

20120313 SLAC SEECS meeting

Minutes of SLAC SEECS Meeting March 13th, 2012

General

We may not be able to get the MATLAB Toolkit for Tulip. Dr. Anjum proposed that we can use Mathematica instead of MATLAB. He also proposed that he will try to find out some MS students who might take this conversion as his MS thesis. Students have to finalize their thesis by 6th April. By the second week of April, we will get to know if we have any student doing this.

IPV6 - Anjum and Ghulam

Les has the passwords for monitor and the IPv6 host at SEECS. He has successfully logged on. When he tries to execute a traceroute to ipv6.google.com he gets:

```
Executing exec(traceroute6, -m 30 -q 3, 2404:6800:4009:800::1014, 140) traceroute to 2404:6800:4009:800::1014
(2404:6800:4009:800::1014), 30 hops max, 140 byte packets
1  2001:4538:101:3::14 (2001:4538:101:3::14)  3021.987 ms !H  3021.983 ms !H  3021.971 ms !H
```

Les is unclear if he should see an external (to SEECS) host and has sent email. He has had no response.

Les looked at pinger2.pl, it verifies the address is IPv4 4 octets. In addition one will need a copy of pinger.xml with IPv6 hosts and their addresses. There are 2 subroutines in traceroute.pl that may be useful in making pinger2.pl support IPv6. They are valid_ip and gethostbyname6.

Next we need to make pingtable.pl and getdata.pl IPv6 capable, again this could use the 2 subroutines. Since Ghulam and Sadia are working on these it would be good for them to add this.

pinger2.pl

Joun has installed the latest version pinger2.pl that is not supposed to lose <BeaconsList> at 3 sites (cae.seecs.edu.pk, maggie2.seecs.edu.pk, aup.seecs.edu.pk). A month later (march 7, 2012) after install the <BeaconsList> appears OK on all three, i.e. for the <LastChecked> variable showed march 1st, 2012 for aup.seecs.edu.pk and maggie2.seecs.edu.pk, and showed March 7th, 2012 for cae.seecs.edu.pk. Joun may want to extend the deployment of the new pinger2.pl.

Joun has deployed pinger2.pl to 22 hosts, he will be sending us the list of these 22 hosts.

PingER Updates

Collection status webpage generates an email if a node does not work for more than 10 days. Non Pakistani monitors that we cannot gather data from include 1 in Nepal, 1 in Burkina Faso and one at Acme Security in Brazil. Les has sent email to Brazil. Amber will contact the other two.

Update the script to report bug that for first n days of a month we do not get report.

Amber is working on creating smokeping plots for today's data.

HEC Report - Anjum, Amber and Imdad

The report is in Case studies at [PERN Six monthly report](#).

Imdadullah will be writing an executive summary of the report which needs to be submitted before the end of this month. Now the report goes to other universities as well as to the head of HEC. We now need to be very careful in writing the report.

The report was supposed to be submitted last week of February. A few changes were made in the conclusions and an executive summary. Updates?

Pakistani Hosts

Joun and Ghulam have arranged to archive the status of Pakistani hosts in [<http://pinger.seecs.edu.pk/daily-report>]. There is a link from the PingER web site (<http://www-iepm.slac.stanford.edu/pinger/site.html>)

The IPs which start with 111.68.96.xxx are routed in such a way that the data can go out but it cannot come back to Pakistan. They will change the IP of all POP nodes. This arrangement was done for NCP. Only 0.5 MB of that link is provided to other universities. Any node that is having the IP as 111.68.96.xxx will be given a new IP.

There were 12 nodes with this problem which will be solved in the next week. SEECS will give them the IP addresses, and then HEC will change the IP of these nodes.

FSBD and MTN POP have high unreachability values, which is not acceptable. They are looking into it. Backhaul network is currently leased from PTCL however in 3-4 months they will replace it with their own network. There would be no commercial traffic on it. As a result it is expected that RTT and losses will improve drastically. So next 6 months are important for observing the network performance.

HEC is changing the DNS names of these hosts, we expect to hear from them before next meeting.

PingER Archive Site - Ghulam

Current Schema : see [here]IEPM:Pinger+PerfSonar schema].

