

OPENVMS ON NEW INTEGRITY SERVERS

Clarete Riana, OpenVMS Engineering

©2011 Hewlett-Packard Development Company, L.P.
The information contained herein is subject to change without notice



Agenda

New Integrity Servers Overview

Differentiating features

Server configurations and field upgrades

Tips & Tricks

Integrity Roadmap



i2 server family

Introduced rx2800 i2 rack-mounted server



OpenVMS V8.4

Integrity Server Blades

World's first scale-up blades built on the industry's #1 blade infrastructure



Integrity 2 Socket Rack Server



8-core scalability with 3x improved density without sacrificing RAS


New!




NEW INTEGRITY SERVERS: OVERVIEW



BL860c i2 Base Blade - Top View



Rx2800 i2



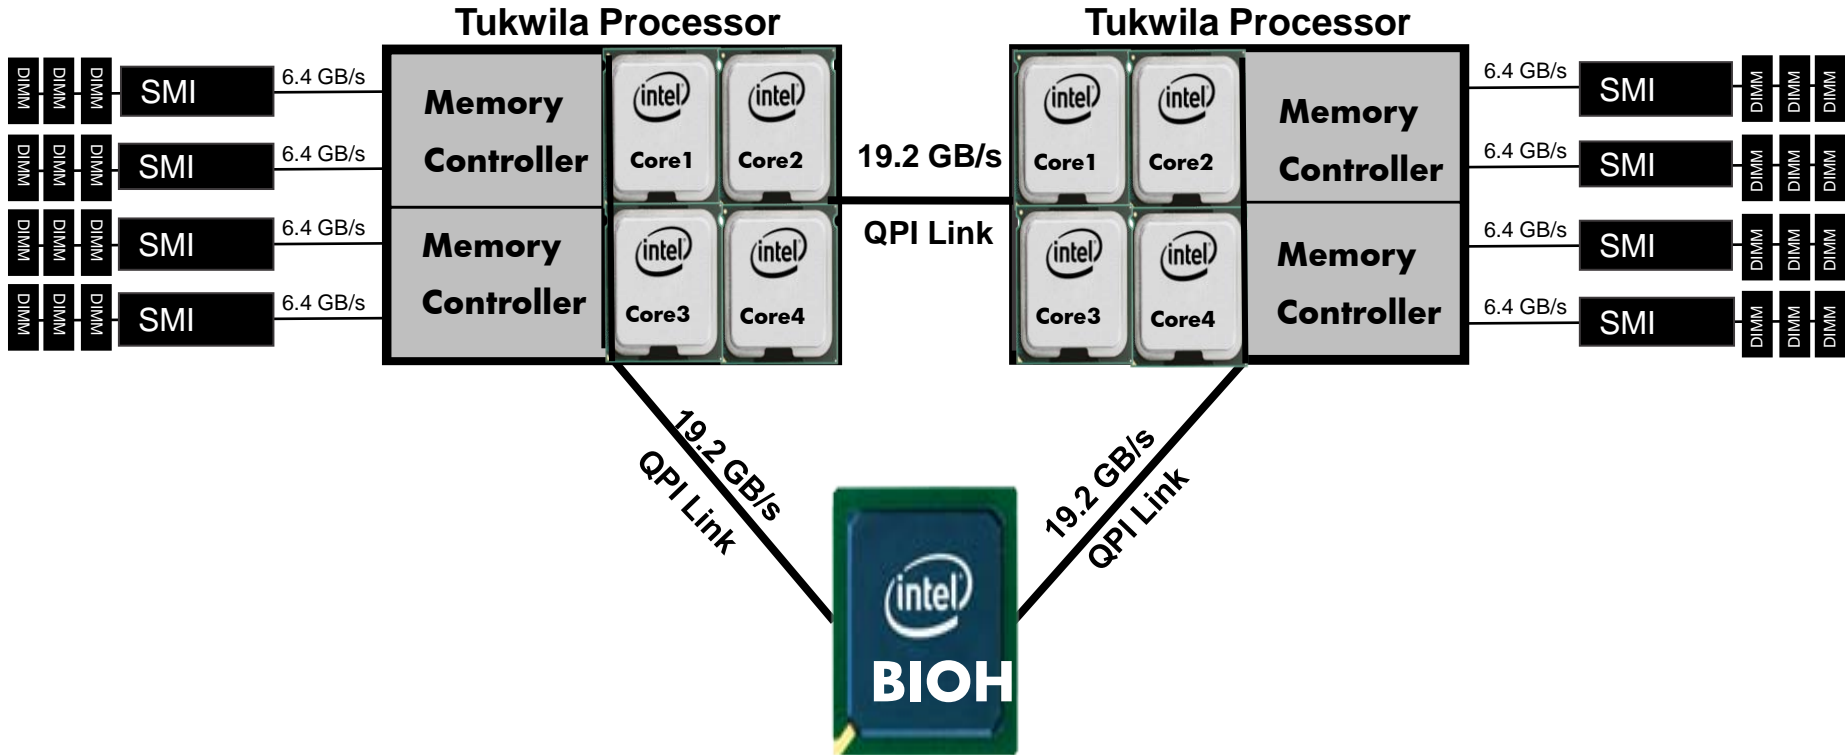
New!



