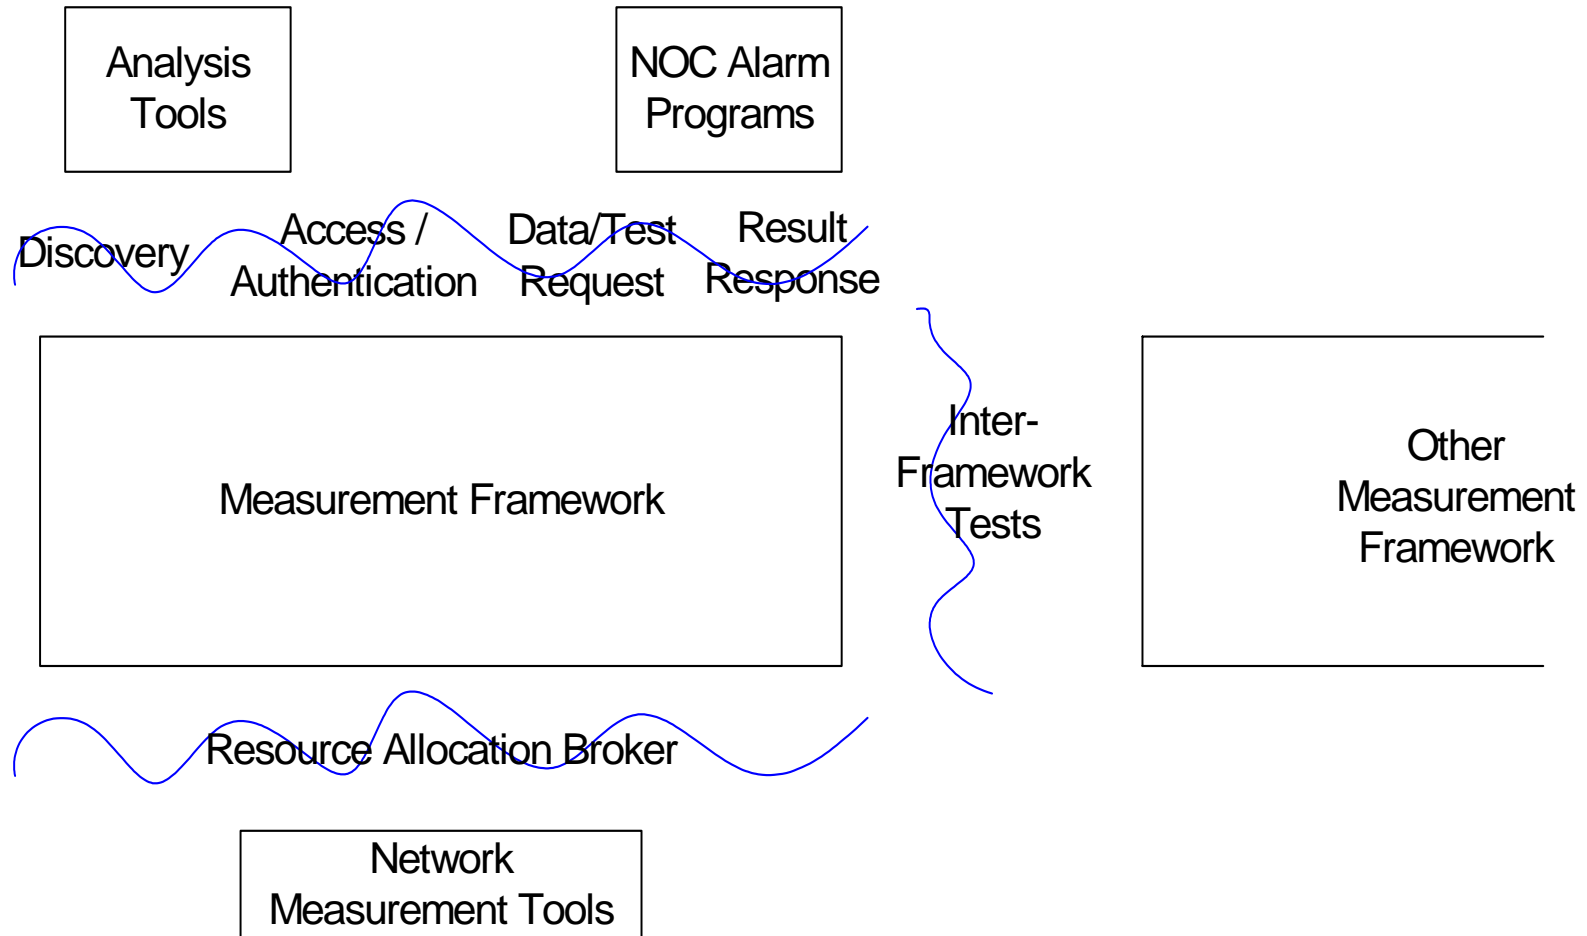
- People want to work together, but they also want autonomy (don't constrain possible solutions by architecture)
- Want ways to share results, enable cross-platform testing (but not forced)
- Tools folks would like to use existing infrastructures
- ...want “Interoperable Measurement Frameworks”

■ Common Themes

- Discovery (domains, boxes along path)
- Interfaces (requests, response)
- Authorization, Authentication
- Calibration, Meta-data
- Infra. Operational issues
- Algebra for metrics
- How know physical topology [not common, but big discussion!]



Measurement Infrastructure Federation Interfaces



Reference for Workshop

- <http://e2epi.internet2.edu/WK03/>
- But... we are still working on a report.



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



Working to Create Interoperable Federations

- DANTE GN2 JRA1, EGEE JRA4
http://people.internet2.edu/~eboyd/transatlantic_workshop.html
http://people.internet2.edu/~eboyd/ucl_workshop.html
- Planning with DoE (MAGGIE)

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 heterogeneous tools

Acknowledgements

■ NSF

- This material is based in part on work supported by the National Science Foundation under Grant No. ANI-0314723. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

■ San Diego Supercomputer Center

■ CAIDA

■ Internet2, especially the End-to-End Performance Initiative

■ All the participants!



www.internet2.edu

- Abilene Observatory, piPEs in Abilene
 - Ongoing measurements, openly accessible; on-demand with piPEs for throughput (bwctl) or latency (owamp)
- TeraGrid
 - Ongoing measurements
- ESnet
 - On-demand for ESnet sites
- (Ongoing work on GEANT)

Wide-ranging Projects

- pinger, iepm-bw
 - Ongoing active
- AMP
 - Ongoing active
- NIMI
 - Active, but focussed on experiments
- (pipEs)
 - Ongoing active

Drawing conclusions

- Advisor (based on iperf, Web100)
 - Intelligence at ends
- Rich Carlson's NDT
 - Intelligence at middle (packet streams, Web100 at Server)
- piPEs provide data to others,
Rich Carlson is working on the piPEs
analysis component

New tools

- Most of the projects that integrate use existing tools
 - Ping, traceroute, iperf (or similar)
 - SNMP (some)
- Some new tools arriving
 - Capacity discovery, lower-overhead throughput testing
 - Web100-style host instrumentation

- NPAD – New tool(set) under development at PSC, Matt Mathis
 - Capitalize on Web100, but look for flaws in link layers
 - Tweak TCP on short runs, detect problems that would show up on long runs [drops, reordering]: hits network harder than UDP tests, but doesn't drive into congestion
 - Find bugs in campus infrastructure
 - Way to bench-test applications via packet reflection



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)



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