

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

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

1st March 2004

TO: Distribution
FROM: David Lapsley
SUBJECT: 23 February 2004 e-VLBI telecon summary

Attendees:

Bill Fink, Lee Foster, Pat Gary, Paul Lang, Mary Shugrue – GSFC
Dennis Baron – MIT
Tom Lehman – ISI-E
Terry Gibbons – Lincoln Laboratory
Charles Yun – Internet2

Kevin Dudevoir, Hans Hinteregger, David Lapsley, Arthur Neil, Alan Whitney – Haystack Observatory

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 Gownet, Bossnet, ISI-E, SuperNet, Max and GSFC/HECN.

ACTION ITEMS ARE HIGHLIGHTED IN RED.

Glownet/Bossnet

Tom Lehman: Jerry is working on Bossnet this week. Making plans to move the Bossnet rack. Then Jerry would hook in the DRAGON fiber. On the DRAGON fiber, I think Jerry is about a month and a half away from hooking it in. Terry can fill us in on whether or not Bossnet is being moved

Terry Gibbons: Qwest is doing a fiber survey today in Eckington. Our rack, which is currently in the transport area scheduled to be moved to co-lo area on the 5th of March. It would be adjacent to the ATDnet and DRAGON racks. Interconnection between DRAGON and MAX still TBD.

Tom: Guessing that time to get DRAGON fiber connected is a month to a month and a half away.

Alan Whitney: This is required in order to get Haystack reconnected to the network?

Tom: Could come up with an intermediate solution. How hard should we work towards a quicker, intermediate solution?

Alan: A month and a half is a bit further out by a few weeks than what we were expecting

Tom: Need to talk with Jerry about contractual details etc.

Alan: Until this is finished, we can't get anywhere on Bossnet including to Goddard?

Tom: We need to do some things to make this happen.

Alan: We are interested in the short term in getting a connection to Goddard - to GGAO in particular.

Tom: Will talk to Jerry and Terry to see if there is anything we can do for intermediate connection while waiting for DRAGON fiber. Can report back to Alan prior to next meeting to see if there is something we can do for a reasonable amount of effort.

Alan: The reason we are interested in GGAO is that we are in competition with the Europeans to do real-time correlation. We would like to do it before they do. We think we could do it within a couple of weeks if we had the fiber available. We are gearing up the software here and making the connection to Westford to demonstrate a one antenna VLBI experiment. Doesn't cut the mustard as far as VLBI. Want GGAO so that we can be first. Bandwidth requirements are 256 Mbps dropping down to 128 Mbps if we have to. Demonstration of real-time rather than high bandwidths. Need a gigE connection with at least 128 Mbps of data throughput reliably.

Terry: At this point a contractual issue. Technically feasible once Bossnet rack relocated

Alan: Also want to start the intensives. Can manage with Mark5 at ISI-E for the time being.

Pat Gary: No e-VLBI news. One good indirect thing: Goddard interested in working with NLR to get to the west coast Scripps Institute of Oceanology. We have written an internal research and development proposal. Some funding has been awarded. Funding will be tied up in one of two things: membership fees for joining NLR, hope it can be reduced with a package deal including NASA Ames, JPL etc... Goddard needs to get to McClean to the Level3 POP. Hoping to use DRAGON network to do this. Have some money that they can invest in DRAGON network to get to NLR.

Alan: Congratulations on this. Wanted to talk about some information dug up by some detective work by Hans.

Hans Hinteregger: Keeping track of the progress of commodity technology. There are now a number of companies that make DWDM SFPs (Small Form-factor Pluggable) modules that will go into cheap switches, such as Dell's. According to the advertisements, the pricing is at the same level as for CWDM(< \$1000 per transceiver). Can get all the ITU wavelengths. If you need to be compatible with commodity equipment with existing DWDM multiplexers, this is an inexpensive way to do it. It looks like the technology is just maturing - may be some technical risk. Still worth considering as an inexpensive way to light up fibers that you have access to.

Alan: Some of you may have been aware already, but interesting news to us.

Hans: Luminent OIC, TriQuint, Finisar are offering these products. Luminent offers 40 channels on the ITU 100 GHz grid. Finisar demonstrated DWDM XFP 10 Gbps device last year with 40 km reach. Not commercial yet. Can get XFP with 10 km reach for \$500.

