



CplantTM Overview

Lee Ward

Scalable Computing Systems

Sandia National Laboratories

lee@sandia.gov

<http://www.cs.sandia.gov/cplant>

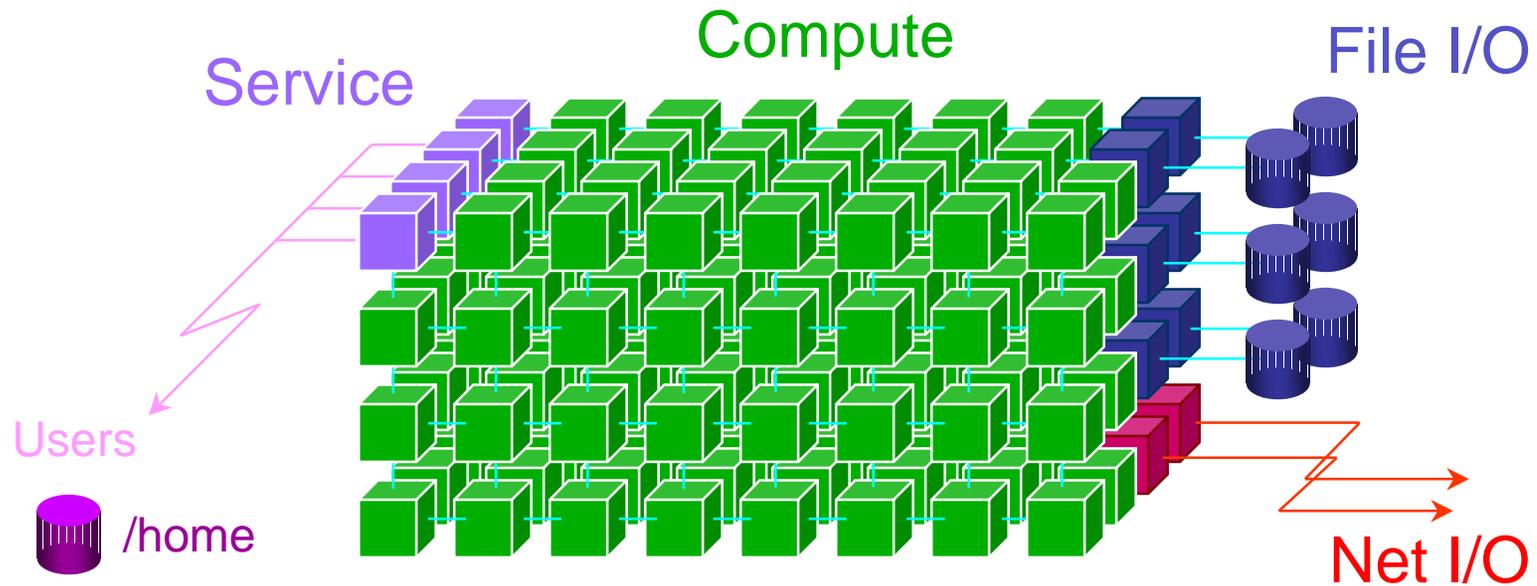


Computational Plant (Cplant™)

- **Hybrid approach combining commodity cluster technology with MPP technology**
- **Emulate the Intel TeraFLOPS environment**
 - Partition model (functional decomposition)
 - Space sharing (reduce turnaround time)
 - Scalable services (allocator, loader, launcher)
 - Complete compute node resource dedication
- **Use Existing Software when possible**
 - Red Hat distribution, Linux/Alpha
 - Software developed for Intel TeraFLOPS