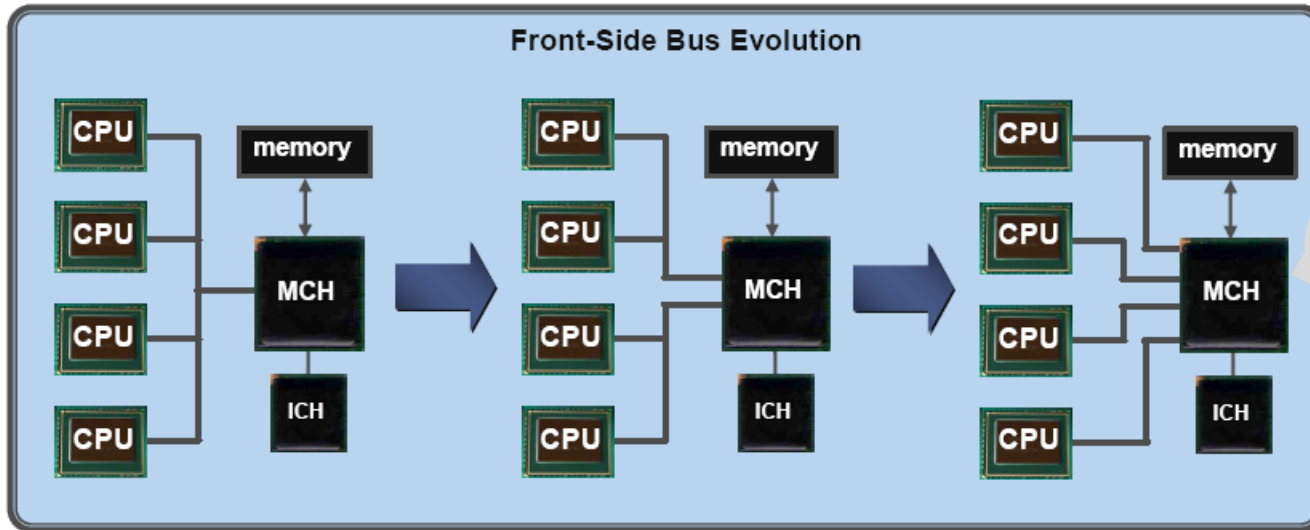
NEW INTEGRITY SERVERS: DIFFERENTIATING FACTORS



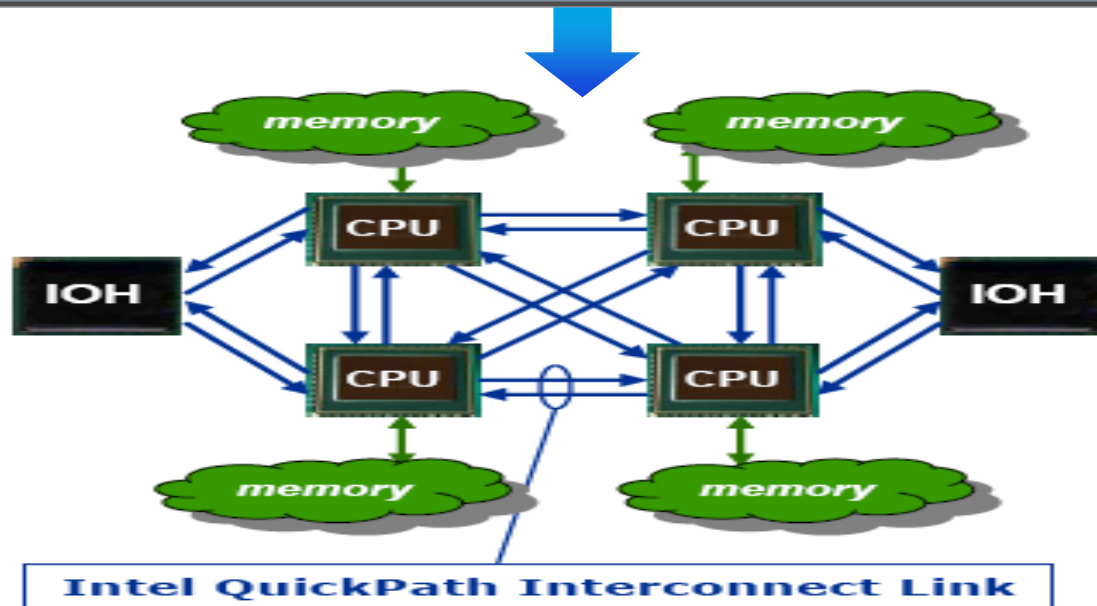
9300 Processor & Memory



QuickPath Interconnect



All CPU hog same MC; Memory Controller Bottleneck



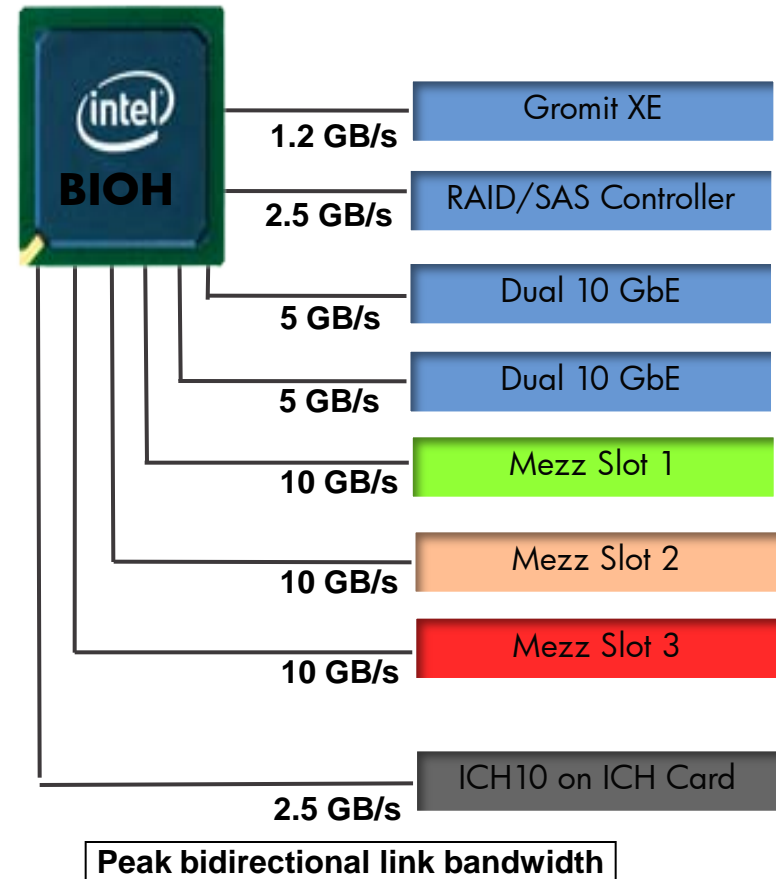
Intel® E7500 IOH & ICH10 south bridge

I/O Hub (E7500 Chipset)

- Connects to local CPUs via QPI links
- Provides 36 PCIe Gen 2 lanes
 - Order of magnitude peak IO bandwidth increase over previous generation in BL870c i2
- Hosts major IO functions
 - p410i RAID/SAS controller
 - Two dual-port 10GbE Flex-10 NICs
 - Three x8-provisioned mezzanine slots
 - Gromit XE (iLO3) mgmt controller
 - ICH10 I/O Controller Hub (SouthBridge)

ICH10 utilization

- x4 PCIe Gen 1 link for partner blade support
- Support VGA controller, USB controller



Flex-10 LAN-on-Motherboard


Each BL8x0c motherboard integrates two Broadcom 57711E dual port 10GbE controllers (total of four physical 10GbE ports per motherboard)

With a Flex-10 Virtual Connect module (and VC 3.0 FW), the resulting built-in FlexNIC counts are:

- 16 FlexNICs in a BL860c i2 server
- 32 FlexNICs in a BL870c i2 server
- 64 FlexNICs in a BL890c i2 server

Need more? Add Flex-10 capable mezz cards (like the NC532m)

- Up to 128 FlexNICs supported in a BL890c i2 !



