

Minutes of T11.1 HIPPI Ad Hoc Working Group

June 8-9, 1999
Minneapolis, MN

1.0 Opening remarks and introductions

The T11.1 Chairman, Roger Ronald of Power Micro Research, opened this meeting at 9 AM and thanked Robert Cornelius of Ancor for hosting this meeting.

This group is constituted as both the HIPPI Working Group under T11.1, and the High Performance Networking Forum (HNF) - Technical Committee (TC).

Roger lead a round of introductions. The list of attendees is at the end of these minutes. Roger volunteered to produce these minutes.

2.0 Review / modify the draft agenda

Draft agendas were distributed via the www.hippi.org web site before the meeting.

Steph Bailey had made two proposals for additions to the ST document and these were added to the agenda and the previous item 7.4 (document processing plans) was moved to be item 7.6:

- 1) Add a version code to ST (item 7.4)
- 2) Change ST to support a variable Block size (item 7.5)

These minutes reflect the approved agenda.

3.0 Review minutes of previous meeting

The minutes of the May 13-14, 1999, HIPPI Working Meeting in Milwaukee, WI were reviewed.

It was noted that the minutes needed to have the location for the August meeting changed from Minneapolis, MN to Rochester, MN.

Jeff Young moved and Don Woelz seconded, to approve the May 13-14, 1999 working meeting minutes as amended. Motion passed unanimously.

4.0 Review old action items

The action items from the May 13-14, 1999 meeting were reviewed.

1. John Gibbon to consider development and documentation of HIPPI-6400 MIBs. (**carryover**)
2. Jean-Michel Pittet to provide the IETF RFC number for "ARP and IP broadcast over HIPPI-800 (GSN)" to Don Tolmie and Roger Ronald. (**carryover**)
3. Marck Doppke to process the HIPPI-800 switch MIB through the IETF. (**carryover**)
4. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400. (**carryover**)
5. Greg Chesson to collect information for an "ST FAQ" html document. (**carryover**)
6. Greg Chesson to draft a top-level structure for the ST-API document. (**carryover**)
7. Don Tolmie to make the edits to the Scheduled Transfer document (rev 3.0) agreed to at this May, 1999 meeting. (**done**)
8. Steph Bailey to update the SCSI on ST text with the changes agreed to at this May, 1999 meeting. (**done**)
9. Eric Salo to update the level 0 ST API text with the changes agreed to at this May 1999 meeting. (**carryover**)
10. Ed Cady to provide the committee with information on the PCI version of the FCI-Berg Micropax connector. (**done, but more information is requested by the August meeting**)
11. Eric Salo to specify field sizes in the ST API in accordance with the approach selected at the May, 1999 meeting. (**carryover**)
12. The Chair (Roger Ronald) to summarize the VC2 size issue on the HIPPI reflector and ask for comments and/or proposals for change in this area. (**done**)

13. The Chair (Roger Ronald) to contact Berg and SGI to collect status on the cable and connector problems (EMI/RFI and PCI form factor compliance). (**done**)
14. The Chair (Roger Ronald) to prepare a project proposal for the SCSI on ST Mapping, including a proposal for liaison with ANSI NCITS T10. (**done**)
15. Arie Van Praag to investigate why exactly eight bytes of zeros are required in the ST on DIX Ethernet mapping versus six bytes of zeros or a length field. (**done**)
16. Ian Philp to contact Eric Salo by email to discuss the open issues in the ST API. (**done**)
17. Don Tolmie to add text on VC usage recommendations to the HIPPI-6400-PH amendment. (**done**)
18. The Chair (Roger Ronald) to produce the minutes for the May, 1999 meeting. (**done**)

5.0 HIPPI-6400-PH (ref: Rev 2.3)

The draft amendment provided by Don Tolmie is waiting on resolution of cable EMI and PCI form factor compatibility issues prior to completion.

A call to ask for any further input on the project proposal calling for a an amendment to change HIPPI-6400-PH was made.

5.1 VC2 Usage

During discussions in April, a possible additional change for HIPPI-6400-PH was identified. Because HIPPI-6400 is flow controlled, a slow host talking to a fast host results in both sides operating at the slowest common denominator for the VC in use. If only a single VC can be used, this means that an entire system is only as fast as the slowest client.

