

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

HAYSTACK OBSERVATORY

WESTFORD, MASSACHUSETTS 01886

Telephone: 978-692-4764

Fax: 781-981-0590

22 August 2003

TO: Distribution
FROM: David Lapsley
SUBJECT: 18 August 2003 e-VLBI telecon summary

Attendees:

Bill Fink, Pat Gary, Andy Germain, Kevin Kranacs, Paul Lang, Bill Wildes – GSFC

Steve Bernstein – MIT Lincoln Lab

Tom Lehman – ISI-E

Charles Yun – Internet2

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

This telecon is one of an ongoing series of telecons to prepare for 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.

Glownet/Bossnet

Steve reported that he is still very optimistic that the Bossnet project will continue.

Tom reported Bossnet is now running two wavelengths, GigE and OC-48; only the GigE is currently being used. OC-48 is undergoing further testing; Tom expects OC-48 would switch over to GigE wavelength when OC-48 path is cleared.

10 GigE plans are on hold until OC-48 is working.

NASA/GSFC

No infrastructure changes at GSFC. Pat is very interested in understanding e-VLBI needs for coming few years. One concern is mandated firewalls. Can easily handle 1 Gbps, maybe 2 Gbps, but significant challenges at even higher data rates; need to put serious plans to together in near future. Pat suggests pushing for infrastructure changes to support up to ~10 Gbps so that the necessary infrastructure will be in place when it is needed.

10GigE Intel NIC cards

Bill Fink reported on back-to-back tests of new 10GigE Intel NIC cards. Uni-directional 3.2Gbps. Bi-directional 4.5Gbps (sum of two directions). Not clear where bottleneck is. PC Platform was dual Intel 3.0 GHz Xeon with 3 x 133Mhz PCI-X buses with Supermicro X5DL8-GG. Managed to get 5 Gbps using quad GigE cards (4 Gbps in one direction, 1 Gbps in the other direction). Looking at

getting S2IO cards which should be able to get twice the performance.

International Connections

Kevin Dudevoir reported on his conversations with people in Germany regarding E3 (34 Mbps) connection to Wetzell telescope site in October.

David Lapsley reported on his conversations with people from the University of Hawaii and the Australian Academic Research Network regarding possible 10GE links from the East coast of Australia to the West coast of US (via Hawaii). If approved, timescale for installation would likely be 1st quarter of 2004.

Performance Monitoring/Testing

Andy Germaine completed his action item to mail out his software package and writeup. Tests are continuing to run with no significant changes in performance. Best performance is from Goddard to ISI-E in the 300 Mbps range. From Goddard to MIT, 3 nodes at MIT got under 100 Mbps. This is due to them using a different server. Only Goddard "Wilson" nodes can get to the good server (~250 Mbps throughput), other nodes are limited at 100 Mbps. ISI-E is a well connected node (~350 Mbps from Goddard). Transitioning to a new monitoring website. A demonstration is available at: <http://ensight.eos.nasa.gov>. Traceroute information is currently being collected, but only hop counts are displayed.

David Lapsley reported that some more e-VLBI test servers had been added: one on the island of Hawaii at the Smithsonian Institute (currently DS3 limited, but may be connected at 10 GE in the future - depending on approval of Australia to US links); another two new servers in Japan located in Tokyo off APAN backbone. Installed new High Speed file transfer tools TSUNAMI and SABUL. Both use TCP control stream for keeping track of error statistics on link and acknowledgements, also use a rate-limited UDP data stream for data transfer so they are able to make better use of available bandwidth on high delay, high bandwidth pipes. More information available at: <http://web.haystack.mit.edu/staff/dlapsley>.

David reported that he had started automated testing from Tom's server at ISI-E between various e-VLBI sites around the world, using scripts based on Charles Yun's performance script. More information and preliminary results available from web site above. Pat had asked would it be possible to use Andy's script in parallel on some of the test machines (for example, the host at Hawaii). The only potential limitation is that some of the hosts would not have administrative access.

Action Item: David and Andy to look at installing Andy's tool on some e-VLBI performance servers.

David reported that Internet2 performance people suggested that SLAC tools may be useful for us. They have been collecting data around the world for many years, have a large database of results and wide deployment of their tools (17 monitoring sites in 10 countries, 300 remote sites in 27 countries). Major tool is PingER tool (but also have other tools, e.g. bulk transfer, etc.). It would be nice if we could hook into their network as this would provide more detailed statistics covering a wider area. Could run their tool or Andy's tool exporting data into a format that their servers can download.

David reported that Haystack had received a binary version of the FAST TCP kernel. He has not yet had a chance to test but will have it up and running this week. David will put an iperf server on it and will open it up for testing to other people. Currently, kernel is distributed as a binary that is compiled on an Intel/Redhat platform.

Action Item: David to check with Caltech to see if it's ok to distribute FAST kernel to other members

of the group.