*NEW INTEGRITY SERVERS:
PERFORMANCE & RAS
FEATURES*



Keep Your Business Running


Processing

- Error correction codes (ECC), parity protection
- Soft Error hardened latches
- Cache Safe technology
- MCA



Memory

- Double-chip sparing & EC
- SDDC & DDDC
- Pro-active memory scrubbing



QPI & I/O

- PCIe link cyclic redundancy check
- Isolated I/O buses for error containment
- Intelligent Error Management



Components

- Color coded latches for faster upgrades and repairs
- Redundant, hot-swap power supplies and fans



Faster Performance


Processing

- Quad core Enhanced Thread-Level Parallelism (TLP)
- Intel® Turbo Boost Technology – Performance on Demand
- Intel® VT-i2 Introduced -
- Data TLB support for 8K and 16K pages



Memory

- 2 integrated memory controllers, peak memory band-width up to 34 GB/s (6x)
- Capability to support more than 1TB with DDR3 DIMMNS
- Directory-based Cache Coherency – Reduces Snoop traffic and contention



QPI

- New Intel® QuickPath Interconnect Technology -
- 4 full-width Intel QuickPath Interconnect links and 2 half-width links
- Peak processor-to-processor and processor-to-I/O communications up to 96 GB/s (9x)

I/O

- Pci-e gen 2 lanes
- 2x faster than gen 1 lanes
- Faster IO devices



NEW INTEGRITY SERVERS: CONFIGURATIONS



Server Configurations

Linear scalability with industry's first 2-4-8 socket server blades

Blade Link combines multiple blades into a single, scalable system

Scale Up, Out and Within

Scale More Only 8-Socket blade in industry standard blade enclosure

CPU	2s/8c	4s/16c	8s/32c
Memory	384GB	X 2 = 768GB	X 2 = >1.5 TB
LAN	4 x 10GbE	8 x 10GbE	16 x 10GbE
HDDs	2 Slots	4 Slots	8 Slots

Scale Linear System resources grow evenly across CPU, memory, I/O, and etc

8 socket system at 2x the performance in half the footprint

Improve Business Response Time With Compact Scaling

Run a range of mission-critical applications in a 2U footprint

From dual-core, 2- and 4-socket Integrity servers...

rx2660



rx3600



rx6600



To one quad-core, 2-socket, 2U Integrity server that can do more

rx2800 i2



- Per socket performance increases
- Quad-core Intel® Itanium® 9300 series processors
- Up to 96 GB memory *
- Up to 2.4 TB storage *
- 6 I/O slots
- 4 integrated 1GbE LAN
- SATA DVD

Supported configurations for OpenVMS

Supported



BL860c i2, BL870c i2, BL890c i2, Rx2800 i2

LAN, FC pass thru and switches

c3000, c7000 enclosures

Core I/O SAS disks (RAID mode)

Network NICs

- 10 GigE LOM, MEZZ
- 1 Gbps & dual-port (NC364m, NC360m)

Fibre Channel HBA

- 8 Gbps dual-port FC (Q-logic)

External SAS – P700m

MDS600, P2000G3, MSA2000G2

OpenVMS guest

vMedia, DVD (internal, USB)

8GB, 16GB DIMMs with BL8x0c i2

Virtual Connect, Flex-10

Power Management through ILO

Not yet Supported

8GB & 16GB DIMMs on Rx2800 i2- Coming soon!

450GB and 600GB disks on Rx2800 i2 – Coming soon!

FireMV – Coming soon !

Blade link and field upgrades



BL860c i2



BL870c i2



BL860c i2

Or



BL870c i2



BL890c i2



TIPS & TRICKS



Tips & Tricks

Shadowing across controllers and caching



NUMA, RAD & Fastpath settings



DPR



Partial Dump, Console and excluding devices during configuration



Hardware Setup

SAS controller in RAID mode



Port mapping



vMedia Installation & Boot



CPU & Memory Loading Rules



Leadership in I/O and Storage on i2 architecture

High performance, reliable and scalable

SAS provides a point-to-point connection to each HDD

Parallel SCSI with rx7640 has a shared bus
Ultra 160 SCSI

Provides four p410i RAID controllers (one per blade) on BL890c i2,
One P410i RAID on rx2800 i2

Configured as RAID 0/1 or HBA mode [Future]

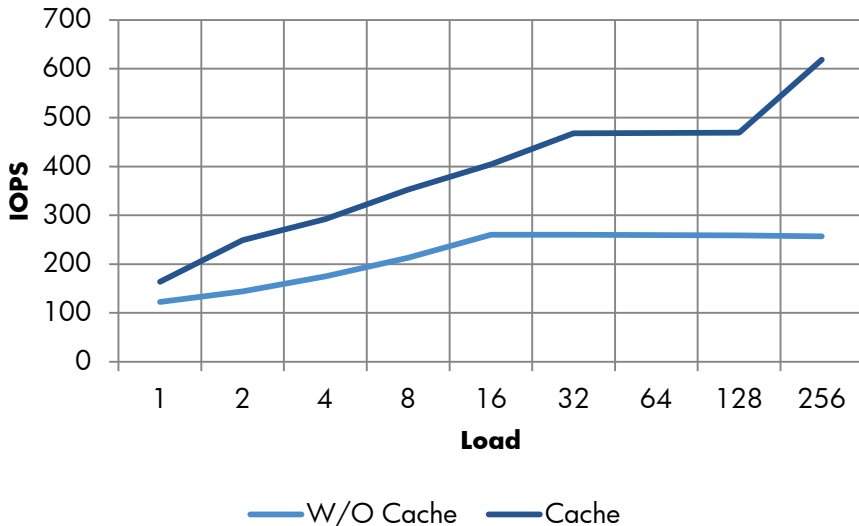
Stripe data across multiple p410i RAID controllers (OpenVMS Shadowing)