Since VC2 is not designated as a channel for any specific traffic and overlaps with VC0 in its size range, most in the group agree that it is underutilized. The proposal that surfaced during the April meeting would modify the Message size limitations of VC2 and allow larger Messages (perhaps as large as the VC3 limit of 4GB).

However, during discussions at the May meeting, a downside to allowing a larger mix of Message sizes became visible (inadvertent priority effects and potential blocking of Admin responses). These concerns convinced the group that quick action

was unwise. (see the May, 1999 minutes for a more complete summary).

In the absence of a consensus to proceed, the chair accepted an action item to disseminate this information to the HIPPI e-mail reflector and to call for any further proposals in this area. This call was made and there was no response.

In response to this flow control concern, Don Tolmie added recommendation text on VC usage and buffering. This addition to the draft amendment was reviewed and accepted.

5.2 EMI and Cable problems

During the April and May meetings, no input on cable testing was available from SGI. The chair accepted an action item to collect status from SGI on this topic.

Joel Darnauer responded by email with information that testing was still proceeding and that earlier approaches had failed. His reply is reproduced herein:

1) The current Gore cable construction (double-sided mylar foil has inadequate shielding. Tests were performed with a faraday cage built around the front portion of our system and bonded 360 to the cable shields at the cable exit. This idealized system had only 1dB of margin for CISPR at a number of frequencies related to the null traffic on the link. A second cable with a separately added outer layer of shielding was constructed and replaced the standard cable. This doubly-shielded cable passed by about 15dB, which gives some margin which would allow the faraday cage to be removed. We are proceeding with Gore and Berg to develop a doubly-shielded cable. Samples will be arriving here for test this Friday and we should have some results next week. I have also prompted Berg to allocate some time to testing of the raw cable shielding effectiveness.

2) A second experiment shows that the bulk of the emissions in the previous measurements are due to the cable and not due to the system. A special "fake" cable was constructed by placing loopback paddleboards into the backshells of a standard cable, and connecting the backshells with a 5m length of cable whose signal lines are not terminated. The backshells connect to the outer shield of this "fake" cable (which was the standard Gore material). This cable, which allows traffic to pass into the backshells, but not down the length of the

cable is very quiet, and passes our requirements with large margins. This also shows that the emissions are leaking out of the cable assy, and are not simply system noise that is being conducted out to the antenna along the cable shield.

3) A third experiment shows that the termination method is as important as the cable shield effectiveness. The specially constructed "doubly-shielded" cable was tested in two configurations. In the first configuration, it was tested with the backshell removed and the outer shield bonded to chassis at its exit from the special faraday cage built around our system. The inner shield for this test was tied to the board ground plane. This cable passed with large margins. In the second configuration, the cable was tested with the faraday cage removed, and the standard backshells re-applied. It was discovered earlier that isolating the two cable shields was necessary for the best performance, so this configuration had the inner shield taped to the shell of the plug and the outer shell bonded to the backshell. This "triaxial" configuration is maintained through the board connector. This cable is marginally passing, which indicates that the current backshell design is not as effective as the faraday cage.

4) Further measurements showed that the source of much of the noise was common-mode coupling to the pcb micropax connector. If the connector is excited, this noise travels on the outside of a singly-shielded cable and chassis if the chassis is connected to the cable. The noise also radiates from the jackscrews that make contact with the threads of the connector. The doubly shielded cable isolates this noise from the outside to the extent that the shields are isolated. The faraday cage pulls the outer shield away from the hot pcb connector, providing superior isolation. The common mode coupling to the connector is caused by three factors:

a) unequal rise and fall times for the true and complement outputs of the SUMAC driver create common-mode noise even on a perfectly balanced differential circuit,

b) the true and complement legs of the wires in each connector are unequal lengths, producing preferential coupling,

c) the capacitance between the connector body and each pin is on the order of 5pF. For all 46 pins in the interface, the total capacitance is ~230pf, which is a reactance of only 7 ohms at 100MHz. It

is easy for high-frequency common-mode noise on the pins to show up at the connector.

In retrospect, the choice of a right angle connector without good common-mode shielding/grounding severely limits our ability to manage EMC.

5) An additional issue with the connector was discovered. There are two metal alignment tongues that are on the plug side of the connector. There is a clearance between these tongues and the matching grooves on the receptacle side that allows this contact to be intermittent. This contact is needed to form a low impedance path back to board ground for the signal wires. Without this path, these return currents will be forced to flow on the outer cable shield or chassis.

This is by no means a 100% complete picture of what is happening.

