

# 20120111 SLAC SEECS Meeting

## Minutes for SLAC-SEECS Meeting January 11th, 2012.

### General

Ghulam has filled out and returned a PDF of the DoE form. Sadia has sent it to the SLAC User Organization. **The concerned person at SLAC is on holidays till thursday.**

### IPv6 - Anjum and Ghulam (this has been de-prioritized until new database PingER is working)

IPv6 machine is working fine. Ghulam installed pinger2 on it and tried to collect data. But it was unable to resolve the IPv6 address. Seems like we need to do some changes in the whole architecture to make it workable with IPv6. Les looked at pinger2.pl, it verifies the address is IPv4 4 octets. He made a suggestion to use valid\_ip.pl to verify both ipv4 and ipv6 addresses. Ghulam are there other parts of pinger2.pl that need modifying. It already can access ping6. In addition one will need a copy of pinger.xml with IPv6 hosts and their addresses.

A possible project would be to make traceroute.pl work on a dual stack IPv6 host (say to traceroute to ipv6.google.com). Will need to look at gethostbyname etc. Is there any interest?

Another is to make pingtable.pl and getdata.pl IPv6 capable, again this could use valid\_ip.pl. Since Ghulam is working on these it would be good for him to add this.

### pinger2.pl

There is a bug in pinger2.pl that results in the <BeaconList> being empty. Les has made a fix that may help. **Sadia has created a new tar file /afs/slac/g/www/www-iepm/pinger/tools/pinger-2.0.3.tar.gz. It has been sent to Joun who will install to ensure it works.**

### Traceroute Archive - Sadia

Traceroute Archive documentation is required which can be referred in the ICFA reports . Sadia? **Done, There is a documentation available at [here](#).**

New Improvements in Traceroute Archive: [Click here](#)

Previously pulldown menu was showing default value even after selecting the new date and nodes. **Done, Now drop down menu's are being updated with the choice selected .**

Nodes name need to be sorted alphabetically? **Done**

**To have a better understanding of changes observed in three traceroute , coloring of each change observed in IP name or IP address is done. [See here](#)**

### ICFA Report - Les and Amber

- Amber is making sure the graphs use a common color scheme and formatting, removing anomalous Russian data from the throughput trendlines. **Done.**
- Amber is adding the African moves from terrestrial to satellite links. **Done.**
- Amber is working on the Pakistan case study for 2011.

### HEC Report - Anjum and Amber

**HEC report will have three sections:**

1. POP to POP Analysis: This will further have two parts
  - a. Inter-Regional POP to POP analysis: This will include analysis from ISL FJWU POP to all the POPs in different regions of Pakistan.
  - b. Intra-Regional POP to POP analysis: This will include POP to POP performance within a region. Only three regions have more than one POP nodes, Karachi, Lahore and Islamabad. Analysis will be carried out for these regions only.
2. PoP to Non POP Analysis: This will include two parts
  - a. Inter Regional POP to Non POP Analysis: A POP node in each of the 5 regions will be chosen and the analysis from these POP nodes to all the Non POP nodes of other regions will be carried out.
  - b. Intra-Regional POP to Non POP Analysis: A POP node in each of the 5 regions will be chosen and the analysis from these POP nodes to their own regions will be carried out.
3. Outliers: This section will include some interesting POP to POP and PoP to Non POP outliers with their behavior analyzation using traceroutes.
  - Overall network link performance will be considered as good or bad based on Standard metric values for VOIP as mentioned in Pinger Tutorial.
  - Amber and Imdad had a meeting at 8:30pm on 9th January. Imdad is now working on PoP to POP analysis section with ISL FJWU as monitoring POP instead of ISL HEC POP. **He has sent this to Amber which she will be adding in report.**
  - Amber is working on Inter-Region PoP to Non POP analysis.

### Pinger Explorer - Amber

- Amber has put together an email list of Pinger contacts to send the video to. We will send email about Explorer to the list.

