



ACCELERATE

outcomes with convergence

HP **TECHFORUM** 2010

CONVERGE

INNOVATE

TRANSFORM



COMMAND

your infrastructure like a
visionary, not a functionary.

OpenVMS 8.4 on Tukwila Quad Core

9300 – Blazing Performance

Prashanth K E, Technical Architect, OpenVMS Lab

24th June 2010

k.e.prashanth@hp.com



Agenda

- Today's Mission-Critical challenges
- Mission-Critical Converged Infrastructure
- HP's Server Strategy
- Introduction to Intel ® Itanium ® 9300-series processor based servers
- BL8x0c i2 Configurations Supported by OpenVMS
- BL8x0c i2 Performance Features and OpenVMS Test Results
- OpenVMS Support for RAS features on BL8x0c i2
- OpenVMS Power Management Support
- Q & A

Mission-Critical Customer Challenges

Financial Services

Every minute of downtime = a minute of lost revenue!



Manufacturing and Distribution

Production comes to grinding halt



Healthcare

Patient outcomes depend on 24x7 access to data



Public Sector and CME

Customer retention and fraud detection at risk



No tolerance for downtime


Increasing SLAs with decreasing budgets

Islands of legacy apps and monolithic systems



The First Mission-Critical Converged Infrastructure

New Integrity systems optimized for the converged infrastructure



A common, modular architecture that simplifies, consolidates, and automates everything

A mission-critical infrastructure delivering the highest levels of reliability and flexibility