Despite this partial information, we are moving ahead with the following design strategy:

1) Gore will provide cable in a double-shielded configuration. Two configurations are being built for testing:

* inner quad - foil - insulator - foil - braid - jacket

* inner quad - thin braid - insulator - foil - braid - jacket

The two shields will be insulated and separately terminated (triax) through to the board. The isolation of these shields, as well as the quality of the joints at the cable, shell, connector, and board interfaces is critical.

2) Berg is looking at changes to the inside of the current backshells that will allow better isolation between the connector shell and the backshell. This change would be totally transparent to most users. In addition they are exploring ways to improve the connection between the two parts of the connector.

3) SGI is re-designing our bulkheads and connecting hardware to provide:

+isolation of the jackscrew from the pcb connector

+robust gasketing of the backshell to the chassis

We are working under the constraint that any new cable design be mechanically compatible with existing equipment and that any changes to our

board hardware be mechanically compatible with existing cables, with a possible need to change out a jackscrew. We are also going to use the current backshell as a base design. We should have rough proposal drawings by May 26th. The mechanical design should be complete and in fabrication by June 4th. It would be possible for someone at SGI to present our plans at the June 8th meeting, but without verification, they would only be plans. It will be several weeks to fab parts, with testing occurring in late June.

We are still a substantial ways from a concrete solution. I apologize for the slow pace of this progress, but it is a long loop to close when changes to the cables are needed, and the work is very difficult and unfortunately involves a lot of trial and error.

Looking forward, there are a number of decisions that were made very early in the design process that make the grounding and shielding problem worse. Future designs/standards should consider the following:

- 1) Explicit limits on common-mode power on the IO driver. Alternatively, we could constrain skew and edge rates on each pair.
- 2) Large voltage swing outputs have higher power and are correspondingly more difficult to shield.
- 3) Edge rates that are very fast cause high-frequency harmonics to form that are out-of-band for the data being sent. Currently we have large harmonics in the 650MHz and 850MHz range that does not necessarily contribute to signal integrity.
- 4) Consideration of both differential and common mode termination in the termination scheme design.
- 5) Consideration of both far-end and near-end termination in the termination scheme design.
- 6) Explicit of cable shielding effectiveness as part of the specification.

Joel obviously has been working hard on this topic. Unfortunately, Joel also announced that he would be leaving SGI May 28th.

It was also noted in previous meetings that a PCI capable connector/cable assembly was needed and that Berg was supposed to be working on one. Joel Darnauer of SGI noted that SGI was not giving priority to a PCI compliant cable/connector assem-

bly and that other companies should be driving this area.

Bob Newhall was present and talked about modifications to increase RF insulation in the bulkhead area. The general driving concept resolved around triaxial shielding with the backshell tied to chassis ground and the inner shield tied to logic ground.

Don Woelz strongly stated that everyone needed to be on board for building an EMI compliant connector that worked with the PCI form factor.

Ed Cady of FCI-Berg provide information that Berg would be doing a PCI version of the connector. It would take 2-3 weeks for drawings to be completed and the total cycle for connector modifications would be about 28 weeks.

Bob Newhall took an action item to list the specific tests that they are trying to pass on the www.hippi.org reflector.

5.3 Draft Amendment Processing Plans

Since the cable issue was not resolved, we will not be able to forward the amendment at the June meeting. August seems a more reasonable date, but even August will not be achieved without an increased pace of progress.

6.0 HIPPI-6400-SC (reference Rev 2.5)

6.1 Processing status

HIPPI-6400-SC Rev 2.5 passed its second public review on November 24, 1998, with no comments. Unfortunately, it missed the NCITS meeting cycle and had to wait for the March 30, 1999 NCITS meeting to complete its management review.

The management review is now complete (it passed) and an NCITS letter ballot will complete on May 13. No new information was available at this meeting. Hopefully, the final procedural steps in the process will complete quickly and HIPPI-6400-SC will become a standard within the next month or two.

7.0 Scheduled Transfer (ref. Rev 3.1)

7.1 Review general document changes

Changes were reviewed and accepted with a few minor edits. Roger Ronald recorded these changes

to give to Don Tolmie for inclusion in the next revision of the document.

There were also a few inconsistencies found when using this document (for SCSI on ST) as the example for determining how acronyms are handled. These nits will also be passed on to Don.

7.2 Other Proposed Changes

There were no other proposed changes without a section heading in the agenda and in these minutes.

