
VSI OpenVMS on -i4 servers

Why you should upgrade to VSI OpenVMS on
the HP Integrity -i4 servers


Colin Butcher
XDelta Limited

www.xdelta.co.uk
+44 117 904 8209

XDelta – who we are

- VSI Professional Services Alliance (Europe)
- Independent consultants since 1996:
 - UK based with international reach
 - Over 30 years experience with OpenVMS
- We design and implement solutions:
 - Mission critical systems
 - Cross-sector experience
 - Engineering background
- Gartner (2009):
 - Identified XDelta as one of few companies world-wide capable of OpenVMS migration related projects

Business Partner



HP Integrity -i4 servers - highlights

- “Poulson” 2.53GHz 8 core processor with shared L3 cache
- Re-architected CPU design
 - Around 30% per core greater throughput
 - Reduced NUMA effects for same core count
 - Better memory latency and bandwidth
 - Improved floating point and integer performance
 - Improved hyperthreading – less stall time
- bl870c-i4 (32 cores) about 1.3x better than bl890c-i2

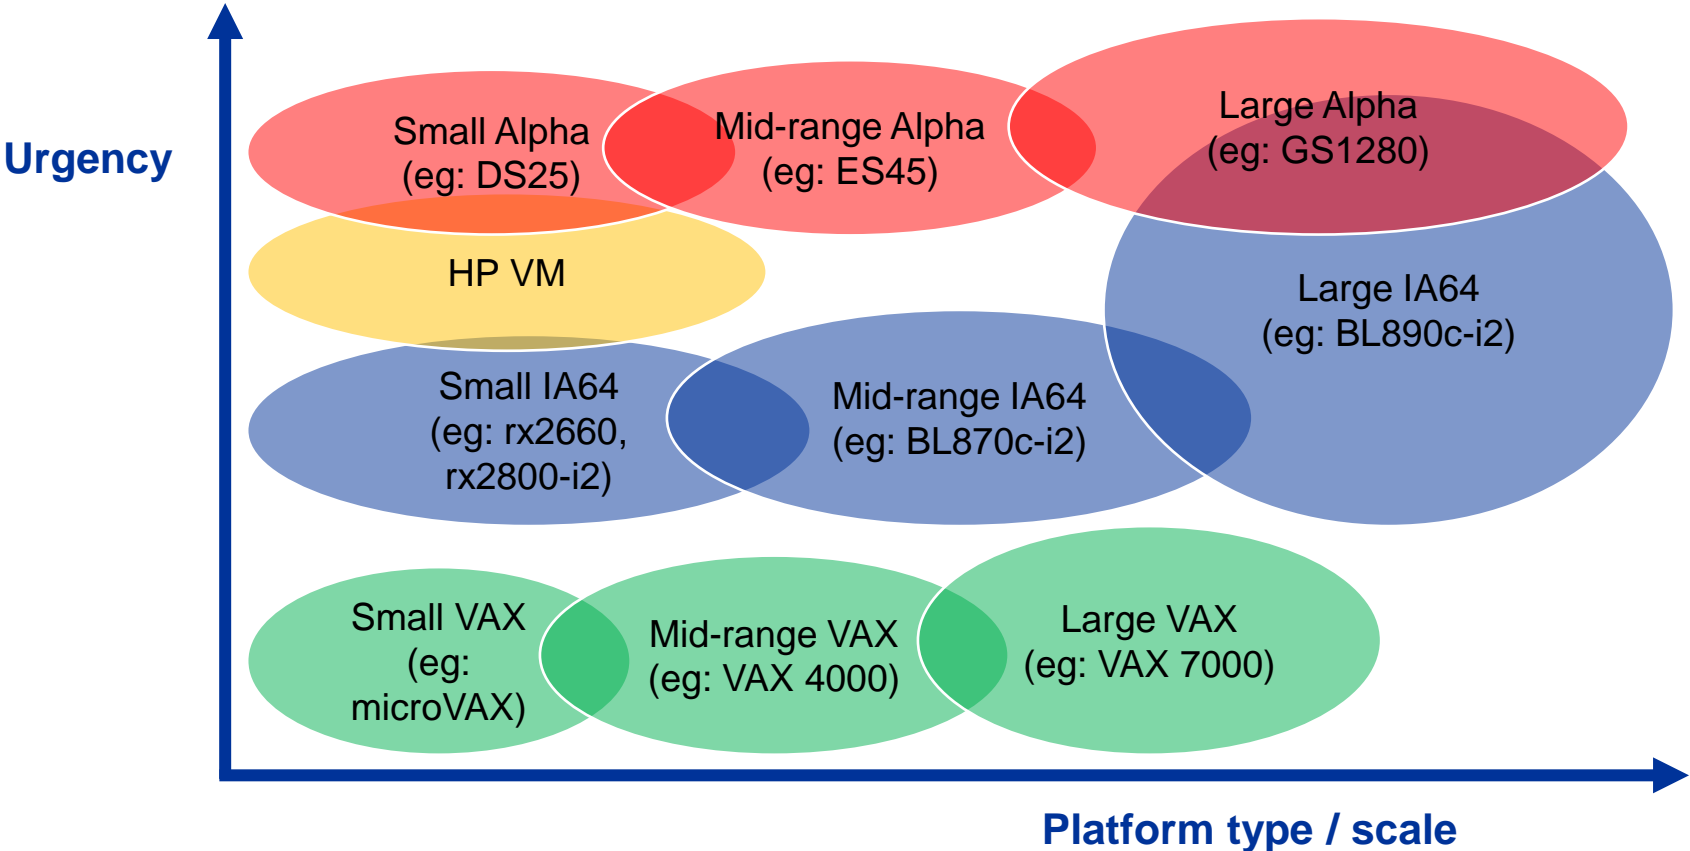