Striping across controllers provides no SPOF storage

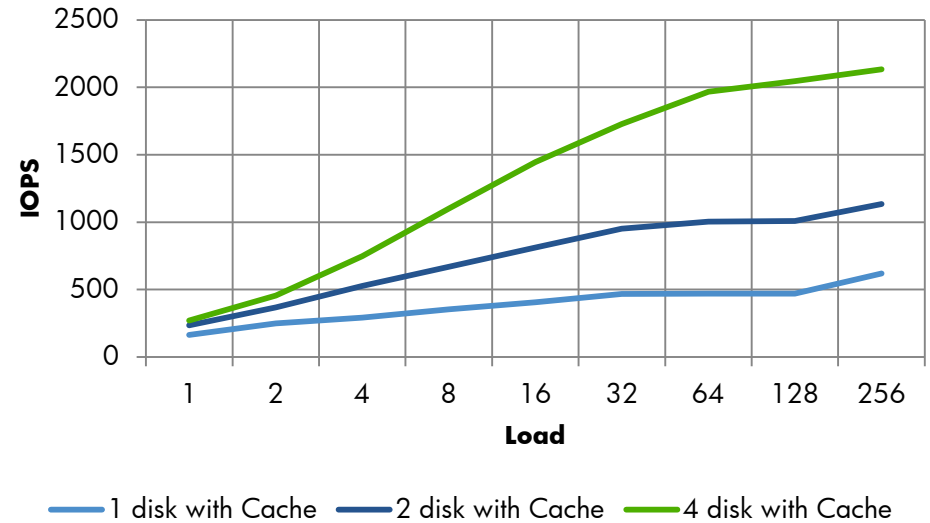
Each BL890c/rx2800 supports eight SFF SAS HDD, up to 4.8TB capacity

Rack Mounted rx2800 i2 – Core SAS


rx2800 i2 - Core SAS Caching



rx2800 i2 - SAS Logical Disk (Striping)




- Rx2800 i2 server comes with p410i SAS Controller as Core SAS I/F
 - Small Block Random Tests were run on same sized disk
 - Logical Volumes were spread across multiple disks to show the I/O striping effect
- p410i with cache exponentially boosts the performance W/O (upto 2.5x)
- Increased number of disks in a Logical Volume increases performance