Pat: On the fiber between Goddard and the fiber at MAX at College Park we have CWDM equipment from a manufacturer called LuxN. We can then hook up different inputs and their gear will put them

on a particular lambda. Could also buy Cisco GBICs that are already tuned to CWDM lambdas - same frequency used by LuxN. If we light up both ends using the optical OCM modules (Cisco), reduces cost from \$12K using LuxN to \$6k using Cisco. Could operate Cisco wavelengths over LuxN transport. Is that what you are talking about?

Hans: Not sure about the commodity packages.

SFPs are like mini-GBICs.

Hans: As of a year ago, could get standard short range at 850 nm, intermediate range at 1310 nm and CWDM at this time last year with 20 nm spacing from 1310 nm to 1610 nm. Up to 16 wavelengths across 300 nm span. DWDM at this time was much more expensive. As of last month, DWDM available from a number of companies with thermal stabilization built-in and have dropped price down comparable to CWDM. Major objection to DWDM has been removed and will probably continue to go in this direction over the next couple of years.

Paul Lang: This is great. Best way to go, rather than having to buy extra optics. Cuts out expensive intermediate steps that need to be done.

Hans: Can go over cheap switches for short reaches without amplification. Very interchangeable, just buy what you need.

Paul: Some vendors may complain about using wavelengths from other vendors. E.g. on the 6509 over at MAX Dan Magorian had got an error message complaining that the GBIC serial number did not match. It appears that Cisco may be running certain checks to make sure that you are using certain XFPs.

Hans: If the interface is at the WDM multiplexer, you don't have to go through a router.

Paul: This is SFPs for the 6509 itself. Should be ok if you are using an inexpensive switch. Hopefully these switches will not check serial number.

Bill Fink: I was the one that discovered this. That was Cisco running on one of the 6509 NGIX switch. It did a check and determined that the serial numbers were the same on a couple of different ports and refused to bring those interfaces up. The thought was that if you used some other vendors GBICs then it might complain that their serial numbers matched and not work.

Paul: Dan said that they were Cisco SFPs they were using.

Hans: Finisar demonstrated a 10 Gbps SFP module about a year ago. Haven't seen an announcement that it is commercial yet. 10 Gbps intermediate range XFPs for \$500 being produced.

Bill: For 10G most of them use Xenpack technology. Force10 might use XFP.

<http://www.finisar.com/home/>

<http://www.luminentoic.com/>

<http://www.triquint.com/>

Alan: Pat had raised subject of increasing bandwidth to GGAO. I spoke with Bill Wildes about this a couple of weeks ago. I didn't get a yes or no. Bill said he will look into it. I'm not sure if he has

contacted you.

Pat: Not yet.

Alan: Maybe I should ding him. I've raised that subject with him.

Bill: One possibility is to use the Marconi DWDM equipment from another group to get 2 Gbps (2 Lambdas) to GGAO. This group just bought Marconi DWDM gear and don't have an immediate use for it. Could also use the SFP route you were talking about.

Hans: Very convenient because it lets you use commodity devices that are very cheap. Can defer the aggregation question until you get to the nearest pop and stick with gig E.

Bill: With Extremes and round robin equipment should be able to get up to 8 Gbps.

Paul: Extreme has a summit 400 that supports a router card with 2 x Xenpacks. Should be fairly reasonably priced. \$26K list price for 2 x 10 Gigs and 48 x 1 Gigs.

<http://www.extremenetworks.com/homepage.asp>

Performance Testing

David Lapsley: At the last telecon I mentioned the Internet2 tool bwctl for scheduling and authorization of tests and intermediate hop testing to servers in the Abilene backbone. Got that installed this morning and had a play around with it. Very nice. Quite easy to get up and running. At the moment, not using the authorization functionality. Seems like an easy way to do this. I've been talking with some of the Abilene folks to see if we can get access to some of their internal servers so that we can do some intermediate hop testing across the backbone. Should be good. I've scheduled some testing on Wednesday to do some of that. On Friday we will start testing with Onsala in Sweden. This is a telescope in Sweden that is connected at 1 Gbps, so we will do some bandwidth to see what the network looks like in preparation for an experiment that we have coming up on the 25th of Mark. We would also like to start some testing to the Mark5 down at GGAO. Performance testing in preparation for some of the real-time experiments that Alan had mentioned. Can't do that from here at the moment, but at least from ISI-E can do some testing there. Try out some new transport protocols to see what sort of performance we can get.

