SLAC Internet Measurement Data

Les Cottrell, Jerrod Williams, Connie Logg, Paola Grosso
SLAC, for the

ISMA Workshop, SDSC June, 2004

www.slac.stanford.edu/grp/scs/net/talk03/isma-jun04.ppt

Partially funded by DOE/MICS Field Work Proposal on Internet End-to-end Performance Monitoring (IEPM), also supported by IUPAP

Main interest as end-user
•Active probes, E2E
•Passive border: characterization & security

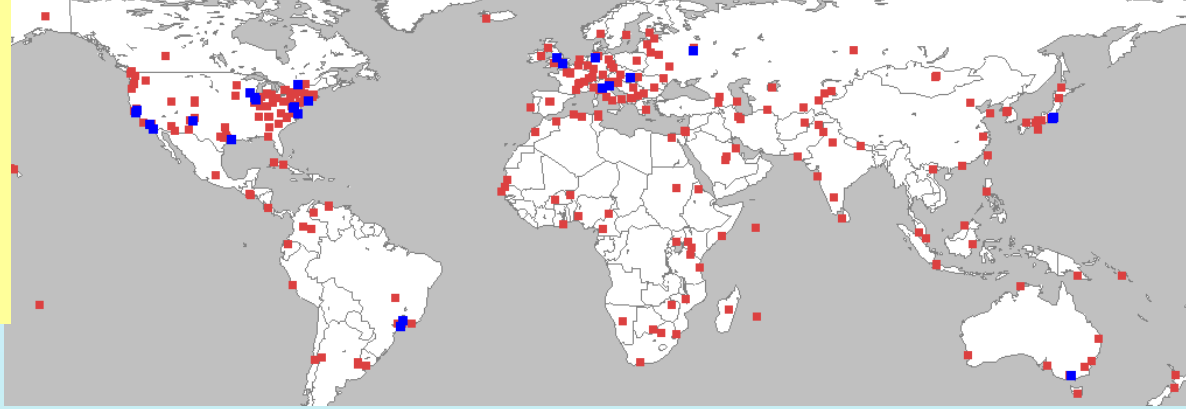

- PingER:
  - 7 years of data, > 100 countries, ~35 monitoring sites, ~550 remote sites, lightweight, good for developing countries
  - pings every 30 mins growing number of sce-dest pairs (~3700 currently)
  - Monitor site collects 0.5MB/pair/month
  - Two archives: SLAC & FNAL
    - Gather data from monitor sites at regular intervals
    - Kept in flat files at SLAC
    - Adding to Oracle database for recent data, and web services access following NMWG schemata, e.g.
      - path.delay.roundTrip ms (min/avg/max + RTTs),

IEPM-BW

- Measurements for hi-perf paths with multi & single-stream iperf, bbcp, bbftp, GridFTP, ping
- Ten monitoring sites, ~60 remote hosts (9 countries)
- Measurements ~ 90 mins intervals, ~ 10-20 s per measurement
- Kept in flat files on monitor host, no regular central gathering
- Network intensive, requires scheduling
- Also available via web services with Oracle back-end, e.g.
Characteristic	Toolname
path.bandwidth.achievable.TCP	iperf
path.bandwidth.achievable.TCP.multiStream	Iperf,bbftp, bbcp, GridFTP
  - Used by MonALISA (so WSDL changes need coordination)

IEPM-LITE

- Currently about 40 sites, expect to expand
- ABwE measurements every 3 mins
  - Provides capacity, X-traffic, available bandwidth, RTT
- Traceroutes every 10 mins
- Network low impact (ABwE 20 packets / direction), no scheduling needed
- Kept in flat files, also web services, e.g.
  -

Characteristic	Toolname
path.bandwidth.capacity	ABwE
path.bandwidth.utilization	ABwE

  - Working (with Warren Matthews/GATech/I2) on defining / providing access to traceroutes for AMP & IEPM-LITE

Data types

- Raw measurements
  - Maybe saved in flat files or in an SQL dB
  - Flexibility in querying vs. speed of access
- Analyzed data
- Plots, Tables
  - Some on demand (CGI scripts) in particular PingER
    - Takes longer to get information for user
  - Others generated daily and saved (IEPM-BW & LITE)
    - Faster access for user, but more storage
- Data kept in network file systems (AFS/NFS)
  - Allow access from monitor host
  - Web servers
  - Can be reliability problems