Conceptual Partition Model

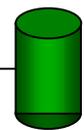




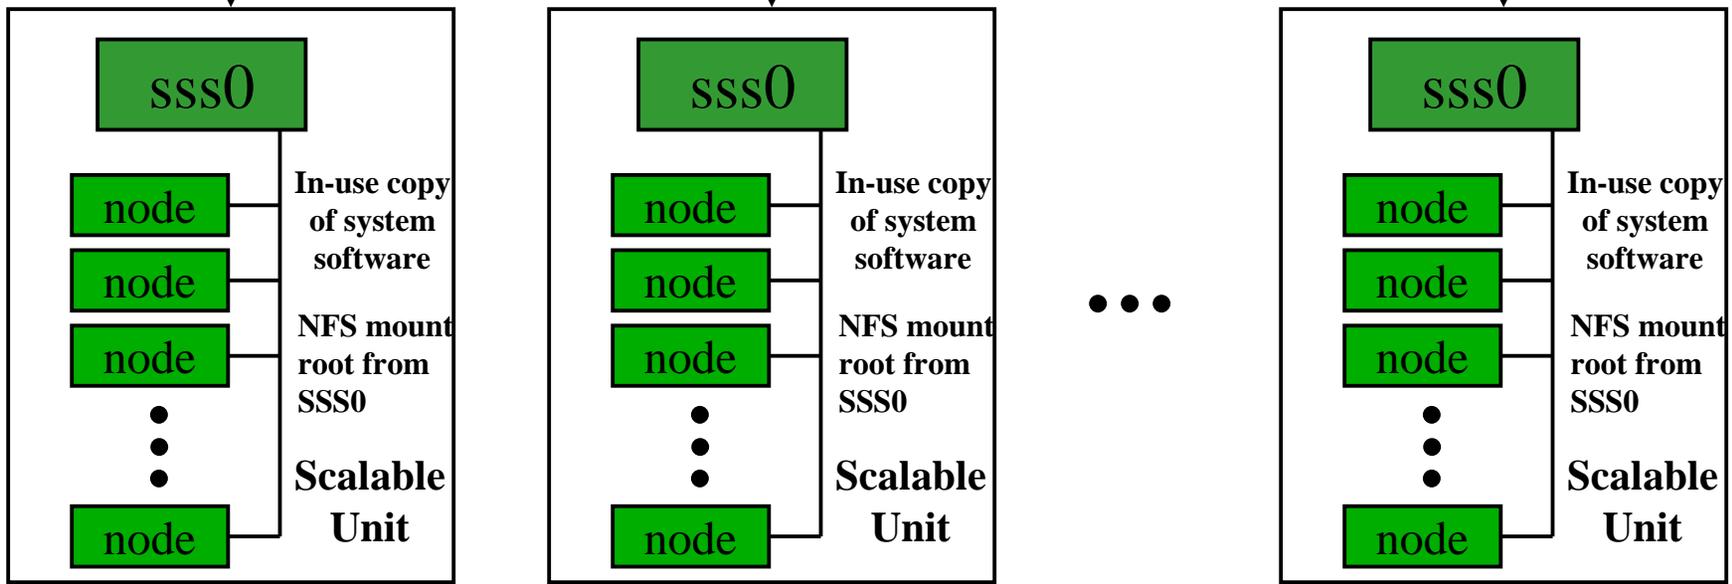
System Support Hierarchy

Admin access

sss1



Master copy of system software





1999 - Phase III- Production (Siberia)

- 624 Compaq XP1000 (Monet)
- 500 MHz Alpha 21264 CPU
- 4 MB L3 Cache
- 256 MB ECC SDRAM
- 16-port Myrinet SAN/LAN switch
- 64-bit, 33 MHz LANai-7 NIC
- 1.73 TB disk I/O
- Integrated by Compaq and Abba Technologies





2000- Phase IV – Development (Antarctica)