Alan: David and I attended an NSF PI meeting last week (as were Tom and Jerry). A number of interesting talks. Tom organized this meeting. He can summarize.

Tom: Shared Cyber-Infrastructure at NSF. NSF recently organized their CISE division into multiple divisions. One of these is SCI. Initiative for dealing with e-Science applications over network infrastructure. PI meeting last week in Crystal City. Brought together various projects, applications and network. There is a URL and slides will be up on the web later in the week:

<http://hpn.east.isi.edu/nsf-sci/agenda.htm>

Alan: A number of presentations on new protocols and testing of new protocols. Which was interesting. A lot of things going on, including gathering momentum for data rate limiting, rate-based protocols. Growing interest in this and some discussion on that.

Tom: Also a lot of interest in lambda-grids, still need some type of protocols for these two. Interaction between IP networks, lambda grids and rate-based protocols. Interesting to see where this all ends up.

Alan: A lot of NSF people attending, including Mari Maeda, who originally gave us some seed money when she was with DARPA. Unfortunately, Mari's term is up at NSF in March. She will be leaving to go back to her research. Which institution was she from?

Tom: Originally CNRI - Corporation for National Research Initiatives. Bob Kahn's organization.

Alan: A new person brought in as head of SCI.

Tom: Sang Kim, division director of SCI. Most recently in industry for a large pharmaceutical firm. Research history, cyber-infrastructure and managing large research projects. Should be a good fit for NSF's desire to apply cyber-infrastructure to their large e-Science applications and networks.

Alan: Unfortunately, I don't think there has been any substantial money allocated yet.

Tom: I heard rumors that CI would get a lot of money, but haven't heard results of that. Don't know the results.

Alan: I spoke with Doug Gatchell. He said that it is stuck at \$20 million per year (the Atkins report recommended 1 billion dollars per year).

e-VLBI Testing

Kevin Dudevoir: Same status with Kokee. Unexplained asymmetry. Haven't done much work, because target bandwidth is 30 Mbps. Can get 34 Mbps to Germany. Using Tsunami and UDT protocols able to reliably get 30-40 Mbps using Tsunami. This is to Kame at ISI. One of the geodetic applications is called an "Intensive". Run 4 times per week between Wetzell and Kokee. Run for an hour with about 30 minutes of observing time. They run at 128 Mbps for 30 minutes of actual observation time. About 30 GB per experiment per hour. Brought some scans back to Kame. Brought them to Haystack over regular network. Correlated the data. On the point of transferring entire experiments here and correlating them once BOSSnet is up and running, or could transfer them to the Mark5 at ISI in the short term.

Alan: We talked with Tom and Jerry about using the Mark5 as a temporary termination point for data that needs to be correlated at USNO. Need to have a shuttle service to ferry disc packs to and from USNO and ISI-E. Appreciative of Tom and ISI-E people for allowing us to do that.

Tom: We have had a couple of emails go back and forth. I have Kerry's information here, so I will make arrangements for Kerry to be able to access this data as he needs to.

Alan: We will co-ordinate with Kerry to start moving that data to ISI as soon as those arrangements are made.

Tom: When do you think the first data sets will be available for pick up?

Kevin: In a couple of weeks. You might want to wait until after the real-time experiments are finished. We could possibly start within a week. It would be nice if we could have some additional storage

capacity available at Kame.

Tom: I have quite a few hard drives there that are just not visible to your home directory, I can change that.

Kevin: We have 30GB of data per station per experiment. We could probably transfer into Kame and then transfer into the Mark5. We can also transfer directly into the Mark5. That is a bit restrictive as we don't have access to all of the transport protocols we want to experiment with initially. Preferable to transfer to Kame initially, until David has additional protocols installed on the Mark5. We need 60 GB of storage per experiment. Is that a SCSI subsystem you have? We can buy discs if you need it.

Tom: It has those hot swappable drive slots. Let me send you the specs for this. I think I have some additional space you can start using.

Bill: You could use an external fire wire drive.

