



CU G 2008 HELSINKI • MAY 5–8, 2008
CROSSING THE BOUNDARIES

Cray Operating System and I/O Road Map Charlie Carroll

Cray Operating Systems Focus

■ Performance

- ⚙️ Maximize compute cycles delivered to applications while also providing necessary services
 - ▶ Lightweight kernel on compute node
 - ▶ Standard Linux environment on service nodes
- ⚙️ Optimize network performance through close interaction with hardware

■ System stability

- ⚙️ Correct defects which impact stability
- ⚙️ Develop and implement features to increase system robustness

■ Scalability

- ⚙️ Scale to increase performance without compromising stability
- ⚙️ Provide better system management tools to manage larger systems

Cray Operating Systems and I/O

■ Compute node kernels

- ✿ XT CNL
- ✿ XT Catamount
- ✿ X2 CNL
- ✿ XMT

■ Service node kernel

- ✿ Supports all compute nodes

■ File systems

- ✿ Lustre
- ✿ DVS (Data Virtualization Service)

■ Networking

- ✿ Portals
- ✿ TCP/IP
- ✿ uGNI and DMAPP

■ Operating system services

- ✿ Checkpoint / restart
- ✿ Node health daemon
- ✿ CSA (Comprehensive System Accounting)

■ System management

- ✿ Interface to system data
- ✿ ALPS (Application-Level Placement Scheduler)
 - ▶ Interfaces to PBS Pro, Moab/Torque and LSF
- ✿ Command interface

Cray OSIO Themes

■ System stability

- ❁ Failover
 - ▶ Lustre
 - ▶ Service nodes
- ❁ Portals & Lustre
 - ▶ Significant effort to improve robustness, defect corrections, and increased testing
- ❁ Node health check
 - ▶ More and better tools to evaluate compute node health

■ Performance

- ❁ Tension between lightweight kernel and features
 - ▶ We'll hold the line on features
- ❁ Huge page support
- ❁ Analyze efficacy of topology-aware job placement

Cray OSIO Themes, continued

■ System management

- Unify the interface to system management data
- “Play nicely” with customers' existing data center infrastructure
- Look ahead to increasing scale

■ Support hardware

- Ideal is to release software in advance of the hardware
 - ▶ OS for Quad-core, PCIe, and NUMA support went well

■ Lustre

- Work with Sun to build their test capability
- Continue to improve our troubleshooting tools
 - ▶ Nic Henke talk at 9:15am on Thursday
- Possibly become more selective about taking Lustre features

Cray OSIO Themes, continued

■ System size and scalability

- ✿ Portals working to run Global Arrays across some very big systems

■ Internal infrastructure

- ✿ Become more like Linux in our build and delivery infrastructure
- ✿ Better mechanics, such as kernel source release

OSIO Release Process

- Moving to two significant releases per year
 - ✿ GA in roughly Q2 and Q4
 - ✿ LA release one quarter earlier
 - ▶ LA to GA requires testing at large (40+ cabinets) scale
- Mid-release hardware may be supported with a product-specific release
 - ✿ XT5 will require v2.1HD release
 - ✿ Goal is to minimize risk to the v2.1 customer base
- Maintenance releases will be consolidated and scheduled
- Moving toward having the ability to release service node software independently of the compute nodes

Features in Amazon (XT V2.1, GA in 3Q08)

- Lustre 1.6
 - ✦ Performance improvements
 - ✦ New configuration methodology
- DVS (Data Virtualization Service)
 - ✦ Ability to project NFS to compute nodes
- SLES10 SP1
 - ✦ Kernel and user space
 - ✦ Automated site data migration tools for software upgrades
- SIO node reboot
 - ✦ Increased system uptime
- Node health, phase 1
 - ✦ Ping nodes of jobs which terminate abnormally
 - ✦ Admin-downs the nodes that do not respond