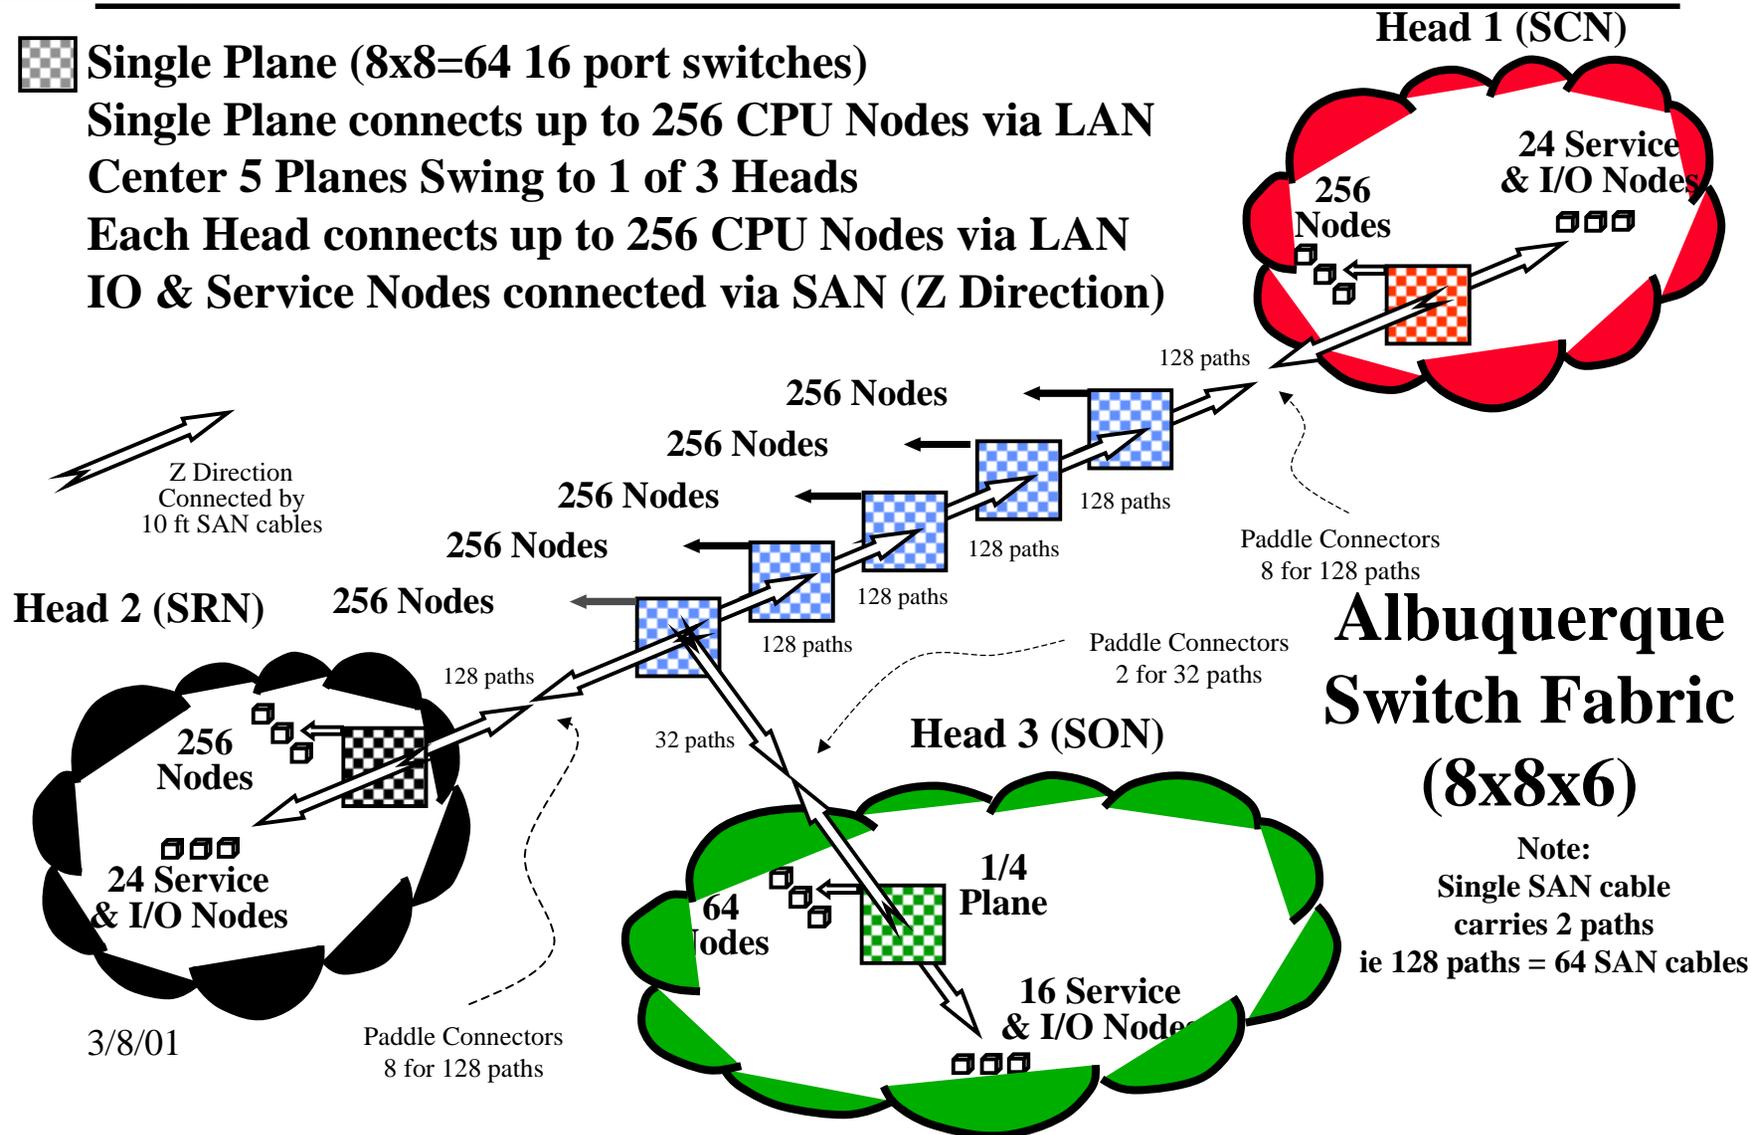
- **~1760 DS10 Slates (NM+CA)**
- **466MHz EV6, 256 MB RAM**
- **Myrinet 33MHz 64bit LANai 7.x and 9.x**
- **Will be combined with Siberia for a 2000+-node system**
- **Red, yellow, green switchable**