Tom: Does linux 2.4 support firewire?

Bill: Yes. I use it at home.

Tom: I don't have a firewire port on it.

Bill: As a matter of interest, I just saw an announcement that LaCie has a firewire drive that is 1 TB for \$1200. 5.25" external. 7200 RPM drives with firewire2 800 Mbps. But need a machine that can support firewire 2, otherwise can only support 400 Mbps.

Tom: I will email Kevin if we need to add anything.

David: Tom, on the Mark5, there is an experiment running now, that will be finished by Wednesday. That will be ready for pickup then.

Tom: Send me an email when it is ok to move it.

Alan: Other antennas. David has been looking into tests with Onsala.

David: Start testing on Friday to Sweden. There are two ways to go: routed over GEANT (European equivalent of Abilene) then through ISI-E or through Chicago through StarLight and setting up a lambda from Europe, through SURFnet, to Chicago. May give us better quality. See what performance is like routed and then decide.

Alan: Any other e-VLBI?

David: Last week we had another experiment, T2026 geodetic experiment, finished last Wednesday, transferring data between Kashima, Japan, correlator in Bonn and Chile. This was the first time they participated in an experiment. Koyama-san at CRL got fringes over the weekend using their software correlator. A nice result.

Miscellaneous

Alan: Two weeks ago, David and I were in an international VLBI meeting in Ottawa. David gave a nice presentation on among other things the new proposed VSI-E standard. This is the RTP protocol designed specifically for e-VLBI. We had a warm feeling coming away from the meeting. Next step is

a formal review of the draft specification by the international committee. Get back comments and make a review of the comments and see if any changes are required. It has been feather in David's cap because he put this together basically by himself and has been working with John Wroclawski at MIT on the specification in such a way that we hope the extensions will be accepted by the IETF as part of the RTP protocol. This is ongoing, we hope to wrap that up in the next few months, maybe even sooner, but probably not. Any other subjects? Dennis, do you have any items you want to raise?

Dennis Baron: No, not at this time.

Alan: Have you heard anything about an IEEEAF donation from New York to Boston?

Dennis: I think I am meeting with them on Friday.

Alan: We would be very interested to hear about them. I was button-holed by Ed Fontagrossi in Hawaii 4 weeks ago at the joint techs meeting. He told me had a donation of fiber from New York to Boston and would be working with NOX to see how it would be utilized.

Dennis: The meeting on Friday is a NOX meeting.

Alan: If you get any information, send out an email, that would be much appreciated. Any other subjects?

David: I was wondering if it would be possible to arrange a meeting within the next couple of weeks to discuss the details of the new wavelength on BOSSnet.

Alan: With the people at Lincoln. Terry are you going to be around?

Terry: Completely free on Wednesday 25th.

Alan: 9:00 am meeting on Wednesday. We will give you a call at 9:00 am on Wednesday.

Pat: I've mentioned previously joint testing we have done with Storage Area Networking technologies. Over the weekend, one of the gentlemen we are working with, Hugh Thompson, has completed a paper that has been accepted at a conference: 21st IEEE Conference on Meta-storage systems and technology, April 13-16th. In that paper, we have captured a lot of the work that Bill, Paul have provided as network support, where we have used these Nissan boxes for connecting SANs as well as some from Sandcastle. Not all possible tests run, but revealing to see that some other proprietary efforts to do high performance using UDP with your own packet sequencing, or trying to spoof TCP, that distance (RTT) can have a significant effect on throughput. We are still getting credible, even useful results testing between Goddard, and NCSA, Urbana Champagne with Sandcastle and San Diego Super Computer with the Nissans. I can't give you the report today, but it will be made public in a couple of months. One other different kind of animal: you've heard of our interest in 10 Gbps technology. Paul Lang made some arrangements with Force10. We already own one small switch - E300 with a couple of 10 GE blades. We have borrowed a mirror image of this, so we can do back to back testing or even between here and College Park or on campus. Can deploy a second unit with 10 GE capability on some fibers we have access to and do some testing of that technology. And Tom we can factor you in so that we can do some sort of "scheming" where we can.

Next telecon

Next telecon is scheduled for Monday, 22nd March 2004 at 2 pm EST.

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
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
Arthur Niell, Haystack
Joe Salah, Haystack