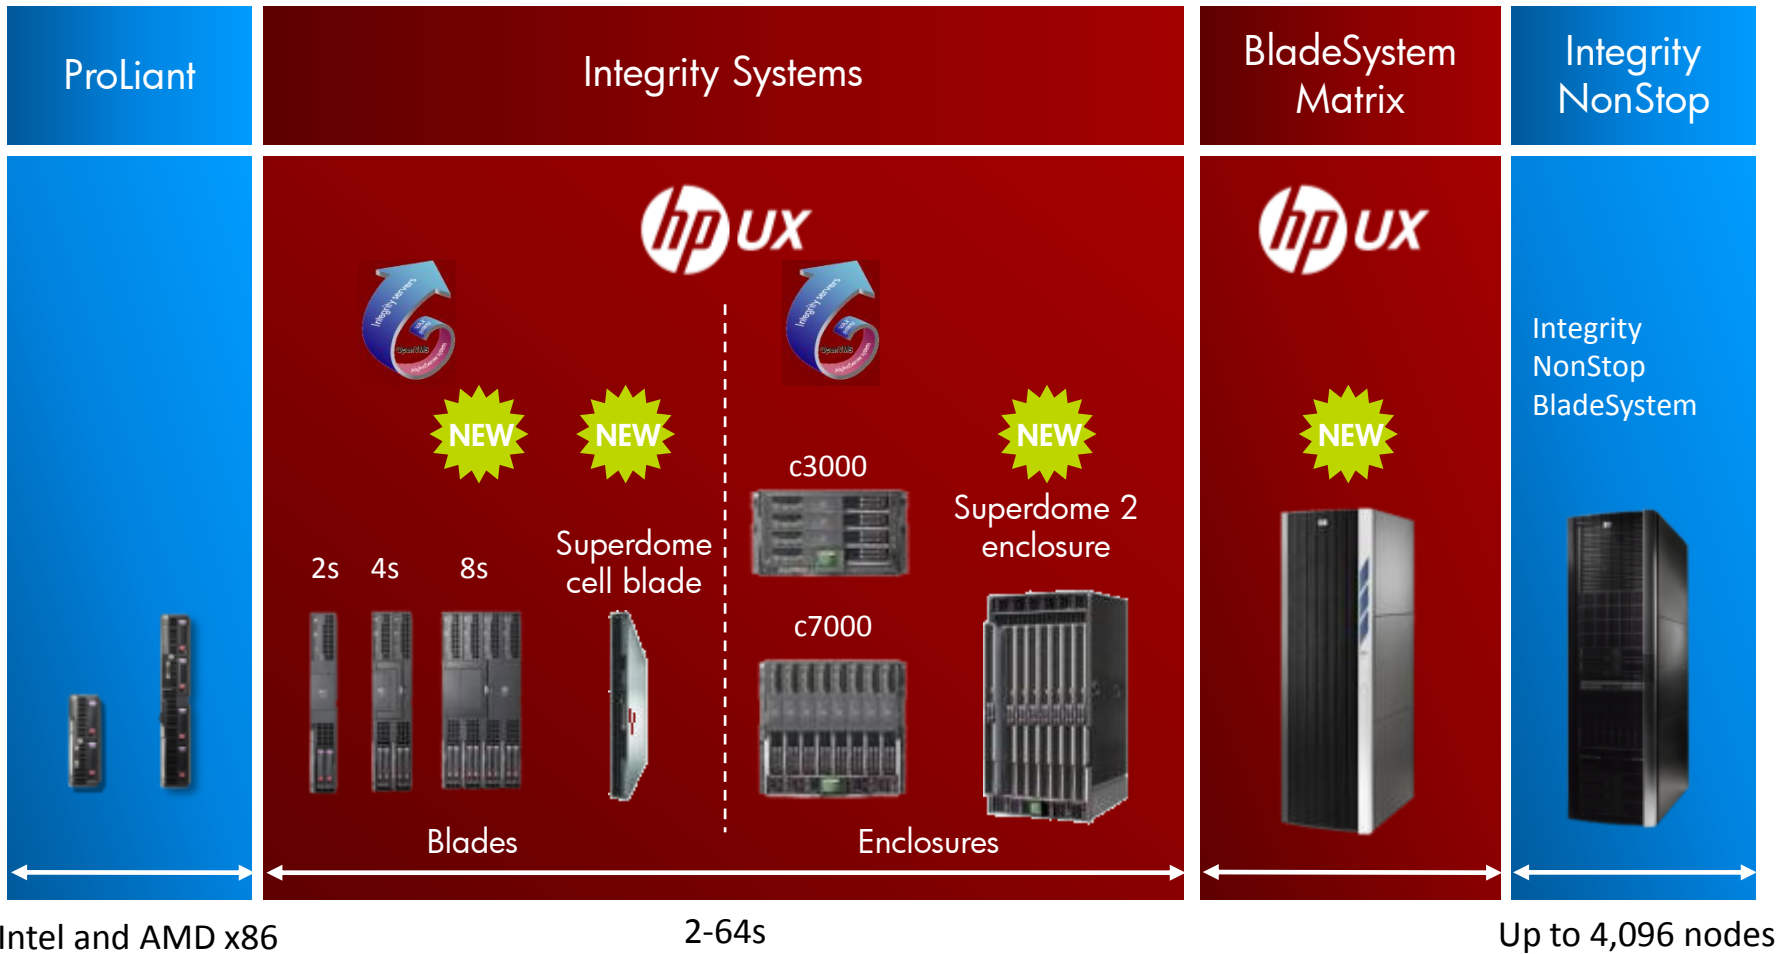
Introducing the Revolutionary Blade Scale Architecture

The first mission-critical Converged Infrastructure on the industry's #1 blade platform



