



Introducing Project Carlsbad

Technology for a New Era in Computing



This presentation contains forward-looking statements regarding the SGI Altix® XE server family and roadmap, other SGI® technologies, and third-party technologies that are subject to risks and uncertainties. These risks and uncertainties could cause actual results to differ materially from those described in such statements. The viewer is cautioned not to rely unduly on these forward-looking statements, which are not a guarantee of future or current performance. Such risks and uncertainties include long-term program commitments, the performance of third parties, the sustained performance of current and future products, financing risks, the impact of competitive markets, the ability to integrate and support a complex technology solution involving multiple providers and users, the acceptance of applicable technologies by markets and customers, and other risks. These forward-looking statements are subject to risks and uncertainties as set forth in the company's Forms 8K dated September 8, 2006, and most recent SEC reports on Form 10-Q and Form 10-K. Silicon Graphics is under no obligation to publicly update or revise any forward-looking statements, whether to reflect new information, future events or otherwise.

©2006 Silicon Graphics, Inc. All rights reserved. Silicon Graphics, SGI, SGI Altix, the SGI logo and the SGI cube are registered trademarks. SGI ProPack, Performance Co-Pilot, and Innovation for Results are trademarks of Silicon Graphics, Inc., in the United States and/or other countries worldwide. Linux is a registered trademark of Linus Torvalds in several countries. Linux penguin logo created by Larry Ewing. Itanium and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries trademarks. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries. Windows is a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. All other trademarks mentioned herein are the property of their respective owners. (11/06). Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Images courtesy of Stony Brook University, NASA Ames, gsiCom, Accelrys, Landmark, and Leonard Wikberg III.

Product plans, descriptions, and dates are estimates only and subject to change without notice. SGI may choose not to announce or make generally available any products or programs discussed in this presentation. Therefore, you should not make changes in your business operations on the basis of the information presented here.

Carlsbad, Dixon, Ultraviolet, Santa Fe, Oro Valley, Chama and Taos are internal project code names.



SGI Today

Industry Leading Innovation

- More than 1600 employees
- 800+ Customer-facing employees
- 300+ Engineers to continue innovation
- More than \$500m in annual revenues
- 6000+ Customers in over 50 countries around the world




Core technology leadership in

- Advancement of Linux® OS into HPC market
- Scalable system architecture
- Global shared memory
- File systems and shared storage
- Consulting and services

SGI Unique Capabilities

20+ Years of expertise in solving the most demanding compute and data-intensive problems

- Unified server, cluster and storage architecture
- Wide use of Open Standards, including Linux® OS
- Largest and fastest storage systems
 - Global memory address to over 100TB
 - Filesystems over 100TB and 12GB/s disk-to-SAN
- Renown for deep vertical expertise of employees
 - More than 200 employees with security clearances
- World class customer service organization



Project Carlsbad

Next-generation integrated blade platform, with breakthrough performance density and reliability.


DENSITY

POWER

RELIABILITY

Project Carlsbad: Technology for a New Era in Computing


- **Next generation blade platform for breakthrough scalability and price/performance.**
 - Modularity to add & update resources independently for perfectly right-sized systems: memory, storage, processors
 - Packaged for best overall price/performance – 512 Intel Xeon Processor cores per rack, easily scales to thousands of processors.
- **Integrated blade platform reduces complexity, simplifies management, and lowers total cost of ownership.**
 - 50% less space (based on Tflops/rack versus ‘rack’ or ‘box’ competitors)
 - Fewer blade components reduce potential points of failure.
 - Leading energy efficiency: average \$100K in annual savings for 10 Tflops compute power.
 - Enhanced serviceability, blade based platform that is monitored and managed at the blade, chassis, rack levels.
 - Fully redundant system components, hot-swappable blades.
 - SGI Platform Manager (name TBD) provides multi-level management across complete Carlsbad system.
- **SGI “out of the box” deployment, backed by SGI world-class support and service, for immediate productivity.**
 - 10 Tflp of compute power “up and running” user apps in a day
 - Standards-based – Intel Xeon Processor technology, certified Linux, Microsoft Windows CCS
 - Fully integrated, includes SGI Platform Solution (name TBD), a complete software solution stack.




System Hardware Overview

Project Carlsbad Compute Blade

- Intel® 5000X Chipset (Greencreek)
- (2) Intel® Xeon® DP SKU Processors
 - Dual-core Woodcrest
 - Quad-core Clovertown
- (8) Fully buffered memory DIMM slots per blade
 - 1GB, 2GB, 4GB DIMMs
 - 32GB Memory Support
- (2) x4 DDR IB ports on embedded HCAs
- No on-board storage
- Power: 487W at 12VDC (high-bin processor SKU & (8) 2GB FB-DIMMs (2GB/core))



10U 16-Node Individual Rack Unit (IRU)




Top View

(4) 4x DDR IB Switch Blades Shown

(2) 24-Port IB switch ASICs per blade
(6) 4X IB + (1) 4X IB external
Cable connections per blade

