

# Copy of 20140305 SLAC SEECS and UNIMAS Meeting

## Time & date

Wednesday February March 19, 2014 9:00pm Pacific Standard Time, Thursday March 20th 2014 10:00am Pakistan time, Thursday March 6th, 2014 1:00pm Malaysian time, Thursday March 6th, 2014 3:00am Rio Standard Time.

## Attendees

### Invitees:

Anjum, Hassaan Khaliq, Kashif, Raja, Johari, Nara, Abdullah, Badrul, Ridzuan, Ibrahim, Hanan, Saqib, Adib, Les, Renan, Bebo

+ Confirmed attendance

- Responded but Unable to attend:

### Actual attendees:

## Administration

- **Hassaan has been added to the pinger-my email list.**
- UUM are interested to host a PingER workshop at UUM in 2014. However, they have two issues, one is the budget (they need to run this workshop on cost recovery basis); second is their commitment to MYIGF which is hosted by UUM (btw 12-13 august). There was no update 3/5/2014.

## Renan

Renan realized that testing on Virtuoso is going to take longer than he thought because of the big number of triples. He did not have enough time to continue his tests.

Renan completed writing the paper on Linked Open data access to PingER data. Bebo 7 less have reviewed and commented. It was submitted to the Stanford conference March 1st, 2014. The deadline for the submission was extended from March 1st to March 15. The notification is May 1st, the final submission is May 9th, conference is May 27-31. Poster submissions are due April 4th, we might submit a poster just in case. The acceptance rate is very low. Information on the conference is at <http://www.scienceengineering.org/ase/conference/2014/bigdata/sanjose/website/paper-submission/>.

Les will try and set up a Skype meeting with Renan and Bebo to go over strategy and also to discuss Ibrahim's work (see below).

Les will send another email to Vinicius and Renan to see if we can get the UFRJ host running again (the web server is not working).

It would be good if somebody else was able to work with Renan to move PingER LOD forward. However we think the performance issue has to be resolved first.

## UM

UM appears to be experiencing very high packet losses (>4%). We need help in understanding these. It does not look like simple congestion since the jitter and RTT is staying low. Saqib needs assistance from Badrul as part of his case study of Malaysia. Badrul writes: "UM is now currently changing the new policy regarding the bandwidth and connection. They are upgrading the switch and router to a new router - mostly every weekend they shut down the network. therefore we are now having some difficulty." Saqib has identified what appear to be losses starting at a JARING router in KL. Badrul will run some mtr trace/route/pings to various sites. No update 3/5/2014.

Badrul plans for his undergraduate student (Abdulrahim Haroun Ali) to have a draft paper on anomalies in PingER measurements to share for review by the end of semester (January). No update 3/5/2014.

Ridzuan successfully setup hadoop environment for the servers. He is now looking into how best to input the pinger data to the environment and type of analytics that can be done. Ridzuan is working through some issues due to lack of IPv4 addresses and the need to use IPv6.

Ibrahim Abaker is planning to work on a topic initially entitled "leveraging pingER big data with a modified pingtable for event-correlation and clustering", The purpose of this project is to make Pinger more efficient and more scalable to retrieve and analysis pingER raw data. Consequently it may improve pinger data store. He is planning to use key-value store model with high expandability and shorter query response time. The benefit of such project will help in the design of real time analysis and distributed pinger data across multiple servers for analysis. Les has sent him a pointer to documentation on the data flow at SLAC and also listings of pingtable.pl and associated script. He is currently writing a proposal, which will contain a problem statement, objectives, scope and the approach to tackle the problem. He asked: how much PingER data is captured and stored so far (~ 1TByte)? Is there a published paper (see <http://www-iepm.slac.stanford.edu/paperwork/ieee/ieee.pdf>)? Bebo & Les will set up a Skype meeting with Renan to get an update and see about setting up direct contact between Renan and Ibrahim. Ibrahim provided a copy of his proposal (see <https://confluence.slac.stanford.edu/download/attachments/17162/leveraging%20pingER%20big%20data%20with%20a%20modified%20pingtable%20for%20event-correlation%20and%20clustering.docx>).

## UNIMAS

No progress 2/19/2014. Johari was not on the call.

The Raspberry Pi is at the data centre, but due to miscommunication, the public ip which was previously assigned to the rasp pi unit has been assigned to another server. So now Johari has to wait for the ip to be reassigned back to rasp pi unit and give new ip to the other server. They suggested a new ip for the rasp pi unit but Johari wanted to have the old one so that we do not have to update the record at slac for the archival process, etc

The tool to enable synchronizing Malaysian monitors is completed. It provides the ability to add, edit and delete nodes. Raja has tested and after a mod was made to replace a private IP address it is working well. Badrul and Saqib need to try it out.