Features in Amazon (XT V2.1), continued

- CSA (Comprehensive System Accounting)
 - ✿ System management and billing
- Mazama log manager
 - ✿ Centralized log management
 - ✿ Search, filtering and log features
- Virtual Channel 2 (VC2)
 - ✿ Higher throughput in some high-load situations such as all-to-all
- Kernel changes for NUMA
 - ✿ Needed for XT5; base kernel going forward
- EAL3 support
 - ✿ Security validation

Features in XT's Congo Release (GA in 2Q09)

- Node health, phase 2
 - ✿ User configurable for when to run, how to react to errors
 - ✿ More checks: file systems and OS
 - ✿ Initiated locally on each node, that is, scalable
- Attribute management
 - ✿ Single, documented interface to system information
- SLES10 SP2
- Build split
 - ✿ Mostly internal, RPMs put in locations more like Linux
 - ✿ Source updates easier

Features in XT's Congo Release, continued

- Checkpoint / restart
 - ✱ Mitigates job failures
 - ✱ Support MPI and Shmem applications
- Portals changes for XT5
 - ✱ Better network performance in XT5's NUMA architecture
 - ✱ More consistent performance
- SDB node failover
 - ✱ Aids system resiliency
- LDAP integration into CSA
 - ✱ Eliminates need for separate user database for CSA
- DVS
 - ✱ See Dave Wallace's talk at 8:45am on Thursday

Features in XT's Congo Release, continued

- Package manifests
 - ✿ Smoother installation process
- Open Fabric Enterprise Distribution (OFED) / Infiniband support
 - ✿ Enabler for external Lustre

- Catamount not supported in Congo and later releases

Features Being Discussed

■ External Lustre

- ✿ In 2008 we will provide IB cable to connect to customer-provided Lustre servers
- ✿ Broader involvement under discussion

■ External login nodes

■ Dynamic libraries

■ Resiliency features

- ✿ We will do more for system and application resiliency. Exact features are under discussion.