**(7+1) 1625W 12VDC Output
Front-End Power Supplies**

**(16) 2-Socket Nodes
(Supports (8) 4-Socket Nodes)**

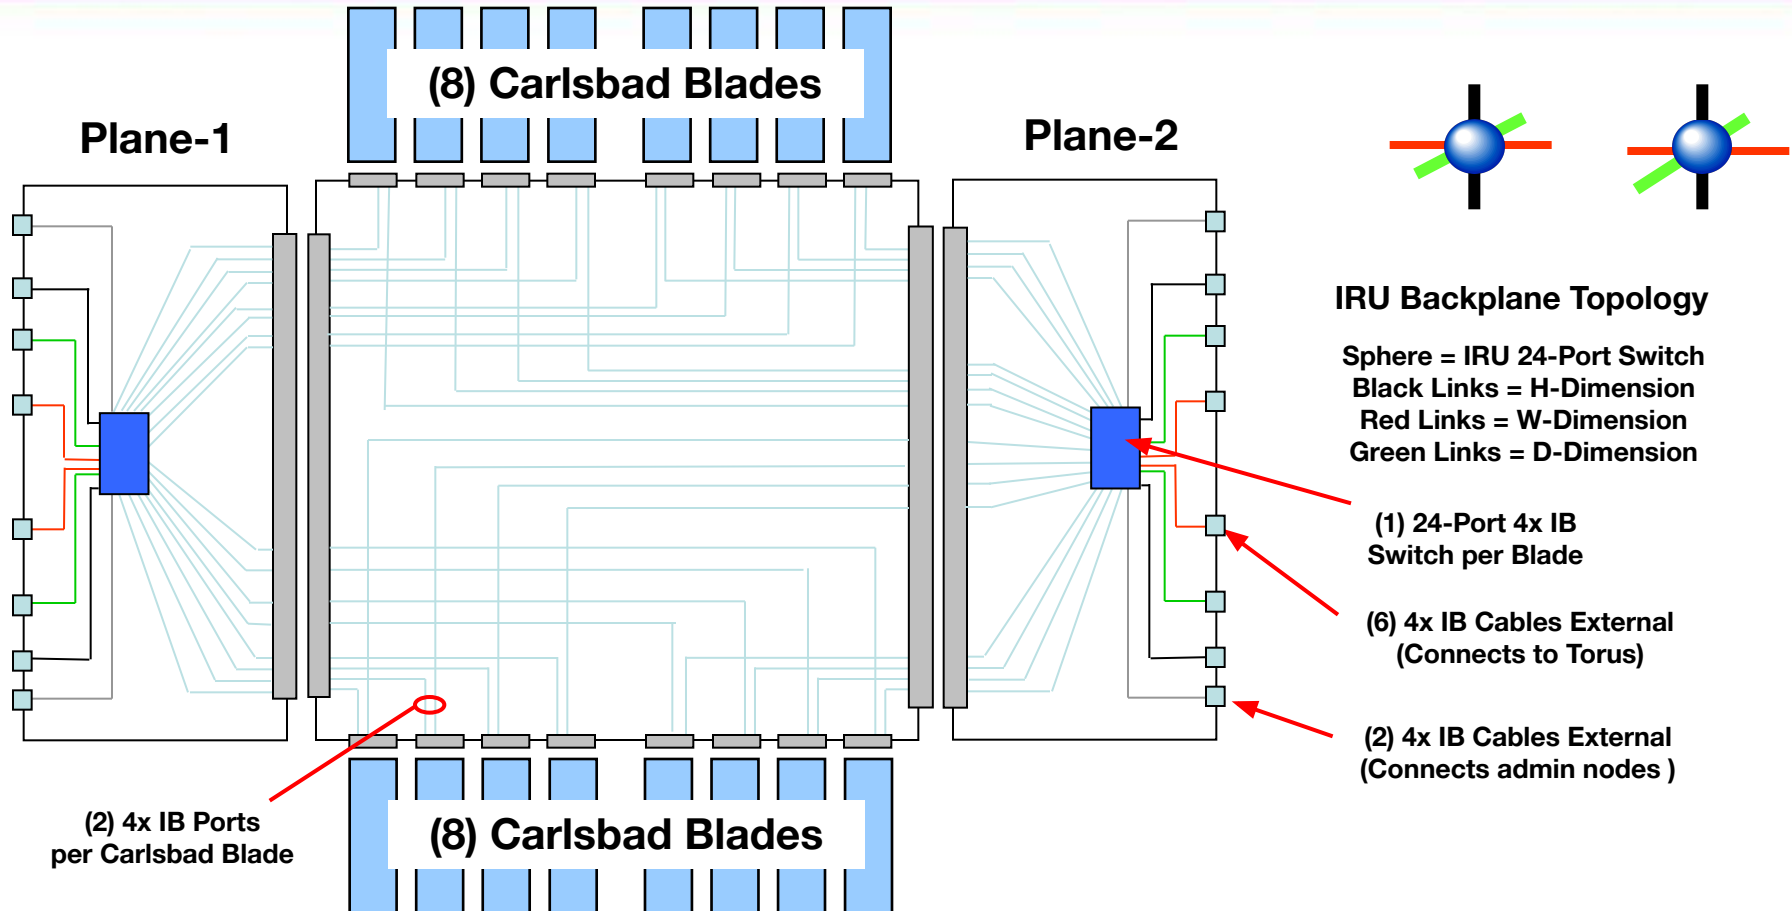


Front View

**10U 24-inch EIA Form Factor
(17.50-in H x 22.5-in W x 32-in
D)**


**(1) Chassis
Management
Controller**

Project Carlsbad 4x DDR IB Backplane Topology

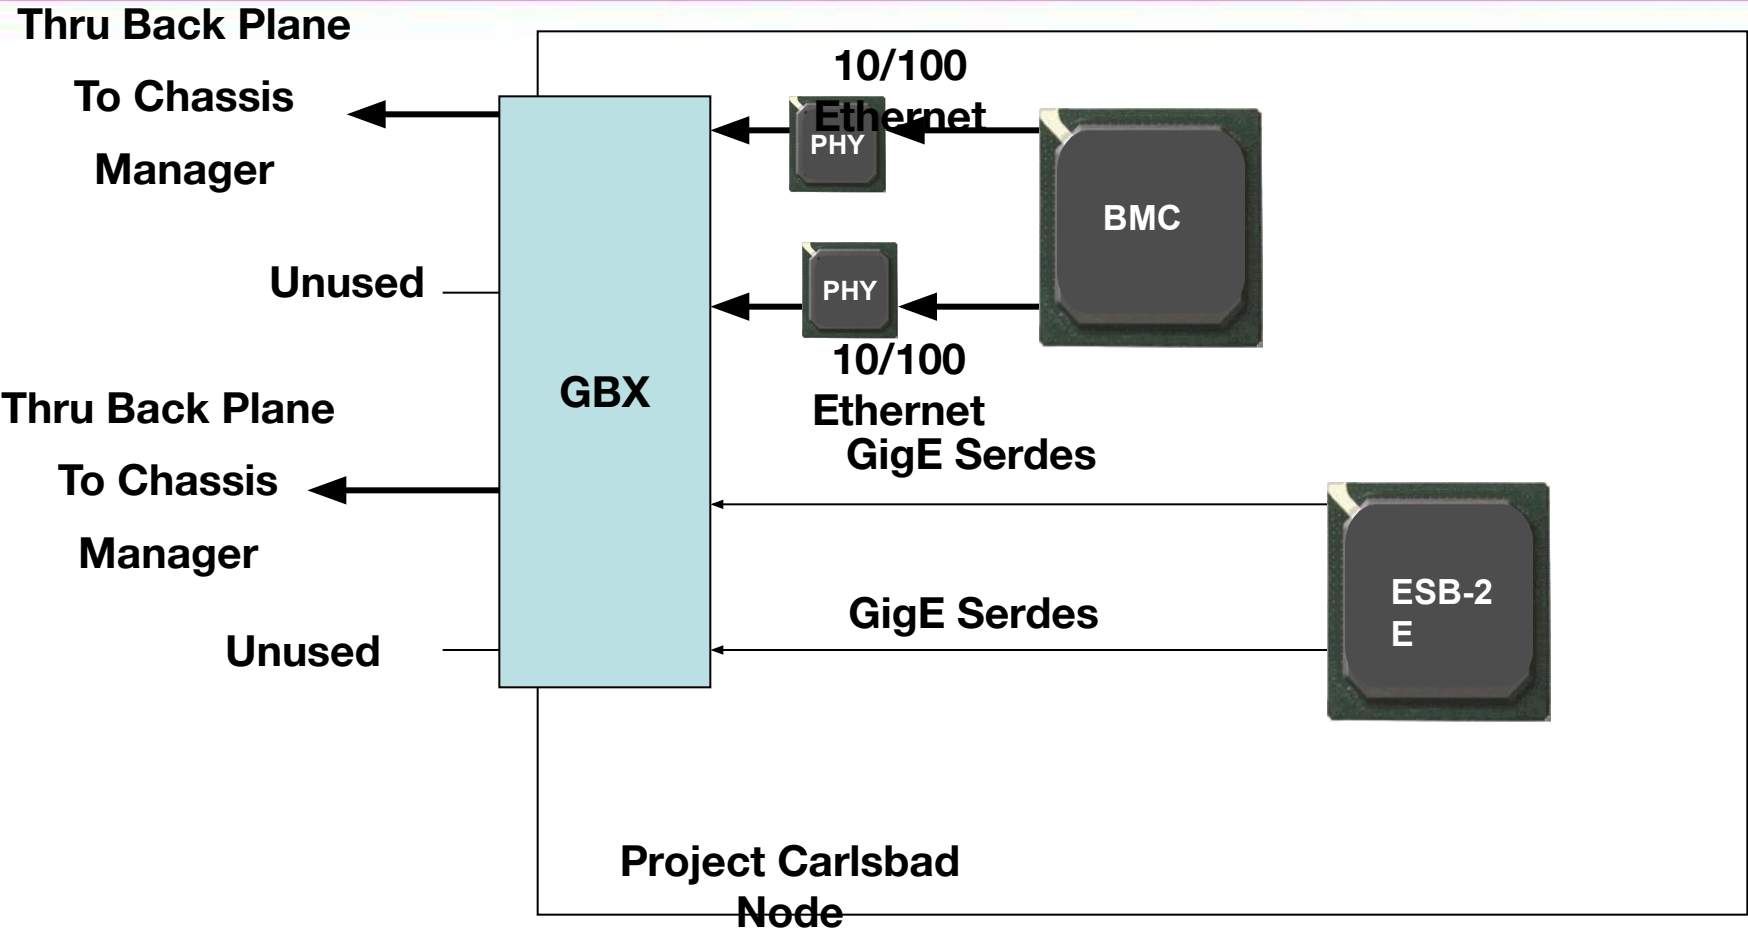


Blue Links = 4x DDR IB (Atoka to IRU Switch Cards)
 Gray Links = 4x DDR IB (IRU to external admin nodes)
 Black Links = 4x DDR IB (H-Dimension Torus)
 Red Links = 4x DDR IB (W-Dimension Torus)
 Green Links = 4x DDR IB (D-Dimension Torus)

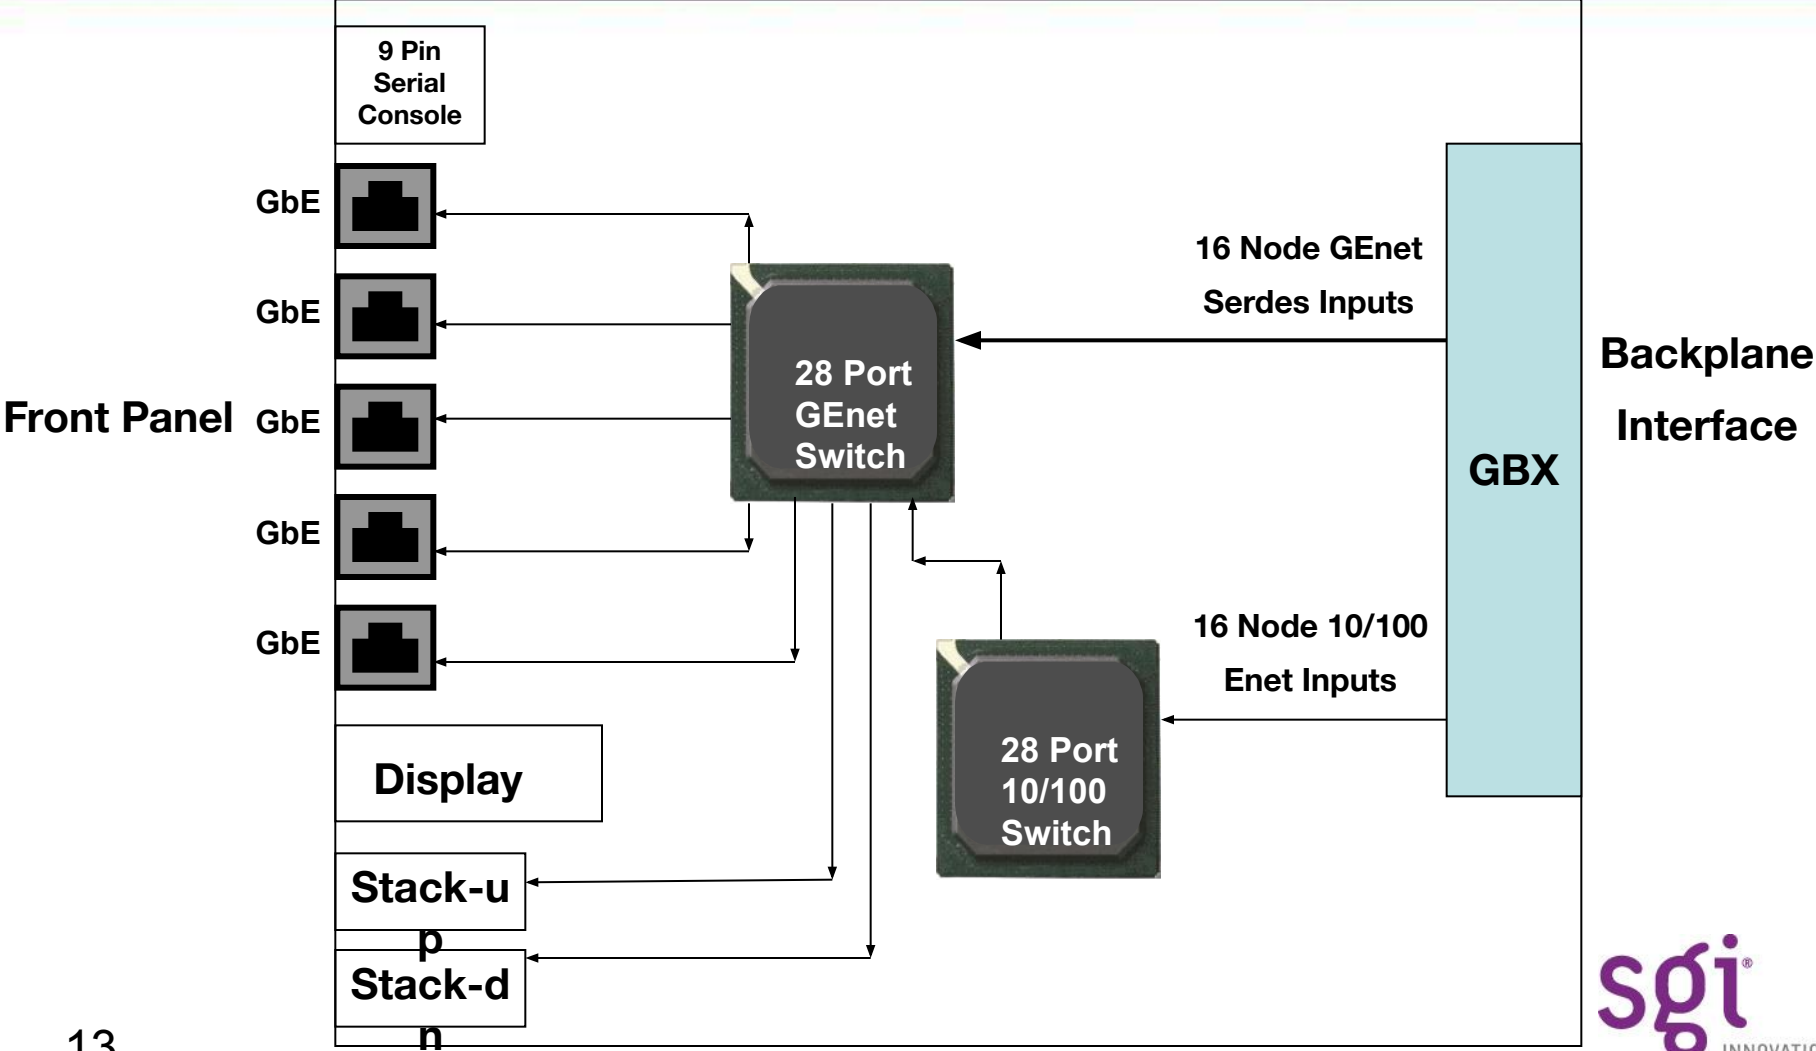
Project Carlsbad 4x DDR IB Backplane Topology



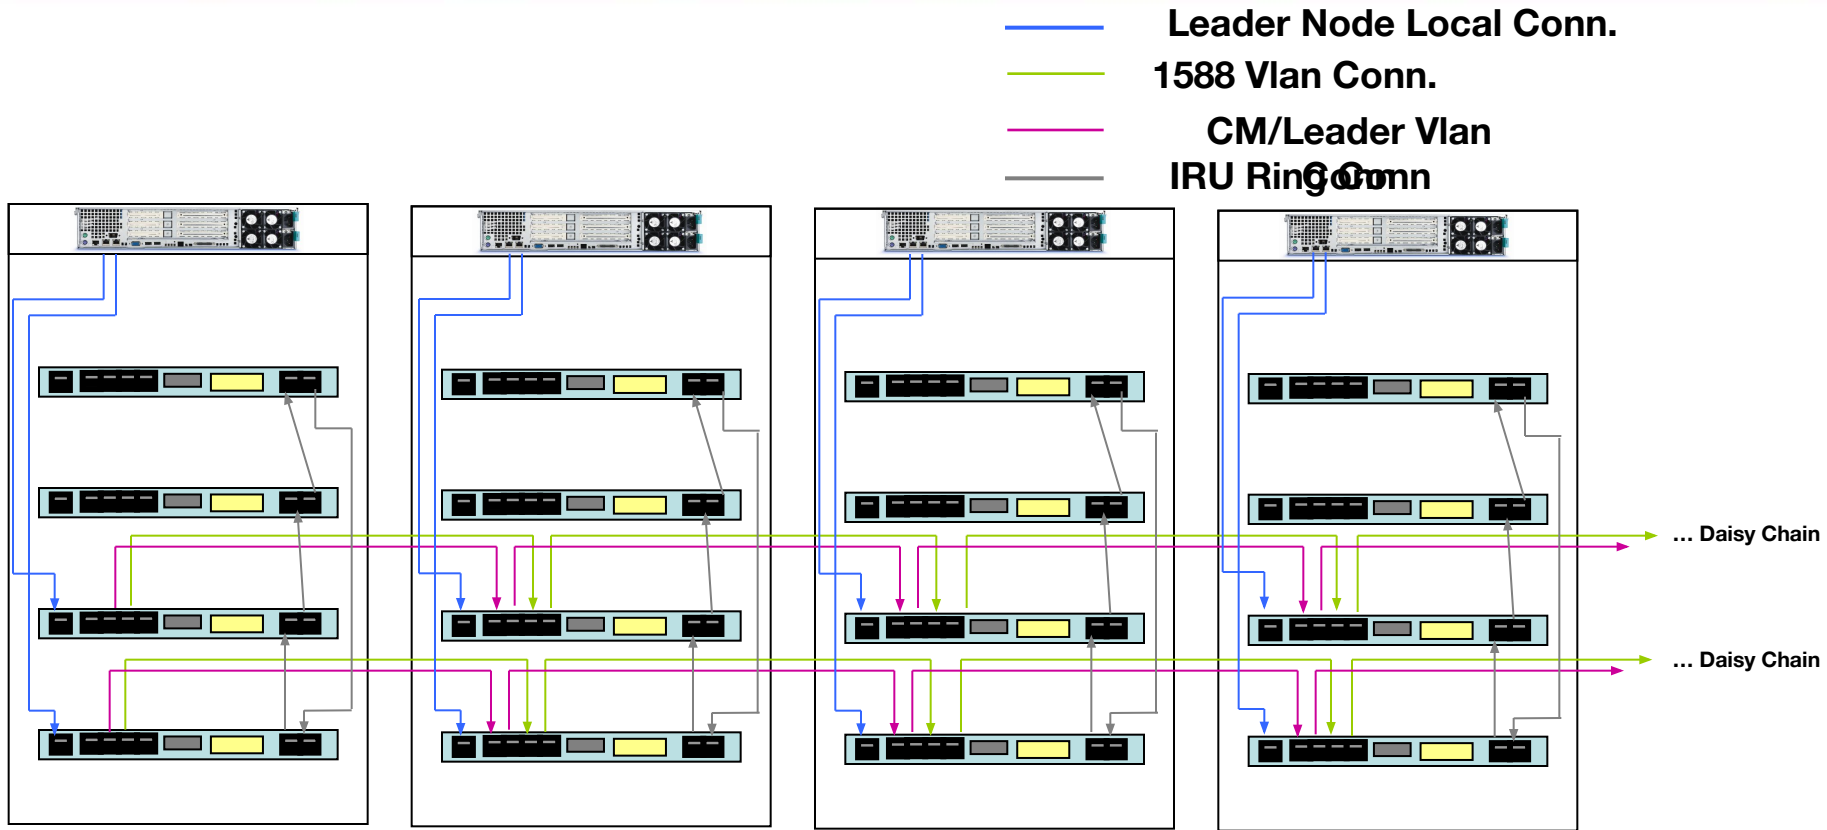
Project Carlsbad Blade




Chassis Manager




Rack Chassis Manager Cabling Topology



10U 16-Node Individual Rack Unit (IRU)