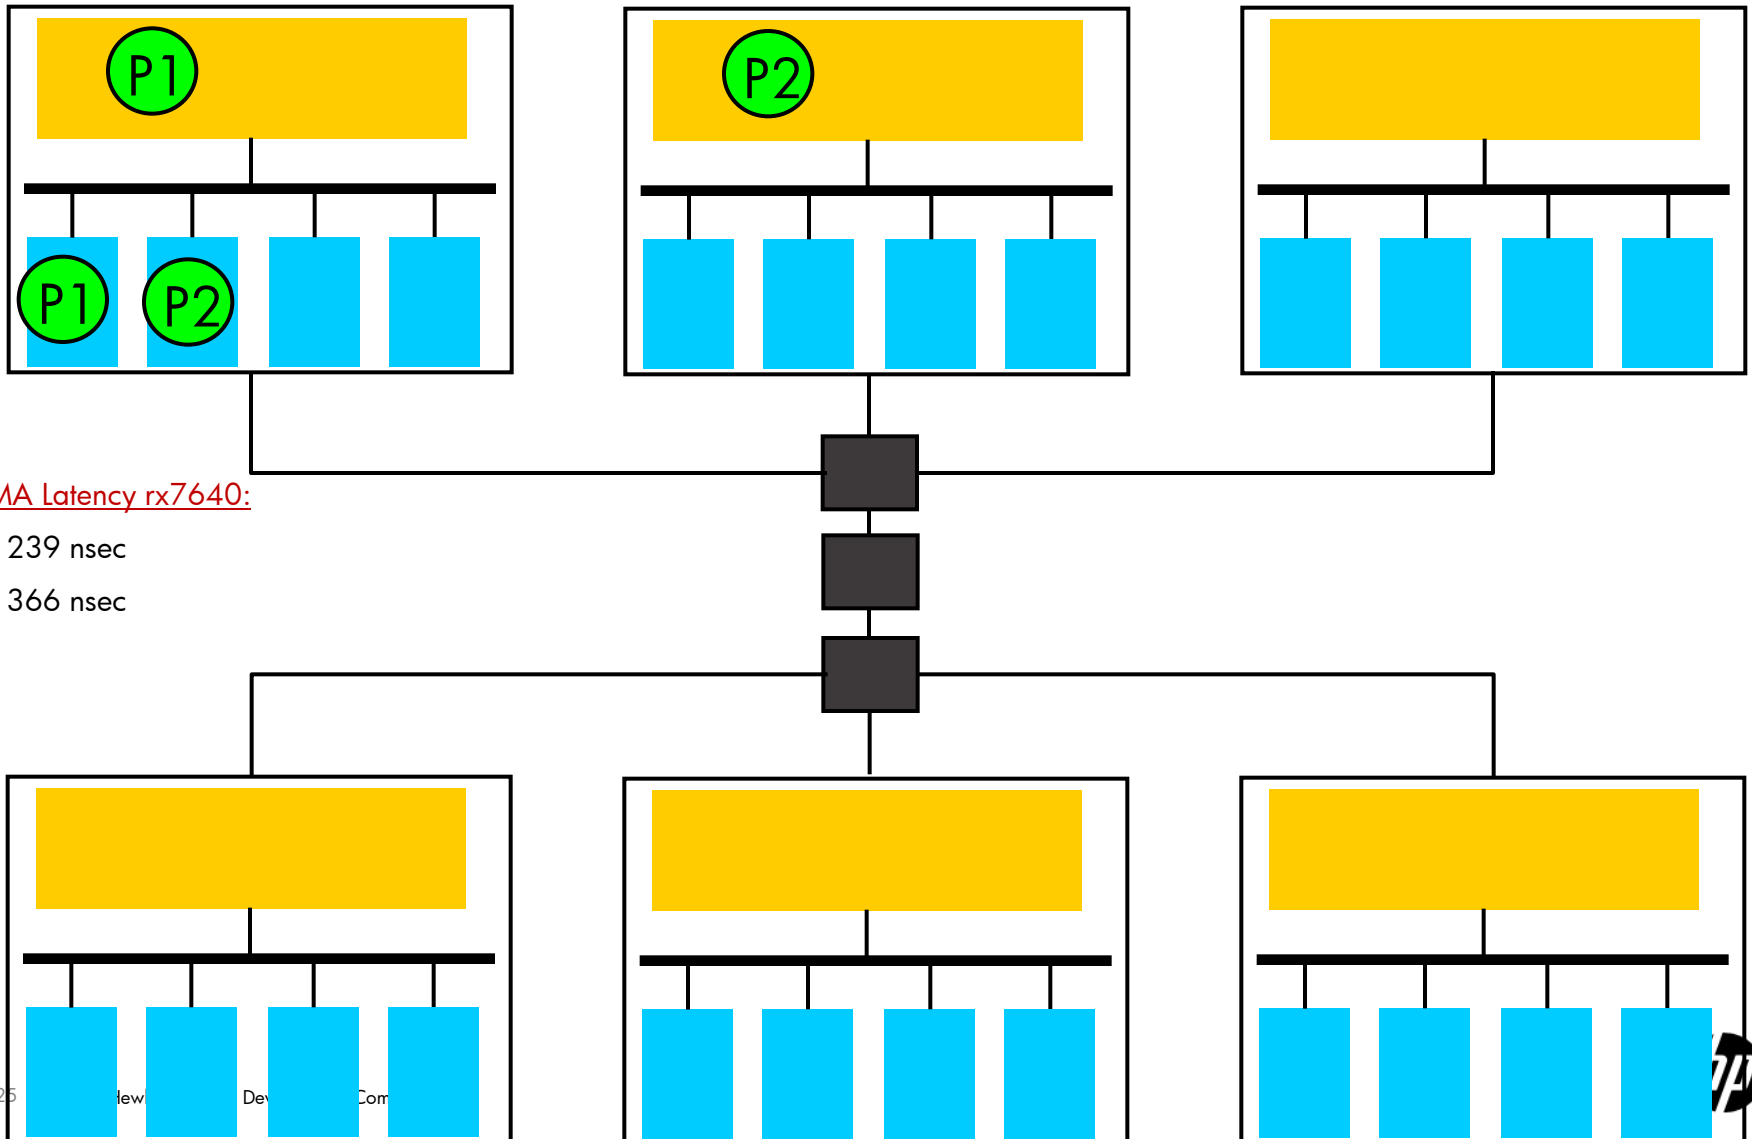
Hardware Evolution

Traditional Scale Out System - Uniform Memory Access (UMA)



NUMA Introduction

Existing Scale Out System – Non-Uniform Memory Access (NUMA)



NUMA Latency rx7640:

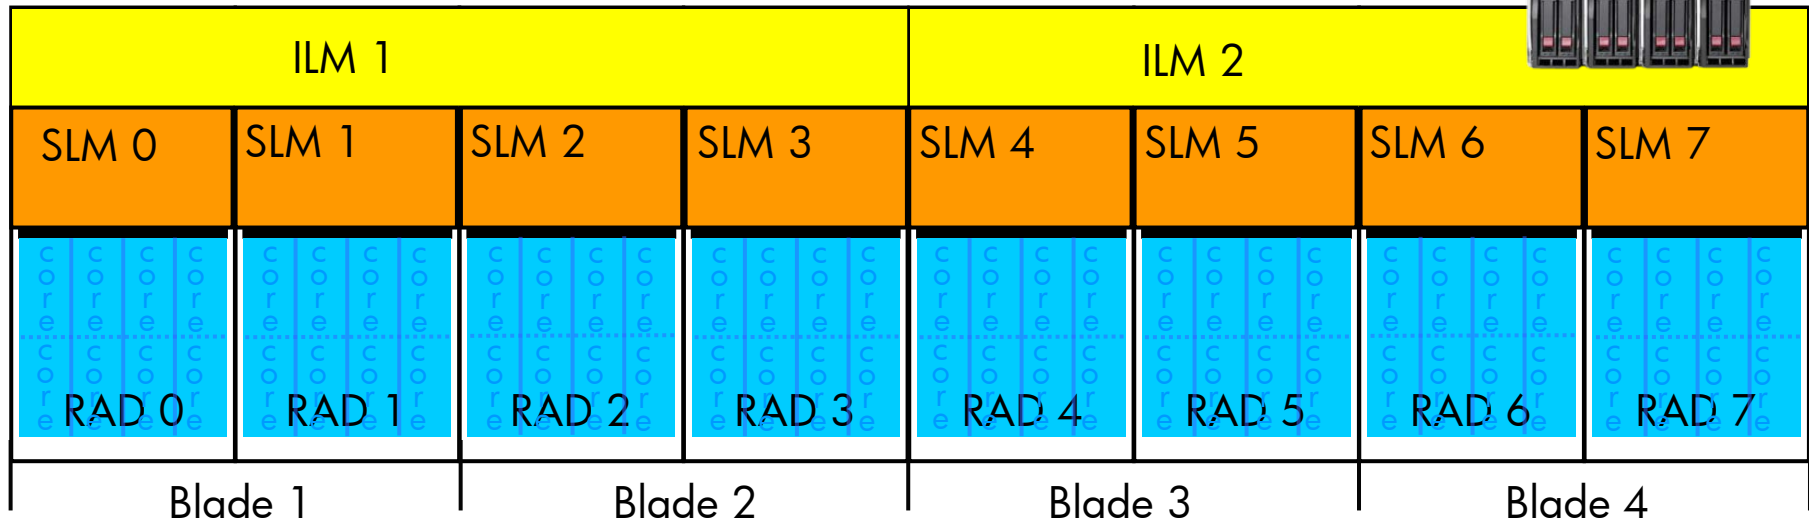
P1 – 239 nsec

P2 – 366 nsec



BL890c i2 NUMA

SLM and ILM in i2 Blades




- BL890c i2 is 4 blade server connected through Blade Link (QPI)
- BL890c i2 comes with 8 Sockets
 - Each Blade will have 2 Sockets
 - Each Socket comes with Memory Controllers and DIMM's, so the SLM
 - 5 different memory combinations/configurations are supported on i2 server
 - Default configuration shown above will have , 8 SLM and 2 ILM (total of 10 RAD's)



rx2800 i2 NUMA

– SLM and ILM



rx2800 i2 comes with total of 2 Sockets (Quad-Core/Dual-Core

- Each Socket comes with Memory Controllers and DIMM's, so the SLM
- 5 different memory combinations/configurations are supported on i2 server
- Default configuration shown above will have , 2 SLM and 1 ILM (so total of 3 RAD's)

Configuring Memory on NUMA systems

- Integrity Cell Based Systems
- BL8x0c i2 server blades
 - EFI shell provides the memconfig –mi command

Memory Option	ILM	SLM
MaxUMA	8/8	0/8
MostlyUMA	7/8	1/8
Balanced	4/8	4/8
MostlyNUMA (Default)	1/8	7/8
MaxNUMA	0/8	8/8



RAD Performance Guidelines

A NUMA system should perform better when memory configured with a combination of CLM and ILM


Magnitude of performance improvement depends on memory access patterns of the application