Unified Blade Architecture from x86 to Superdome

Simplify by consolidating applications on a common platform



Train once, certify once, deploy once

FlexFabric

Scale resources to meet any workload demand

Within the system



Blade Link

Linear scaling from 2-8 sockets



Crossbar Fabric

Fault-tolerant flexible scaling

Across the network



Virtual Connect

Wire-once, change-ready connectivity

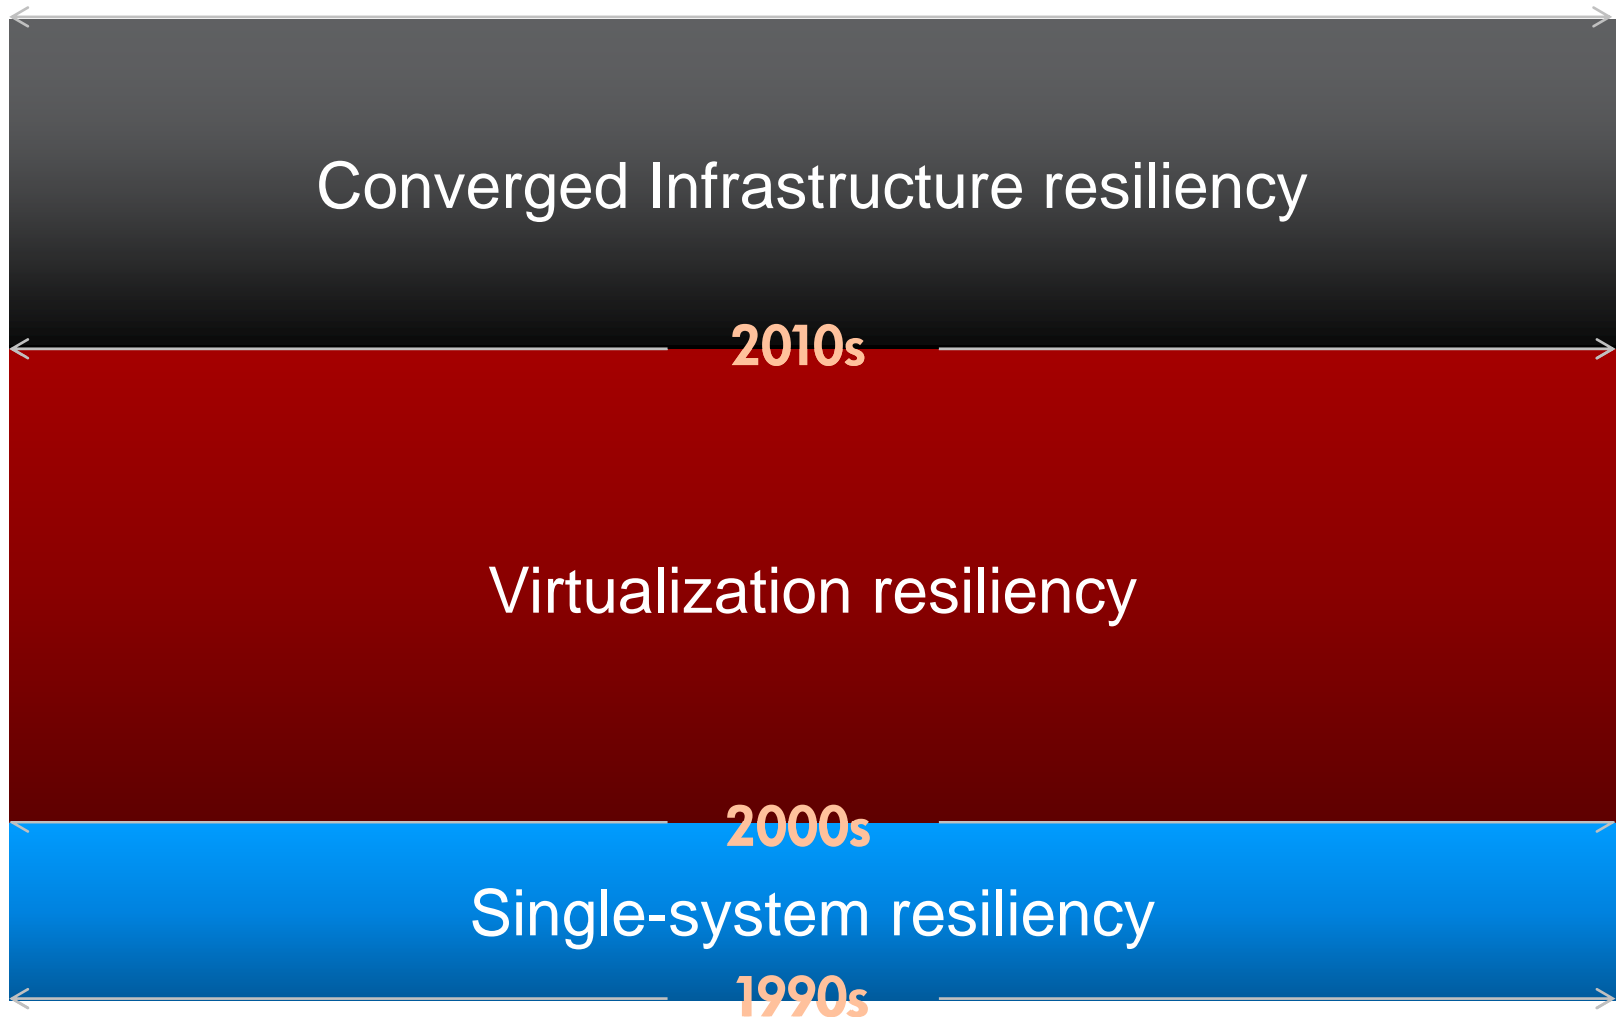
Dynamically managed



Matrix operating environment, delivered by Insight Dynamics

Always-on Resiliency

100+ new innovations to ensure global business continuity



Matrix Operating Environment

Common management delivered by HP Insight Dynamics

**Provision
infrastructure
in minutes**



New for HP-UX:

Automated orchestration
for physical blades


**Optimize workloads
confidently**



Only for HP-UX:

Most comprehensive automated
workload management

**Reduce data center
and energy costs**



HP ranked #1

Optimized energy usage without
compromising performance or
reliability

HP's Server Strategy

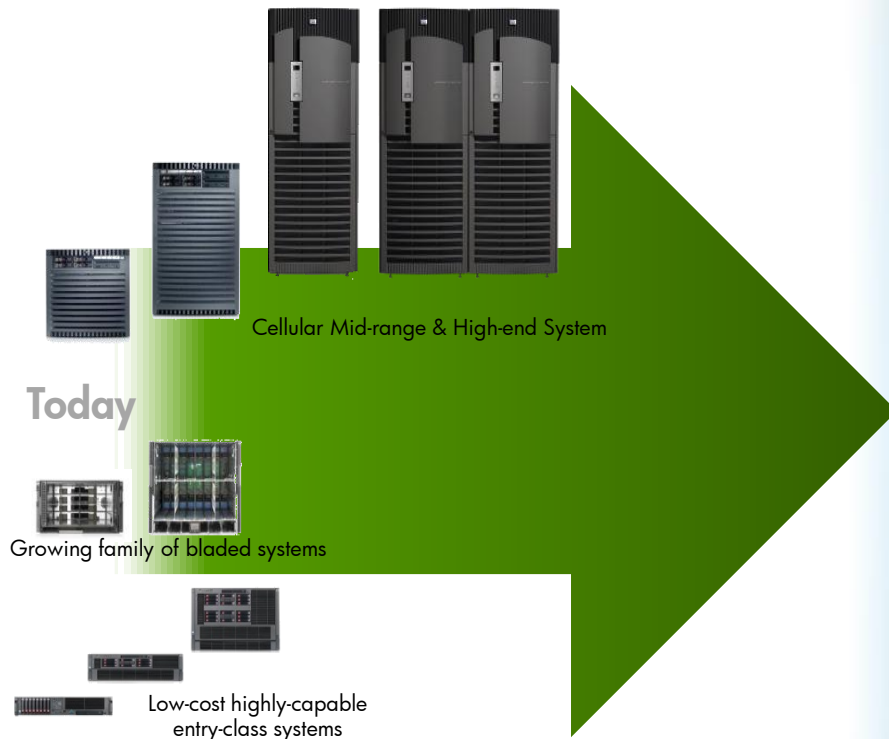


Integrity value redefines the data center with...

2010

Future

A future business critical infrastructure



More reliability and business criticality with:

More modular design

More common components

More flexible and customizable architecture

More performance

Enhanced robustness and fully virtualized

Delivering improved data center economics

Integrity value redefines the data center with...