(4) IB Switch Blades



IB Backplane

SGI PROPRIETARY AND CONFIDENTIAL



**(7+1) 1625W 12VDC Output
Front-End Power Supplies**


sgi
INNOVATION
FOR RESULTS™

IRU Rear Blower Assembly

(7+1) 175mm Blowers
(Reused Altix 4700)




10U (17.50-in H x 12-in D)
24-inch EIA Form Factor



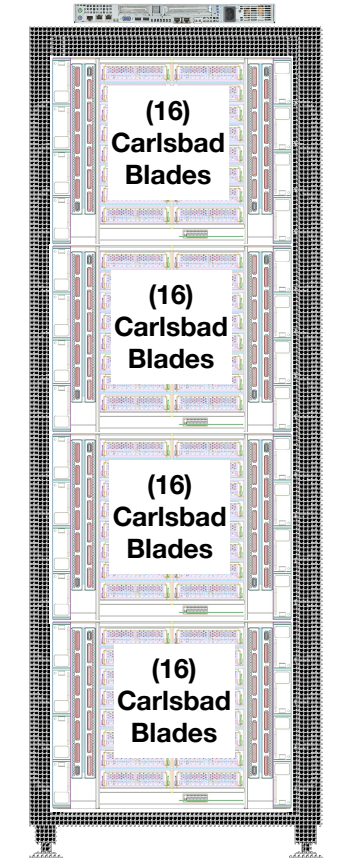
Rear View

Project Carlsbad IRU Assembly Exploded View




Single Project Carlsbad Rack


- Each 42U rack (30" W x 40" D) rack has:
 - (4) IRUs with (16) 2-Socket Carlsbad Nodes each
 - (128) DP Xeon sockets
 - DDR IB ports on (4) backplanes for torus
 - (48) 4x DDR IB
 - 2U Space at Top of Rack Contains
1U SGI Altix XE210 Leader Node (1 per Rack)
- SGI offers optional chilled water-cooled units for use in large system configurations
- 39.5kW (high-bin SKUs + (4) FB DIMMs /socket)
 - 31.6kW (assuming 80% system-level derate)
- Rack weight ~ 2050 Lb (246 Lb/ft² footprint)



19" Standard Rack Also Supported...



- Selected 2 IRUs based on ~20kW std. rack
- 3 IRUs => ~30kW class rack



Project Carlsbad Administrative Support Node

- SGI® Altix® XE240 (Default)
 - Used to provision and manage the cluster using Cluster Management SW
 - Network connections: GigE to leader nodes, communications to/from CMC & compute nodes administratively restricted.
 - Quantity: one per cluster



SGI® Altix XE® 240 Administrative Node

2U Server Board

- Dual Intel® Xeon® Processors (Woodcrest or Cl
- Intel 5000P chipset (Blackford)
- 8 fully buffered DIMMs
 - Quad Channel DDR-2
 - Memory Sparing, Mirroring
- Optional expansion modules (SAS or Dual GigE)
- Dual Gigabit Ethernet ports
- Integrated graphics (ATI ES1000 w/ 16MB)
- 5 slots, hot-swap drives (SAS/SATA) with HW RAID 0, 1, 5, 10
- Up to 3 PCI-X & 4 PCIe
- Optional redundant power



Leader Support Node

- Provisioned & functioned by the administrative support node
- Runs fabric management software
- Monitors, manages & pulls data from IRUs and compute nodes within the rack
- Consolidates and forwards upon request data from IRUs & compute nodes to the administrative node
- Provides shared read-only kernel/initrd (~40MB) & root fs (~1.6GB) images for rack's compute nodes
- Provides non-shared read-write system storage (~64MB /var, /etc) & minimal swap space (256MB) for rack's compute nodes
- Can be combined with fabric management support node
- Quantity: 1 per rack
- Network connections: GigE to other leader nodes & to first IRU within the rack, IB to whole cluster



Additional Nodes

Login service node

- User's login here to create/compile programs, etc.
- Quantity: 1 or more per cluster, commonly combined with batch and gateway service nodes

Batch service node

- Runs batch scheduler (PBS/LSF). User's login or connect here to submit jobs to the compute nodes.
- Quantity: 1 or more per cluster, commonly combined with login and gateway service nodes

Gateway service node

- Acts as a gateway from IB to various kinds of services such as storage (direct attached, fiber channel, etc.)
- Quantity: 1 or more per cluster, commonly combined with login and batch service nodes


Storage service node

- A NAS appliance bundle that provides shared, IB attached, filesystems for the cluster
 - Quantity: 1 or more per cluster
- A storage appliance that provides node private, IB-connected, scratch storage for the cluster
 - Quantity: 1 or more per cluster



Fabric management support node

- Provisioned & functioned by the admin node
- Runs fabric management software, monitors & manages the IB fabric
- Forwards upon request fabric status to the admin node
- Quantity: 1 or more per system, commonly combined with one or more leader nodes in the cluster®

42U Project Carlsbad 24-inch EIA Rack



42U Project Carlsbad 24-inch EIA Rack (Empty)



Concerns about Facility (Space, Weight, Power)



Emerson Network Power Presents Industry Survey Results That Project 96 Percent of Today's Data Centers Will Run Out of Capacity by 2011

11/16/2006 ...

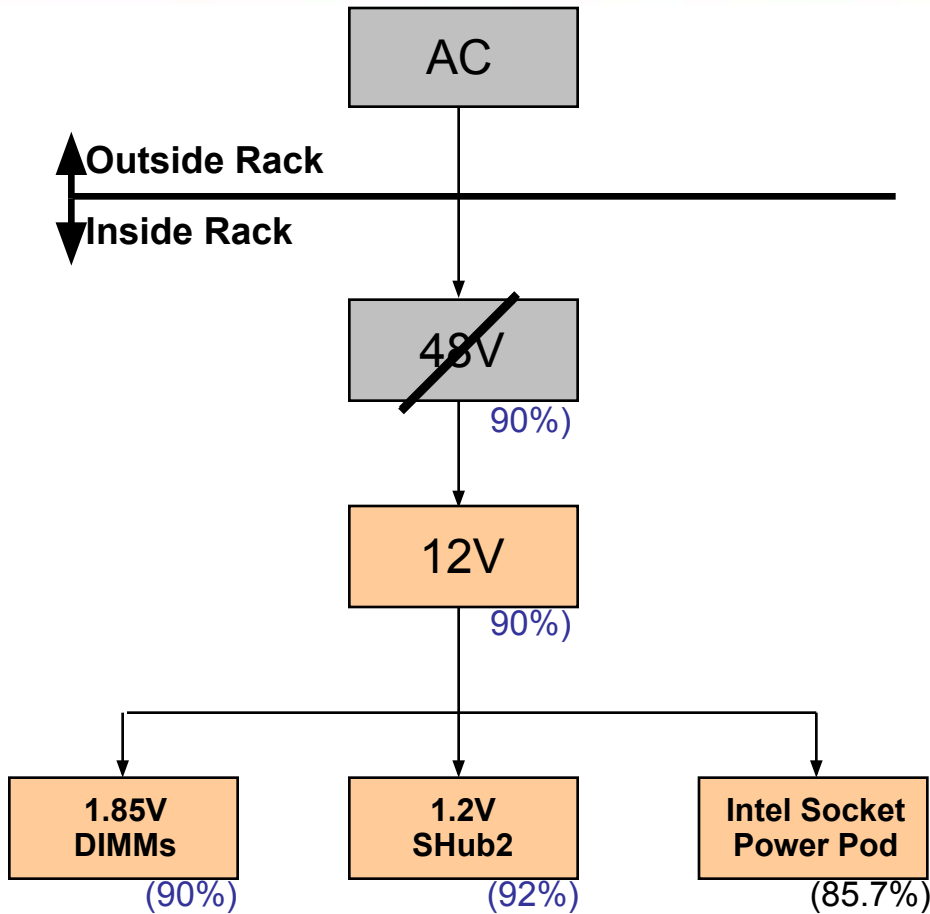
Data Center Users' Group survey confirms that increased densities, consolidation and energy efficiency are driving change in the industry

Why?

Computer Cycle: 2-5 yrs

Facility Cycle : 10-25 yrs

Facility (Power) : Energy Efficiency of Altix 4K & Carlsbad



AC DC conversion counted

(a) Remove 48VDC step

(b) Pick best supplier here



(c) Apply trade-secret

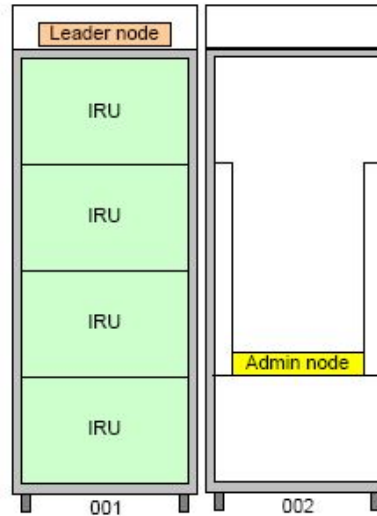
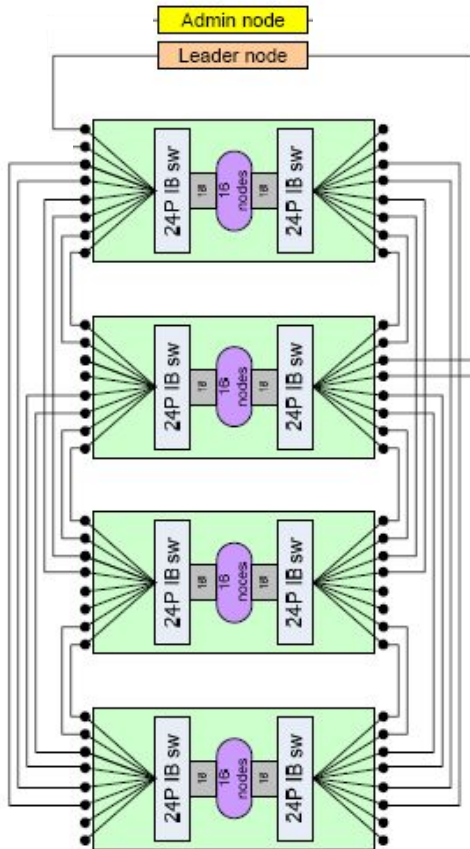
SGI Energy Efficiency

- SGI® Altix® 4700 server delivers a world-class power solution
 - High efficiency, high reliability, high density, remotely manageable
 - Standards-based
 - Over 90% efficiency on 12VDC front-end power supply
 - Up to 87% efficiency on compute blades
 - Up to 76% efficiency at rack-level
- Project Carlsbad design leverages the SGI Altix 4700 power architecture
 - 3rd generation water cooled solution
 - For systems above 15.0kW per rack, SGI strongly recommends the customer uses water cooled solution
- SGI remains committed to evolving high-efficiency power architectures for current and future products

Topology Overview


Example Project Carlsbad Configuration

Single rack topology (64 blades, 128 soc.)