■ SLES11

- ✿ We expect to track Novell's releases

Baker-Gemini Features

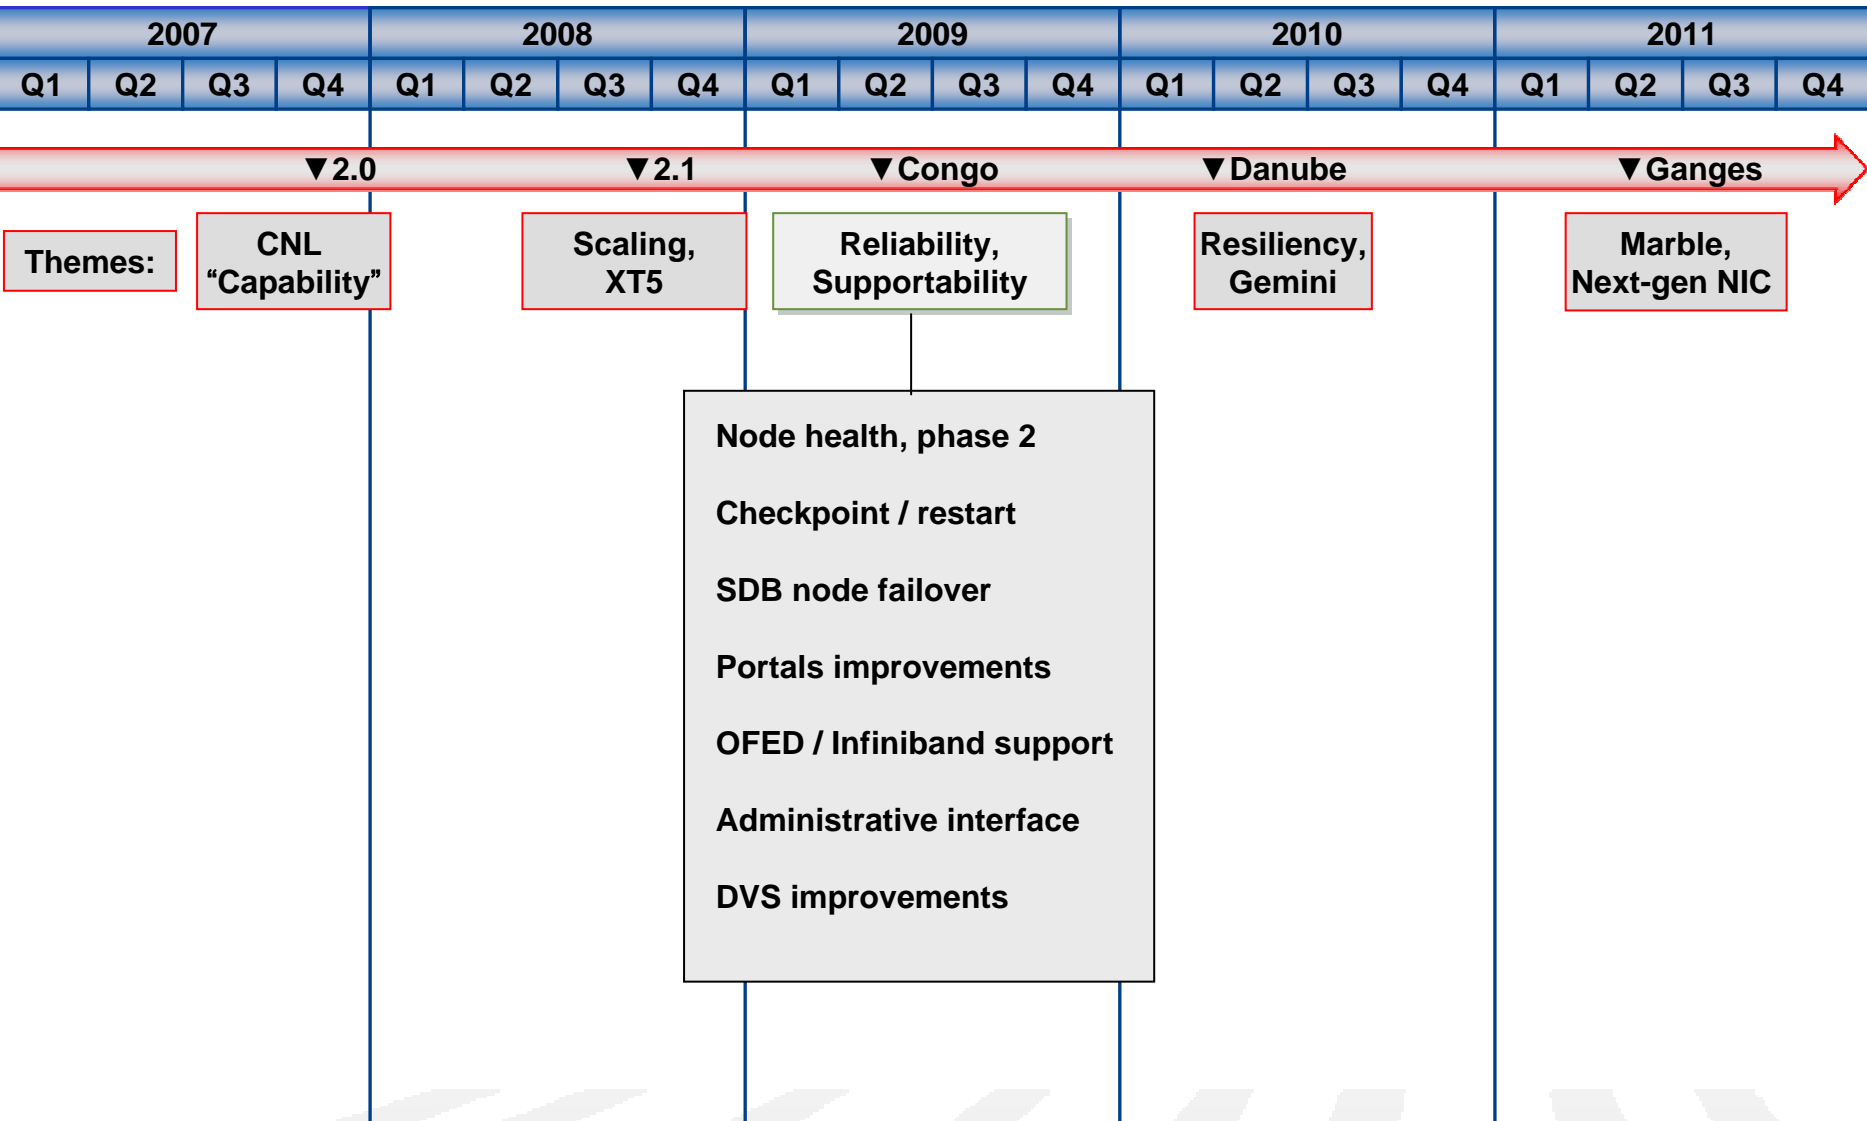
■ Support for Gemini

- ✿ Support for MPI applications via User-level Gemini Network Interface API (uGNI API)
- ✿ Support for PGAS languages via Gemini Distributed Memory Applications API (DMAPP API)

