



Getting It All Together

Daryl Grunau
June 19, 1998



Outline


- **A message from our sponsors**
- **Hardware schedule**
- **Software configuration**
- **MPI/HIPPI**
 - **scalability**
 - **concerns**
 - **superheros**
 - **outstanding issues**
- **Machine reliability and stability**
- **A non-subliminal message**



Partnership for Advanced Computing



ASCI: Accelerated Strategic Computing Initiative



CHAMMP: Global Climate Modeling Program

**HPCC
Phase II**

DOE HPCC Phase II Grand Challenges



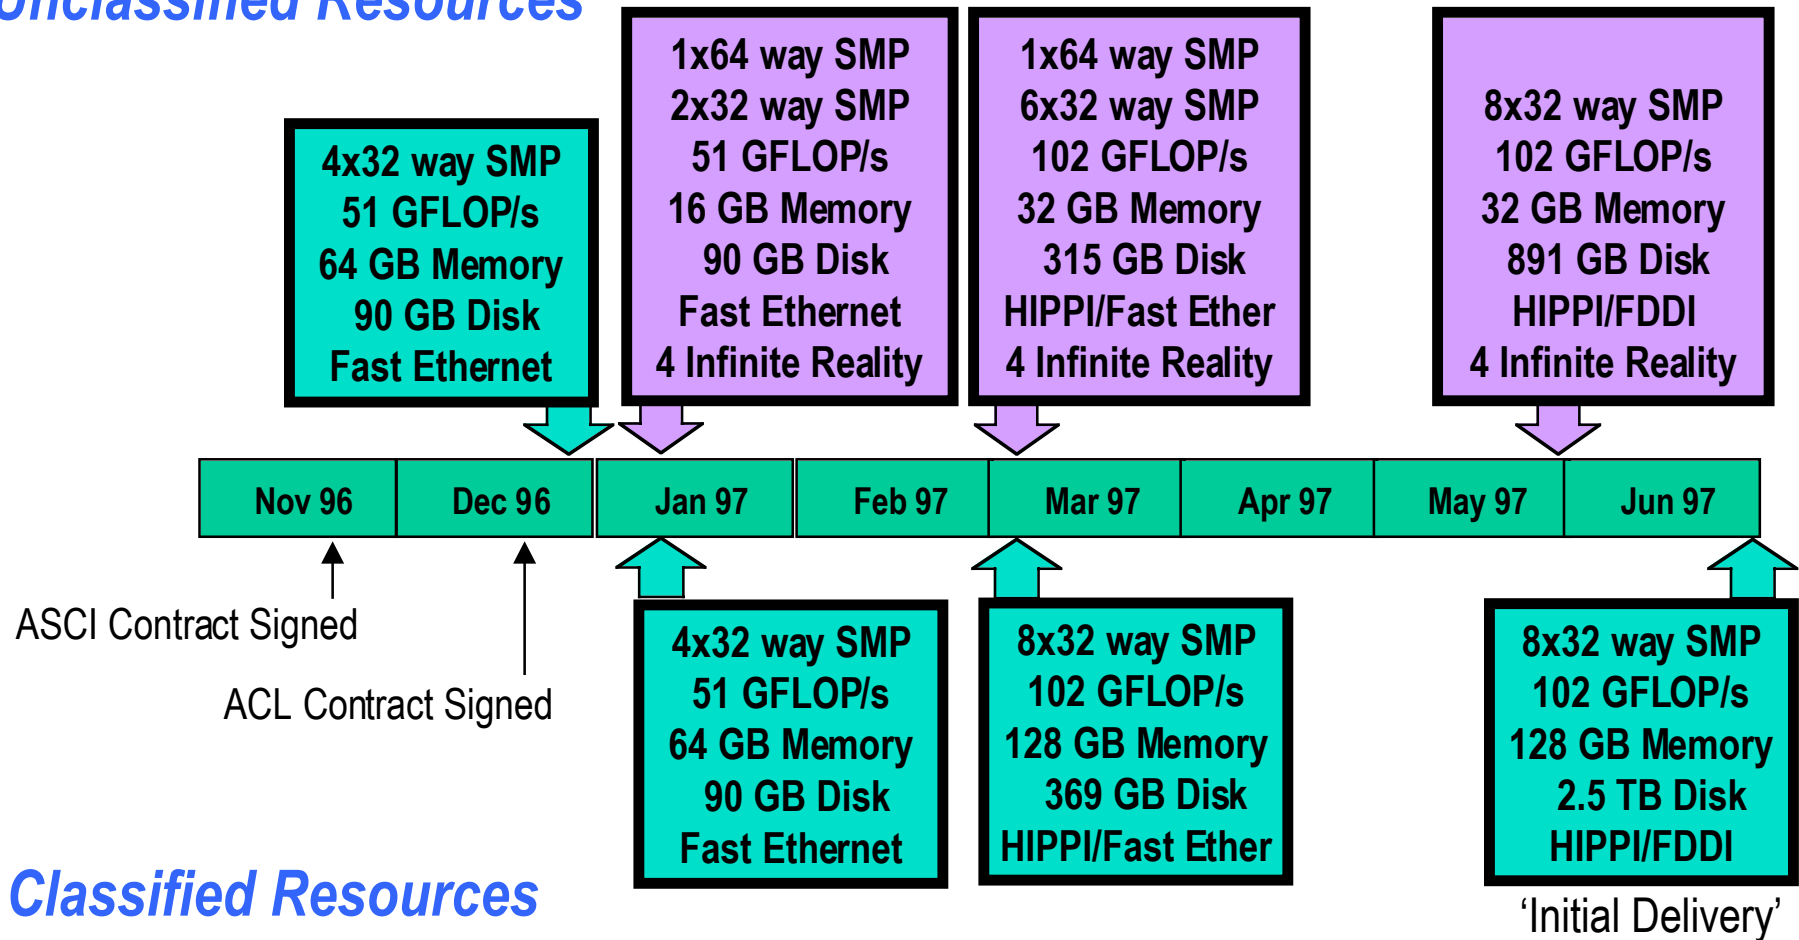
Los Alamos National Laboratory LDRD



Hardware Schedule

Unclassified Resources

'Initial Delivery'



Classified Resources



Hardware Schedule (cont.)