HP Integrity -i4 servers – hardware

- **bl860c-i4:**
single width, 16 cores, 384GB, 4x 10GigE, 3x mezz, 1c2d SAS
- **bl870c-i4:**
double width, 32 cores, 768GB, 8x 10GigE, 6x mezz, 2c4d SAS
- **bl890c-i4:**
quad width, 64 cores, 1.5TB, 16x 10GigE, 12x mezz, 4c8d SAS
 - OpenVMS V8.4-1H1 currently supports a maximum of 32 cores
 - OpenVMS V8.4-1H1 also supports nPARs
- **rx2800-i4:**
2U rack, 16 cores, 384GB, 4x 1GigE, 6x PCIe, 1c8d SAS

Migrating from Alpha to Integrity

- Multi-core processors, NUMA, hyperthreading
- 10GigE network
- 8GigFC SAN
- Blade chassis connectivity for bl8x0c-i4
- EVA to 3PAR storage migration
- bl870c-i4 and bl890c-i4: good for GS1280 migration

Migration to Integrity and beyond



Server hardware differences (-i2 to -i4)

- Re-architected CPU, not just an updated design
- Higher clock rate (2.53GHz v 1.73GHz)
- “Out of order” instruction execution
- Minimal stall time between co-threads with hyperthreading
- Double the core count (8 cores)
- Greater memory capacity
- Reduced memory latency
- Shared on-chip cache