David reported on his survey of Net100. This was in response to a discussion on the last telecon regarding a framework that would allow an application to retrieve historical connectivity to a destination, prior to the start of transmission (would allow the application to automatically tune itself). Tom Lehman had mentioned that Net100 people had created such an application and that it would be good to look into more detail. Net100 is a DoE funded project. PIs from PSC, LBNL, ORNL and UCSB. Net100 project has developed a framework for network monitoring and application performance tuning. Software is currently available. Net100Logger application, automatically logs statistics etc. from any host to a central server. Database framework that automatically accepts statistics and stores it. API for retrieving statistics from database site. Also have a module for forecasting performance across the network. Work very closely with Web100 folks. Some Web100 people also working on Net100. Using Net100/Web100 it is possible to pull out very detailed statistics on TCP from the kernel and store it in a database. Net100 allows applications to automatically tune their window sizes and various TCP parameters to optimize their performance. David contacted people from Net100, they are interested in working with us (will send us Net100 patches to Linux kernel). David will install and test in the lab. If anybody is interested links are available on his web page. Source code will be available soon. Tom Dunigan was the contact at Net100. Very nice framework with a lot of effort and thought put into the design. Currently a lot of different approaches to the same thing. Net100 folks have already started standardizing this in the IETF, would be good to have a uniform approach.

Kevin Dudevair reported on his testing to Germany and Hawaii. Testing to the University of Regensburg (bottleneck capacity of 155 Mbps) achieved throughput of 100 Mbps to the server approximately in both directions. Most of the path is in common with the proposed link at the end of October to Wetzell (at 34 Mbps). Testing to Hawaii still stalled. Chris Blackstad at PMRF has agreed to host a test box at PMRF before the last microwave link to the Antenna. Kevin will have a box configured and shipped out to PMRF for testing within the next week. Rick Larkin attended a DREN meeting last week and had spoken with people from PMRF. Didn't know of any problems with the Microwave link. Performance on the link seems to be good. Currently have one DS3 assigned to us. Currently do UDP tests from 1 Mbps up to 20-30 Mbps, see transmitter/receiver rates are good up to about 20 Mbps. Fixed loss rate of about 4%. Test node may help to isolate problem.

Pat Gary asked about the Mark 5's at Goddard: will testing be done on those ramping up testing between Mark5's and our site? Kevin is looking at improving the performance on those boxes. Currently the maximum rate is 400 Mbps from Streamstore card driving network. Looking at improving this performance to at least 512 Mbps. Is it worth putting a test node there that can be reported by Andy's software? Currently test to Pluto. Need to get regular performance testing to GGAO to see the health of the link.

Action Item: Andy and Kevin to setup automated testing between Haystack and GGAO.

Performance Test Scheduling

Tom Lehman commented that there was a need to co-ordinate all of the testing. Unscheduled testing at the moment should be ok. However, in the future an updated/more advanced web calendar system for iperf and performance tests would be required. David suggested that the schedule could be stored in a database, that applications could access via an API to see if a timeslot is scheduled/reserved. There could also be a GUI for scheduling in advance and an API for applications to be able to check whether or not they can run. It would also help with future e-VLBI experiments to ensure they are not

impacted by performance testing.

Action Item: Tom to survey web calendar software.

Bill Fink reported that he is also looking at scheduling software for nuttcp.

Tom Lehman commented that it would be good to tightly integrate performance measurement, scheduling, real-time debug and isolation. Need to work towards database and storage of performance testing for historical/predictive purposes. Andy currently working on storage retrieval. Andy has based system on Oracle database and has results dating back to 2000. Currently working on Phase I. Phase II has a list of new enhancements.

Action Item: Tom, David, Bill and Andy to continue conversations.

Internet2 Demo

Alan Whitney mentioned that Charles had asked if Haystack would put together a demo at the Internet2 meeting in the fall in Indianapolis. Potential sites for live demonstrations: Hawaii (Mauna Kea), Arecibo, JIVE, Kashima, GGAO at NASA. Everybody positive. Arecibo connection at risk (155 Mbps currently) shared with University of Puerto Rico, UPR not sure if they will maintain it. Could be cutoff within this month. Trying to preserve this connection.

Charles Yun said that it may be useful for Astronomers to get in contact with Cornell to encourage Arecibo to keep the link up by identifying users of the link and bandwidth. Connectivity to be driven by usage of link. Charles has spoken with some people, but believes that every one should talk at once. Major user is at Berkeley. Major users are currently mostly non-VLBI. Future usage for VLBI would be important. Haystack would be the only VLBI user.

Action Item: Alan will look into sending a note regarding future VLBI use of this link to Arun Venkataraman (Arecibo) and Guy Cormier (University of Puerto Rico).

Next telecon

Next telecon is scheduled for Mon, 8 September 2003 at 2 pm EDT

cc:	Steve Bernstein, LL	Tom Lehman, ISI-E
	Jim Calvin, LL	Jerry Sobieski, MAX
	Rick Larkin, LL	Guy Almes, Internet2
	Lorraine Prior, LL	Charles Yun, Internet2
	Peter Schulz, LL	Richard Crowley, Haystack
	Leslie Weiner, LL	Kevin Dudevoir, Haystack
	Herbert Durbeck, GSFC	Hans Hinteregger, Haystack
	Bill Fink, GSFC	David Lapsley, Haystack
	Lee Foster, GSFC	Arthur Niell, Haystack
	Pat Gary, GSFC	Joe Salah, Haystack
	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	