The traceroute server at <http://pinger2.unimas.my/cgi-bin/traceroute.pl> has the same problem as before. They know (sort of) the problem but haven't got the chance to rectify it (mapping NAT address, needs to be added). There is no progress 12/4/2013, 1/8/2013, 1.22.2014, 2/5/2014. Now that the historical traceroutes are working for UM (see below) there is an extra incentive to get the reverse traceroute working at UTM and UNIMAS

Custom iso: He can get as far as the boot screen, but is unable to get to the desktop. No progress 2/5/2014.

Johari has created a shell script to automate the installation of pinger package in Ubuntu/Linux distro and has added Centos 6.2 and earlier versions as well. Next he will test on Fedora 19 and 20. He is using SLAC repo version 2.3 with a virtual box and ubuntu server 13.04 for testing purpose. Johari has added a page at [pinger.unimas.my/pinger](http://pinger.unimas.my/pinger) website on the usage of a shell script to automate the installation steps for pinger package. It is available at <http://pinger.unimas.my/pinger/install-tutorial.php>. Johari will check with Kashif to see the version(s) needed for Pakistan to ensure it works for them. Then Kashif can try it out.

Johari has a research student who finalized a proposal in order to officially apply for his masters. He will start in February. He is currently working on threshold/anomaly detection, and will extend to correlating performance over multiple routes. He will share the proposal with Les and others in the next 2 weeks to get reviews.

## UTM

No progress 2/19/2014

Saqib received an email regarding the successful selection of research assistant-ship at UTM. However, he is waiting for an official offer letter from research management center.

Saqib is working with the IT department in UTM to solve the problems of delays in traceroute. They have a new ISP. Now the pinger server is working with single static IP (161.139.68.188) both for intranet and internet. Previous internet IP 161.139.146.158 is removed. Again given the new historical traceroute facility (see below), it is important to get this working.

- No load balancing is involved as UTM is working with single ISP.
- Still the problem of the delay in traceroute is not solved. It appears to be due to the DNS lookup delay.

Saqib has started a case study of Malaysian hosts seen from Malaysian Monitors. The case study is identifying the need for the Malaysian monitors to collect traceroutes to their target hosts daily. SLAC has a script to facilitate this, for routes seen from SLAC, that could be adapted for other sites. The latest copy of the word file was shared with Badrul and Johari. No update 2/5/2014.

## UUM

Adib was unable to attend the meeting. He provided an update by email: "The PingER monitor will be up soon"

## NUST

The PingER monitor at SEECS/NUST (maggie1) is being rebuilt on a different computer following a crash. There are 2 approaches that can be tried in parallel:

1. Rebuild from scratch. In the initial case Umar took 2-3 weeks. He wrote documentation, so this will be a test of the documentation. This could be valuable should Malaysia or another site decide to set up an archive. Kashif and Joun are ready to start on this.
2. Dump and restore the database. This is much easier. However it requires the database password. Les will contact Umar to see if he can provide the password. Umar cannot logon to maggie1 since it is blocked by the firewall.

The BeaconLists at NUST and SLAC are out of sync. Raja will compare the two pinger.xml files. It turns out NUST was not using the HostList for the non Beacons and so had an expanded BeaconList which was updated from SEECS. Raja has explained this to NUST and believe it will be easy to synchronize.

Hassaan reports that they have established contact with people in Bahawalpur. The contact person forwarded the request to higher authorities for approval of PingER deployment. Same status 3/5/2014

The installation of PingER at Sahiwal is complete. It just needs the Beacon List.

We were gathering data from Pingerisl-air (this is continuing to work), and Quest (stopped working Feb 13)

We are still unable to gather data from airuniversity, buitms, cae, duhs, pingerisl-qau, [sau.seecs.edu.pk](http://sau.seecs.edu.pk), uaf.

Monitors	Status
Airuniversity	New installation will be install in a week
Buitms	Problem in hardware, they are trying to fix it. <b>System Replacement Required</b>
Cae	Problem with IP, they are trying to resolve it
Duhs	They are purchasing new servers in two week time

[pinger.pern.edu.pk](http://pinger.pern.edu.pk) Fixed Feb 19.

Pingerisl-qau Contact person is still on leave. He is now busy in other pending works.

[Quest.seecs.edu.pk](http://Quest.seecs.edu.pk) down since February 20th.

[sau.seecs.edu.pk](http://sau.seecs.edu.pk) for all of January. Contact person is not attending

Uaf Problem in system, visit approval for Faisalabad University in progress. Approved for 12th Feb

Anjum pointed out that there are working PingER monitors that SLAC appears unable to gather data from. An example is the Comsats host. Les takes this to mean [pinger.comsats.edu.pk](http://pinger.comsats.edu.pk) and will investigate.

