Early X1 Experiences at Boeing

Jim Glidewell
Information Technology Services
Boeing Shared Services Group

Early X1 Experiences at Boeing

- HPC computing environment
- X1 configuration
- Hardware and OS
- Applications
- Support
- Summary

Current HPC systems

- Two Cray T-90's
- A 384 CPU Origin 3800
- Three 256-CPU Linux clusters
- Cray X1

BOEING

Cray X1

- Fully populated liquid-cooled chassis
  - 64 MSPs
  - 512 GB of memory, 32GB per node
- Additional Java Server
- Total of 26 terabytes of LSI RAID disk
  - managed by ADIC StorNext software
- X1 is partitioned into two systems
  - 14 nodes production partition
  - 2 node test partition
  - Allows testing of weekly OS updates
  - Added additional complexity to network

Timeline

- Early 2003 - Continuing discussions with Cray on X1
- August 2003 - Detailed plan for X1 transition
- September 19, 2003 - Final approval for X1 acquisition
- November, 2003 - Factory Visit
- January 2, 2004 - System Delivered
- January 15, 2004 - Early User Access
- March 1, 2004 - Limited X1 Production
- March 8, 2004 - Full X1 Production
- April 24, 2004 - StorNext managed SAN put into production

Hardware experiences

- First system shipped with 32GB per node
  - Memory errors seen at factory
  - 1GB DIMMs were out of timing spec
  - Cray developed new memory test to validate DIMMs
  - Resulting in a one month delay in ship

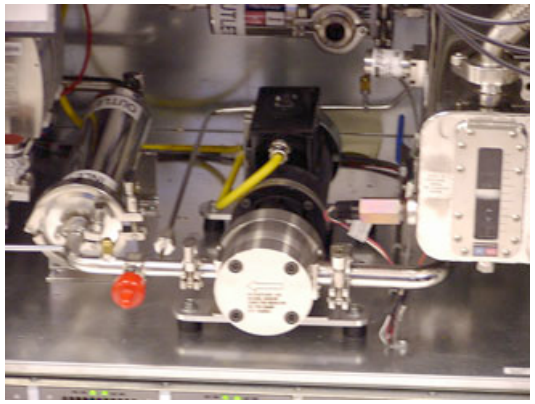

- Fluorinert pumps
  - Four of four pumps replaced
  - One pump replaced twice, due to install error which ran it dry
  - Bad batch has been identified by Cray

- CPU and memory reliability have been excellent

UNICOS/mp experiences

- System stability
- Data compare errors
- Administration and Operations

System Stability

- Kernel panics
  - We've seen very few since start of production in March
  - Application migration is still disabled
  - Job aborts due to node memory oversubscription
- Overall the OS has been stable
  - Exceeded our expectations of one kernel panic per week
  - Improvements continue

Data compare errors

- Problems arose during acceptance testing of disks
- Errors were extremely infrequent (one failure every 72 hours of testing)
- Concerns about data integrity
- Cray dedicated multiple systems to replicating the problem
- Root cause was lack of "I/O cache coherency"
  - I/O started before data was flushed from cache to memory
  - DMA picked up stale data
  - Problem resolved by ensuring all data flushed before writes
  - Similar timing windows closed in other I/O related code

Administration and Operations

- Experience with IRIX really helpful
- Disk configuration & backups
- Network
- PBSPro
- Accounting

Disk configuration

- All disks are LSI RAID
- Complex
- Leaned heavily on Cray support
- Limited guidance on performance tuning
- Delayed StorNext managed SAN for months
- Backups
  - Weekly backups to tape, daily disk to disk copies
  - Long term plan to rely heavily on StorNext

Network

- Lots of components - X1, CNS, CPES, X1-JS, backup CNS
- Very complicated - hope nothing breaks!