## Status of Pakistani PingER hosts - Amber

1. There is a discrepancy between the PERN monitors reported by Joun as not working and those SLAC is able to gather data from (reported in [chec kdata](#), [here](#)). It was reported a couple of weeks ago. Kashif were Joun are looking at. For example we are unable to gather data from pinger.pern.edu.pk, there is something strange with the wget not working (see [here](#)). This maybe tied to some improvements to security installed on Dec 14th. **Les confirmed there are still problems affecting 4 PERN POP nodes and 12 other Pakistani monitors. See [Hosts pingable from some regions but not others](#) for more details.**
  - a. Currently pinger.pern.edu.pk is not pingable so we are unable to test further.
  - b. All nodes deployed at PERN PoPs are being upgraded to enhance security. Anjum had discussed with Umar as to what security features should be implemented on these nodes.
2. In addition some PERN POP monitoring hosts (pingerisl-fjwu.pern.edu.pk and pinger.pern.edu.pk ) are only pingable from Pakistan and Jordan. This may or may not be related. This needs to be resolved. It is probably related to Amber's observation that nodes such as nukhimain.seecs.edu.pk ; nuisb.seecs.edu.pk and pingerisl-fjwu.pern.edu.pk are pingable from SEECS but not SLAC. Using reflector.pl to ping nukhimain.seecs.edu.pk and also www.cern.ch, the number of landmarks able to ping nukhimain was 26, while for cern it was 106. It appears only landmarks in Pakistan, Algeria, India, Brazil, and Russia can ping nukhimain. Kashif and Joun are looking at.
3. Also <http://pingermtm.pern.edu.pk/cgi-bin/traceroute.pl?function=ping&target=www-wanmon.slac.stanford.edu> is not responding, it appears the web server or CGI script may be down.
4. Kashif reports we need a system for air university because they have a shortage of systems.

As updated on 01/10/2011.

Responsible person: Joun Muhammad

**HEC is sending out letters to the contact persons (who are non-cooperative) after which the nodes will be more reliable. 2-3 weeks will show much more stable nodes.**

Node	Status	Description
pinger.ustb.edu.pk	UP	Pinging by another IP. Data not collected. Issue will be resolved soon.
pinger.giki.edu.pk	Down	Off because of network issues.
hu.seecs.edu.pk	Down	Pinging but not fetching data. Trouble shooting in progress.
pinger.uar.edu.pk	Down	Network issue, will be up soon.
airuniversity.seecs.edu.pk	Down	System issue, will be up soon.

## PingER Archive Site - Ghulam

Ghulam has rebuilt the database.

- Zafar has raised concerns about using the perfSONAR schema. As a result we propose to extend the data and meta table's columns (to eliminate the need for joins).
- If we use the SEECS schema we have concerns over the number of rows exceeding 1 million and performance issues. SLAC proposes using monthly shards for the data. **We are unclear how SEECS would address the > 1 million rows. Have any tests been made?**
- SLAC is unclear whether the data at SEECS is stored for each hour or each 30 minute measurement. If the former then how are the SEQs and RTTs stored (need for out of order and consecutive packet loss probability).
- perfSONAR sets the interval between groups of pings to be 5 mins with a flat random distribution of 1 minute. The 5 minute setting is configured in Measurement Point Settings.
- Ghulam will work on getdata.pl to modify the scripts to use the database rather than flat files and to use parallelization for speed (**do we need parallelization, since today getdata.pl at SLAC takes about only an hour and we thought most of the time goes on the wget so it should not get much worse wusing the database as opposed to flat files**).
- If we do need parallelization then we will need to use parallel loops rather than parallel threads. Ghulam with the help of Zafar is removing threads and adding parallel loops. This would help in fixing the sorting issues of pingtable-db.pl. **Ghulam is working on it.**
- Pingtable.pl using database, aggregates data on fly for results.
- How long pingtable.pl takes to display the results for 12 months? Ghulam said that they did a testing few months back. Then pingtable.pl for 100 pakistani nodes to pakistani nodes took less than a minute to display the results. It was suggested to do a testing proof for that.
- Dr Les is inclined to the idea of PerfSonar database schema as it is well structured and a lot of support is present. It would not go worse than what is deployed at seeecs.
- Sir Umer suggested to use the idea of views in database which are logical tables having fields from different tables.
- Finally it is decided: we should go with PerfSonar schema using shards/views.
- Brainstorm the answers of the following few questions and other possible questions if ghulam or any one finds interesting, kindly do share with all:

1. What are typical queries for pingtable.pl
2. Which query has highest frequency
3. Do we need to shard the tables. e.g if we shard into monthly. it means there would be 1.8M rows for 1300 pairs with 48 pings a day in 30 days. It would be doubled if we use different packet size.
4. Which way is best to shard the data in terms of

- Time
- Region
- Sites
- Metrics

- Months etc

5. Put a document having

- New schema
- Queries which are run
- Time queries took to run
- two/three lines of content regarding the query and results
- graph would be a good thing to add up

Sadia should (Pending since schema gets finalized)

1. Document the database Schema
2. Data transfer from flat files to new database
3. See PerfSonar scripts to do the aggregation of data like IPDV etc

The new Schema proposed by Ghulam is placed [here](#) , you can view the changes in the schema on page 5 of the document. Page 6 has answers to the questions about most used queries etc. There is still a confusion whether we can merge location\_location and host as there is no physical connection between the two. This will help in reducing one of the joins. If we can merge the two tables, then the schema would be similar to SEECS schema with only little differences.