2010

Future

A future business critical infrastructure



"Superdome-Gen2"
2-64s SMP server



New Processor: 9300 series

2-8S SMP Server Blades



Supported in C7000 & C3000 chassis

2s/2U rack mount server



New Processor: Poulson

New Processor: Kitson

Delivering improved data center economics

Introduction to Intel® Itanium® 9300-series processor based servers



BL8x0c i2 Overview

Processors and chipset

- Intel® Itanium® 9300-series processors
- Intel® E7500 Scalable Memory Buffer
- Intel® E7500 IOH
- Intel® ICH10 south bridge

Form factor

- Full Height c-Class form factor
 - Single wide BL860c i2 – 2s
 - Double wide BL870c i2 – 4s
 - Quad wide BL890c i2 – 8s
- Supported in c3000 and c7000 enclosures

Memory (per blade)

- ❖ ~6X memory BW increase over previous generation
- 24 PC3-8500 DIMM sockets
- 192 GB capacity per blade with 8GB DIMMs

I/O subsystem (per blade)

- ❖ 5X IO BW vs. previous generation
- Integrated p410i RAID controller
- 2 dual-port 10GbE Flex-10 NICs; VC support
- 3 PCIe G2 mezz slots

Additional IO options

- Two hot-plug SFF SAS HDDs per blade
- Partner blade support

High availability

- Enhanced processor RAS features
- Memory double chip spare
- Internal SAS RAID
- Enhanced interconnect RAS

Power Management

- Enhanced Demand based switching
- Turbo boost

Management

- Integrated Lights Out (iLO3)
- Integrated VGA console
- iLO 3 Advanced Pack firmware
- c-Class Onboard Administrator

Operating Systems Support

- HP-UX 11i v3
- OpenVMS 8.4 (Aug '10)
- Windows 2008 R2 (to follow)



Scaleable Blade Link

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	96GB	X 2 = 192GB	X 2 = 384GB
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

Migration from Current Integrity Server Products



rx2800 i2 – 2p/8c RACKMOUNT

Value and Key Features

Coming soon !!!

rx2800 i2	
Processor	Up to two Dual-Core or Quad-Core Intel® Itanium® processors
Memory	Industry standard DDR3 technology 24 PC3-8500 DIMM sockets 192GB max (with 8GB DIMMs)
Internal Storage	8 Hot-Plug SFF Serial Attached SCSI HDDs 1 CD+RW or DVD+RW
Networking	Dual Integrated Gigabit Ethernet ports Manageability LAN
I/O Slots	6 PCI-E slots (2 x8, 4 x4 slots)
Management	Management Processor with iLO3 (Integrated Lights Out functionality)
Form Factor	2U Rack Mount Server (office tower conversion kit available)



Key Points:

- Affordable, 8-core scalable entry-level non-x86 server
- High Density Compute Server (2U footprint)
- Excellent memory capacity
- Continued innovations in RAS
- Ease to deploy into today's racked environments
- Office Friendly pedestal option
- N+N redundant power and cooling

BL8x0c i2
configurations
supported by
OpenVMS



Supported configurations for OpenVMS

Supported

- BL860c i2, BL870c i2, BL890c i2
- LAN, FC pass thru and switches
- c3000, c7000 enclosures
- Core I/O SAS disks (RAID mode)
- Network NICs
 - 10 GigE LOM
 - 10 Gbps mezz
 - 1 Gbps quad-port (NC364m)
 - 1 Gbps dual-port (NC360m)
- Fibre Channel HBA
 - 8 Gbps dual-port FC (Q-logic)
- 3 Gbps external SAS – P700m

Supported

- OpenVMS guest, HPVM V4.2 PK1
- HP Insight Control 6.1
- vMedia, DVD (internal, USB)
- MDS600, P2000G3, MSA2000G2

Not Supported in Aug '10 release

- 8GB DIMMs
- 8 Gbps FC HBA from Emulex
- Core I/O SAS HBA mode
- Virtual Connect, Flex10
- FCoE
- vFlash, USB flash disk



BL8x0c i2 Performance Features



Performance Enhancing Factors on BL8x0c

Enhanced thread level parallelism

- Double the number of cores
- Enhanced hyper thread management

Directory-Based Cache Coherency

- 9100 series processors used “snooping” mechanism for cache coherence
 - Increase in coherency traffic with increase in number of processors
 - Contention for bandwidth
- 9300 series uses directory based cache coherency
 - Home agent for each memory controller
 - Track owners and shares of a given cache line

