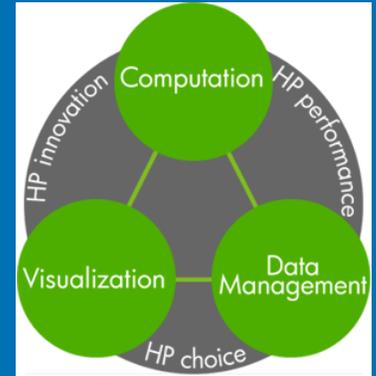


SRNWP Workshop

HP Solutions and Activities in Climate & Weather Research

Michael Riedmann
European Performance Center



Agenda

- A bit of marketing: HP Solutions for HPC
- A few words about recent Met deals
- Met Application Support at HP

HP Solutions for HPC

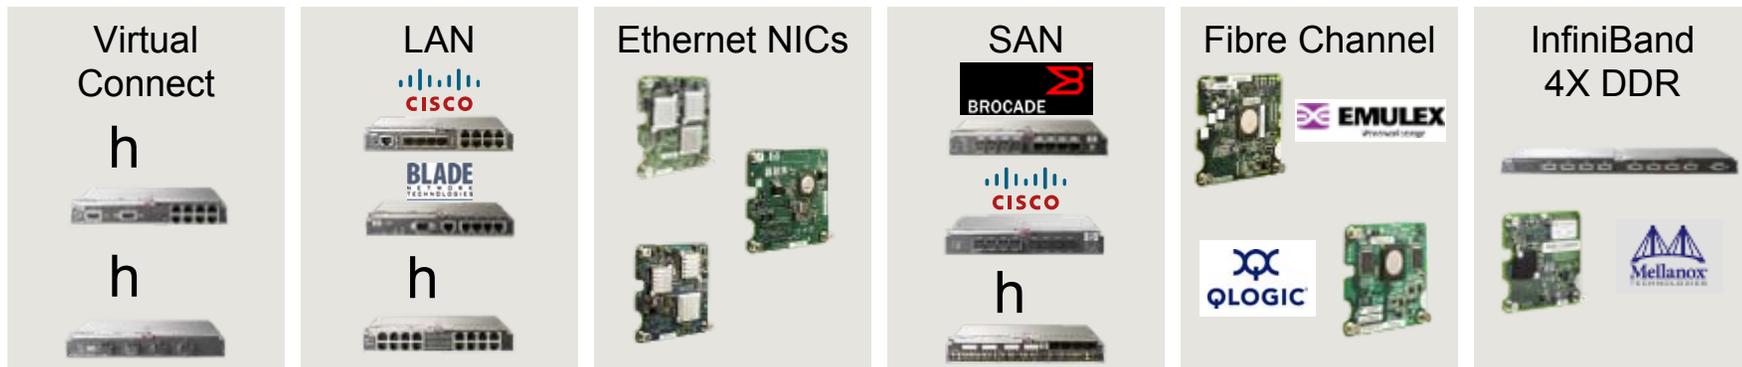
HP BladeSystem c-Class Portfolio



A Full Range of 2P and 4P Blades



Interconnect choices for LAN, SAN, and Scale-Out Clusters



Unified Management



Choice of Power

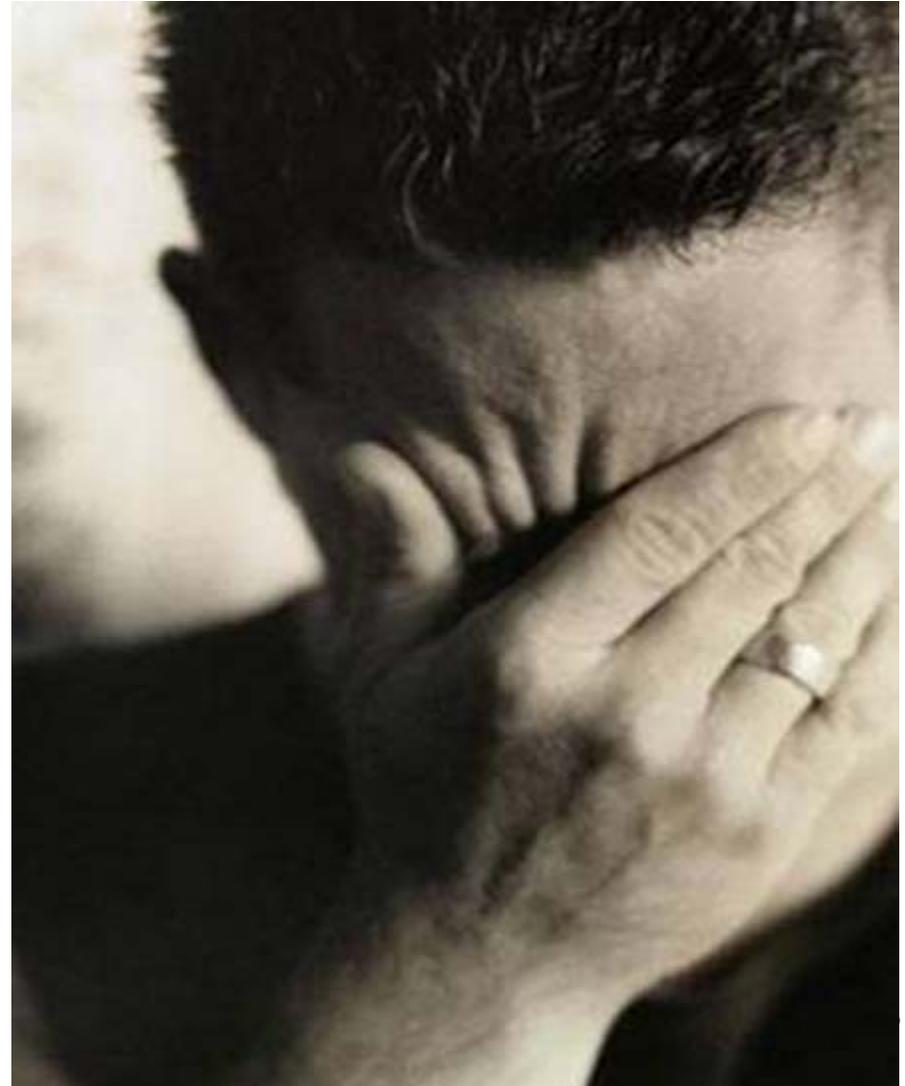


Complete Services

Assessment
Implementation
Support

Cluster Implementation Challenges

- Manageability
- Scalability
- Integration of Data Mgmt & Visualization
- Interconnect/Network Complexity
- Version Control
- Application Availability



XC System Software

- A complete HP-supported Linux cluster stack for operation and management of HP clusters
- An integrated and tested collection of best-in-class software technologies
 - open source where appropriate
 - easily portable across servers and interconnects
- HP-developed installation and configuration procedures and defaults, management commands, scripts and plug-ins