7.3 DIX Ethernet

Arie Van Praag from CERN had previously presented a proposal for carrying ST traffic over DIX Ethernet.

The reception from the group in May, 1999 was positive and it was agreed that directly documenting ST on DIX Ethernet was a much better approach than documenting ST on DIX Ethernet on ATM LANE.

There was consensus that all of the ATM mappings (token ring, 802.3, and DIX) were unneeded and these were removed from the ST document.

The only downside for ST on DIX Ethernet was the absence of a length field. Without a length field, fill at the end of a payload is indistinguishable from payload data. Arie Van Praag accepted an action item to investigate whether the eight bytes of zero fill replacing the LLC/SNAP header could be used to carry a length field.

Arie completed the action item by responding on the reflector. This response stirred up a major flurry of activity.

During the meeting, it was determined that the length field in **all** common (802.3 and DIX) Ethernet implementations was not guaranteed to correctly represent the length of the next high layer protocol. Therefore, to run correctly on these lower layers, ST (just like IP) **must** carry its own length field. The place where lengths are likely to be incorrect are where Ethernet is bridged to some non-Ethernet and non-GSN media, and back again. There was a high degree of group confidence that this could add a byte or more of fill.

There was even some discussion about changing the ST format in all cases, including on HIPPI-

6400. But this change would invalidate current practice/hardware and would not make any substantial functional difference (it would make bridging a bit more straightforward if all the formats were the same). So it was agreed that GSN would have a ST format without an included length field (Ether-type x'8181') and that other media would use a ST format with 4 bytes of fill and a 4-byte length field (Ether-type x'8183'). The use of separate Ether-Types would allow implementations to deal with one or both of these formats, as appropriate. Despite the fact that x'8183' format could also be used on HIPPI-6400-PH, it was agreed that x'8181' use should be strongly encouraged to support interoperability.

7.4 Version Code Proposal

Steph Bailey proposed via email that ST should include a version code in the connection setup sequence.

Tom Gilbert pointed out that he had proposed this earlier and it had been voted down.

It appeared that this proposal was headed for adoption, but in the end, it was determined that a major new protocol revision should have a new Ether-type and that a new Ether-type resolved the problem neatly.

So, no change was made.

7.5 Variable Block Sizes

Steph Bailey proposed via email that ST should support variable Block sizes.

There was substantial debate on this topic on the reflector. Many objected to such a major change so late in the game. But Steph did a very good job of defending his proposal.

During the meeting this trend continued. Finally, the chair asked if perhaps an optional payload could be used to modify a CTS to request less than a full Block. Greg Chesson elaborated on this proposal with some modification.

There was general consensus that this sort of change was acceptable, but SGI wanted to go off and check it against their implementation.

Steph Bailey accepted an action to describe the latest proposal by Greg Chesson on the HIPPI-

reflector and to document his comments. This mail should be a starting point for further discussions.

7.6 Processing plans

ST had seemed almost finished and Don Tolmie had suggested that ST should be forwarded this month. However, this was prior to Steph Bailey's email comments and the changes related to the mappings onto Ethernet.

The requested changes this month will prevent the document's forwarding in June. October is the next date that seems likely.

8.0 HIPPI-LE (reference Rev 3.7)

8.1 Processing Status

The main gating item before forwarding HIPPI-LE is inclusion of the RFC number for the "ARP and IP Broadcast over HIPPI-800" RFC in clauses 4, 6.1.1 twice, and Annex B. According to Jean-Michel Pit-tet, the IETF RFC number should have been provided by this meeting, but it was not. Therefore, there was no progress made.

8.2 Other Proposed Changes

A call was made for anyone with requirements to change HIPPI-LE to come forward. No change proposals were made.

9.0 SCSI on ST Mapping

9.1 Project Proposal

After some discussions about the lack of progress in moving forward towards a document (during the Level 0 ST API talk at the May meeting), the group agreed that SCSI on ST ought to be a separate project. The chair accepted an action item to write a project proposal.

Roger Ronald presented a project proposal for a SCSI on ST standard and noted that there would be a vote in the plenary session to forward the Project Proposal to T11 for further processing. (This vote later failed in the plenary due to lack of sufficient members voting).

The project proposal was reviewed and accepted with two very minor modifications.

- The word "amendment" was inadvertently left in two lines and these were removed
- The acronym for the project was changed from SoST to SST.