Best suited memory configuration has to be selected

- The OS has features to help get the most out of RADs
- Memory allocation attempts to obtain process private memory from the memory associated with the RAD
- The scheduler attempts to execute the process on a CPU that is part of the RAD
- Almost any memory allocation routine and process creation type allows a RAD to be specified
- Application should ensure processes sharing common data runs in the same RAD and the common data also lives in the same RAD



RAD Guidelines for I/O



Keep I/O
Devices close to
process which is
heavily accessing
it


Make use of
FASTPATH
efficiently

Statically load
balance the
devices across
multiple RAD's

Make use to SET
PROC/AFFINITY
to bind processes

vMedia – Virtual media

- vMedia logic allows DVD drive in PC/laptop to be configured as a direct attached USB drive on server
- System Manager can do firmware updates and OS installs/upgrades on the server from media/ISO files on the PC/laptop DVD drive



hp **Integrated Lights-Out 3 Advanced**
HP Integrity rx2800 i2



iLO Hostname: sent11
Current User : Administrator
[Home](#) | [Sign Out](#)

Virtual Front Panel

LEDs: Locator UID System Health System Power

State: Boot complete

Navigation Menu

- System Status
 - Status Summary
 - System Health
 - System Event Log
 - Forward Progress Log
 - System Inventory
 - iLO Health
 - iLO Event Log
- Remote Console
 - Remote Serial Console
 - Integrated Remote Console
 - Virtual Media**
- Power Management
 - Power & Reset
 - Power Meter Readings
 - Power Regulator
- Administration
 - Firmware Upgrade
 - Licensing
 - User Administration
 - Local Accounts
 - Group Accounts
 - Settings

Virtual Media USB Key for EFI Only **Launch** ?

CD/DVD-ROM and USB Key Support ✔ Status: vMedia is available

Connect the CD/DVD-ROM or USB key physical devices or images on your client system to the remote server so they appear local to the server during system boot or while the operating system is available.


Virtual Media OS Support
Before launching, check [OS support](#) for Virtual Media.

Advanced License
Advanced Pack license is required to use this feature. The iLO comes with a full-featured permanent default Advanced Pack Licensing Key built into the firmware.

Virtual Media Right
The Virtual Media right is needed to use this feature. This right can be enabled from the User Administration pages by a user with the User Administration right.

Status Message Update
The status message will update every 10 seconds.

Integrity rx2800 i2



Change P410i SAS Controller Mode

- P410i controller mode – HBA to RAID using SAUPDATE.EFI
- [SAUPDATE.EFI](#) utility is available with the I/O firmware bundle
- Download the tar.gz file and SAUPDATE.EFI is part of it
- SAUPDATE.EFI on USB Pen Drive accessed from SHELL

– To get the mode of the controller

– *fs0:\saupdate get_mode all*

– To set the mode of controller to RAID

fs0:\saupdate set_mode p410i raid

– Execute “RECONNECT –R”

fs0:\reconnect –R

The boot disk for the system will need to be created from EFI

MSA\$UTIL can be used after boot to create rest of the units



DPR

DPR (Processor Indictment)

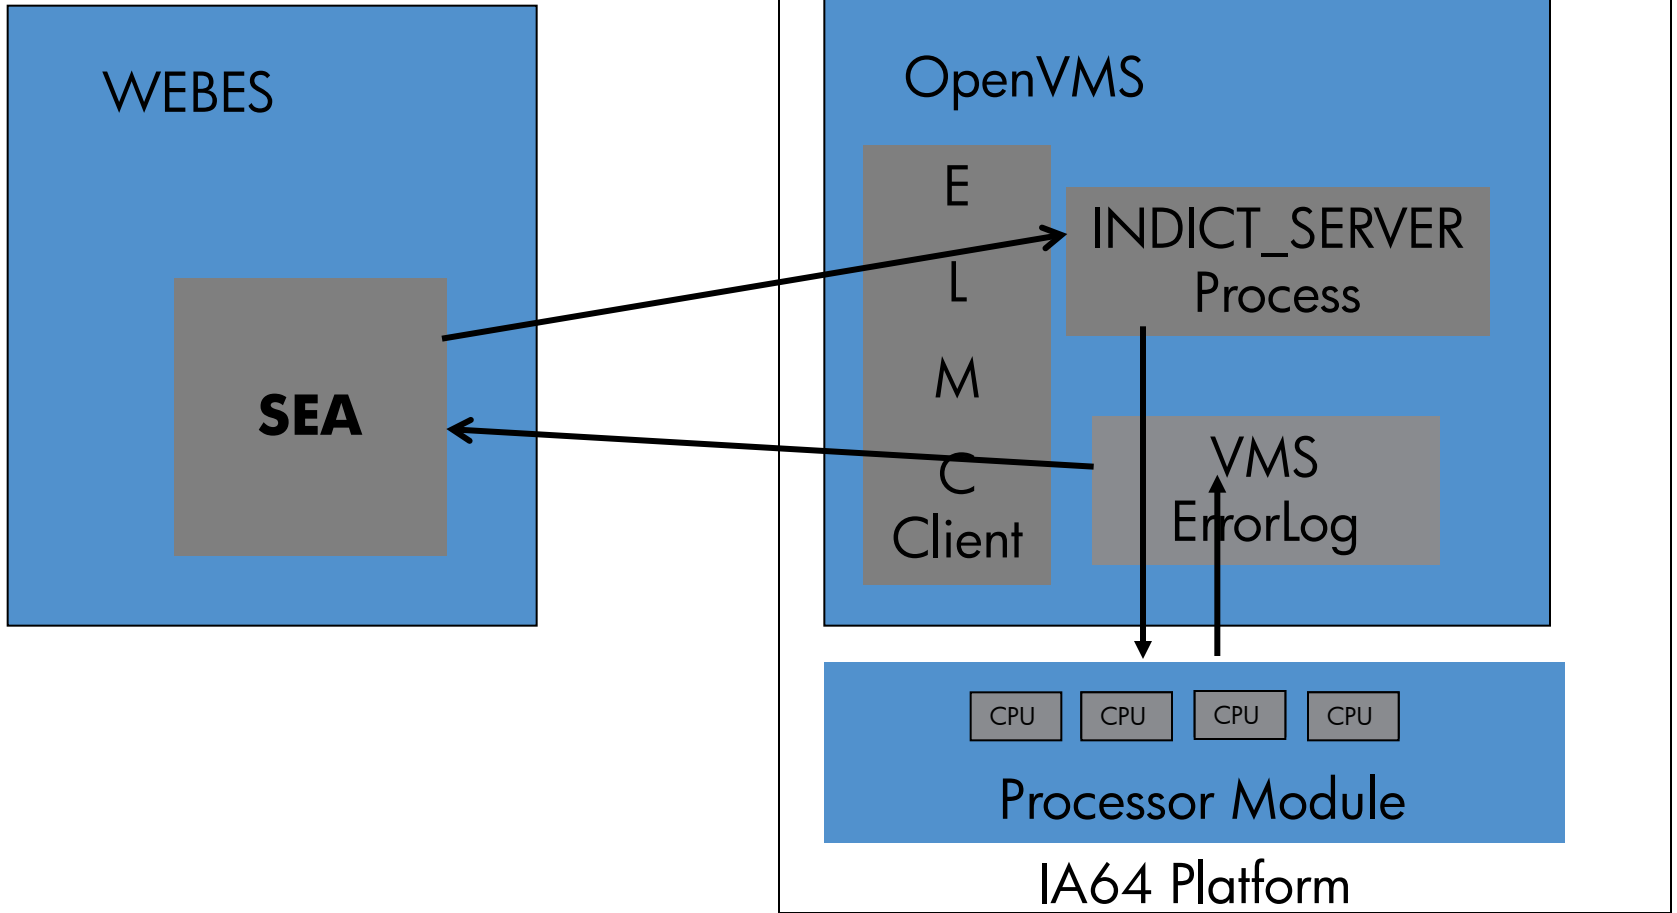
- Recognize degrading processors
- Isolate them from further usage