Data Requests

- Big analyses (e.g. 7 years of PingER RTT & Loss data)
  - Tar and zip data and FTP (few requests/year)
- Recent data (e.g. for Grid application steering)
  - Web services (MonALISA for IEPM-BW)
  - Currently real-time PingER data not available, i.e. one day old, we are working on this with NIIT
- Intermediate term available from web pages in TSV format for Excel etc., easily automated
  - PingER: roughly 40 hits/day
- PingER data NOT anonymized, IEPM host name hidden (network name visible)

Challenges 1/2

- Keeping remote sites accessible (port/protocol blocking, hardware failures, changes in address or name or hardware …)
  - Result in holes in the data, or new host/site replacing old
- Collecting data from monitoring hosts
- Recovering “lost” data and rippling it back into the analysis chain.
- WSDL
  - Complexity, steep learning curve, tools currently limited
  - Schema definition stability inhibits deployment

Challenges 2/2

- Running continuous measurements, collecting data etc. is hard

More Information

- PingER
  - http://www-iepm.slac.stanford.edu/pinger/
- IEPM
  - http://www-iepm.slac.stanford.edu/bw/
- Web services access to IEPM & PingER
  - http:www-iepm.slac.stanford.edu/tools/web_services/
- Example SOAP client for IEPM-BW
  - www-iepm.slac.stanford.edu/tools/soap/IEPM_client.html

Access mechanisms

Web Services

- See http://www-iepm.slac.stanford.edu/tools/web_services/
- Working for: RTT, loss, capacity, available bandwidth, achievable throughput
- No schema defined for traceroute (hop-list)
- PingER
  - Definition WSDL
  - http://www-iepm.slac.stanford.edu/tools/soap/wsdl/PINGER_profile.wsdl
    - path.delay.roundTrip ms (min/avg/max + RTTs),
    - path.loss.roundTrip
    - IPDV(ms),
    - <definitions name="PINGER" targetNamespace="http://www-iepm.slac.stanford.edu/tools/soap/wsdl/PINGER_profile.wsdl">
    - <message name="GetPathDelayRoundTripInput">
    - <part name="startTime" type="xsd:string"/>
    - <part name="endTime" type="xsd:string"/>
    - <part name="destination" type="xsd:string"/>
    - </message>
    - Also dups, out of order, IPDV, TCP thru estimate
    - Require to provide packet size, units, timestamp, sce, dst
  - path.bandwidth.available, path.bandwidth.utilized, path.bandwidth.capacity
- Mainly for recent data, need to make real time data accessible
- Used by MonALISA so need coordination to change definitions

Perl access to PingER

```
#!/usr/local/bin/perl -w

use SOAP::Lite;
use Data::Dumper;

my $startDate = "2004-05-18";
my $endDate = "2004-05-19";
my $destination = "fnal.fnal.gov";


my $measurement = SOAP::Lite
    ->service('http://www-iepm.slac.stanford.edu/tools/soap/wsdl/PINGER_profile.wsdl')
    ->GetPathDelayRoundTrip($startDate,$endDate,$destination);


print Dumper ($measurement);

exit;
```

PingER WSDL

```
<?xml version="1.0" encoding="UTF-8" ?>
- <definitions name="PINGER" targetNamespace="http://www-
  iepm.slac.stanford.edu/tools/soap/wsdl/PINGER_profile.wsdl" xmlns="http://schemas.xmlsoap.org/wsdl/"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/" xmlns:tns="PINGER"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsd1="http://www-
  iepm.slac.stanford.edu/tools/soap/wsdl/response_Jan2004/nm_root.xsd">
 - <message name="GetPathDelayRoundTripInput">
    <part name="startTime" type="xsd:string" />
    <part name="endTime" type="xsd:string" />
    <part name="destination" type="xsd:string" />
   </message>
 - <message name="GetPathDelayRoundTripOutput">
    <part name="networkMeasurementSet" type="xsd1:networkMeasurementSet" />
   </message>
 - <portType name="PINGERPortType">
   - <operation name="GetPathDelayRoundTrip">
      <input message="tns:GetPathDelayRoundTripInput" />
      <output message="tns:GetPathDelayRoundTripOutput" />
     </operation>
   </portType>
 - <binding name="PINGERBinding" type="xmlns:PINGERPortType">
    <soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http" />
   - <operation name="GetPathDelayRoundTrip">
      <soap:operation soapAction="urn:PINGER#GetPathDelayRoundTrip" />
     - <input>
        <soap:body encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" use="encoded"
         namespace="urn:PINGER" />
       </input>
     - <output>
        <soap:body encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" use="encoded"
         namespace="urn:PINGER" />
       </output>
     </operation>
   </binding>
 - <service name="PINGER">
    <documentation>SLAC PINGER Web Services</documentation>
   - <port name="PINGERPort" binding="PINGER:PINGERBinding">
      <soap:address location="http://www-iepm.slac.stanford.edu/cgi-wrap/pinger.cgi" />
     </port>
   </service>
 </definitions>
```