Boeing X1 Nets – PAM, RAID, StorNext

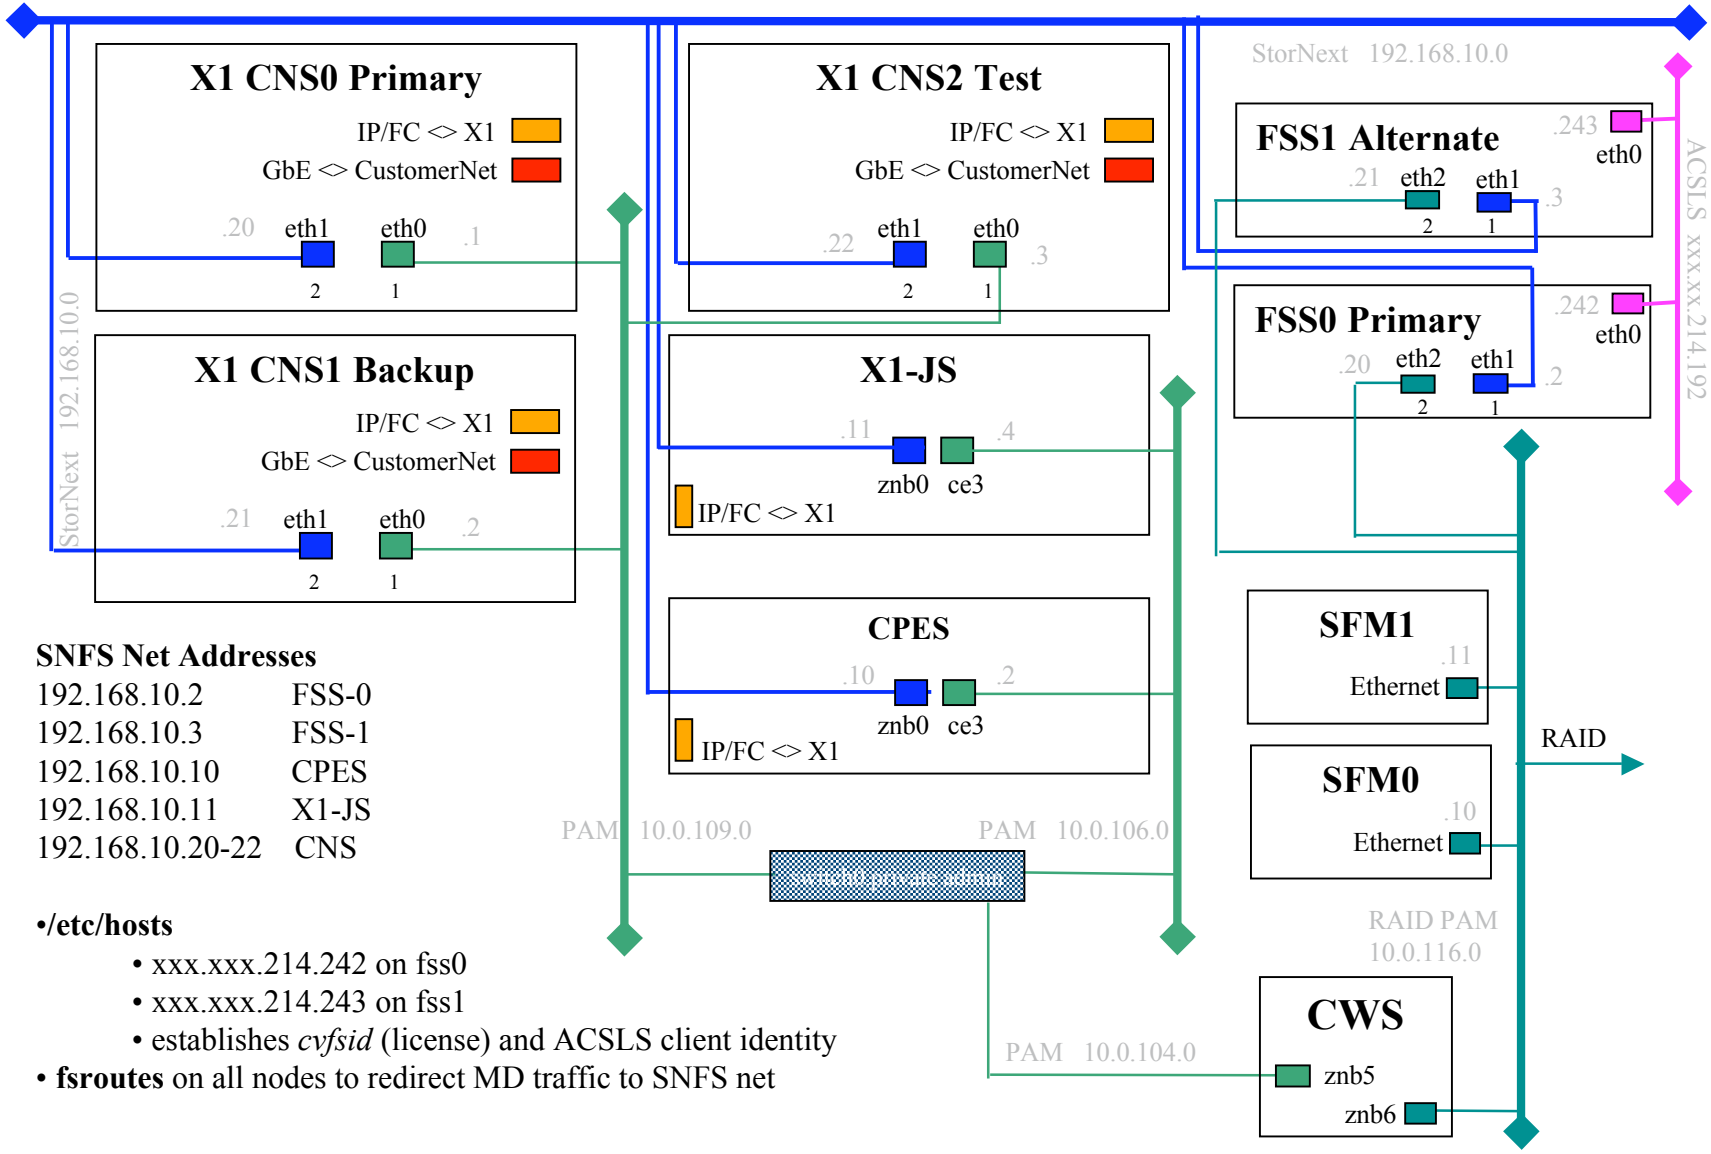

SNFS Net Addresses

192.168.10.2	FSS-0
192.168.10.3	FSS-1
192.168.10.10	CPES
192.168.10.11	X1-JS
192.168.10.20-22	CNS

- **/etc/hosts**
  - xxx.xxx.214.242 on fss0
  - xxx.xxx.214.243 on fss1
  - establishes *cvfsid* (license) and ACSLS client identity
- **fsroutes** on all nodes to redirect MD traffic to SNFS net

PBSPro

- Solid and reliable for us
- Excellent documentation
- User documentation needed (What's an MPPE?)
- Good interaction with psched
  - But PBS knows nothing about flexible vs. accelerated mode
  - Disabled migration can leave applications stuck in posted queue
  - Expect improvements at 2.4

Accounting

- No Cray System Accounting...
- But project accounting is supported
- Using a locally written program to sum usage by user, project
- Session id will be included in UNICOS 2.4 process records

Applications and Tuning

- Compilation time is still an issue
- Cray has provided significant help in getting our key applications performing well on the X1
- Overflow
- Tranair
- Other applications

Overflow

- CFD code developed at NASA Ames
  - Parallel CFD code using MLP (Multi-Level Parallelism)
  - Scales well to large number of nodes (256 or more)
- Runs very well on our NUMA-based system
- Initial performance was roughly 16X that of a 400Mhz CPU on our existing Overflow system
- After tuning by Cray personnel, performance is now 25X
- Code changes have been integrated to NASA's version

Tranair

- Boeing developed CFD code
  - Adaptive grid single-CPU FORTRAN code
  - Out-of-core solver - made heavy use of SSD on T-90s
  - Highly optimized for T-90 series
- Large memory requirement - 7-12GB per job, desire larger
- Forced to use MSP mode due to memory needs
- Initial speed ratio was 0.9-1.18 relative to T-90, both MSP and SSP
- Current MSP speed ratio is 2.0 times the speed of T-90

Support

- Cray has provided excellent support throughout the process
  - Initial hardware and software setup
  - Local technical support
  - Technical folks in Chippewa Falls and Mendota
  - Hardware resources to reproduce and debug problems
- Excellent training & documentation
- Cray's assistance was essential in the success of our installation on a very aggressive schedule

Summary

- Significant hardware and software issues were encountered and resolved
  - Memory DIMM problems
  - Data compare errors
  - Kernel panics
  - StorNext teething pains
- Cray has provided great support throughout the process
- Users are happy with the turnaround and overall reliability of the X1
- The X1 is already a key part of our CFD design process

Coming soon…