9.2 Document Review

Don Woelz read through the latest version of SCSI on ST. This version (1.0) was the first one in ANSI style format.

Since this was the first read-thru of the ANSI formatted document, there were many changes. The read-thru occupied most of the June 9th meeting (~9:30 AM --> 3PM).

These changes were almost all at a formatting or style level.

10.0 SGI Low Level API Proposal.

10.1 Document Review

Eric Salo was not present to work on the Level 0 API document. Therefore, no work was done.

10.2 Document Editing Plans

It's not clear if anyone is willing to do what is needed to make the Level 0 ST API into a standard. No work was done in this area.

Both Jim Pinkerton and Greg Chesson have a boilerplate document in Framemaker format.

11.0 Other HIPPI items

11.1 GSN Lite

The group talked briefly about GSN-Lite, a half-speed variation of HIPPI-6400. There was mention of other competitive technologies at a lower price point (e.g. Myrinet).

This group is interested in GSN-Lite, but the decision to proceed is basically pending on a direction from companies interested in pursuing the technol-

ogy. This direction is something the group believed would need to come out of the HNF meetings.

11.2 Bridging to HIPPI-800

At the April meeting, Nicolas Droux provided an electronic copy of his early work on HIPPI-800 bridging. The document is available on the web site.

Nicolas did not provide an update of his document and did not attend the meeting so there is nothing further to report.

11.3 ARP and IP Broadcast over HIPPI-800

Jean-Michel forwarded this document to the IETF in December and again in February. We are waiting to get an RFC number for this RFC, so that it can be added to the HIPPI-LE document as a reference. Although an RFC number was expected by this time, it has not been received. Nothing further at this meeting.

11.4 IP and ARP over HIPPI-6400 (GSN)

An electronic copy of this document is still needed for posting on the web site.

During the April, 1999 meeting, Jean-Michel was present to defend his current HARP approach and successfully defended keeping the documents unchanged; at least until they moved forward in the IETF process. He noted that the offending "extra" packet formats disliked by some (requiring HIPPI-6400 hosts to know HIPPI-800 formats) could easily be removed per IETF rules at a later date, but that adding the format back would be difficult if it were really needed.

Nothing further at this meeting.

11.5 MIBS

The chair noted that if no one was interested in MIBs, there was no good reason to keep them on the agenda.

11.5.1 HIPPI end-point MIB

If we want to pursue this MIB further then we need a champion for it. Nothing new at this meeting.

11.5.2 HIPPI switch MIB

Marck Doppke of Essential Communications has a draft document out for comment. Michael McGo-
wen previously offered Marck's services for pro-

cessing the HIPPI-800 switch MIB through IETF, and Marck has acknowledged this assignment. Nothing new at this meeting.

11.5.3 HIPPI-6400 MIB

Von Welch of NCSA has a draft document, based on HIPPI-6400-PH Rev 1.4, out for comment. Von was not at this meeting and nothing new was reported.

11.6 Tutorial for HIPPI-6400 ULA use

Greg Chesson has obtained the format material from Bob Snively of Sun. Drafting of the actual text is pending. Greg verified that this item is not holding up forwarding ST.

12.0 Future meeting schedule

12.1 Attendance

Don Tolmie expressed frustration with the lack of progress towards the goal of a standards document. Spotty attendance and participation by various people who have been involved in the effort has reduced progress even when the meeting time has not been completely utilized. It's hoped that the pending change back to bi-monthly meetings will reduce the standards "attendance fatigue" and encourage participants to treat each meeting as important.

12.2 Interim meeting, July 13-14 in Seattle, WA

The next (and perhaps last) interim working meeting will be hosted by Sean Rieb and Tera Computer in Seattle, Washington. The meeting information is linked to the www.hippi.org web site. The interim meeting schedule will be:

Tuesday 7/13 1 PM - 9 PM HIPPI ad hoc

Wednesday 7/14 9 AM - 5 PM HIPPI ad hoc

12.3 Plenary week, August 2-6, Rochester, MN

The August Plenary week location is in Rochester, MN. The hotel information is available on the www.hippi.org web site. The tentative schedule is as follows:

Tuesday 8/3 9 AM - 11 AM HIPPI-6400-OPT

Tuesday 8/3 11 AM - 9 PM HIPPI ad hoc

Wednesday 8/4 9 PM - 6 PM HIPPI ad hoc

Wednesday 8/4 6 PM - 8 PM T11.1 Plenary

Note that the optical working group meeting has been moved to occur on Tuesday morning.