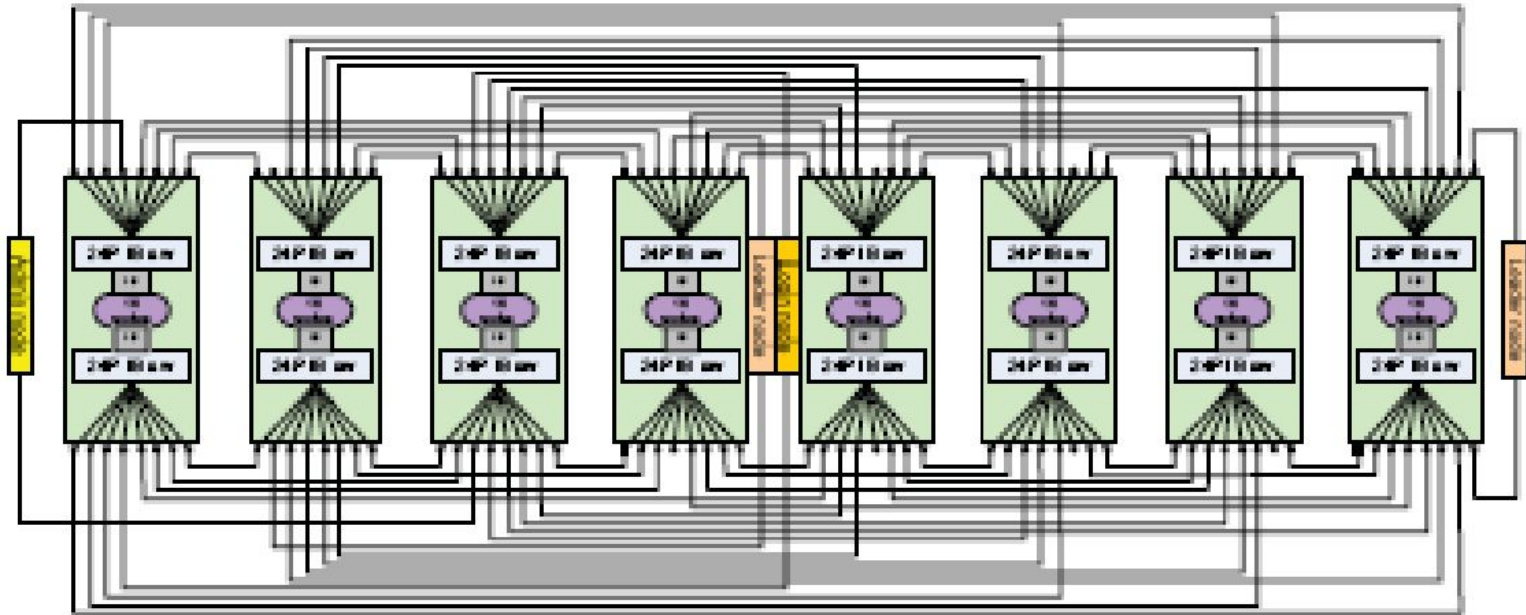
Maximum hop count = 3
 Bisection Bandwidth = $((16 * 4GB) / 64 = 1.0$
 GB/s/node)
 InfiniBand DDR switches

Logical drawing
 (Represents one side of the IRU)



Example Project Carlsbad Configuration

128 blade – Dual-plane Torus Topology (2 Racks)



Logical drawing
(Represents one side of the BRU)


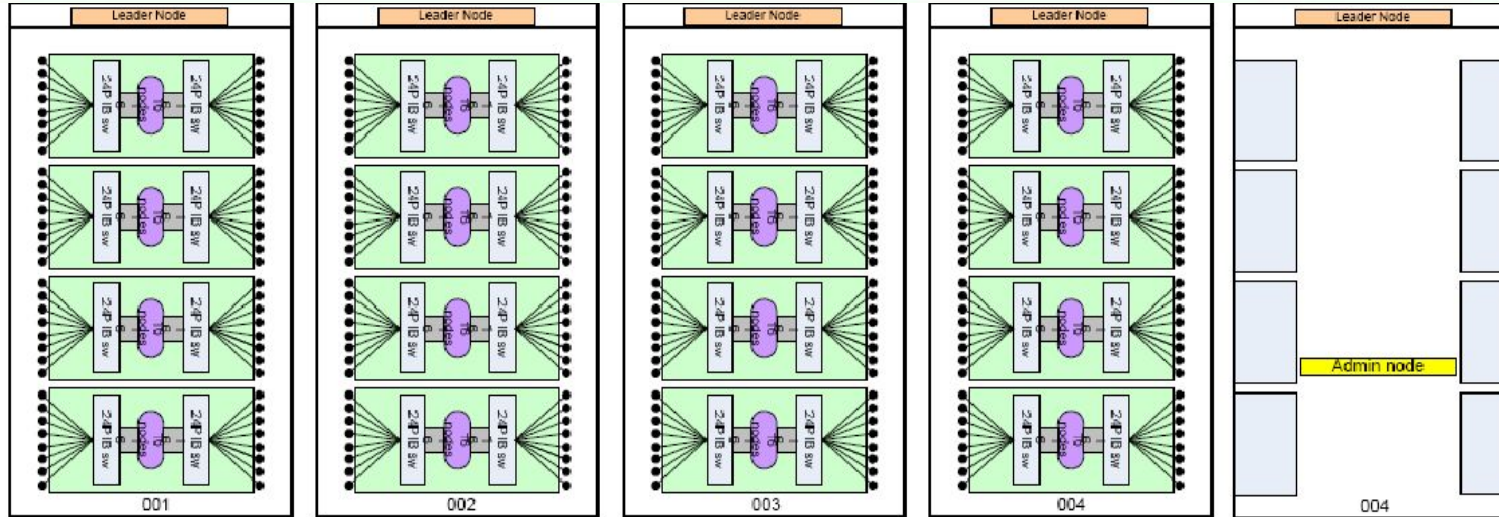
ATION
RESULTS

Maximum hop count = 3
 Block on bandwidth = $(32 * 400) / 128 = 1.0$
 GB/second)

InfiniBand DDR switches

Example Project Carlsbad Configuration

4-rack topology (256 blades)




Maximum hop count = 6
 Bisection Bandwidth = $((16 * 4GB) / 256 = 250 \text{ MB/s/node})$
 InfiniBand DDR switches

The logical cabling diagram shows the IRU interconnection for all of the cabling.


8192-Socket Project Carlsbad 4x DDR IB

4H x 8D x 8W Torus (1-rack Group x 8 x 8 = 64 Rack)

Only Torus Connections Shown
(Node Fan-In / Fan-Out
are Additional Hops)



1-Rack Group
(Contains H-Dimension)



Red Links (Interleaved down the ranks)

4H x 8D x 8W = 256 Switches / Plane
512 Total Switches
(256) 1.0m Cables H
(256) 2.0m Cables H
(128) 5.0m Cables D
(384) 8.0m Cables D
(64) 2.0m Cables W
(256) 3.0m Cables W
(192) 4.0m Cables W
1536 Total Cables

(256 IRUs)

Bisection:
256 links H-Dim
128 links W-Dim
128 links D-Dim

Green Links
(Interleaved across
the aisles)

● = 1-Rack Group

$$(128 \text{ links}) * (4 \text{ GB/s}) = 512 \text{ GB/s Bisection} / 8192 \text{ sockets} = 0.0625 \text{ GB/s/socket}$$

$$12\text{-hops} = [(2) * (160 \text{ nS node fan in/out} + 4.95 \text{ nS for 30-in pcb}) + (1+2+3+4+4+4+5+8+8+8 \text{ m network cables}) * (4.3 \text{ nS/m}) + (10\text{-hops torus network}) * (160 \text{ nS} + 3.3 \text{ nS for 20-in pcb})] = 2,165.0 \text{ nS 1-Way Longest Path Latency}$$

Project Carlsbad Topology Summary (4X DDR IB)

Carlsbad Torus 4x IB							
Sockets	24-Port 4x IB Switch ASICs	Cables	Hops	Longest Path Latency (nS)	MPI Latency (nS)	Bisection (GB/s /socket)	Total Bisection (GB/s)
64	4	12	3	498	3,438	0.75	48
128	8	24	3	506	3,446	0.50	64
256	16	48	5	846	3,786	0.25	64
512	32	96	6	1,026	3,966	0.125	64
1024	64	192	7	1,211	4,151	0.125	128
2048	128	384	8	1,408	4,348	0.125	256
4096	256	768	10	1,770	4,710	0.0625	256
8192	512	1536	12	2,165	5,105	0.0625	512
16384	1024	3072	14	2,530	5,470	0.0625	1024
32768	2048	6336	18	3,313	6,253	0.0430	1408

8192-Socket 4x DDR IB (4H x 8D x 8W Torus)

- 62.5 MB/s/socket Bisection
- 5,105 nS MPI Latency (2,165 nS 1-Way Longest Path Latency)

Software Overview

Complete, Factory Integrated Solution Stack

Linux[®] Operating System

Performance Optimization

SGI[®] ProPack[™] 5 software for Linux OS

Cluster Management

SGI developed and branded (based on OSCAR)

Workload Manager

Factory Integrated and Tested

MPI

Intel MPI Runtime

IB Fabric and Subnet Management

SGI InfiniBand Fabric Subnet Management (based on OFED and OpenSM)

Development Tools

Intel[®] C++ and Fortran compilers, VTune, Math Kernel Library

- Complete cluster solution stack
- Cost-effective, standards-based
- Optimized for ease of use

SGI and Linux®

Open Standards Industry Leadership

- SGI Linux leadership:
 - Unmatched in the industry, major contributor to Linux standard
 - Expertise to resolve kernel-level issues quickly, efficiently
- 100% Linux - scalable, robust, standards-based
 - Industry standard SUSE® Linux® Enterprise Server 10
 - Red Hat® Enterprise Linux® 5 (avail. Q4 CY07)
 - SGI® Propack™ Toolkit combines essential tools for workflow optimization
 - SGI® InfiniteStorage delivers complete data lifecycle management solution.
- Superior reliability, availability, serviceability:
 - Comprehensive RAS roadmap, ease of service with blades



SGI® ProPack™ Benefits for Project Carlsbad

- Dramatically enhanced performance:
 - FFIO: Accelerated I/O bandwidth
 - CPUSETS, NUMATOOLS: Fine tuning for processors, memory
- Simplified system administration:
 - Performance Co-Pilot™, ESP, Cluster Manager
 - Storage administration tools

SGI kernel-level Linux® expertise unmatched in the industry – to resolve customer issues in-house, fast, effectively.

Booting and Configuring OS

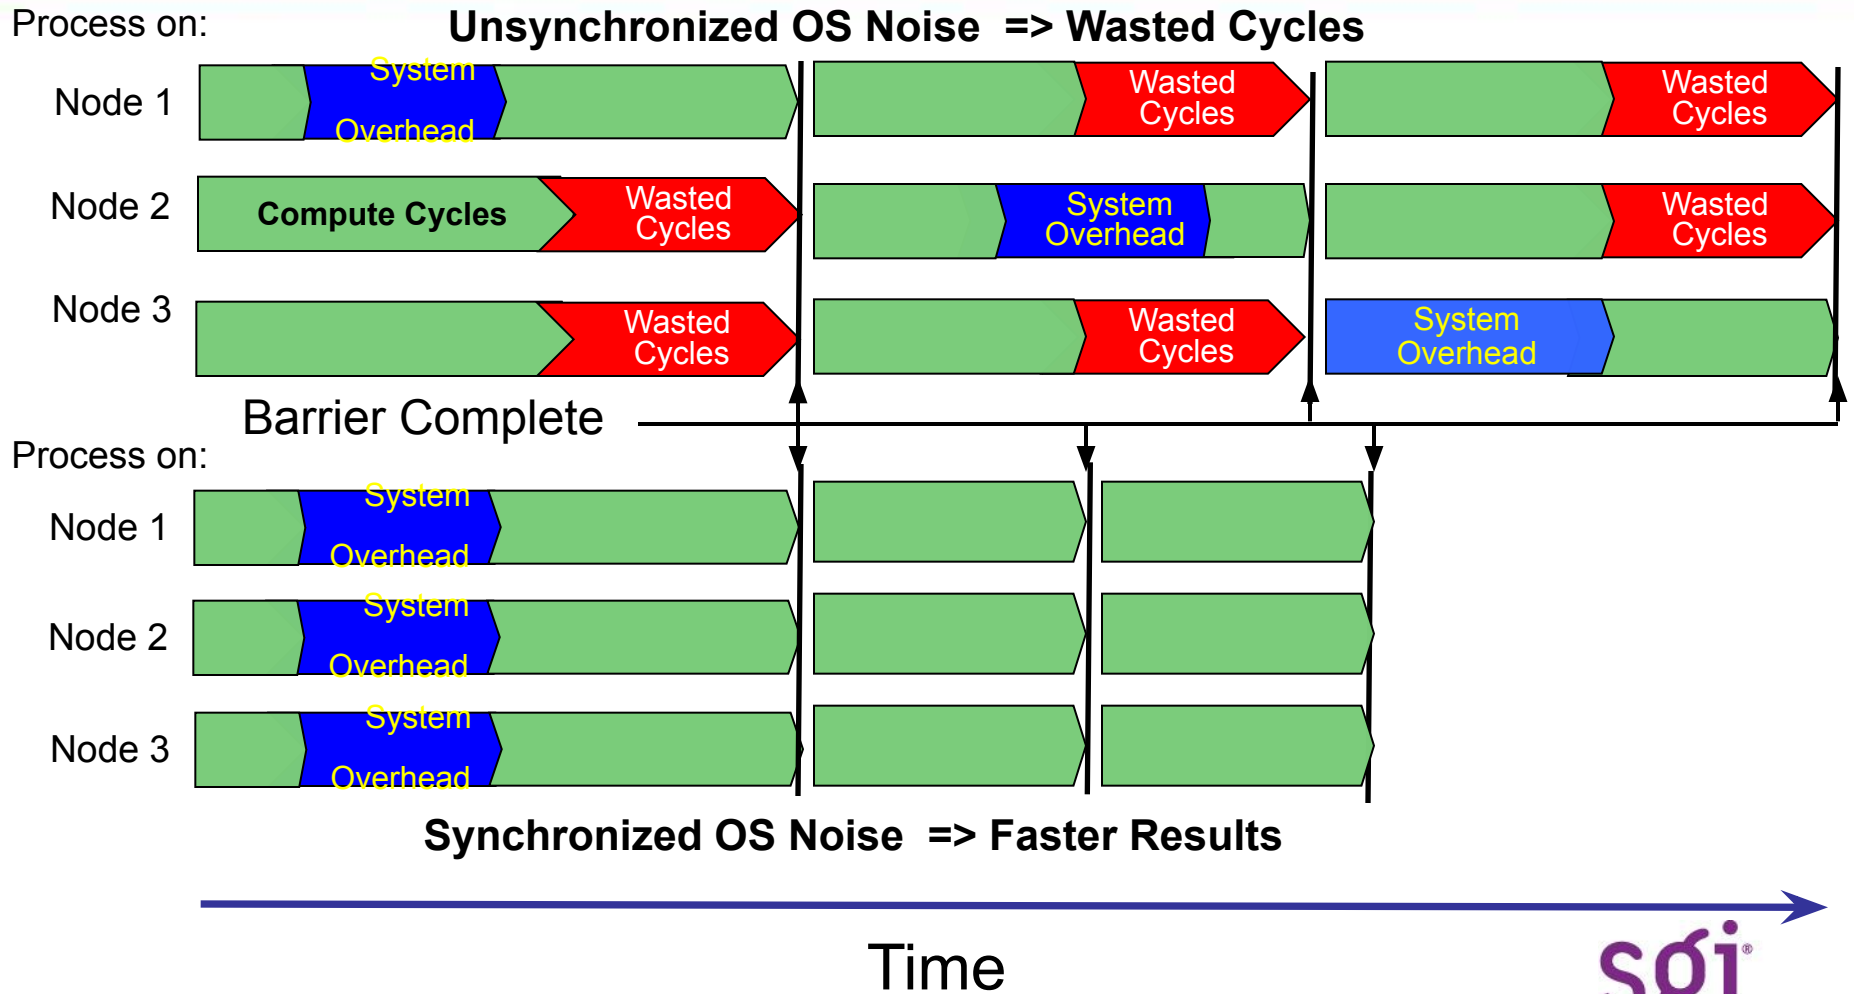
- Boot Services provided by rack's leader node
 - (1) Leader Node services all (64) diskless AtokaP nodes in a given rack
 - Linux® OS images received over administrative GigE network
 - Enables scalability (leader nodes are single point of control in each rack)
 - No BIOS modifications necessary
- File System Services
 - Root images mounted via InfiniBand (using NFS from rack's leader node)
 - Root images can be shared by all blades in a rack
 - Use InfiniBand native storage (otherwise NFS)



