

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764
Fax: 781-981-0590

June 11, 2004

TO: Distribution
FROM : David Lapsley
SUBJECT: 24 May 2004 e-VLBI telecon summary

Attendees:

Bill Fink, Lee Foster, Paul Lang, Mary Shugrue, Bill Wildes – GSFC
Terry Gibbons – Lincoln Laboratory
Jerry Sobieski – University of Maryland, Mid-Atlantic Cross-roads
Charles Yun – Internet2
Kevin Dudevoir, Hans Hinteregger, David Lapsley, Alan Whitney – Haystack

This telecon is one of an ongoing series of telecons to prepare for 10 gigabit/sec e-VLBI demonstrations between NASA GSFC and MIT Haystack Observatory using a combination of network facilities including all or part of Glownet, Bossnet, ISI-E, SuperNet, Max and GSFC/HECN.

ACTION ITEMS ARE HIGHLIGHTED IN RED.

Bossnet

Alan Whitney: OK, let's start. Let's go through the block diagram. I'll let David talk about what we are doing up here. We just committed to spend some money.

David Lapsley: That's right, the Purchase order for upgrade equipment went from MIT to Movaz Networks last Thursday. Now waiting for delivery of equipment. Should be delivered anywhere from 2–4 weeks.

Alan Whitney: Terry, once equipment arrives we need to co-ordinate with you to install equipment.

Terry Gibbons: Yes.

Alan: We should co-ordinate with both Terry and Jerry. We will keep you up-to-date on this.

Jerry Sobieski: David, you might want to check with Henry about whether Richard Solis (Movaz engineer) can come out and help with that. That might be a day's work or so.

David: I will talk to Henry about that.

Jerry: Would be helpful to have someone who is familiar with the equipment to help install it.

David: Terry, is there any other information that we need to tell Movaz? Power levels etc.?

Jerry: One issue with DC power for those devices. Need to make sure that install locations have DC power.

Alan: Terry, we need to co-ordinate with Russ for equipment to be installed here.

Terry Gibbons: Yes. There is also a piece at Russ's side for Glownet where wavelength is converted from 1500 nm down to 1310 nm and then forwarded to me. Russ will have to co-ordinate this.

David: Sounds good.

Alan: Anything new at Lincoln?

Terry: Not really. I have identified a couple of DWDM MUX/DEMUXes out of some on-ramp switches that I think I can scavenge and put into place. We are not going to be getting this optical add drop from NRL.

Jerry: You're not going to get it at all or is it delayed until some future date?

Terry: I'm not holding my breath. I have MUX and DEMUXES that cover this wavelength plus 4/5 other ATDnet wavelength. That should take care of the immediate need.

Alan: We thank you Terry and Jerry for all of the co-ordination that had to be done to make this happen.

Jerry: I have a high confidence level, but there is a fair amount of work that needs to be done to make sure it all works. On the south end at Eckington, we will be installing a Ray Express there. We can take the hand-off from ATDnet from the switch. The fiber is going in as we speak on the Qwest ring. I've asked them to expedite segment from Eckington to ISI-E which would allow us to progress the Haystack wave to Tom's place as early as possible. We can get that back into shape and begin testing down to ISI-E and via conventional IP to Goddard. Second phase once the whole ring is up. The lateral for the University of Maryland, where all the College Park activity occurs needs to be constructed. Won't be spliced in until the end of June. Might be moved up, but wouldn't depend on it. Can't easily pull it back to College Park and then over to Goddard until the ring is up. I am hoping that the ring will be up and operational by the fourth of July.

Alan: Sounds good. How is the installation of other things, such as your optical switches going?

Jerry: We have been doing a lot of work in the last month. Paul, the fire brigade run that goes from College Park over to Goddard?

Paul Lang: That is supposed to be in this month. We'll check on the status.

Jerry: Let me know. That is currently the long pole on moving the Ray Express out to Goddard. At least as far as I am aware. If there is something else about building locations at Goddard, we will have to work on those as well. We have to get that up in the priority queue. The boys out on the DRAGON project have successfully gotten GMPLS extensions to RSVP to work. We are actually doing some dynamic signaling from server to server over the optical cloud. This is very good. A few more days to polish it up and as soon as fiber out to Goddard is ready, we will be moving the Ray Express out to Goddard. Finally things are starting to come together. I may be talking to David in the next couple of weeks about how we start integrating the stuff into a real application.