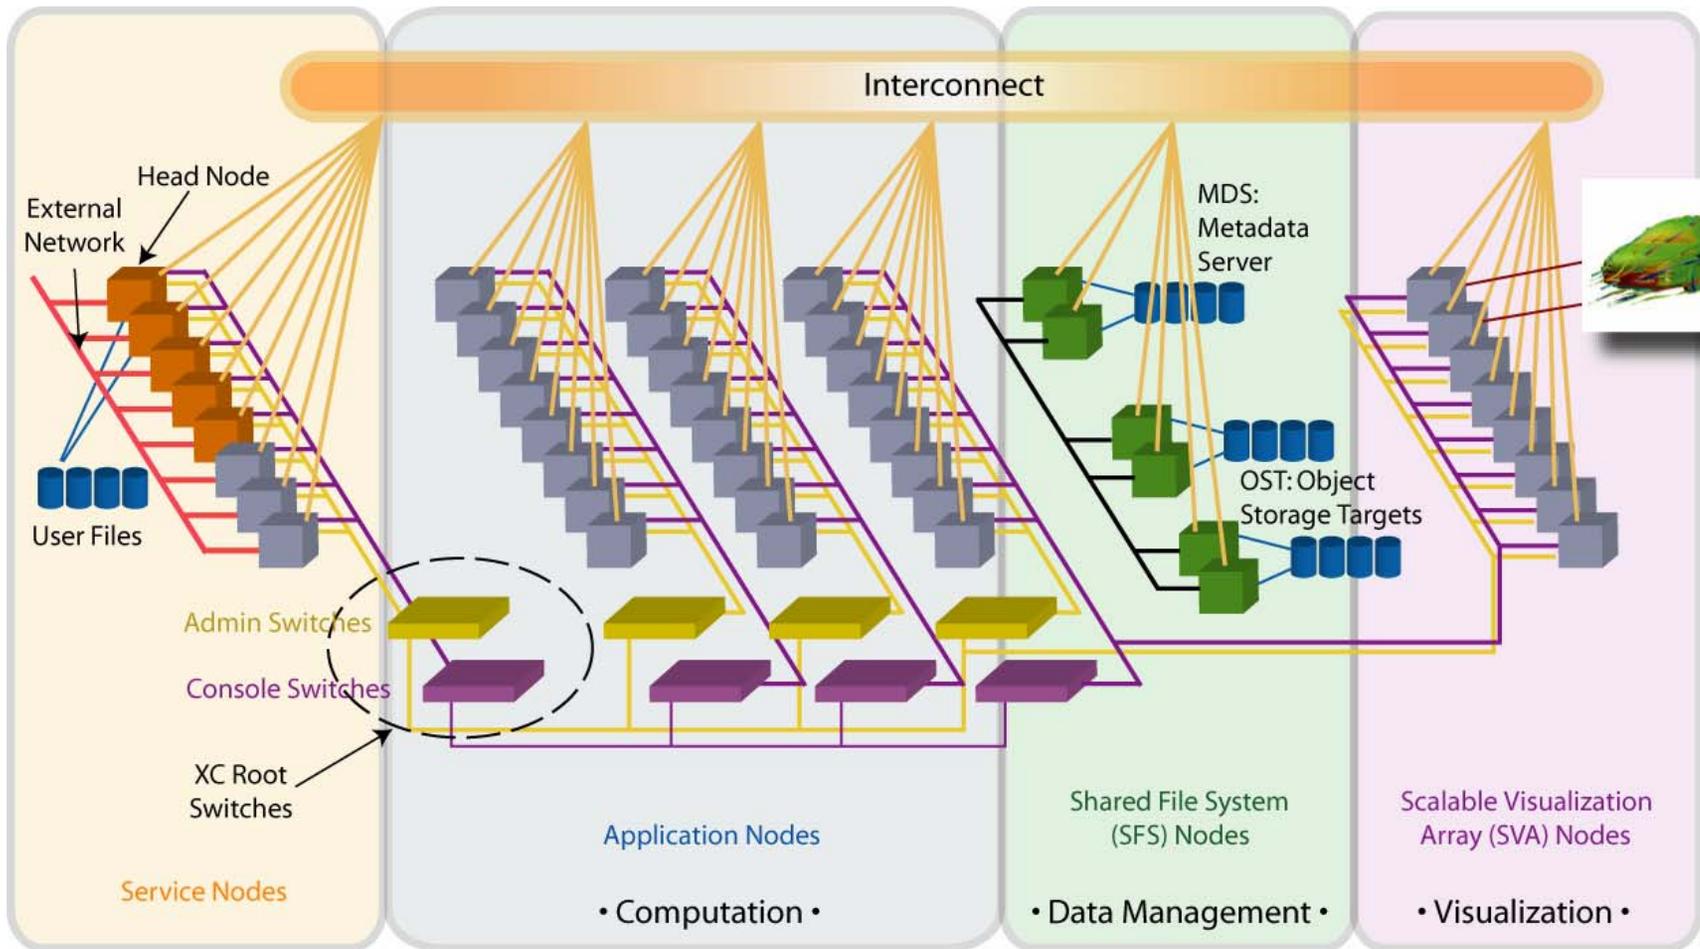
XC System Software V3.0

- Available for Xeon, Opteron and Itanium HP servers and optimized for CP
- Worldwide HP support
- Designed to be compatible with Red Hat EL 4.0 U2
- Includes:
 - Suite of system mgt and monitoring utilities
 - LSF workload mgt (choice of LSF HPC or LSF standard)
 - Interconnect libraries
 - HP MPI (MPI-2 compliant)
 - Modules for version selection of tools
 - Lustre client support
 - Preconfigured firewall
 - Preconfigured NAT router services
- LVS for single system login
- Some key availability features
- **Designed for scalability**
- Can be integrated with high performance viz
- Suite of qualified ISV apps

XC V3: Key Software Technologies

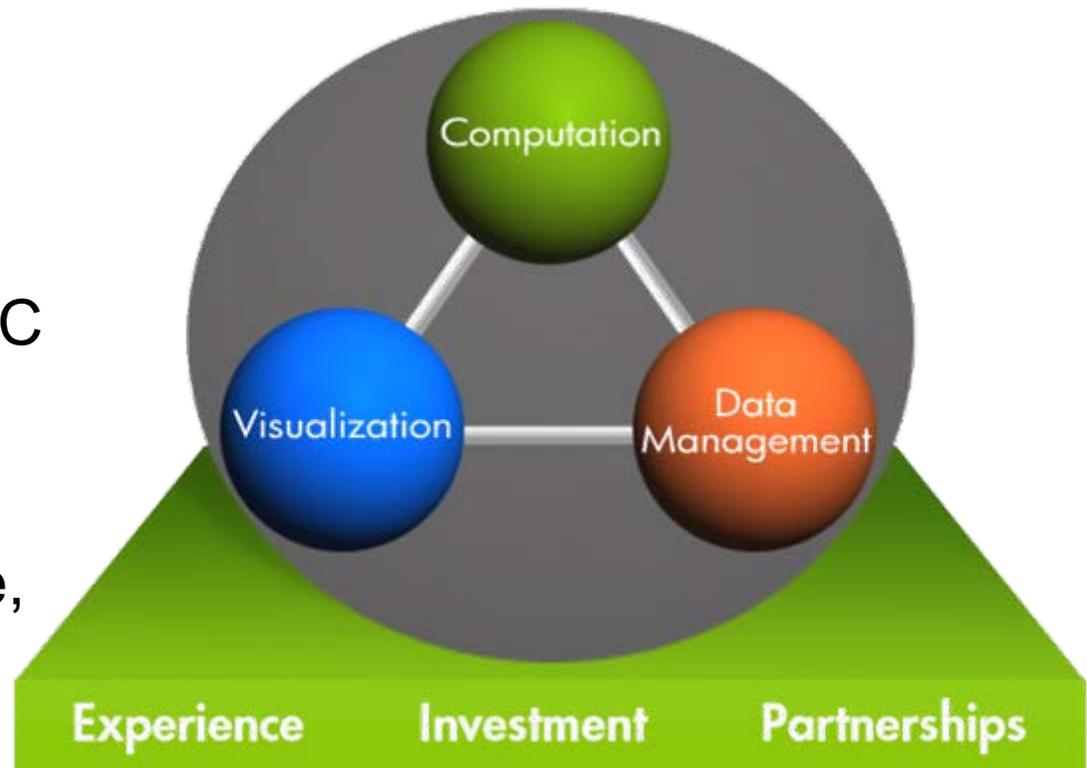
Function	Technology	Features and Benefits
Distribution and Kernel	RHEL 4	Red Hat Current shipping product, Posix enhancements, support for latest Opteron and Core (“Woodcrest”), ISV support
Inbound Network / Cluster Alias	LVS	Linux Virtual Server High availability virtual server project for managing incoming requests, with load balancing
Batch	LSF 6.x	Platform LSF HPC Premier scheduler, policy driven, allocation controls, MAUI support. Provides migration for AlphaserverSC customers
Resource Management	SLURM	Simple Linux Utility for Resource Management Fault tolerant, highly scalable, uses standard kernel
MPI	HP-MPI 2.x	HP’s Message Passing Interface Provides standard interface for multiple interconnects, MPICH compatible, support for MPI-2 functionality
System Files Management	SystemImager Configuration tools Cluster database	SystemImager Automates Linux installs, software distribution, and production deployment. Supports complete, bootable image; can use multicast; used at PNNL and Sandia
Console	HPLS_PowerD [power mgmt] Telnet based console commands	Power control Adaptable for HP integrated management processors – no need for terminal servers, reduced wiring IPMI, ILO server interfaces to low level console controls CMF content management framework
Monitoring	Nagios SuperMON	Nagios Browser based, robust host, service and network monitor from open source. SuperMon supports high speed, high sample rates, low perturbation monitoring for clusters.
High Perf I/O	Lustre 1.2.x	Lustre Parallel File System High performance parallel file system – efficient, robust, scalable