The T11.1 (i.e., HIPPI), Plenary meeting will still be on Wednesday evening of the T11 Plenary week, following the HIPPI working meetings. The HIPPI and T11.1 meeting days are not specified within the Plenary weeks.

12.4 Future meeting dates and locations

The group continues to feel that interim meetings after July will not be required.

1999

Oct 4-8 Plenary Miami Beach, FL Jaycor

Dec 6-10 Plenary Reno, NV Solution

2000

Feb 7-11 Plenary Huntington Beach, CA Qlogic

Apr 3-7 Plenary Palm Springs, CA Brocade

Jun 5-9 Plenary Boise, ID HP

Aug 7-11 Plenary ** open ** ?

Oct 2-6 Plenary Manchester, NH Hitachi

Dec 4-8 Plenary Austin, TX Crossroads

2001

Feb 5-9 Plenary So. California Qlogic

Apr 9-13 Plenary Palm Springs, CA Brocade

13.0 Review action items

The following action items were carryovers or were added at this meeting.

1. John Gibbon to consider development and documentation of HIPPI-6400 MIBs.
2. Jean-Michel Pittet to provide the IETF RFC number for "ARP and IP broadcast over HIPPI-800 (GSN)" to Don Tolmie and Roger Ronald.
3. Marck Doppke to process the HIPPI-800 switch MIB through the IETF.

4. Greg Chesson to draft an IEEE tutorial on HIPPI-6400 ULA usage and the ULAs special to HIPPI-6400.
5. Greg Chesson to collect information for an "ST FAQ" html document.
6. Greg Chesson to draft a top-level structure for the ST-API document.
7. The chair to provide the updated DIX format drawings to Don Tolmie for inclusion in ST, version 3.2.
8. Don Tolmie to make the edits to the Scheduled Transfer document (rev 3.1) agreed to at the June, 1999 meeting.
9. Don Woelz to update SCSI on ST (version 1.0) with the changes agreed to at the June, 1999 meeting.
10. Eric Salo to update the level 0 ST API text with the changes agreed to at the May 1999 meeting.
11. Eric Salo to specify field sizes in the ST API in accordance with the approach selected at the May, 1999 meeting.
12. Ed Cady to provide the committee with information on progress towards a PCI version of the FCI-Berg Micropax connector.
13. Bob Newhall to provide a summary email on the HIPPI reflector detailing the exact FCC testing that SGI is working to pass.
14. Steph Bailey to summarize the proposal for modifications to the ST CTS operation described by Greg Chesson and to provide comments on its impact.
15. The chair to attempt to improve attendance among key participants by sending an email about the lack of progress in key areas due to the lack of attendance.
16. The Chair (Roger Ronald) to produce the minutes for the May, 1999 meeting.

14.0 Adjournment

Motion to adjourn by Jeff Young. Seconded by Don Woelz.

The meeting was adjourned at 3:15 PM on the 9th of June. There were no objections.

Attendance

Name	Status	Company	Phone	E_mail
David Hyer	A	Compaq Computer Corp.	978-493-6139	david.hyer@digital.com
Harry Levinson	O	Computerm Corporation	412-394-1084	hlevinson@computerm.com
John Gibbon	O	Essential / ODS	505-344-0080 x338	jgibbon@esscom.com
Michael McGowen	O	Essential / ODS	505-344-0080	mikemc@esscom.com
Ed Cady	P	FCI-Berg	503-359-4556	edcady@aol.com
Stephen Bailey	O	Genroco	414-644-8700	steph@genroco.com
Donald Woelz	A	Genroco, Inc.	414-644-2505	don@genroco.com
Ian Philp	P	Los Alamos National Lab	505-667-4305	philp@lanl.gov
Roger Ronald	P	Power Micro Research	972-437-9461	rronald@pmr.com
Walter R. Smith	P	Raytheon Systems Co.	972-205-5378	wrsmith@gar.esys.com
Greg Chesson	P	SGI	650-933-3496	greg@sgi.com
Jeff Young	A	SGI/Cray Research	651-683-5536	jsy@cray.com
Albert F. Kelley	P	Tensolite Co.	904-829-5600 x281	akelley@tensolite.com
Sean M Rieb	P	Tera Computer Co.	206-701-2164	srieb@tera.com