David: Sounds good.

Alan: Very good, thanks Jerry. Anything at Goddard to report, Paul, Bill, Mary, Lee?

Paul: We were trying to check on some stuff to blend in to the DRAGON project. We had a discussion with Jerry on Friday. We're looking at running DWDM XFP into the external wavelength module and filter/OADM MUX/DEMUX to get it into the DRAGON network. We have found two different possibilities for DWDM XFPs: Optnix and Finisar. We have some particulars to nail up, which wavelength we should use. We have problems with our procurement system, so we need to get purchase request in as soon as possible. One of those vendors, Optnix, makes DWDM Xenpacks as well. That was our initial approach, using a Xenpack in an Extreme Networks summit 400, but it looks like we can run it through the Force10 instead, which is our preferred path. We were just going to use the Extreme as a converter which is a little on the expensive side.

Hans Hinterreger: Does the Force10 only accept XFPs?

Paul: Right. They will have that. I have been asking Extreme to do an XFP module for the Summit 400. It's a daughter card. I was hoping they could spin up an XFP daughter card. They were saying that the XFPs were more expensive and less mature. It appears that the difference has been caught up, we have found the XFPs to be less expensive than the Xenpacks.

Hans: They are designed to be.

Paul: Yes, and they are a lot smaller too. Xenpacks are quite big: an inch high, couple of inches wide, 8 inches long.

Hans: I was encouraged to hear that you got two sources for DWDM XFPs. They were first announced a year ago.

Paul: Yes. Neither one is currently released yet. We are trying to get pre-release versions of them. I have sales contacts if anyone is interested. We have some particulars to try and get this hooked into DRAGON. We have been working with Jerry and Henry on this.

Alan: Good, thanks. Anybody have anything to work on status updates or changes in the network. Anything to report on 10 Gig testing at Goddard? Paul, Bill?

Paul: Bill did a test where he sent a petabyte worth of data.

Bill Fink: We took our existing Force10 E300 and had a loaner E300 that we took over to the MAX at College Park. We did some 10 Gig testing across it. We verified that we could do full line rate. We sent a petabyte of data across the link and sustained that. Took 6 days to send that. Sustained the full line rate for that period.

Alan: What sort of error rate did you see?

Bill: It was negligible. We are doing TCP so if there were any drops I would not have been able to get line rate. On the UDP there were a few drops (on the order of less than 100) for the entire 6 day period, doing a UDP test.

Alan: That sounds very impressive. Anything more to report on that. Anyone else have anything to report on testing, 10 GigE or otherwise?

Bill: Nothing to report yet, but relating to the DRAGON applications, planning to make modifications to nuttcp to enable it to use RSVP to setup a light path before sending the data across the network.

Paul: Bill got that request from Jerry.

Hans: Did you say 1 petabyte in 6 days?

Paul: We only had the one 10 GigE link and we were doing bi-directional.

Bill: Half a petabyte in both directions.

Paul: We did a local test where we did 20 Gbps between two different buildings, so it only took 3 days there.

Alan: Let's talk about some of the e-VLBI work. Both David and Kevin have been very busy and have things to report.

e-VLBI Experiments

Kevin Dudevoir: Status of Kokee link same as last meeting. Roger Hall was finishing his testing and reconfiguration trying to get us more bandwidth and solve our problem. Ended up with symmetric performance for UDP and we can sustain 60–80 Mbps. TCP still suffers from packet loss, anywhere from 20% to 0.1%. So it really cripples our hopes of using TCP at all. We think we can use one of the UDP-based protocols, Tsunami and/or UDP. We are just ramping up to start our test application which is transferring data from Wetzell and Kokee 4/5 times a week. They have an experiment that records at 128 Mbps and we figure we can transfer that at 30 Mbps from both sites. Should take 3 hours for one hour's worth of data. Hopefully by the next meeting we should be deep into production.

David: I've been working on a few different things. One of the things is FAST TCP, the Caltech folks released the source code for that a couple of weeks ago, so now you can download the patches and patch a standard Linux kernel, compile it and get it up and running. We've swapped out a Mark5 down at ISI-E and we put FAST on that. This is the first Mark5 we have got that is able to make use of FAST. I did some quick tests over the weekend, just some very quick and dirty tests, testing it out of the box without any optimization to the FAST protocol and was able to get approximately 330 Mbps over a 170 ms round trip time path going to servers that are in Japan. That was pretty good because when I compared that to regular TCP performance that we get, it is about 6.7 Mbps for the same path. So that's quite an improvement. So we are looking forward to deploying that on a few more Mark5's. This morning I downloaded the patched source to the Mark5 at GGAO and I am looking at getting that running there as well.