XC System Architecture



Delivering a Complete HPC Solution

- Innovation based on standards
- Broadest choice of customer-focused HPC solutions
- Affordable, accessible, supercomputing performance



A few words about recent and current Met deals

DWD (1)

- Proposals submitted by HP, NEC and IBM
- HP Solution for the compute servers was based on BL685c blades (4 Sockets, 16-Cores) with AMD Barcelona Quad-Core processors.
- Power & Cooling requirements were met.
- Suspend-Restart Scenario was tested successfully.

Anyway ... NEC made the best offer ☺

DWD (2)

What went well

- Benchmarking: One single application,
Code is well known,
Performance requirements were reasonable.

What went not so well

- Sizing of DB-Server and DB-Storage caused some headaches.
- AMD and Intel processor roadmaps are unreliable
Difficult to make long-term performance commitments

UK Met Office

... ongoing activity ... no comments.

Qing Dao, China

Focus on Ocean Modelling

- HP won this deal and is delivering a full-blown SuperDome system (128-way Itanium SMP)
- OS is HP-UX due to superior compiler

Met Application Support

Working on Top10 Met Applications

The most important applications are assigned to dedicated engineers within a WW team.

HIRLAM, Aladin, ARPEGE/IFS, LM_RAPS, HRM, MM5, WRF, GFS, FMS, MOM4, ETA

Activities:

- Performance characterization
- Migration to new OS and compiler platforms
- Optimisation and benchmarking

Performance Characterisation with LM_RAPS

Single node comparison



Test case: LMK_Small 100x100x50, ts=30 sec, Time for 1 hour forecast

LMK_Small is a downsized model which shows the same behaviour as full size LMK model.
The data set per core is roughly the same as in the full size model.

Hardware	Compiler	Cores / Node	1 Core	4 Cores	8 Cores
HP BladeSystem BL 460c Xeon 5160 DC 3.00 GHz 4 MB Xeon 5150 DC 2.66 GHz 4 MB	Intel 10.0 beta	4 4	431 sec 468 sec	186 sec 191 sec	
HP BladeSystem BL460c Xeon 5355 QC 2.66 GHz 4 MB	Intel 10.0 beta	8	470 sec	192 sec	190 sec
HP BladeSystem BL 465c Opteron revF DC 2.6 GHz 1 MB	PGI 7.0 AMD64e	4	529 sec	179 sec	
HP BladeSystem BL 685c Opteron revF DC 2.8 GHz 1 MB	PGI 7.0 AMD64e	8	524 sec	166 sec	91 sec
AMD Whitebox Barcelona QC 2.0 GHz Proto	PGI 7.0 AMD64e	8	560 sec	166 sec	114 sec
HP Integrity RX2660 Itanium2 DC 1.6 GHz 9 MB	Intel 10.0 beta HP-UX f90 v3.1.1	4 4	395 sec 417 sec	178 sec 165 sec	

Thank You !