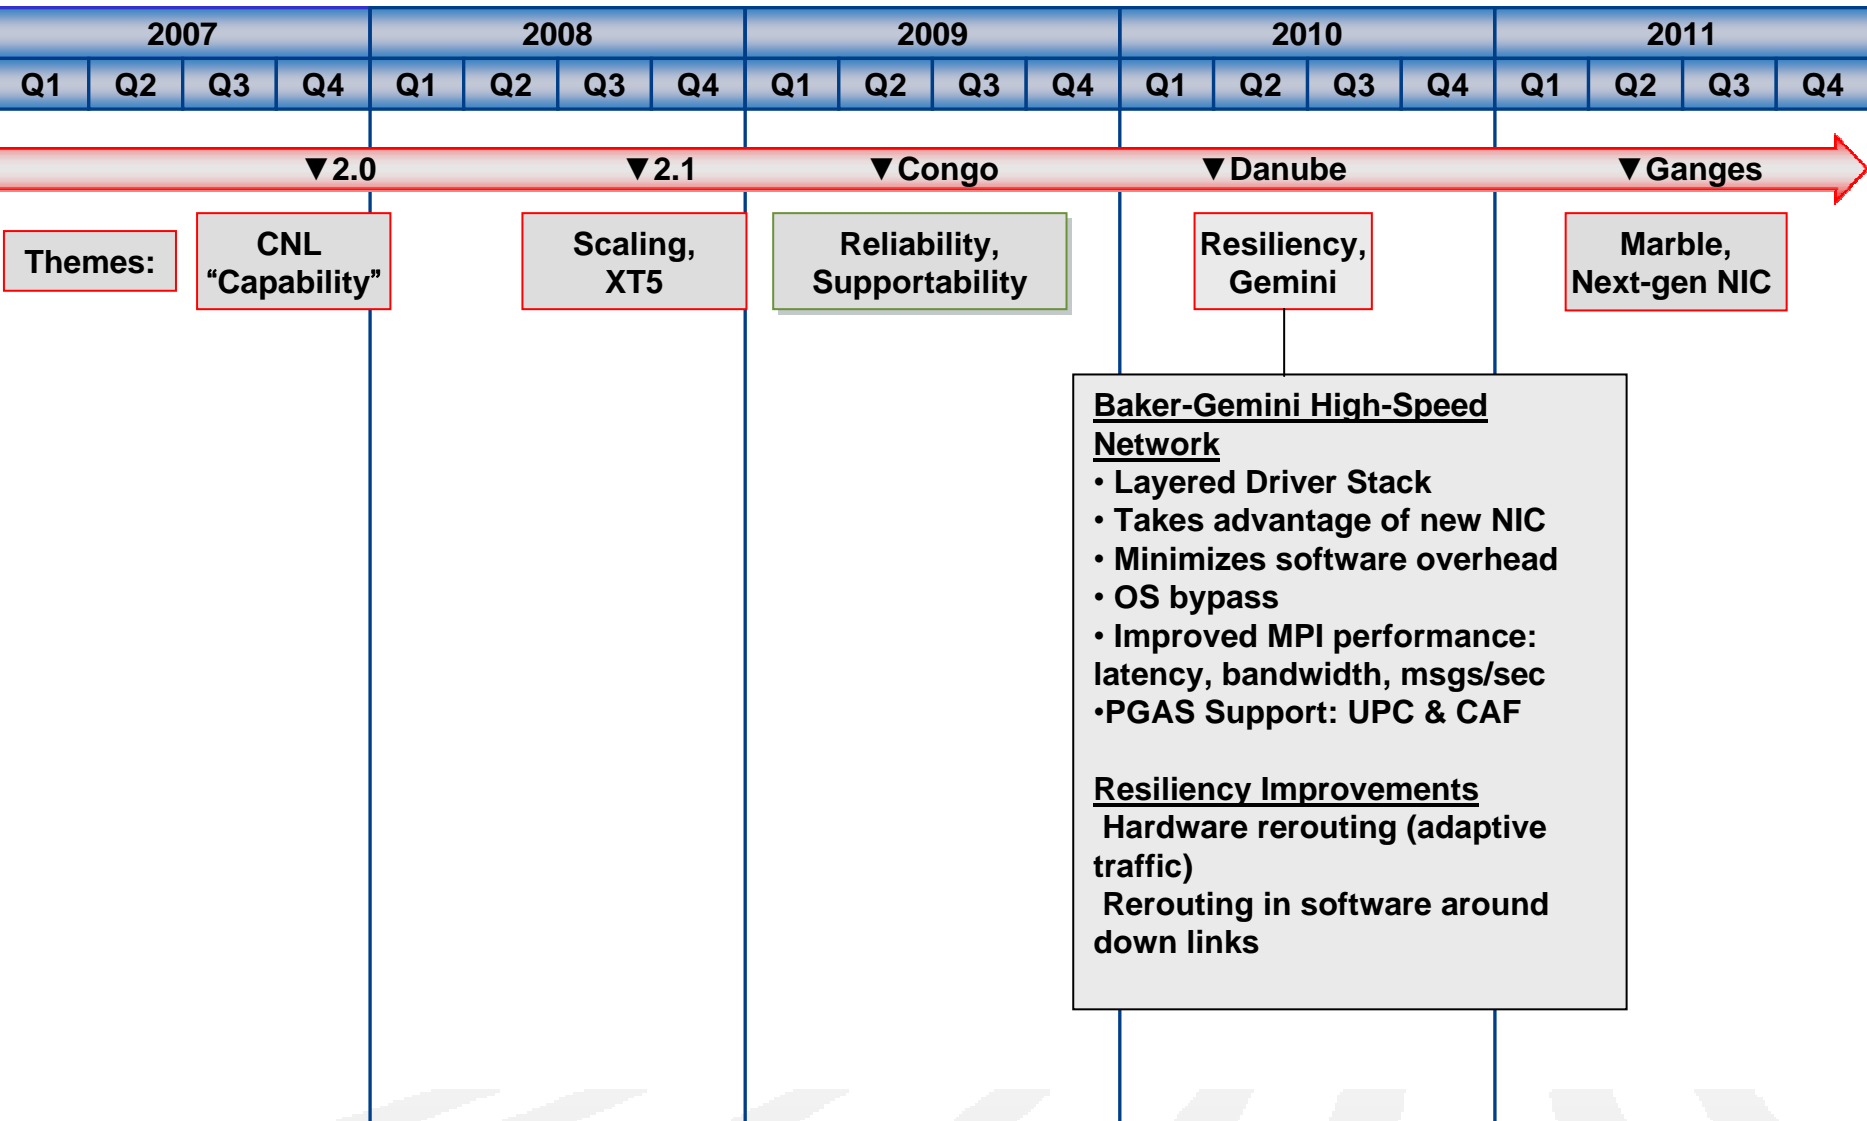
■ Link resiliency

- ✿ Baker/Gemini will provide the capability to ride through many types of link errors. A single hardware link failure will not take down the entire system, although some applications may be terminated.

Cray Linux Environment (CLE) Congo



Cray Linux Environment (CLE) Danube



Cray Linux Environment (CLE) Ganges