Performance Enhancing Factors on BL8x0c

Memory

- Memory controllers integrated with processor (2 per processor)
 - rx6600 - 1 memory controller for 4 sockets
 - BL870c i2 – 8 dual port memory controllers for 4 sockets
- DDR3 memory
- Support for larger memory configurations

QuickPath Interconnect

- Point-to-Point connections
- Higher bandwidth between processors and IO

Performance Enhancing Factors on BL8x0c

IOH (Intel E7500)

- Connects CPU to PCIe gen2 IO (core and mezz)
- Larger number of PCIe gen2 lanes increase IO throughput

Data TLB support for 8K and 16K pages

- Faster translations via the TLB


Resource Affinity Domains

- 5 standard RAD configurations
- Provides users the flexibility to define how memory layout optimize for your application

OpenVMS Performance Test Results



Price performance migration path



New Integrity Blades are more scalable than previous generation

- Double the number of cores
- More memory
- More I/Os

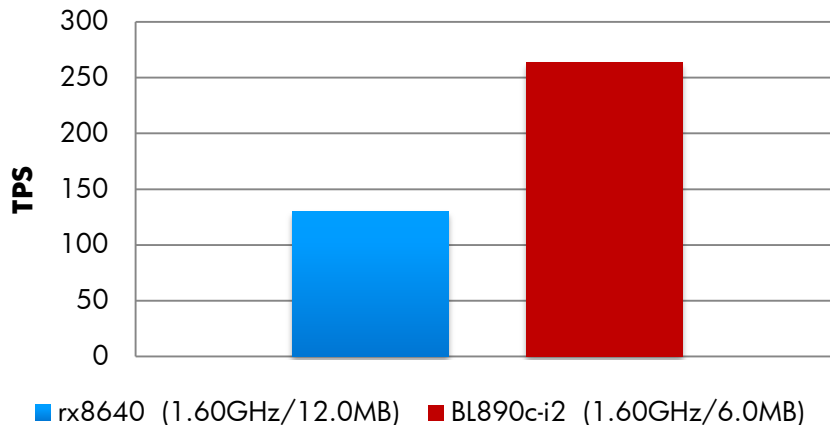
Rdb Tests

Rdb Performance Load Tests – OpenVMS V8.4

2X performance improvement with HT enabled

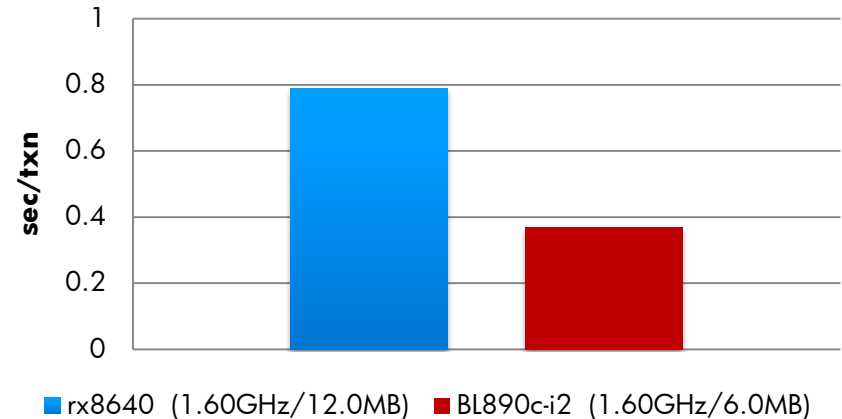
Rdb Performance

(More is Better)



Rdb Performance

(Less is Better)



- BL890c i2 performs 2x better than rx8640
 - The Transaction Per Second (TPS) is 2x
 - The time taken per transaction is reduced to half on i2 server
- Rdb takes advantage of hyper-threading on i2 servers




Apache Performance

2X performance improvement


Apache Bench Tests on OpenVMS V8.4

Bandwidth
(More is better)




■ BL860c (1.59GHz/9.0MB)
■ BL860c-i2 (1.73GHz/6.0MB)