At the Connect Asia Pacific Summit in Bangkok in January and seeing the project "Mapping the pan Asia Pacific information Superhighway and closing gaps in infrastructure connectivity" Shahryar found that very much related to the work in the PingER project. So Shahryar sent email to a UN agency for a possible collaboration with them on PingER project. He has heard nothing so he will write a detailed proposal and then should contact them again. No update 2/5/2014, 3/5/2014.

## TULIP - Raja

Please try out <http://www-wanmon.slac.stanford.edu/cgi-wrap/reflector.cgi?function=vtrace> It is a visual traceroute. It uses TULIP to locate the routers. It is in beta test and needs testing by friendly users. We are looking to add a link to it from the perfSONAR home page.

Raja has added 56 more perfSONAR landmarks, most of them are in North America. Currently we have 298 active landmarks (from total of 918).

The Visual traceroute and TULIP have been modified to support IPv6. Also Raja has improved the laundering.

Raja is working on a paper on TULIP implementation and on the Visual traceroute that uses TULIP.

Raja is adding a web page to select a landmark to request it to ping a target.

- Tool for pinging from PlanetLab landmarks <http://www-wanmon.slac.stanford.edu/cgi-wrap/reflector.cgi?function=plping>
- TULIP landmark list with Ping and Traceroute URLs <http://www-wanmon.slac.stanford.edu/cgi-wrap/reflector.cgi?function=landmarks>
- TULIP landmark map with Ping and Traceroute URLs <http://www.slac.stanford.edu/comp/net/wan-mon/viper/tulipmap.html>

We have sent a reminder email to Yahoo giving them an update on TULIP and how it might be extended to meet their needs. They have responded, they appear to be interested. We are trying to interest them in providing landmarks. No update 2/18/2013.

## Pinger at SLAC

We put together a slide presentation to ICFA in Geneva, to be given by Harvey Newman. It is at: [http://www.slac.stanford.edu/xorg/icfa/icfa-net-paper-jan14/ICFASCIIPresentation\\_MonitoringGroupSlides20140205.pptx](http://www.slac.stanford.edu/xorg/icfa/icfa-net-paper-jan14/ICFASCIIPresentation_MonitoringGroupSlides20140205.pptx)

Raja has also re-written the historical traceroute facility that takes nightly traceroutes and allows one to retrieve and compare them. He then extended it to use the reverse traceroute facility to provide histories for NUST, CERN, UM as well as SLAC. When the UNIMAS and UTM traceroutes work, we can include them. See for example <http://www-wanmon.slac.stanford.edu/cgi-wrap/traceroutearchive.cgi>

There is an empty pinger/xml file at ICTP. Les & Raja are studying. The Pinger hosts at FIU in Florida is now restored. Les is contacting the Jordanian PingER monitor contact to restore it.

## Old Items

### Linked Open Data

Renan finished the new pingerlod web site. The new thing is that it should be much easier now to modify the info texts. What Renan did was to put the texts into a separate [file](#). The new version has been loaded on the server and some text added to describe how to use the map. However there is a bug that prevents it from executing the map. Renan reports that the bugs should be easy to fix. He has talked to his professor who suggested trying RDF Olink, it should have faster responses to queries. Renan will research this. It will probably mean reloading the PingER data so is a lot of work, hopefully this will improve performance. Before the rebuild he will make the fixes and provide a new WAR for us to load on pingerlod.slac.stanford.edu. He is also working on documentation (he has finished the ontology and has a nice interactive tool for visualizing it, since the ontology is the core of the data model of our semantic solution, this will be very helpful for anyone who uses our system, both a developer of the system and a possible user) and his thesis. Bebo pointed out that to get publicity and for people to know about the data, we will need to add pingerlod to lod.org.

Things he will soon do regarding documentation:

1. A task/process flow writing all java classes involved on all those batch jobs;
2. A Javadoc <<http://www.oracle.com/technetwork/java/javase/documentation/index-jsp-135444.html>> which will explain all classes and how they are used.

For the Linked Open Data / RDF which is in pre-alpha days, you can go to <http://pingerlod.slac.stanford.edu>. As can be seen this page is not ready for prime time. However the demos work as long as one carefully elects what to look at:

- Click on Visualizations, there are two choices:
  - Multiple Network Metrics: Click on the image: gives a form, choose from Node pinger.slac.stanford.edu pinging to www.ihep.ac.cn, time parameters yearly, 2006 2012, metrics throughput, Average RTT Packet loss and display format Plot graph, then click on submit. In a few seconds time series graph should come up. Mouse over to see details of values at each x value (year).
  - A mashup of network metrics x university metrics Click on image: gives another form, pinging from pinger.slac.stanford.edu, School metric number of students, time metric years 2006 2012, display format plot graph, click on submit. Longer wait, after about 35 seconds a google map should show up. Click on "Click for help." Area of dots = number of students, darkness of dots = throughput (lighter is better), inscribing circle color gives university type (public, private etc.) Click on circle for information on university etc.
- Renan will be working on providing documentation on the programs, in particular the install guide for the repository and web site etc. This will assist the person who takes this over.