I've also been talking to the guys at Netlab about doing some other work. We are going to have a conference call tomorrow to discuss some other work that we can collaborate on.

I've also been talking to the SLAC guys as well, they are adding that Mark5 to their performance testbed. They have a testbed where they test a whole bunch of different protocols around the US. So it should be interesting to see what results they get from there.

I've also been working on the RTP software. That is the VSI-E software that I mentioned on the last call. That is getting very close. It took a while to get the performance up with all of the processing that it was having to do to encapsulate/decapsulate all of the VLBI data. At the moment it is running at peak rate at about 500 Mbps which I think is good enough for us to be able to make use of it over the long distance international and domestic links.

Bill: On the FAST TCP, is that something that is dynamically loadable/unloadable?

David: No, it patches the TCP software itself. So, it is actually compiled into the kernel. But I think there are flags that you can set within the /proc filesystem that you can use to turn it on and off. I haven't looked into that yet, but I will look into it.

Alan: Any comments or questions? Anyone have anything else to report along these lines? We are slowly making progress here. David hopes to have a first release of the RTP software on the first of June for deployment for testing on a wider scale.

Kevin: Was the RTP testing on Mark5's?

David: That was actually on Turtle and Enterprise. The reason for the throughput limitation is CPU bound. So, at the moment the testbed is a Pentium III 1.1 GHz and Pentium 4 Dual Xeon 2.4 GHz. The receiver is the limiting factor, so I have been running the software with the Dual Xeon as the receiver. It is running in user space which is one of the problems trying to get the performance given that it is running in user space.

Miscellaneous

Alan: Next item is the performance testing and monitoring. Does anyone else have anything else to say on that? Charles do you have anything to say on that?

Charles: I must apologize. I have been on the road for the last month. I have no real updates, but I do have a question for David on some work that I think Guy Almes may have mentioned to you about splitting up long links with a repeater box?

David: No, he hadn't actually mentioned that to me. That sounds like an interesting idea.

Charles: I will take that offline with you, if you have minutes after the call.

David: Sure.

Alan: Anyone else have any comments?

Paul: Yes, one other comment. I checked with the Cisco sales guy and apparently they do key their Xenpacks, which means that you can't use another vendor's Xenpacks. Which means that since they don't have DWDM Xenpacks right now, you are out of luck until they come up with them. Apparently, they check serial numbers on the SFPs, the GigE ones. Some people we know have already run into problems with that with the ones they bought from Cisco, which have some glitch on how the serial numbers got set.

Alan: Well, they will probably end up shooting themselves in the foot. Anybody else have any questions, comments, reports on that? Anybody else have any other subjects they would like to bring up? Well, maybe we can make this meeting short like the last one. Date for the next telecon? I won't be here for the 14th, but after the 16th I will be here.

Bill: How about the 21st?

Alan: Sounds good. 21st of June for the next meeting.

Next telecon is scheduled for Monday, 21st June 2004 at 2 pm EDT.

cc: Steve Bernstein, LL
Jim Calvin, LL
Rick Larkin, LL
Lorraine Prior, LL
Peter Schulz, LL
Leslie Weiner, LL
Herbert Durbeck, GSFC
Bill Fink, GSFC
Lee Foster, GSFC
Pat Gary, GSFC
Andy Germain, GSFC
Kerry Kingham, USNO
Chuck Kodak, GSFC
Kevin Kranacs, GSFC
Paul Lang, GSFC
Aruna Muppalla, GSFC
Mary Shugrue, GSFC/ADNET
Bill Wildes, GSFC
Dan Magorian, UMCP
Tom Lehman, ISI-E
Jerry Sobieski, MAX
Guy Almes, Internet2
Charles Yun, Internet2
Richard Crowley, Haystack
Kevin Dudevoir, Haystack
Hans Hinteregger, Haystack
David Lapsley, Haystack
Jason Soohoo, Haystack
Arthur Niell, Haystack
Joe Salah, Haystack