Throughput
(More is Better)



■ BL860c (1.59GHz/9.0MB)
■ BL860c-i2 (1.73GHz/6.0MB)

Time Taken (sec)
(Less is Better)



■ BL860c (1.59GHz/9.0MB)
■ BL860c-i2 (1.73GHz/6.0MB)

- BL860c-i2 delivered 2x performance compared to BL860c




Java Workload Tests

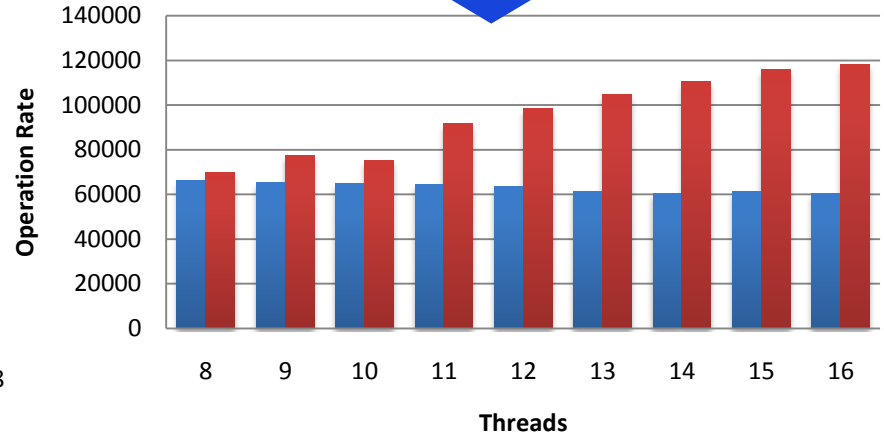
Native Java Tests on OpenVMS V8.4

870c i2 scaled up much better 2X performance improvement at higher workloads

Java Workload



Java



— rx6600 (1.59GHz/12.0MB)

— BL870c i2 (1.60GHz/5.0MB)

■ rx6600 (1.59GHz/12.0MB)

■ BL870c i2 (1.60GHz/5.0MB)


- Java Workloads scale up better on i2 Servers
- Java Workloads are high CPU and Memory Intensive



Memory tests

Throughput

Memory Bandwidth
(more is better)



Single stream test shows 55% improvement in memory bandwidth between rx3600 and BL860c i2

Single Stream Test

- The new BL8x0c i2 server demonstrated 55% improvement in single stream test
- Memory bound applications would benefit from aggregated bandwidth




CPU Ratings

Integer Test Rating

- 1.6 GHz 9300 series cores show 4% improvement over 9100/9000 cores
- 1.73 GHz 9300 series cores show 15% improvement

Ratings
(More is better)



Per core numbers


- These numbers are per Core (within a processor/socket)
- As the frequency increases, we see a increase in rating
- CPU Bound applications should benefit (database queries), specifically integer computational bound applications



CPU Ratings

Whetstone is FP Computation Tests

MWIPS Rating
(more is better)



- 1.6 GHz 9300 series cores show 8% improvement over 9100/9000 cores
- 1.73 GHz 9300 series cores show 17% improvement

Per core numbers

- These numbers are per Core (within a processor/socket)
- Fast response to complex operations; Scientific, Automation and robotic applications should benefit



Performance Tests in Progress

- IO tests
- Oracle tests



OpenVMS V8.4 Performance Improvements

- Integrity RAD Support
- Support for TCP/IP Packet Processing Engine (PPE)
 - Gains of approximately 5%
- Solid State Disks
 - Vendor benchmarks
 - Internal testing with EVA
- Clustering Software Enhancements
 - PE driver
 - Benefits deployments that use multiple channels
 - Improvement of 50% observed in some use cases
 - Dedicated Lock Manager
 - Improved Request buffer handling
 - Results in 2x improvement in some cases
 - Applications like Relational databases are benefitted
 - Shadow Driver
 - Improvement of 10-12% observed in some use cases