Antarctica Topology Layout

-  **Single Plane (8x8=64 16 port switches)**
 Single Plane connects up to 256 CPU Nodes via LAN
 Center 5 Planes Swing to 1 of 3 Heads
 Each Head connects up to 256 CPU Nodes via LAN
 IO & Service Nodes connected via SAN (Z Direction)





Network Fault Attributes

- Each “plane” has multiple, redundant links
- The z-axis is 128 parallel, redundant links
- Each plane, and the z-connections are contained in one physical unit
- Claimed error rate of 2^{-18} holds so long as no node is rebooted
- At reboot, NIC generates noise to the switch
 - Corrupts cross-traffic
 - Deadlocks cross traffic



Runtime Environment Components

- **Yod**
 - Parallel job launcher via fan-out
 - Retries on load failure
- **Bebopd**
 - Compute node allocator, control and status
 - Single point of failure
- **PCT (Process Control Thread)**
 - Compute node daemon
 - Not aggressive about communication with the bebopd
- **Pingd/Showmesh**
 - Compute node status tool
 - Displays the state of the compute nodes



Parallel I/O

- **ENFS filesystem– Extended NFS**
 - **Removes locking semantics from NFS protocol**
 - **Parallel independent I/O to multiple files**
 - **Non-overlapping access to single file**
 - **Internal (compute partition) protocol is related to NFS V2 and remains stateless**
 - **Uses I/O nodes as proxies**
 - **External protocol is pure NFS V2**
 - **Very tolerant of communication related problems**



Next Steps

- **Scalability**
 - To 3000+ nodes in FY01
 - To 5000+ in FY03
- **Performance**
 - Portals 3.0 OS bypass implementation
- **User Tools**
 - TotalView
- **Distributed Resource Management**
 - Globus
- **Establish Cplant™ as a capacity ASCI resource for the tri-labs**



Current Areas of Development

- **Thread support for system libraries**
- **Tools for scalable system management**
- **Tools for automated non-invasive monitoring**
- **Automating system integration**
- **Scalable I/O**
- **System robustness**
 - **Reliability, availability and serviceability**