

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

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

July 12, 2004

TO: Distribution
FROM : David Lapsley
SUBJECT: 21 June 2004 e-VLBI telecon summary

Attendees:

Lee Foster, Pat Gary, Paul Lang, Andrea Macleod – GSFC
Kevin Dudevoir, Hans Hinteregger, David Lapsley, Arthur Niell, 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, DRAGON, MAX and GSFC/HECN.

ACTION ITEMS ARE HIGHLIGHTED IN RED.

Bossnet

David reported on the status of the OC48 implementation on Bossnet for e-VLBI. Optical equipment had been delivered to Haystack. The Lincoln Lab folks had already started to install two of the Movaz Optical Add Drop Multiplexers at Lincoln. Hopefully Westford node will be installed in the next week or so. Then be able to look at testing them and getting them commissioned. Hoping to get Movaz engineer up here when maintenance window is scheduled to help turn up the equipment. Hopefully have equipment lit and tested within the next two or three weeks.

Movaz had generously given Haystack a deeply discounted rate on equipment purchased and donated \$80K worth of equipment. Equipment donated to Haystack has nice expansion capabilities so that future expansion to 10 Gbps without having to purchase additional equipment is possible(just by buying additional modules). Haystack is very grateful to them for this.

The configuration of switches and equipment at Eckington was discussed. In particular, if there was an updated diagram of the equipment at Eckington. It was agreed that this would be a good thing to have.

Connectivity between Haystack and GSFC was discussed. In particular, two options:

- to traverse a lightwave from Eckington to ISI-E and then via a routed network across the “older” infrastructure to GSFC.

-
- to traverse a lightwave directly from Eckington to GSFC.

The issue with this is that addressing etc. for new router/ switch configuration needs to be considered.

The issue of how to terminate OC48 POS from Haystack at GSFC. The possibility of using the Service Interface Module offered by Movaz that takes two gigabit ethernet wavelengths and transports them via an OC48 wavelength was also discussed.

The issue of maintaining connectivity to Abilene while still being able to connect directly between GSFC and Haystack was discussed. This might be a good application of Movaz dynamic wavelength switching. Currently, only one optical switch at College Park until October 1. After October 1st, the second switch would be purchased. Most likely at Eckington. In the third year, scheduled to take delivery of third switch.

It was agreed that these details need to be discussed with Jerry.

Jerry would be talking with Pat about which of the lambdas on the Movaz optical switch can be handled. The current switch was designed to handle 40 lambdas, currently can only handle 15, but only 9 of those handled correctly. Need to know which of these 9 lambdas can be handled for the DRAGON network and which can be nailed down for 10 G lambdas from GSFC. This equipment needs to be purchased from Movaz.

Pat discussed a diagram that he had forwarded to the group. GSFC looking at using NLR to connect to sites West of Goddard: e.g. Ames or Carnegie Mellon. In phase I+, NLR will be connecting to New York. GSFC could get to Goddard Institute of Space Studies. Hopefully NLR would get to Boston. In Phase I, NLR will connect Atlanta, DC, Pittsburgh, Cleveland, Chicago, Denver, Seattle. In phase II, NLR will connect Louisiana, Texas and New Mexico. NLR touches down at the Level 3 POP in McClean (Washington, DC). GSFC will become a member of Mid-Atlantic Tera-scale partnership(MATP): Virginia Tech leads this. This consortium shares an aggregating facility. The aggregating facility (most likely a Cisco 7600) will be used to feed 10 G connections to GSFC and other members of MATP. This can be a connection to a shared 10 GigE lambda or a number of 1 GigE connection to a set of national size VLANs (GigE) coast-to-coast. Also possible to connect through Internet2 HOPI which will allow regional and national networks to experiment. Hopefully Goddard can define certain experiments that Internet2 would be interested in pursuing. GSFC also interested in putting a 10 G router at McClean. Put some G4/G5 Macs with quad GigE cards or 10 GigE NIC cards. Then would be able to test with NLR. This would be an additional site that everybody could test to. Jerry had said he would put a Movaz OADM at McClean. GSFC interested in a couple of 1 GigE connections with normal transponders that Movaz has. Would also look at supporting 2 x 10 GigE connections using Alien/External wavelength module. Transponder was too expensive (\$27K). With external wavelength module with data at correct wavelength can do that for about \$5K connection.