Output from script

```
[3cottrell@hercules:~]$bin/pinger_wsdl.pl
$VAR1 = bless( {
                 'networkMeasurementSet' => {
                                              'networkMeasurement' => {
                                                                        'characteristic' => 'path.delay.roundTrip',
                                                                        'methodology' => {
                                                                                           'parameterSet' => {
                                                                                                               'packetSizeParam' => '100',
                                                                                                               'packetTypeParam' => 'ICMP',
                                                                                                               'numPacketsParam' => '20'
                                                                                                             },
                                                                                           'tool' => {
                                                                                                       'name' => 'ping'
                                                                                                     }
                                                                                         },
                                                                        'subject' => {
                                                                                       'destination' => {
                                                                                                          'address' => {
                                                                                                                         'host' => '131.225.111.1',
                                                                                                                         'version' => '4'
                                                                                                                       },
                                                                                                          'name' => 'fnal.fnal.gov'
                                                                                                        },
                                                                                       'source' => {
                                                                                                     'address' => {
                                                                                                                    'host' => '134.79.240.30',
                                                                                                                    'version' => '4'
                                                                                                                  },
                                                                                                     'name' => 'pinger.slac.stanford.edu'
                                                                                                   }
                                                                                     },
                                                                        'results' => {
                                                                                       'resultSet' => {
                                                                                                        'mean' => '51.19',
                                                                                                        'result0037' => {
                                                                                                                          'timestamp' => '1084928402',
                                                                                                                          'delay' => '50.228'
                                                                                                                        },
                                                                                                        'result0029' => {
                                                                                                                          'timestamp' => '1084914001',
                                                                                                                          'delay' => '50.317'
                                                                                                                        },
                                                                                                        'result0045' => {
                                                                                                                          'timestamp' => '1084942801',
                                                                                                                          'delay' => '50.743'
                                                                                                                        },
                                                                                                        'result0038' => {
                                                                                                                          'timestamp' => '1084930202',
                                                                                                                          'delay' => '50.312'
                                                                                                                        },
                                                                                                        'result0046' => {
                                                                                                                          'timestamp' => '1084944604',
                                                                                                                          'delay' => '51.401'
                                                                                                                        },
                                                                                                        'result0039' => {
                                                                                                                          'timestamp' => '1084932001',
                                                                                                                          'delay' => '50.405'
                                                                                                                        },
                                                                                                        'result0047' => {
                                                                                                                          'timestamp' => '1084946403',
                                                                                                                          'delay' => '51.289'
                                                                                                                        },
                                                                                                        'max' => '59.145',
                                                                                                        'count' => '20',
                                                                                                        'min' => '50.105',
                                                                                                        'result0001' => {
                                                                                                                          'timestamp' => '1084863602',
                                                                                                                          'delay' => '51.19'
                                                                                                                        },
```

Perl AMP traceroute

```
#!/usr/local/bin/perl

use SOAP::Lite;
use Data::Dumper;

my %request = ( methodology => { tool => { name => "traceroute", }, },
                subject => { source => { address => { name => "MONTANA", }, },
                             destination => { address => { name => "GATECH", }, },
                             type => "IPv4", },
                startTime => "now",
                );

my $webService = SOAP::Lite
    -> uri('http://amp.nlanr.net/ATR')
    -> proxy('http://amp.nlanr.net/active/cgi-bin/AtrSoapServer.cgi')
    -> NetworkMeasurementRequest(\%request)
    -> result;

#print $webService->{version},"\n";
print Dumper($webService);
```

AMP traceroute output

```
bin/traceroute_wsdl.pl
$VAR1 = {
          'message' => 'Measurement done.',
          'networkMeasurement' => {
                                    'characteristic' => '',
                                    'subject' => {},
                                    'results' => {
                                                   'resultSet' => {
                                                                    'result' => {
                                                                                  'route' =>
'153.90.170.254 153.90.8.100 198.48.91.21 198.107.150.12 198.107.144.2 198.32.8.50 198.32.8.1
4 198.32.8.80 198.32.8.78 199.77.193.9 199.77.194.41 199.77.194.14 ',
                                                                                  'count' =>
'',
                                                                                }
                                                                  }
                                                 }
                                  },
          'version' => '2004-02-27-Dante-NM-WG-RequestSchema2'
        };
```

Intermediate term access

- Provide access to analyzed data in tables via .tsv format download from web pages.

Bulk Data

- For long term detailed data, we tar and zip the data on demand. Mainly for PingER data.