Booting and Configuring OS

- Use a standard Linux® OS distribution
 - Use a standard kernel and remove all unnecessary RPMs
 - Preserve 3rd party application certification
 - OS and boot support will be based on industry standards to assure compliance with standard data center operations
- Synchronization of OS overhead (OS jitter, OS noise)
 - SGI value added hardware and software will reduce OS overhead effects
 - Enables greater performance on parallel workloads
 - Detailed slide in Back-Up
- SGI® ProPack™ for Linux® OS
 - Combines essential tools for workflow optimization

Carlsbad SW: OS Noise (overhead) Synchronization Significant Speedups for Parallel Workloads



System Management



- Node-level
 - Baseboard Management Controller (BMC) and onboard NICs
 - Utilize industry standard IPMI 2.0 compliant protocols
- Chassis management controller (CMC) in IRU
 - SGI developed CMC
 - Hierarchical design for scalability enabling larger systems
 - Provides dedicated GigE network for all management functions, remote console access, and cluster management
 - Provides dedicated GigE network for synchronization of OS overhead
- System management and monitoring
 - Performed via a common cluster management software tool

Cluster Management

- SGI Developed Solution
 - Based on Open Source Cluster Application Resources (OSCAR) from OpenClusterGroup.org
 - Provides centralized SW and system provisioning, monitoring and cluster-specific management
 - Hierarchical design for scalability enabling larger systems
 - Cluster management features supported include:
 - Software installation (admin, leader, compute, and non-storage service nodes)
 - Software configuration and customization (admin, leader, and compute nodes)
 - Establish, expand and contract the Project Carlsbad cluster
 - Power control
 - Booting/shutdown
 - Console management
 - Monitoring, logging, alarms
 - Project Carlsbad Interconnect Verification Tool (diagnostic tool)
 - Scalable cluster-wide commands (C3)

InfiniBand Fabric Configuration & Management

- SGI developed solution
 - Based on OpenFabrics Enterprise Distribution (OFED) from Open Fabrics Alliance (OpenFabrics.org)
 - Subnet management (SM) based on OpenSM
 - Runs on a leader node
 - Features supported:
 - Automatic fabric configuration
 - Administrative fabric re-configuration (zoning-partitioning)
 - Management of virtual lanes (MPI traffic, Storage traffic)
 - Monitoring, diagnostic testing, SM software updating
 - Redundant SM with fail-over



**Fabric
Manager**

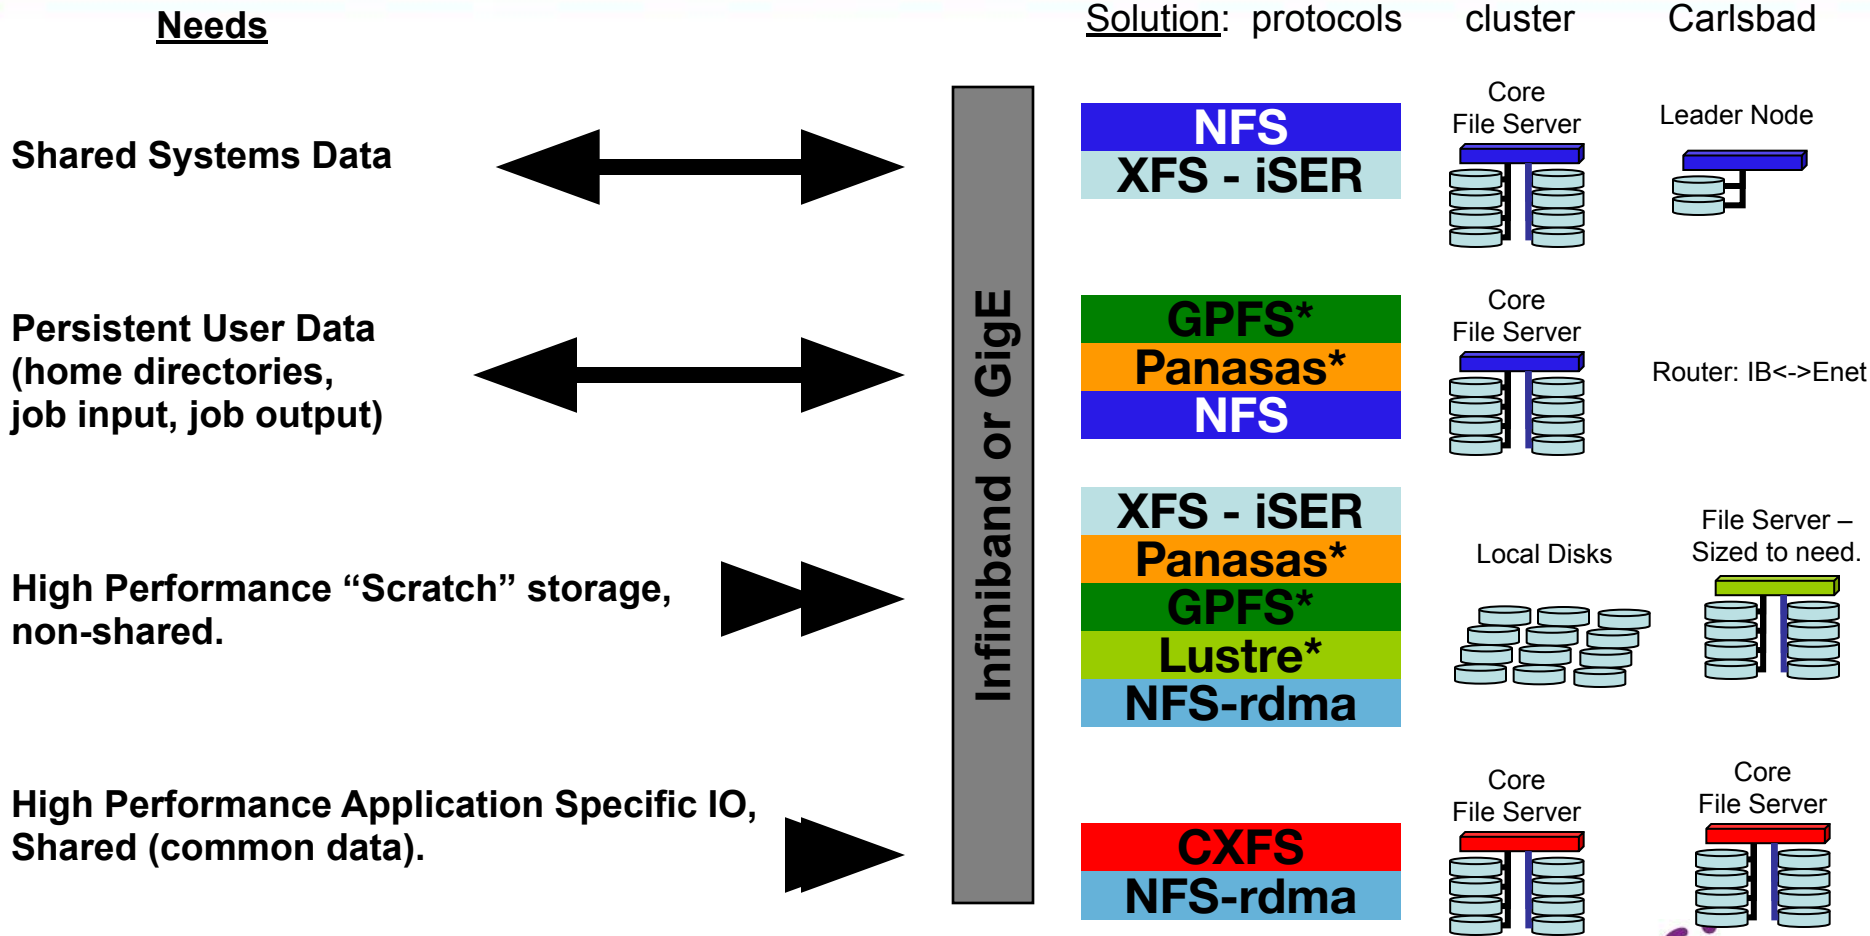
Storage Integration

Storage – Typical needs

Key Types of IO (each with different IO usage patterns)

- Shared Systems Data “installed application” (mostly read-only, low reliability, low performance)
- Persistent User Data “home directories” (read-write, high reliability, low performance)
- High Performance “Scratch” storage, non-shared (read-write, low reliability, high performance)
- High Performance specialized application IO, “shared common data” (read-write, high reliability, high performance)

Storage – provisioning methods






Infiniband or GigE

Storage – Carlsbad options




- Shared Systems Data
 - Place on Leader Nodes – one to serve each rack
- Persistent User Data
 - Use IB to Enet router, for Carlsbad access to facility data
 - Can use multiple routers for bandwidth
- Scratch Storage
 - Configure a fileserver to need, can be 0.
 - disk/node eliminated: saving power, weight, cost.
- High Performance Shared/Common Data
 - Configure specialized fileserving as needed.