At Goddard, connections comes into a GigE switch at building 28. Above GigE switch is Thunderhead cluster, ranked number 6 in top 500. Take 10 CPUs out of 512 CPU cluster, have one GigE up-link into Extreme Networks Summit 400s with 48 x GigE copper ports and two optional 10 GigE up-links. \$6K

for 48 ports, add optional 2 x 10GigE ports for \$8K. Nice for aggregating 1 GigE up-links onto a single 10 GigE network. Thunderhead cluster, SAN, hyperwall displays.. Test stations in building 32. NCCS has other supercomputing computers. Many issues with firewall ports.

In connecting to GGAO will bring out a 1 x GigE connection to GGAO. Look at options for providing more bandwidth for e-VLBI when required.

e-VLBI Experiments

David reported on two additional experimental transfers that had taken place as part of the support of the International VLBI Service (IVS). Last week we started a transfer for a T2030 session: from Kashima in Japan to ISI-E in Washington. Not high bandwidth transfer due to the fact that it is going through a 100 Mbps bottleneck at Japan. Seems to be quite a lot of traffic traversing the link. This transfer is the first to use the FAST protocol. However, because of heavy traffic, not able to see full advantage of using FAST. Japan is in the process of upgrading link to 1 Gbps. Once new link is up, we should be able to see the advantages of using FAST.

After completion of T2030, data from experiment RD0404 will be transferred. RD0404. Data will be transferred from Tsukuba to ISI-E. Tsukuba connected at OC12 622 Mbps. Previous testing has shown rates of up to 300 Mbps using FAST on this link.

Additional two practice experiments coming up for IVS. In the summer, IVS wants to add an additional two stations for weekly intensives. These are short experiments that typically run for two hours of observing time at 50% utilization typically 40–50 GB of data. We are helping to support that. Also working on 4 station experiment. Real-time experiment at high bandwidth: GGAO, Westford, Onsala in Sweden and possibly Tsukuba in Japan. People at SURFnet have been contacted to see if it would be possible to get a dedicated wavelength from Onsala to Chicago. This would give us clear bandwidth to there. Once there, will go on to Abilene. Should not see any loss there. Optimistic about being able to see close to 1 Gbps throughput. For correlation, will be aiming for 512 Mbps.

On performance testing. David had been invited by people at Caltech to participate in some testing with them for a paper they are submitting. Last 3 weeks spent testing between Tokyo XP (termination point for APAN/Transpac connections - one from LA to Tokyo, the other from Chicago to Tokyo). Testing has been from a node connected directly to the border router, which has made a huge difference. Testing has also been done to Pittsburgh Supercomputing Center and to Georgia Tech. Interesting results achieved. High throughput, but also interesting results. Some problems found with the Linux kernel. 990 Mbps was observed from server connected to Tokyo border router over 170 ms round trip time. Almost too good a result because for some tests did not see a single packet drop over 30 minute test interval. The other node connected through a layer 2 infrastructure, lightly loaded, 2 or 3 Foundry switches. Noticed a huge difference in throughput almost 50% drop in throughput between this node and the border node. Even though the layer 2 switch looked lightly loaded (utilization only about 10% from MRTG graphs). Thought to be from the cross traffic. Interesting that this traffic is not shown on the graphs. Typically, these use 5 minute averages. Only see the cross traffic loss during active testing.

Kevin reported on Hawaii and Germany. Wetzell can routinely achieve 80% of available 34 Mbps (26 Mbps). Currently looking at automating scripts for operators. Switched to using Tsunami and UDT for this link since TCP performance poor. Kevin also looking at installing FAST and characterizing loss. Having problems with the software locking up during heavy packet loss. Simply insufficient capacity on links from Kokee. Too much other traffic on link. Packet loss is hard to characterize. Often see very high packet loss: anywhere from 0.05% up to 99%. It is believed that losses are mostly caused by congestive losses due to the 40–50 subscribers on this link.

Paul reported that they had just sent back the E600 today and still waiting to get cluster hooked up to the E300. The Summit 400 was sent back the previous week. Possibility of getting it back again. May have the opportunity to test some new Force10 equipment in the near future. A report on the 10 GigE testing will be forwarded.

Miscellaneous

Next telecon is scheduled for Monday, 12th July 2004 at 2 pm EDT.

cc: Steve Bernstein, LL
Jim Calvin, LL
Rick Larkin, LL
Peter Schulz, LL
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