Building Large-Scale Internet Measurement Infrastructure

Vern Paxson

Network Research Group
Lawrence Berkeley National Laboratory
University of California, Berkeley

vern@ee.lbl.gov

Jamshid Mahdavi, Andrew Adams and Matt Mathis

Pittsburgh Supercomputing Center
Carnegie Mellon University
Pittsburgh, PA 15213

mahdavi@psc.edu, akadams@psc.edu, mathis@psc.edu

June 1998 NANOG Meeting

Dearborn, MI

What is NIMI:

“National Internet Measurement Infrastructure.”

1-year NSF pilot project. Also supported by DOE.

Widespread deployment of measurement “platforms” for:

- diagnosing performance problems
- baseline measurements for how traffic evolves
- assessing performance delivered by IP clouds [*]

This is a talk about <u>architecture</u>, not <u>analysis</u>.

<u>Preliminary.</u>

Deployment model:

Providers might choose to install platforms:

- to aid management of their own network
- to allow their competitors to be measured
- to publicize their own performance [*]
- to facilitate Internet research

And/or: contracts might include provisions for deploying platforms.

Design goals and constraints:

Work in an administratively diverse environment.

$\Rightarrow$ support a wide range of policies & needs

Work in the commercial Internet.

$\Rightarrow$ respect business considerations

Support a wide range of measurements.

- who knows what will be interesting tomorrow?

$\Rightarrow$ notion of “measurement module”

Design goals and constraints, con't:

Active measurements rather than passive.

$\Rightarrow$ greatly simplifies security, privacy concerns

Scale to 1000's of measurement platforms.

$\Rightarrow$ avoid "success disaster"

$\Rightarrow$ require minimal administration, maximal self-configuration.

Solid security and authentication from the beginning.

$\Rightarrow$ prevent abuse of platform

$\Rightarrow$ prevent abuse of measurements

$\Rightarrow$ keep data confidential as necessary

Design goals and constraints, con't:

Platform owners have (potentially) full administrative and policy control over their platforms . . .

. . . but make it easy for platform owners to delegate control

. . . and provide fine-grained control when needed.

NIMI architecture:

Basic model: platforms are servers that service authenticated requests.

Requests might: schedule future measurement; manipulate measurement schedules; retrieve results; delete results; upload new software versions; manipulate the policy table.

Authentication is done by presenting a *credential*.
Credentials have names and associated public keys (RSARef).
You prove you hold a credential by signing your request.

Policy via ACL table:

Access control table expresses a platform's (arbitrary) policy.

Rows are credentials.
Columns are actions.
Table entries yield “yes”, “no”, or “run this script.”

Delete a row: credential no longer has any access.
Delete a column: no one can execute the given action.

Trust boundaries = exactly who holds and shares which credentials.

Delegation:

Can treat subset of ACL table as “subtable” and create an action “add rows to the *Foo* subtable”.

Columns in subtable have predicates associated with them (e.g., “not a `foo.com` measurement target”).

E.g., “hourly TReno mesh measurement” subtable.

Subtables can be nested.

$\Rightarrow$ Mechanism for delegating partial trust.

Easy configuration:

Fundamentally, platforms know just four things:

- its own unique credential
- public key of its configuration “master”
- URLs for locating one or more instances of master
- nascent ACL allowing master to do anything

Configuration, con't:

Platform contacts the master, which downloads additional ACL rows and columns.

Master can redirect platform to another URL for additional configuration. (E.g., “hourly TReno mesh measurement” configuration.) Naturally, it first sets up ACL entries for the additional access.

Those platforms can further redirect.

Implementation status:

On six platforms at five sites (PSC, LBNL, SLAC, FNAL, CERN).

FreeBSD/NetBSD on Pentiums.

Implemented: measurement scheduler, ACL table, encrypted requests, autoconfiguration.

Not yet: ACL subtables; manipulating schedules; auto-maintenance; resource locks; manipulating measurement results.

On-going `traceroute`, `TReno`, and `poip` measurements.

Longer-term research issues:

Result dissemination: multicast?

Distributed database?

Multicast-based inference of path properties.

Large-scale analysis; aggregating results.

The “mapping” problem: self-organization.

Result integrity.

Related efforts:

IPMA: *http://www.merit.edu/ipma/*

Surveyor: *http://io.advanced.org/surveyor/*

Felix: *ftp://ftp.bellcore.com/pub/mwg/felix/index.html*

MINC: multicast-based analysis (w/ UMass, AT&T; postdoc needed!). *http://www.research.att.com/~duffield/minc/*

More information:

INET '98 paper: *http://www-nrg.ee.lbl.gov/nimi-inet98.html*

IEEE Communications paper:

ftp://ftp.ee.lbl.gov/papers/nimi-ieee-comm98.ps.gz

Alpha volunteers? Email to `vern@ee.lbl.gov`.