Roadmap

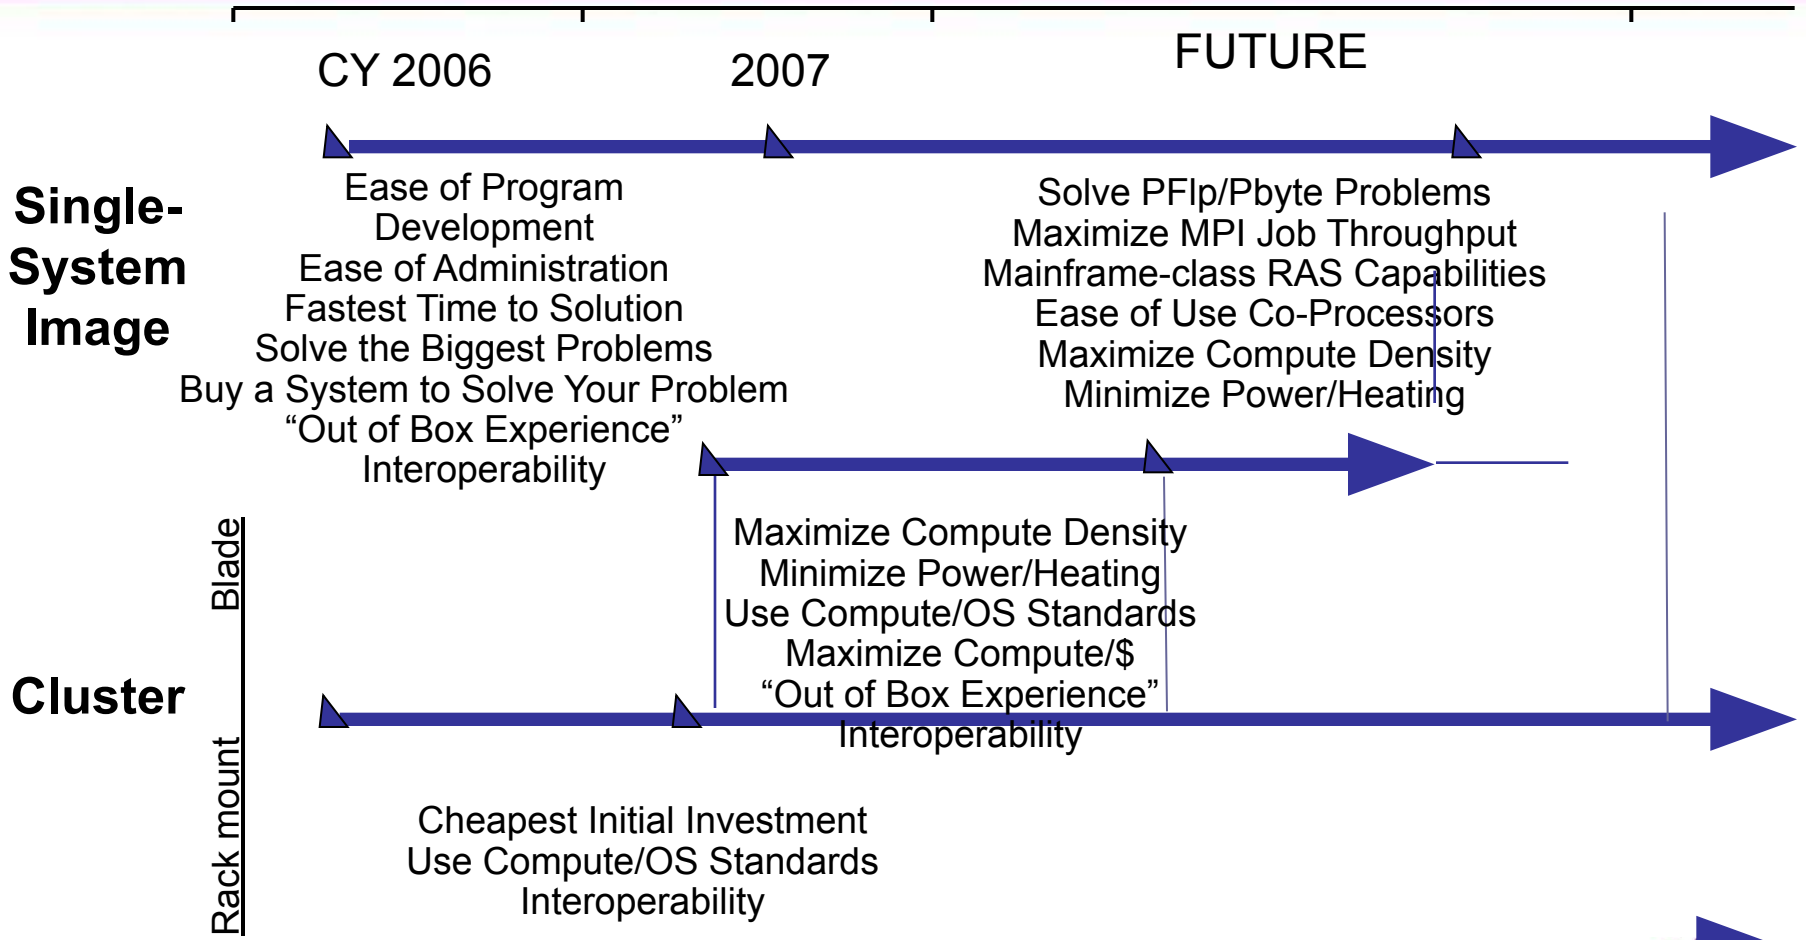
Intel® Server & Workstation Platform Roadmap

<u>Platform Segment</u>		<u>2H'06</u>	<u>1H'07</u>
RISC/Mainframe Replacement Enterprise DP/MP 		Dual-Core Intel® Itanium® 2 Processor 9000 sequence; 24MB L3 Intel® Itanium® 2 Processor 9M / 1.66 GHz / 667 FSB	Dual-Core Intel® Itanium® 2 Processor 9000 sequence; 24MB L3
		Intel® E8870 Chipset (400 FSB) / Enabled Chipsets	Intel® E8870 Chipset (400 FSB) / Enabled Chipsets
RISC Replacement HPC DP/MP 		Dual-Core Intel® Itanium® 2 Processor 9000 sequence Intel® Itanium® 2 Processor (DP only) 3M / 1.60 GHz 400/533 FSB	Dual-Core Intel® Itanium® 2 Processor 9000 sequence
		Intel® E8870 Chipset (400 FSB) / Enabled Chipsets	Intel® E8870 Chipset (400 FSB) / Enabled Chipsets
Enterprise MP 		Dual-Core Intel® Xeon® Processor 7100 Series 667/800 FSB; 16MB L3 Dual-Core Intel® Xeon® Processor 7000 Series 667/800 FSB	Dual-Core Intel® Xeon® Processor 7100 Series 667/800 FSB; 16MB L3
		Intel® E8501 Chipset / Enabled Chipsets	Intel® E8501 Chipset / Enabled Chipsets

Intel® Server & Workstation Platform Roadmap

<u>Platform Segment</u>		<u>2H'06</u>	<u>1H'07</u>
Performance & Volume DP		Quad-Core Intel® Xeon® Processor 5300 Series Dual-Core Intel® Xeon® Processor 5100 Series Dual-Core Intel® Xeon® Processor 5000 Series	Quad-Core Intel® Xeon® Processor 5300 Series Dual-Core Intel® Xeon® Processor 5100 Series Dual-Core Intel® Xeon® Processor 5000 Series
		Intel® 5000P Chipset	Intel® 5000P Chipset
Value DP		Dual-Core Xeon® Processor 5100 Sequence Dual-Core Xeon® Processor 5000 Sequence	Dual-Core Xeon® Processor 5100 Sequence Dual-Core Xeon® Processor 5000 Sequence
		Intel® 5000V Chipset	Intel® 5000V Chipset
Entry UP		Dual-Core Intel® Xeon® Processor 3000 Series Intel® Pentium® D Processor 900 Sequence	Quad-Core Intel® Xeon® Processor 3200 Series Dual-Core Intel® Xeon® Processor 3000 Series Intel® Pentium® D Processor 900 Sequence
		Intel® 3100 Series Chipsets	Intel® 3100 Series Chipsets

Customer Value Roadmap

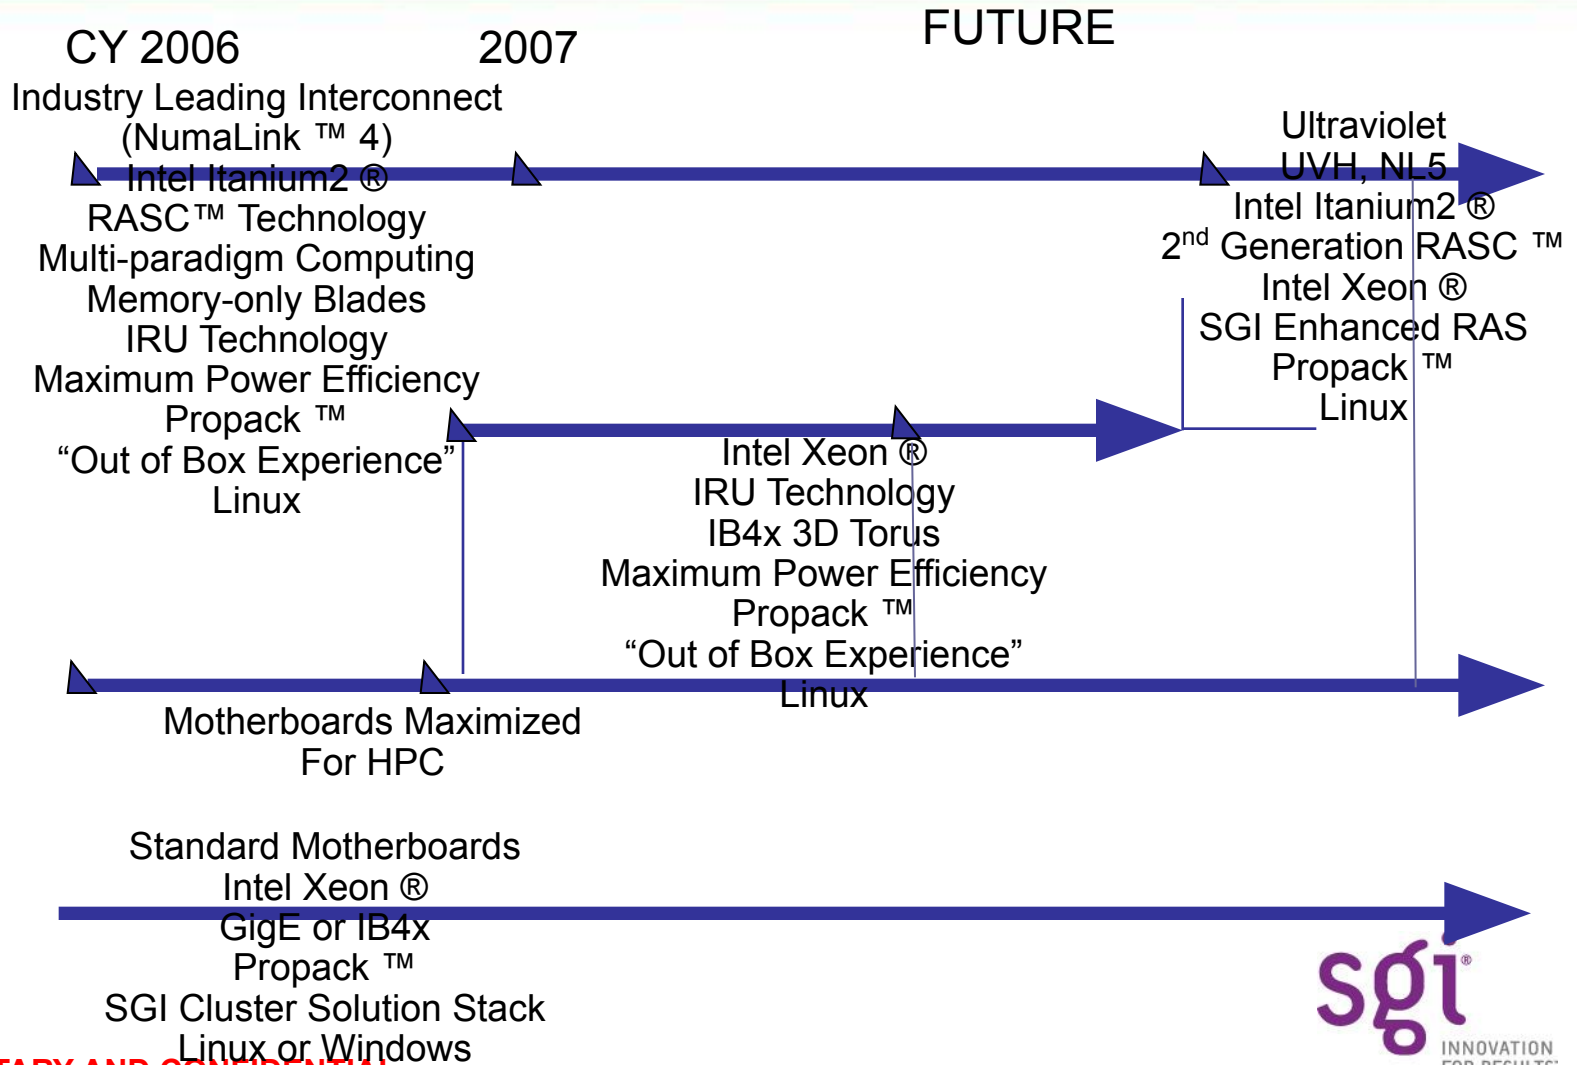


SGI Technologies Roadmap

Big Nodes
to 512S nodes
TB-PB GAM

Cluster
2-4S nodes
10sGB GAM

Blade
Rack mount




Server Roadmap

Big Nodes
to 512S nodes
TB □ PB GAM



Cluster
2-4S nodes
10sGB GAM

Blade
Rack mount



Altix XE 210/240 Woodcrest Santa Fe/Dixon Clovertown Taos – 4S Tigerton

Ongoing advances in: Linux function, RAS, Density, Power/cooling, Easy Deployment



*DPR = Density, Power, Reliability

Summary

Project Carlsbad and Altix® XE1300

Trounce the Competition!