OpenVMS V8.4 Performance Improvements

– General Operating System Performance enhancements

- Enhanced memcmp and strcmp for Integrity systems
- Image Activation
 - Performance gains of 40-50%
 - Application performing lots of image activations gain
- VA tear down
 - Batch jobs, creating and stopping processes improve
- Global Section unmap improvements
- Exception Handling (on par with Alpha now)
- InnerMode semaphore upcalls for Exec and Kernel mode
- System service dispatch enhancements
- Pthread spinlock algorithm changes
- Dynamic enabling of XFC cache for mounted volumes
- PageDyn LALs



OpenVMS Support for RAS features on BL8x0c i2



Intel ® Itanium ® 9300-series RAS Features

Enhanced reliability, availability and serviceability (RAS) over the 9100 series

Extensive capabilities to detect, recover and report errors

RAS features for

- Processor/Socket
- Memory
- Interconnect and Miscellaneous
- IOH and Partitions



Intel® Itanium® 9300-series Processor RAS Features

Error avoidance, detection and correction across all core structures

Soft error hardened latches and registers

- Designed to improve resistance to soft errors by upto 100X

Error Correcting Code (ECC) or parity

- Widely used algorithms implemented in hardware to monitor errors that can occur during transmission



Intel ® Itanium ® 9300-series Processor RAS Features

Intel Cache Safe Technology

- Mapout bad cache lines using heuristics
- Cache Data is automatically scrubbed for single bit errors
- 9300 series covers L2 cache and directory cache as well as opposed to only L3 cache in 9100

Advanced Machine Check

- Many errors that were fatal in 9100 are now correctable
- OpenVMS logs correctable errors in the error logs

Dynamic Processor Resilience (DPR)

- Support for CPU indictment with help of WEBES



Intel ® Itanium ® 9300-series Memory RAS Features

Enhanced ECC protection

- ECC mechanism to detect and fix errors in attached memory components
- Support for Single and Double Device Data Correction

Memory Thermal Protection

- Support for Closed and Open Loop throttling mechanisms to reduce failures due to over heating of components

Memory sparing

- Monitor memory for errors
- On reaching a threshold firmware copies data from DRAM to a spare
- Transparent to the Operating System



Intel ® Itanium ® 9300-series Interconnect RAS Features

QPI Error Detection and Correction

- CRC used to detect errors
- Transactions retried
- Channel Physically reset
- Bad lanes can be mapped out (can impact performance)
- Intelligent error management
 - Single dropped packet can have a cascade effect causing whole lot of errors
 - Sort and analyze the packet to determine the source



OpenVMS Power Management Support



Power Management

Enhanced Demand Based Switching

- Ability to operate at reduced voltage and frequency until more power is needed
- Can run different P-States
 - Predefined voltage and frequency combinations at which a processor can operate correctly

Intel Turbo Boost

- Processor can run at frequencies higher the advertised frequency
- Provided processor package is below rated power, thermal and current limits
- If you have an idle core the active core can use its headroom to get a boost in frequency



OpenVMS Power Management

Highest Performance Mode

- Use Turboboost to maximize performance
- No power savings measures

Dynamic Power Mode

- Use Turboboost when process is executing
- Use C-states (halt states) to save power when a CPU is idle

Low Power Mode

- Use p-states to reduce power usage when process is executing
- Use C-states to save power when a CPU is idle

OS Control Mode

- Enable a system service or sysgen parameter to select among power mode



Q&A



Business Manager (Vivasvan Shastri)

viva@hp.com

Office of Customer Programs

OpenVMS.Programs@hp.com



References and Acknowledgements

- Intel Whitepapers

<http://download.intel.com/products/processor/itanium/323247.pdf>

<http://download.intel.com/products/processor/itanium/318691.pdf>

http://software.intel.com/sites/oss/pdfs/power_mgmt_intel_arch_servers.pdf

<http://www.intel.com/design/itanium/documentation.htm>

- WEBES Documentation

http://h18023.www1.hp.com/support/svctools/webes/index.html?jumpid=reg_R1002_USEN

- HP Whitepapers

[OpenVMS Technical Journal 14 article on OpenVMS Power Management by Prithvi Srihari, Burns Fisher and Veena K](#)



Thank You