Ghulam will modify the script to gather data half an hourly and put the analyze data in a separate table.

Sadia is working on getdata.pl to shift the data from SLAC files to database.

Future concerns:(Will be considered once the performance of above monthly aggregated data is observed)

1. How to store raw data for one year
2. How should it be sharded
3. For how long data should be in database

Sadia :Adding max RTT, MOS and Alpha to pingtable.pl and the analyze scripts

Ghulam there was some problem in pingtable.pl alpha value at SEECs. For some links alpha was having value of 200 . As we know alpha can have maximum of value 2. So there must be something wrong in calculation. **Progress**

Sadia has modified analyze hourly, analyze daily and analyze-monthly for Max RTT, MOS & alpha. She has run the jobs in batch for all days back to 1998. Next step is to copy the new metrics from the new daily and hourly files to the master aggregated files. Then we also need to update pinger.new.cf to add the new metrics for pingtable.

The output looks good. In particular the alpha values are very interesting, we should use these in TULIP (see below). The results may also be useful for identifying pairs of sites in Pakistan that have very indirect connections. We do have groups for KARACHI_REGION etc. that should assist in this.

Ghulam suggests to remove unrequired metrics from raw data files of pingtable.pl. Sadia thinks it wont make any difference unless number of lines being read are huge.

Ghulam thinks Null values in the data table are responsible for getdata.pl not working. We are unable to see any extra data in data table. Ghulam will look into getdata.pl again to get it working. Les will send him alpha subroutine incase he needs it. Les will also be sending Sadia latest version of pingtable.pl which she will send to Ghulam so that he makes it workable with getdata.pl. Ghulam will be sending Sadia his version of getdata.pl.

TULIP - Bilal

With the upgrading of the pinger.slac.stanford.edu host from RHEL4 to RHEL6 the mysql databases got lost. Sadia has worked with Fahad to recover the Tulip data base from the recovered sites.xml file. This is done.

Trying stress testing with reflector instead of reflex. Results are available at [Target Data for reflector tier all.[https://confluence.slac.stanford.edu/display/IEPM/Target+Data+using+reflector+with+Tier+All]] The results also have a comparison of reflex and reflector error in terms of distance. **Repeat this for Europe. ****Europe region stress testing* is completed.**

The result in Europe is not so good due to less number of working landmarks. Only PingER landmarks are working in Europe Region, no Planet Lab and PerfSonar landmarks are working. Also some landmarks are wrongly entered e.g. Pune (India), Germantown (USA), Menlo Park (USA) and Batavia, IL (USA) because these region are not included in Europe region. We know if we want to get more accurate result then we must have more number of landmarks in that region. If we see in Tulip Database then there is no landmark working from Germany, France, Denmark, Netherlands, Spain, Austria and Poland. Due to absence of active landmarks in these countries, the result are bad. Only 6 landmarks are working in whole Europe continent which are given in table below.

| City | Country | IP Address | Type | Tier | Region |
|-------------|----------------|--|--------|------|---------------|
| Vancouver | Canada | andrew.triumf.ca | PingER | 0 | Europe |
| Geneva | Switzerland | pinger.cern.ch | PingER | 0 | Europe |
| Germantown | United States | pinger.ascr.doe.gov | PingER | 0 | North America |
| Abingdon | United Kingdom | icfamon.rl.ac.uk | PingER | 0 | Europe |
| Ottawa | Canada | netmon.physics.carleton.ca | PingER | 0 | Europe |
| Menlo Park | United States | www-wanmon.slac.stanford.edu | PingER | 1 | North America |
| Batavia, IL | United States | pinger.fnal.gov | PingER | 0 | North America |
| Trieste | Italy | pinger.ictp.it | PingER | 0 | Europe |
| Warrington | United Kingdom | icfamon.dl.ac.uk | PingER | 1 | Europe |
| Pune | India | pinger.cdac.in | PingER | 0 | South Asia |

Sadia has updated the TULIP database with the missing PlanetLab landmarks. This adds another 38 landmarks in Europe.

Bilal needs to add the number of landmarks in the region for his Excel Spreadsheets. This will be useful for his paper, i.e. reporting typical accuracy as a function of landmarks in region.