With OpenVMS V8.4

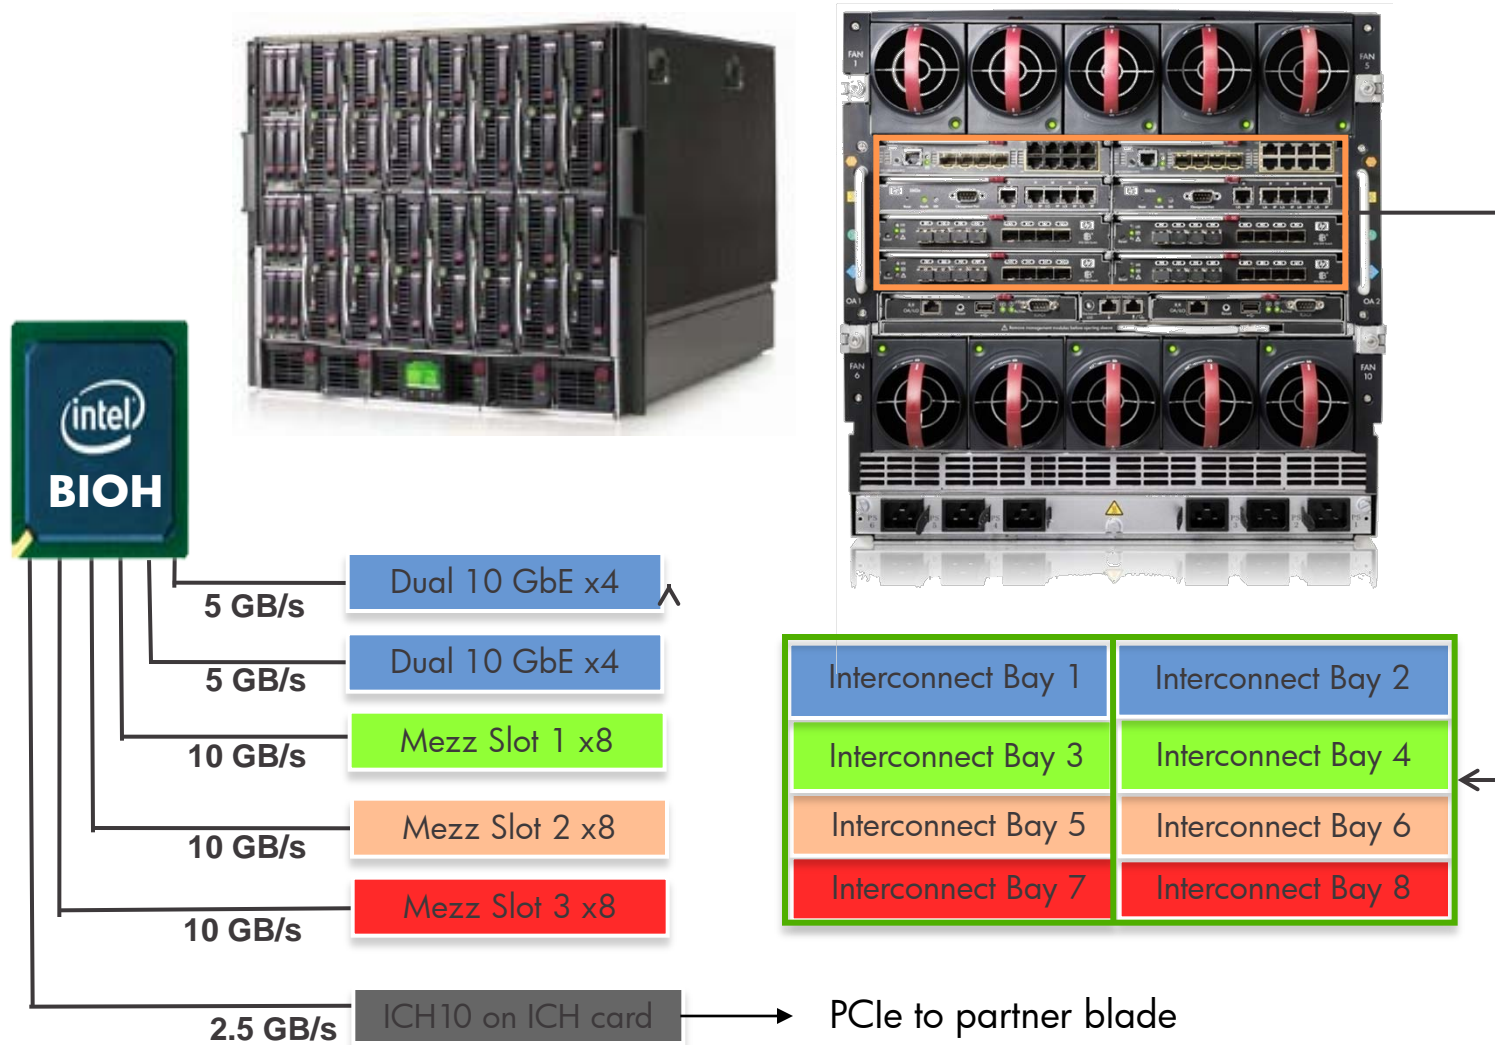
- Mark an indicted CPU as not available for use on the next reboot (deconfigure processor module through ACPI).
- Replace the indicted CPU with iCap spare CPU, if available.