Product	Project Carlsbad	Altix XE1300	Dell PowerEdge 1955	IBM BladeCenter HS21	HP ProLiant BL20p
Form Factor	Blade	Rack	Blade	Blade	Blade
Net \$/Core	\$1,607	\$1,622	\$2,137	\$2,598	\$2,993
TFLOPS/ Rack	3.000	2.766	1.875	2.625	1.875

Product	Project Carlsbad	Altix XE1300	Dell PowerEdge 1950	HP ProLiant DL360	IBM x3550	Sun Fire X4100 M2
Form Factor	Blade	Rack	Rack	Rack	Rack	Rack
Net \$/Core	\$1,607	\$1,622	\$1,894	\$2,040	\$1,965	\$2,658
TFLOPS/ Rack	3.000	2.766	1.594	1.594	1.594	0.744

Product Comparison

	Altix XE210	Altix XE310	Project Carlsbad
Benefits	Excellent Price/Performance Excellent Node Density Head Node Capability Fast drives capability	Superior Price/Performance Superior Node Density Superior Energy Efficiency Superior Bandwidth Alternatives (Native DDR IB option)	Superior Price/Performance Superior Rack Density Superior Energy Efficiency (OS noise synch) Superior Reliability Superior Ease of use Superior Bandwidth (Native DDR EB standard)
Reliability	Industry Standard	Industry Standard	N+1 Hot Swap Power Supplies N+1 Hot Swap Fans Hot Swap Colmpute Blades Redundant Interconnect Topology Cable-less Architecture Disk-less Architecture
Chassis	1U – Single Node	1U – Dual Nodes	10U – 16 Nodes (Blades)
Processor	Dual or Quad-core Intel 2 Sockets	Dual or Quad-core Intel 4 Sockets (2 per node)	Dual or Quad-core Intel 32 Sockets / 10U (2 per Blade)
Hot-swap hard drives	2 x 3.5" SATA or SAS	4 x 3.5" (2 per node) SATA	Diskless Compute Nodes 5 x 3.5" SATA or SAS / Head node
Interconnect	Native Dual GigE Optional PCIe InfiniBand	Native Dual GigE per node Optional Native DDR IB per node	Native multiple GigE per node Native Dual DDR IB per node
Memory	8 DIMM per node 32GB max per node (DDR2 FBD 533/667)	16 DIMM (8 DIMM per node) 32GB max per node (DDR2 FBD 533/667)	8 DIMM per node 32GB Max per node (DDR2 FBD 533/667)
Power	600W (for one 3GHz node)	900W (for two 3GHzz nodes)	N+1 Hot Swap Power Supplies 90% Node Power Efficiency

Project Carlsbad Customer Value

	Project Carlsbad Value
Advanced high-performance platform	<ul style="list-style-type: none">▪ Infiniband 4x DDR in a 3D torus topology▪ Internal GigE backbone for maintenance decreases load on IB for compute▪ Performance features of the quad-core Intel® Xeon® Processor architecture▪ Option to use dual-core CPUs for maximum memory efficiency▪ SGI ProPack™ software including FFIO
Packaging, density saves floor space, reduces power, adds reliability	<ul style="list-style-type: none">▪ Top density, 512 cores per rack▪ Highly efficient air movers, water chilled rack doors and redundant PS, lower thermal dissipation enables greater reliability▪ Cable-free IRU reduces number of failure points▪ Elimination of node-based disk drives removes a frequent point of failure▪ ~ 2X the density of the XE series
Easy to build & deploy	<ul style="list-style-type: none">▪ No external cabling, IRU completely integrated▪ SGI software stack to manage at the IRU and node level▪ Customizable, factory-integrated▪ “Out of Box Experience” in the tradition of the Origin series▪ SGI ProPack tools, for simplified development and administration▪ SGI PS assistance with implementation – deep technology and industry expertise▪ SGI support for hardware and software system components


sggi[®]

Project Carlsbad Water-Cooled Coils

(4) Individual Coils

Target Heat Rejection
95% water / 05% air

Chilled-Water Supply
45°F to 60°F (7.2°C to 15.6°C)
15 psi (103.4kPa) Max.
14.4 gpm (3.3 m³/hr) Max.



Condensate Drain Pan


Branch Feed to
Individual Coil



3/4" (1.91 cm) Coupling

Swivel Coupling to
Supply Hose

Project Carlsbad Water-Cooled Coils




Environmental Operating Windows

Present SGI Operating Window:
 41°F to 95°F (5°C to 35°C)
 10% Rh to 90% Rh
 (non-condensing)

ASHRAE Class 1 Allowable Operating Window*:
 59°F to 90°F (15°C to 32°C)
 20% Rh to 80% Rh
 (62.5°F (17°C) dew point Max)

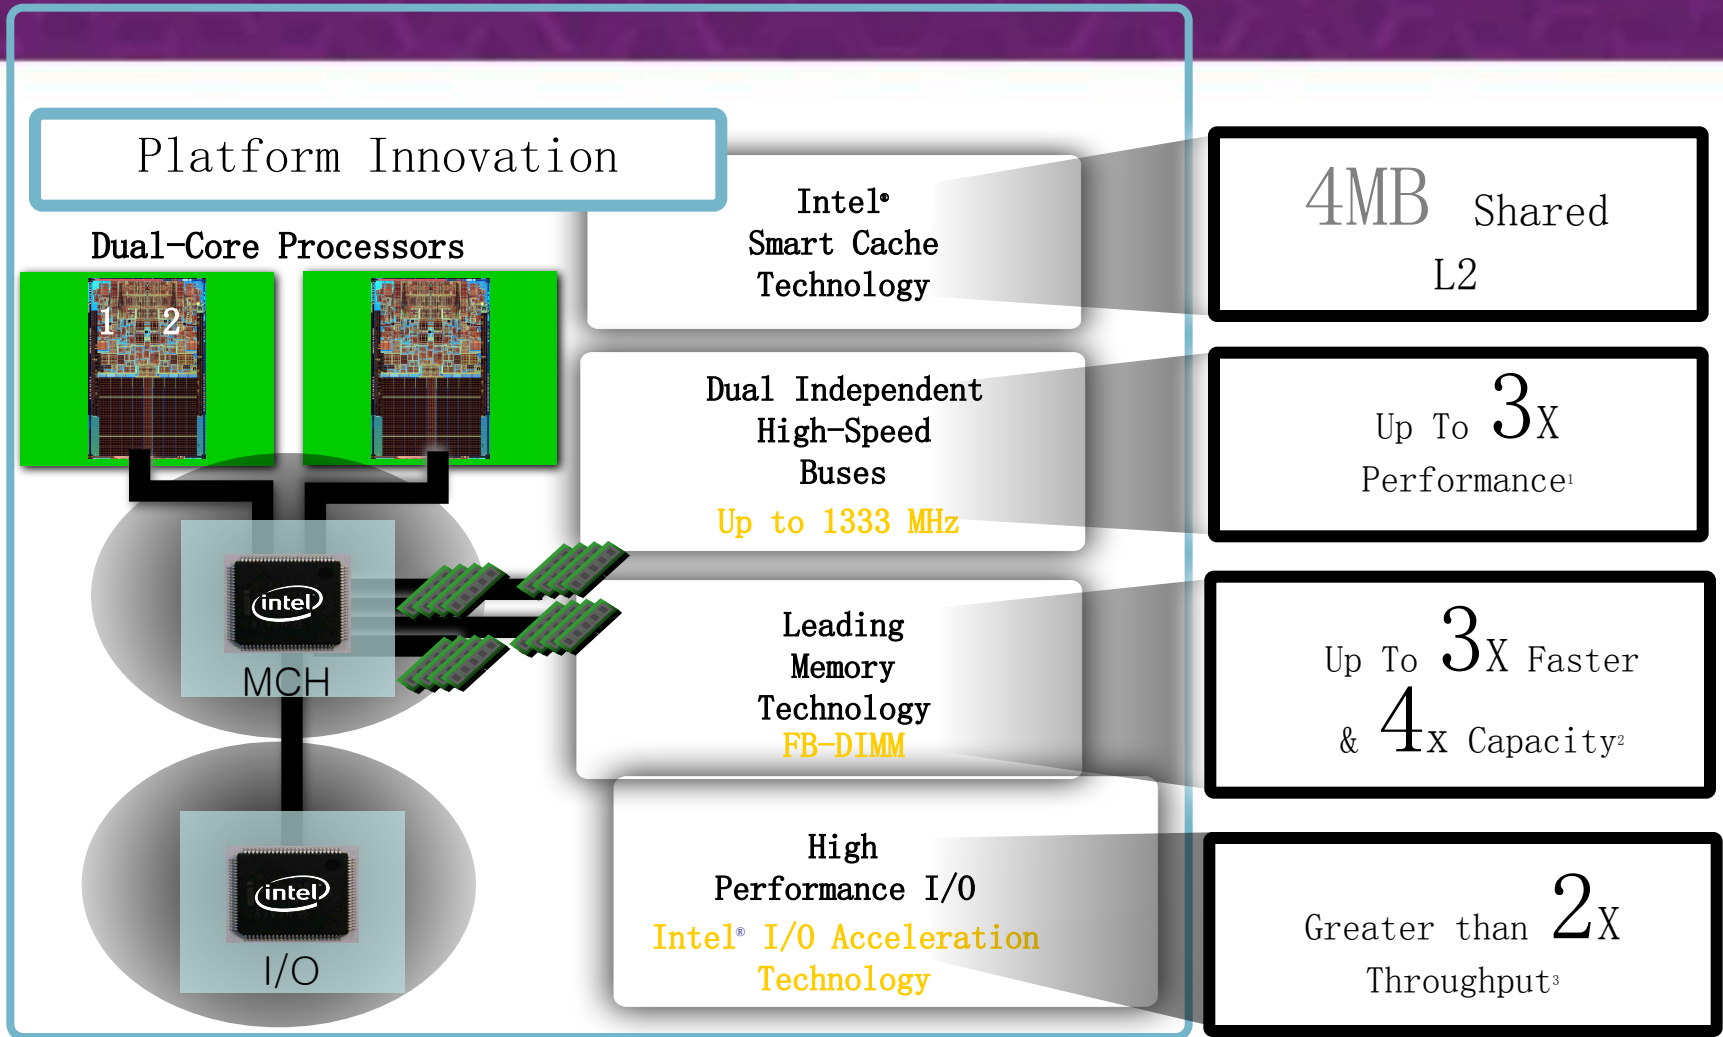
SGI Recommended Operating Window for Water-Cooled Coil:
 68°F to 77°F (20°C to 25°C)
 40% Rh to 50% Rh
 (non-condensing)
 Matches ASHRAE Class 1 Recommended*