Another interesting study would be whether using different alphas (in $\text{distance[km]} = \alpha * \text{min_RTT[ms]} * 100[\text{km/ms}]$) based on the alphas found in PingER for the various regions (see for example <http://www-wanmon.slab.stanford.edu/cgi-wrap/pingtable.pl?file=alpha&by=by-node&size=100&tick=daily&year=2012&month=03&from=United+States&to=United+States&ex=none&only=all&dataset=new&percentage=any>) provides much benefit compared to the single current value of alpha. To facilitate this we have added PingER groups for N.AMERICA, EUROPE, AUSTRALASIA, S. ASIA, S.AMERICA.

Bilal will be sending the tulip draft paper by the end of February. This was delayed by a family emergency.

Sadia is right about SLAC needing to spend ~ \$8K to get the full toolkit needed by CBG. Thus we will not implement at SLAC. Sadia will make sure TULIP at SLAC works for TULUP without CBG. **Progress**

Bilal will rerun the stress testing of North America using new reflex.

Bilal is writing a report with stress testing results of all of the regions. He will be sharing it when it completes.

Possible projects

- There can be a paper about Pinger if we could just find the right conference. MCN, ICC and Globecom do provide network monitoring topics. It could talk of the various metrics and their importance (in particular; MOS, Alpha, max RTT, min RTT), the lessons learnt from running such a worldwide infrastructure, the uses of the data etc.
- We can talk of GEO-Location experiences. For example within Pakistan it works fine, however as we go within regions or continents this gets worse. We can publish some stats on that for example. We can add the impact of changing alpha. We can also indicate the importance of landmark proximity.
- See [<https://confluence.slab.stanford.edu/display/IEPM/Future+Projects>].
- Extend the NODEDETAILS data base to allow entry support for whether the host is currently pingable.
- Improve the PingER2 installation procedures to make it more robust. This might be something for the person(s) in Pakistan who are responsible for installing PingER2 at the Pakistani monitoring sites. They probably have found where the failures occurs. Also look at the FAQ, and ping_data.pl which has been improved to assist in debugging, could it be further improved (e.g. provide access to the httpd.conf file so one can see if it properly configured)? There are 2 students working on the PingER archive. Is this something they could work on?
- [Fix PingER archiving/analysis package to be IPv6 conformant]IEPM:Make PingER IPV6 compliant]. Will build a proposal for an IPv6 testbed. They will try various transition techniques. A proposal has been prepared and that has been submitted to PTA. Adnan is a co PI. It is being evaluated today. A small testbed has been established in SEECs and the plan to shift some of the network to IPv6. Bilal is part of 3 students involved with PingER and they will be involved with IPv6. They are porting the PingER archive site to using a database. They have redeveloped the archive site using Umar's documentation. They have set up a small test archive site. They have gathering, archiving, analysis. They will design a new database. They will also try a port of PingER to IPv6.
- Look at RRD event detection based on thresholds and how to extend, maybe adding plateau algorithm. Umar's algorithm did not work in a predictable manner.
- Provide near realtime plots of current pinger data using getdata_all.pl/wget. It will work as a CGI script with a form to select the host, the ping size, and the time frame to plot. It will use wget or getdata_all.pl to get the relevant data and possibly RRD/smokeying to display the data.

Smokeying Realtime graphs-- Amber

The task involves three steps:

1. Saving the latest data of the request sent. -- This required changes in connectivity.pl and getdata-all.pl
2. Converting this latest data file to RRD file. -- This required changes in SelectConvSrcDest1.pl and PingERtoSmokeying.pm located at /afs/slab.stanford.edu/package/pinger/smokeying. Amber created a new script names PingERtoSmokeying-Amber.pl for this.
3. Generating smokeying graph from RRD file.-- This required changes to graph.pl. Amber created a new file named graph-Amber.pl for this.

Currently, Amber is unable to test if .rrd file generated by this script is in the right format or not. She is looking into it.

Future meeting time - Les

Next meeting on Tuesday 20th March, 2012 at 8:00 pm in US and Wednesday 21st March, 2012 at 8:00am in Pakistan.