Unclassified Resources

'Tech Refresh'

6x128 way SMP
300 GFLOP/s
192 GB Memory
1.7 TB Disk
HIPPI/FDDI
4 Infinite Reality

Contracted 'Final System'

Additional
1 x 1024 way SMP
1 TFLOP/s
256 GB Memory
3.4 TB Disk
HIPPI6400/FDDI
4 Infinite Reality++

CY97

CY98

CY99




16x 64 way SMP
400 GFLOP/s
256 GB Memory
6.7 TB Disk
HIPPI/FDDI

48x128 way SMP
3 TFLOP/s
1.5 TB Memory
76 TB Disk
HIPPI6400/Gb Ether

'Tech Refresh'

'Final System'

Classified Resources



Software Configuration

- Load Sharing Facility (LSF) access to machines
- Kerberos-5 authentication with ticket forward/refresh
- DCE/DFS local storage 
- HPSS archival storage
- Module-based MIPSpro compilers and MPI software
 - Pros:
 - Very nice for providing multiple revs of compilers/libs
 - Can you fit your OS onto a 9G drive w/o them?
 - Cons:
 - Installable/patched without -Vrulesoverride:on?
 - Not all software is 'module aware' (e.g. cvd/workshop)



MPI/HIPPI Scalability

- Major ASCI code groups running across 1000+ nodes
 - lots of medium-sized MPI messages (65 bytes – 16 Kbytes)
- Parameter space study with environment variables

$$4096 * 16K \leq 16K +$$
$$\#hosts * \left(\begin{array}{c} 128 * \$MPI_MSG_PER_HOST \\ + \\ 16K * \$MPI_BUFS_PER_HOST \end{array} \right)$$
$$+$$
$$\#procs * \left(\begin{array}{c} 128 * \$MPI_MSG_PER_PROC \\ + \\ 16K * \$MPI_BUFS_PER_PROC \end{array} \right)$$



MPI/HIPPI Scalability (cont.)

$$4096 * 16K \leq 16K +$$

$$\begin{aligned} & \#hosts * \left(\begin{array}{c} 128 * \$MPI_MSG_PER_HOST \\ + \\ 16K * \$MPI_BUFS_PER_HOST \end{array} \right) \\ & + \\ & \#procs * \left(\begin{array}{c} 128 * \$MPI_MSG_PER_PROC \\ + \\ 16K * \$MPI_BUFS_PER_PROC \end{array} \right) \end{aligned}$$

- **Code team debug cycle:**
 - run with default settings until code crashes:
 - 'ERROR: unable to pin memory on HIPPI' (out of `_BUFS_`)
 - 'ERROR: out of message packets' (out of `_MSG_`)
 - set `_PER_PROC` environment variables to 0
 - minimize `$MPI_MSG_PER_HOST` and maximize `$MPI_BUFS_PER_HOST` subject to the constraint



MPI/HIPPI Concerns

- **What can be done about high latencies w/ HIPPI 800?**
 - Tradeoff for bandwidth in hardware design of the NIC
 - On-box MPI overhead: 10–15 usec (kaf@cray.com)
 - Cross-box MPI overhead: 119 usec (kaf@cray.com)
 - Addressed in the HIPPI 6400 NIC
 - Myrinet NIC shows promise on PC clusters



MPI/HIPPI Concerns (cont.)

- **Should we be concerned about bit error rates?**
 - Connector/cable problems are much more prevalent, but ...
 - HIPPI bypass means 'bypass'
 - LANL code groups now run with CRC checks by default
 - Fixed with HIPPI 6400 having CRC checks in hardware





**'Patch Man', 'App Woman'
and
'Boy Cable'**

in

'The Bit Riddle'



Outstanding MPI/HIPPI Issues

- **Suspect HIPPI firmware problem**
 - MPI error (misrouted headers?) detected on receiving end
 - Manifested when MPI is running out of resources
 - Accomplished by carefully setting ENV variables
- **Sample test code fails once in 100–200 runs**
 - Reproducible on 5 x 63 nodes
 - Cannot be simply placed into an infinite loop
 - Cannot flood the HIPPI devices to generate the failure
- **Estimated 500 – 1000 debug hours already spent**



Reliability and Stability

- **Open and Secure 'Tech Refresh' Systems:**
 - All too frequent hardware fallout
 - All boxes currently undergoing power supply upgrades(4 hr each)
 - Non-fatal memory errors monitored, swapped on multiple occurrences
 - Average between 10 – 15 unscheduled interrupts per week
- **Reassurance from SGI that 'Final Delivery' Systems are more reliable**
- **A plan to help ensure arrival of good hardware on site:**
 - Systems team to run intensive code on each 'phase' of the delivery
 - Met with opposition by SGI but ...
 - Already successful in finding router problems on 1 of 4 x 128



A Non-subliminal Message



- <http://www.hr.lanl.gov/html/jobs/regjobs.html>
Keyword Search : 983176





Daryl Grunau
dwg@lanl.gov