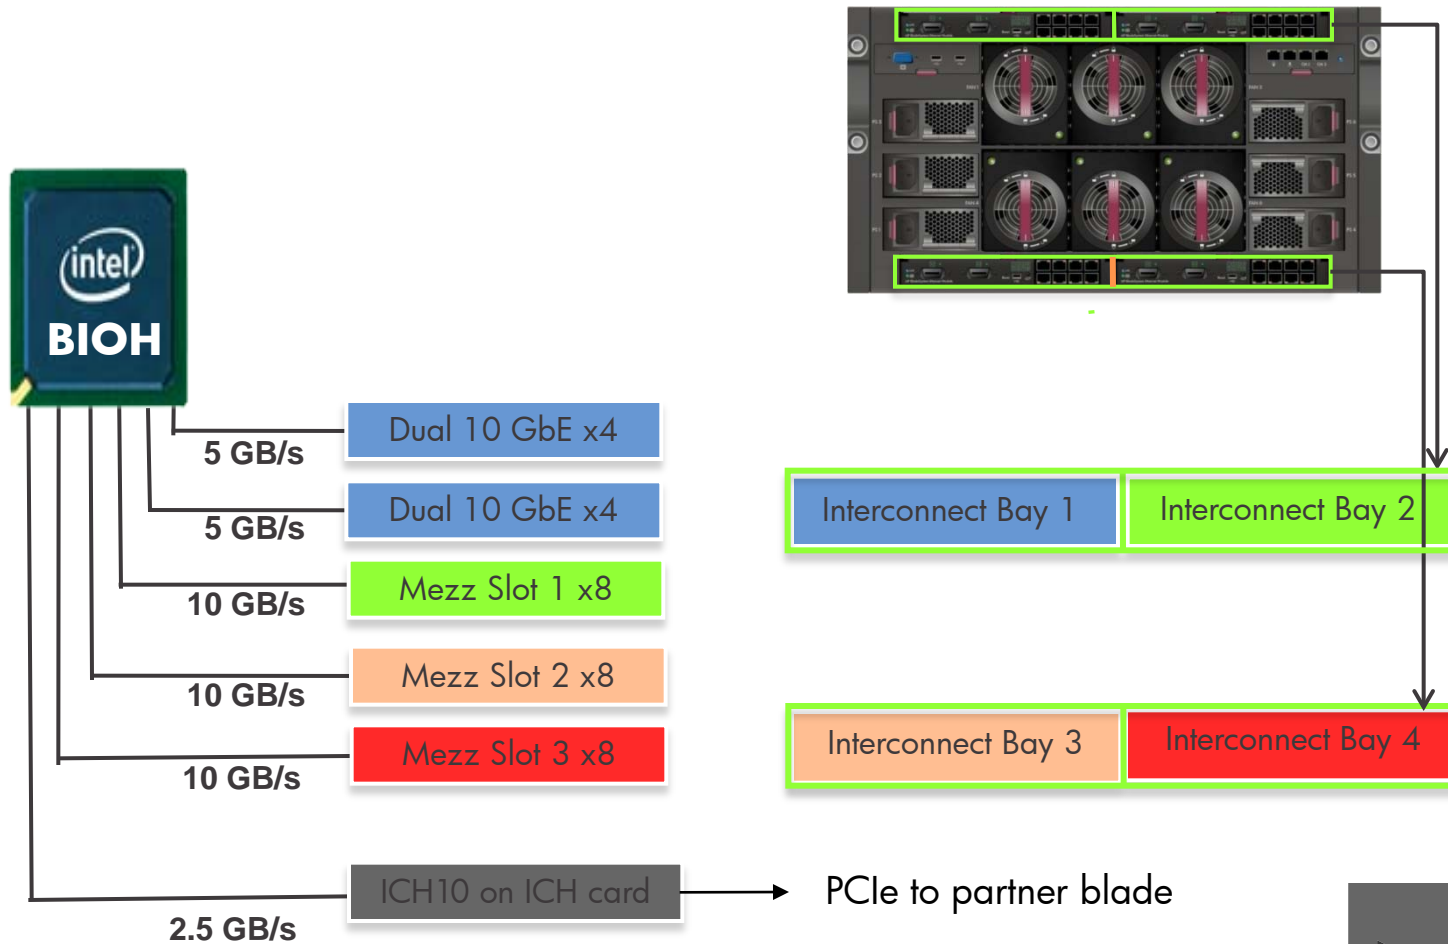
DPR



Interconnect Routing: c7000 Enclosure



Mid-plane Routing: c3000 Enclosure




CPU & Memory Loading Rules

CPU

- Dual core CPUs supported on BL860c i2 and Rx2800i2
- Provision to populate only one socket per blade

Memory

- Load highest density DIMMs first
- Alternate Loading between installed CPUs
- Load DIMM in pairs
- Spread DIMM evenly between CPUs
- mixed DIMM types also supported.



Tips & Tricks

Partial Dump

- Configure large memory systems for partial dumps, so the dumps are faster and quicker to copy
- \$Ana/crash
- SDA>COPY/PARTIAL=KEY SSRVKEY
- SDA>COPY/PARTIAL=PROCESS=NAME=CLUSTER_SERVER SSRVCSP

Console


- Fix for this coming in update 600.

Excluding devices from configuration



- \$mc sysman io set exclude=<devnam>

Enable hyper threading

- Shell>cpuconfig threads on



Integrity so far.....



Dual-Core Integrity servers 2- & 4-socket Integrity

rx2660




rx3600



rx6600



BL860c



BL870c

HP Integrity
rx1620 Server

HP Integrity
rx2620 Server

HP Integrity
rx4640 Server

HP Integrity
rx7640 Server

HP Integrity
rx8640 Server



HP Integrity
Superdome




Poulson

32nm Technology

Socket and Binary
compatibility

Enhanced
performance
features

Advanced RAS via
Instruction replay
Technology



3.1 billion
transistor 12-
Wide-Issue
Itanium Processor

8 cores, with HT
8 more virtual
cores

54MB cache

Q & A



Questions/Comments

Business Manager (Vivasvan Shastri)

Rohini.madhavan@hp.com

Office of Customer Programs

OpenVMS.Programs@hp.com



THANK YOU