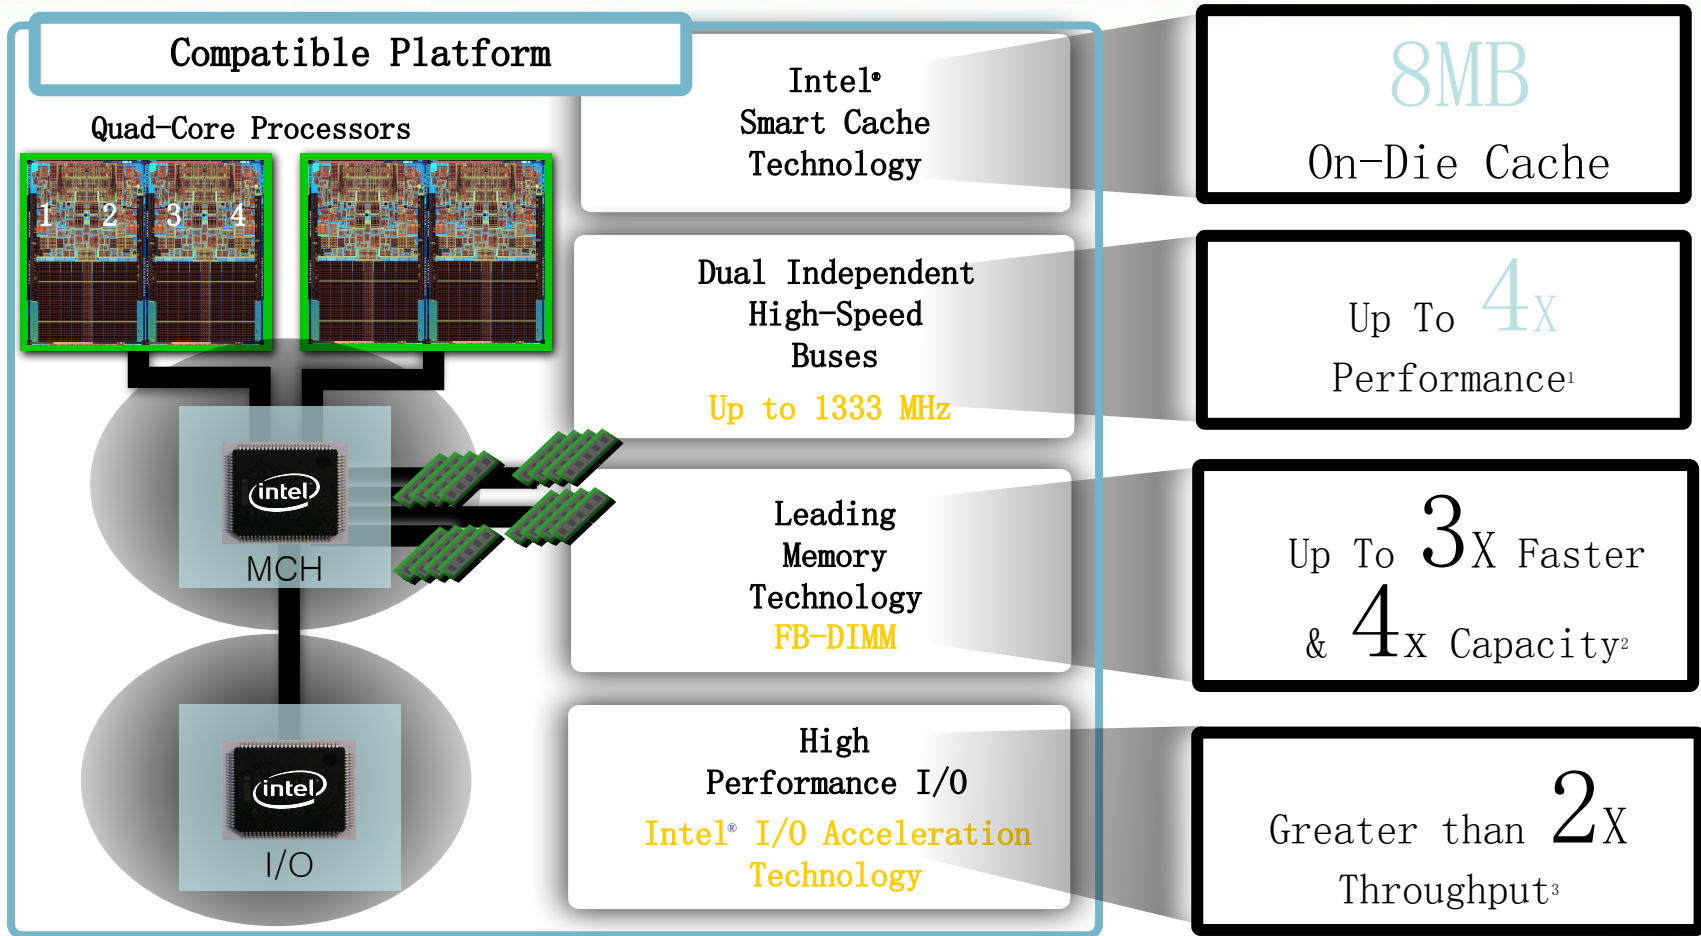
Data Center Energy Use

- Representative breakdown*
 - 59% Computer Loads (33% to 73%)
 - 25% HVAC Pumps & Chiller
 - 10% HVAC Air-Movement
 - 05% UPS Losses
 - 01% Lighting
- 1 kW datacom load ~ 1.7 kW load at facility mains transformer*
 - 1.4 kW to 3.0 kW range

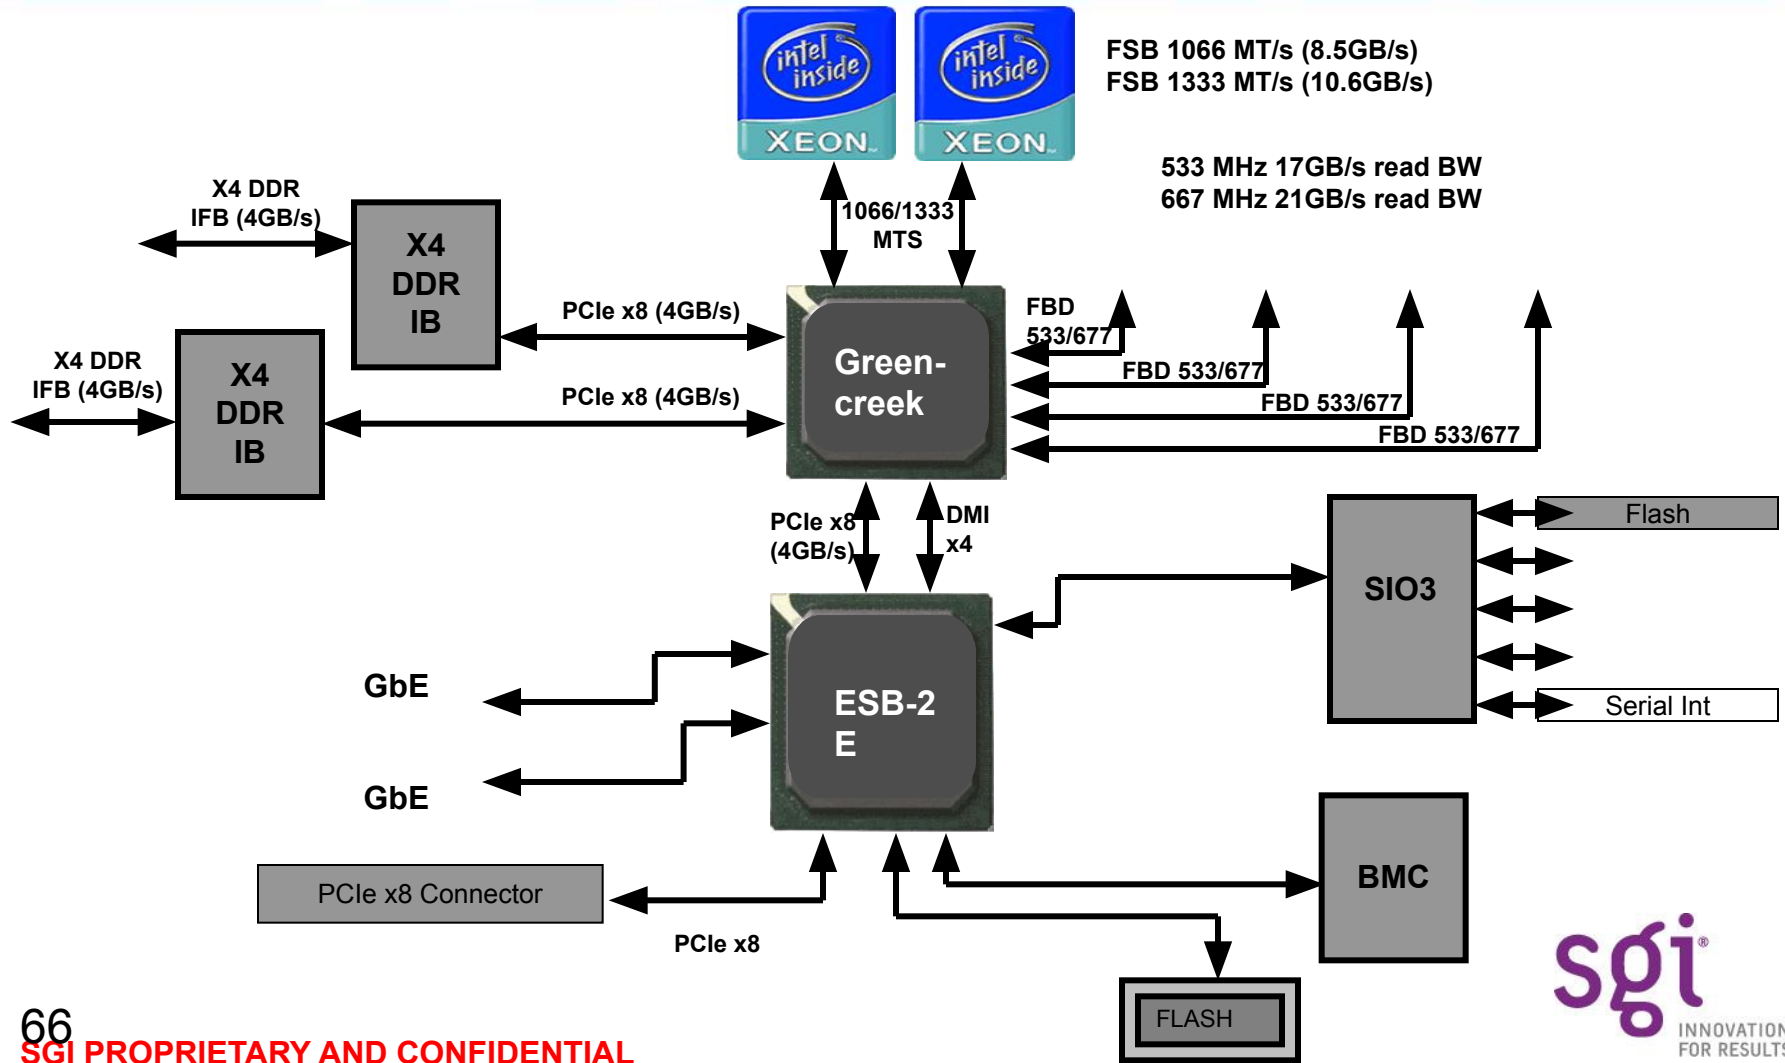
Intel® Xeon® 5100 Series Platform




Intel® Xeon® 5300 Series Platform



Project Carlsbad Blade



Chassis Manager Front Panel



SGI[®] ProPack[™]

Linkless FFIIO	Set as environment variable to accelerate I/O calls. Drives dramatic performance enhancement in I/O intensive cluster configurations.
Intel Runtime Libraries	Current version packaged with ProPack for customer convenience.
CPUSETS	Used directly by cluster workload manager, provides ability to allocate specific CPU for system daemons, etc for improved performance, decreased CPU contention
ESP	Tool used by administrators to monitor system health.
XVM	Provides disk striping, mirroring – makes nodes “CXFS” ready.
NUMATOOLS	Used to specify CPU, memory usage characteristics & fine tuning – accessible by developers, users to tune application execution.
Performance Co-Pilot[™]	System monitoring tool; used to view processor activity, loads, etc.
Storage Administration Tools	Additional tools for managing disk resources – xscsi, udev, LSI commands. Not provided by standard Linux [®] OS.
Failover / Cluster Manager	Basic tool for cluster failover management

SGI® ProPack 5 SP1 Features* and Benefits

Feature	SLES		RHEL5		Benefit
	IPF	x86	IPF	x86	
Linkless FFIO	✓	✓	✓	✓	Drives dramatic performance enhancement in I/O intensive configurations. Can be enabled at site without ISV engagement.
CPUSETS	✓	✓	✓	✓	Gives site administrator ability to control CPU usage for improved performance, decreased CPU contention, boosting job throughput. Some cluster workload managers automatically use.
NUMATOOLS	✓	✓	✓	✓	Developers use to improve application performance, yielding best job throughput. Multi-core CPU systems waste many cycles without NUMATOOLS.
SGI PCP	✓	✓	✓	✓	SGI Performance Co-Pilot™. System administrator monitors instantaneous system behavior and can pull up historical data, in nice graphical format. Zone in on hot spots and candidates for HW upgrade.
ESP	✓	✓	✓	✓	Fire and forget, automatically captures system crash data, reporting to the system admin and to SGI.
Storage Tools	✓	✓	✓	✓	System admin can do a range of low level harddrive parameter setting, firmware updates, etc. Not provided by standard Linux® OS.
Gridstack 4.2	✓	✓	✓	✓	Infiniband Open Fabric package, providing best performance of the IB cluster interconnect. (available separately)
Intel Runtime Libraries	✓	✓	✓		Customer convenience packaging into ProPack
Partitioning	✓		✓		XPMEM, XPNET, and XPC drivers: provide high-speed inter-partition communication within a numalinked Super Cluster environment

SGI® ProPack 5 SP1 Features* and Benefits

Feature	SLES		RHEL5		Benefit
	IPF x86	IPF x86	IPF x86	IPF x86	
SGI MPT	✓		✓		Highest application performance and security: provides MPI, shmem, and global shared memory libraries, optimized for SGI Altix numalink architecture
Intel MPI DAPL	✓		✓		Provides best Intel MPI performance on SGI Altix architecture
Array Services	✓		✓		Launch and stop MPI jobs across a cluster, in a high security environment.
CSA	✓				Accurately track and report on cluster-wide resource usage, to charge-back users.
Realtime	✓	✓			Critical enabling features for customers building realtime applications: SGI REACT®, kernel barriers, guarantee interrupt response time, user level interrupts, external interrupts, Frame Rate Scheduler
XFS	✓	✓			Most popular filesystem in the world: XFS high performance file system, fully supported by Novell SLES.
XVM	✓	✓			Storage virtualization: disk striping, mirroring – makes nodes “CXFS” ready.
DMF / TMF	✓	✓			Data and tape migration facility (available separately)
CXFS™	✓	✓	✓	✓	High Performance, shared file system, provides data sharing, enhanced workflow, and reduced costs in data-intensive environments (available separately as Server or Client)