Renan is using OWLIM as RDF Repository. He is using an evaluation version right now. Renan looked into the price for OWLIM (that excellent RDF Database Management System he told us about). It would cost 1200EUR minimum (~ 1620 USD, according to Google's rate for today) for a one time eternal license. It seems too expensive. No wonder it is so good. Anyhow, he heard about a different free alternative. Just not sure how good it would be for our PingER data. He will try it out and evaluate. He will also get a new evaluation of the free OWLIM lite.

He has also made some modifications on the ontology of the project (under supervision of his professor in Rio) hence he will have to modify the code to load the data accordingly.

Renan has provided a 4 page Appendix on PingERLOD to the ICFA report. This is also available at [PingER LOD Overview](#)

### Raspberry Pi

A quick comparison of the performance of the two hosts (raspberry pi and regular UNIMAS host) without statistical quantification is available at <https://confluence.slac.stanford.edu/display/IEPM/Comparison+of+PinGER+RTTs+from+UNIMAS+monitors+N4+and+RASPBerry>. A page has been created to compare the hardware spec between the [pinger.unimas.my](http://pinger.unimas.my) node (Intel architecture) and the [pinger2.unimas.my](http://pinger2.unimas.my) node (Raspberry Pi ARM architecture), available from the unimas pinger website at <http://pinger.unimas.my/pinger/hardware.php>. There is a link to hardware.php in the [Comparison+of+PinGER+RTTs+from+UNIMAS+monitors+N4+and+RASPBerry](https://confluence.slac.stanford.edu/display/IEPM/Comparison+of+PinGER+RTTs+from+UNIMAS+monitors+N4+and+RASPBerry) web page.

#### Tulip

We are looking at submitting a paper to IMC 2014. April 30, 2014 (see <http://conferences2.sigcomm.org/imc/2014/cfp.html>) is the submission date. This would discuss how TULIP works and share our experiences. Thus we might focus on: managing the appearance/disappearance of landmarks; the selecting of optimum timeouts for speedup plus any other ideas for speedup; where one could go next (e.g. colocation of landmarks on Yahoo, Google, Hotmail ... sites; impacts of heavy use and scaling (e.g. need for cluster for reflector, multiple landmarks at sites to not overwhelm a given landmark); concerns about network utilization); the relation of accuracy to landmark density (big item). Also more on potential uses such as visual traceroute, identifying proxies. Also maybe a bit more on the actual user interface (results from multiple sources besides TULIP) and implementation (parallel threads, tiering and adding new regions). The paper would assume we have a working TULIP as described in the other paper.

#### Follow up from workshop

- Hossein Javedani of UTM is interested in anomalous event detection with PingER data. Information on this is available at <https://confluence.slac.stanford.edu/display/IEPM/Event+Detection>. We have sent him a couple of papers and how to access the PingER data. Hossein and Badrul have been put in contact. Is there an update Badrul?

The Next step in funding is to go for bigger research funding, such as LRGS or eScience. Such proposals must lead to publications in high quality journals. They will need an infrastructure such as the one we are building. We can use the upcoming workshop (1 specific session) to brainstorm and come up with such proposal. We need to do some groundwork before that as well. Johari will take the lead in putting together 1/2 page descriptions of the potential research projects.

1. Need to identify a few key areas of research related to PingER Malaysia Initiative and this can be shared/publicized through the website. These might include using the infrastructure and data for: anomaly detection; correlation of performance across multiple routes; and for GeoLocation. Future projects as Les listed in Confluence here <https://confluence.slac.stanford.edu/display/IEPM/Future+Projects> can also be a good start and also Bebo's suggestion.
2. Need to synchronize and share research proposals so as not to duplicate research works. how to share? Maybe not through the website, or maybe can create a member only section of the website to share sensitive data such as research proposal?

Anjum suggested Saqib, Badrul and Johari put together a paper on user experiences with using the Internet in Malaysia as seen from Malaysian universities. In particular round trip time, losses, jitter, reliability, routing/peering, in particular anomalies, and the impact on VoIP, throughput etc. It would be good to engage someone from MYREN.

#### Potential projects

[See list of Projects](#)

### Future meeting - Les

Next meeting Wednesday March 19th 2014 9:00pm Pacific Standard Time, Thursday March 20th 2014 9:00am Pakistan time, Thursday March 20th, 2014 noon Malaysian time, Thursday March 20th, 2014 02:00am Rio Standard Time.

### Coordinates of team members:

See: <http://pinger.unimas.my/pinger/contact.php>