**Sadia :Adding MOS,max RTT and Alpha to pingtable.pl and the analyze scripts (awaits pingtable-db & getdata-db.pl working first)**

- Schema has been updated by Sadia. The new schema had PerfSonar and pingtable required fields in separate tables. To remove joins we are looking to add extra columns to the perfSONAR table for metrics like throughput, MOS, alpha etc.
- Sadia will work on migrating data from flat files to the new databases. **Sadia is working on it.**
- Analysis scripts to add Mean Opinion Score and Alpha, some things need to be correctly configured. It has been deployed at <http://pinger.seecs.edu.pk/cgi-bin/pingtable.pl> for testing.
- Alpha, max RTT and MOS to be implemented at SLAC site. Sadia will be doing this with the help of Zafar.

## TULIP - Bilal

Following table lists the targets in Europe which are not plotted on maps. For example the first target can be explored [here](#) . This can be compared to a [target which can be plotted on the map](#). Bilal looked into it and found that nodes are plotted using other GeoIP and IP tracking tools.

Country Name	IP Address	Progress
Austria	62.218.39.47	1/10/2012: Fails due to an unreachable host in Germany 212.201.44.81
Austria	212.33.36.188	1/10/2012: Fails due to an unreachable host in Germany 212.201.44.81
Italy	193.206.84.12	1/10/2012: Can be plotted successfully
Ukraine	193.29.220.3	1/10/2012: Can be plotted successfully

Amber looked at the host 212.201.44.81 (in Germany) in Pinger and Tulip database to find out why this host is unreachable. We do not have this host in Pinger Database, however, in Tulip database we only have one information about this host which is its IP. **Amber and Les decided to delete this host from Tulip database. Bilal has to rerun the script to find out if Austria still gets plotted or not.**

Deleting the unreachable host is a temporary solution. We still need to find out a permanent solution for plotting a landmark whose one of the nearest host is not working; Tulip should consider the other nearest landmark to plot it.

## CBG TULIP Integration -- FYP (Bilal)

- Bilal did some stress testing. The landmarks are 331 while the targets will be the ones generated by Sadia. He will compare the results with the 4 month old results with 59 hosts.
- **From the latest results it is apparent that if the landmark is also the target then we can get 0 error. Bilal has modified the tests to filter out such cases. Les has sent him the URL to the Landmarks file so he knows to filter out measurements from a landmark to itself. He will rerun the tests for N. America, Europe, S. Asia, E. Asia and Australia and will send the new results before Sunday.**

## PerfSONAR (Pakistan)

- **PerfSONAR at SEECS:** PerfSONAR throughput and latency nodes are now up and running at SEECS. Hostnames and corresponding IP addresses are:\*\* throughput measurement node: <http://psbw.seecs.edu.pk/> (http:115.186.132.154/toolkit/)
  - latency measurement node: <http://psla.seecs.edu.pk/> (http://115.186.132.155/toolkit/)
- Select options under "Service Graphs" to view throughput or latency graphs. Added 5 Stratum 1 NTP servers to cater for clock delay and everything seems to work fine.
- There are some interesting one-way latency graphs at 115.186.132.155 (SEECS PerfSONAR Latency node). **Dst to Src (e.g MIT to SEECS) latency is less than Src to Dst (e.g SEECS to MIT) latency.** This might uncover some trends in outbound network traffic from Pakistan.
- Bilal and Ghulam will have a meeting with Zafar to know about PerfSonar and to maintain it in future. **Update?**

## Possible projects

- There can be a paper kind of talking on Pinger if we could just find the right conference. MCN, ICC and Globecom do provide network monitoring topics. 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 are yet not ready for Tulip paper.
- See [<https://confluence.slac.stanford.edu/display/IEPM/Future+Projects>].
- Extend the NODEDETAILS data base to allow entry support for whether the host is currently pingable.
- Extend Checkdata to provide emails automatically, see [<https://confluence.slac.stanford.edu/display/IEPM/Extend+checkdata+to+make+it+more+useful>]. Many of the ideas in the script node-contacts.pl are a step in this direction.
- 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.

## Future meeting time - Les

1. Next meeting on Wednesday 18th January, 2012 at 8:00 pm in US and Thursday 19th January, 2012 at 9:00am in Pakistan.