- New 10GigE LoM - LAN only, not FCoE
- Same 8GigFC mezzanine cards

Chassis hardware – c7000 / c3000

- Virtual Connect (GigE, 1/10GigE, 8GigFC)
- Flex10
- LAN side of FlexFabric

- 10GigE chassis based switching
- 10GigE passthrough
- 1GigE passthrough

- 8GigFC chassis switching
- 4GigFC passthrough

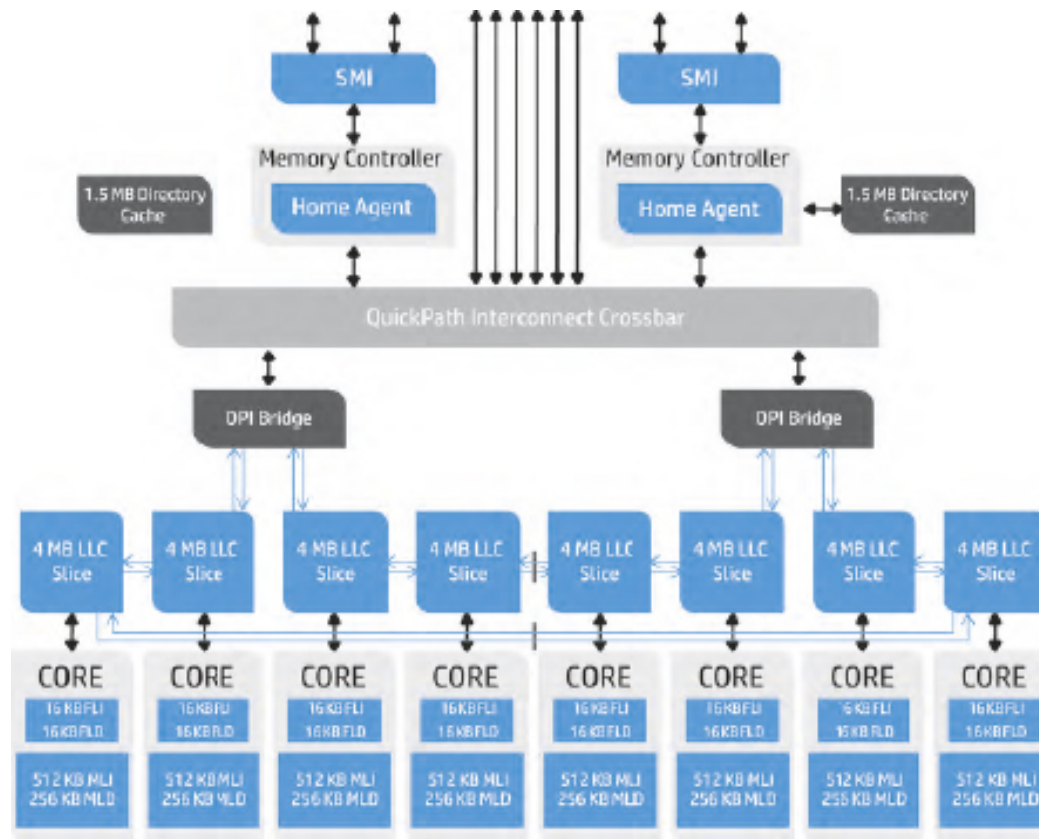
Infrastructure hardware

- 3PAR storage arrays at 8GigFC
- SSD devices for local storage and 3PAR storage arrays
- 8GigFC SAN – HP / Brocade switches
- 10GigE networking – HP Procurve, Cisco

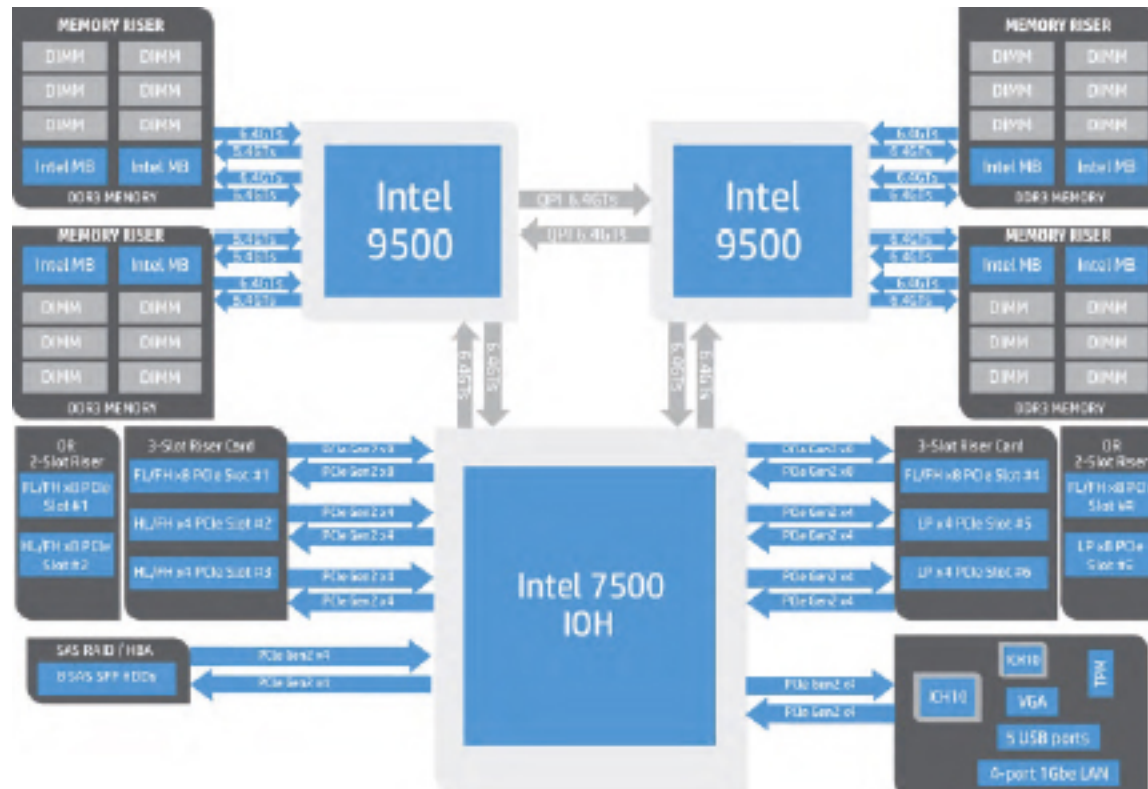
OpenVMS V8.4-1H1 on -i4 servers

- Complete build of base system from sources
- -i4 hardware support (32 cores supported, threads off)
Note: V8.4-2 will support 64 cores, threads off
- New LAN driver for LoM support
- VSI branding
- Patch kits available via HP

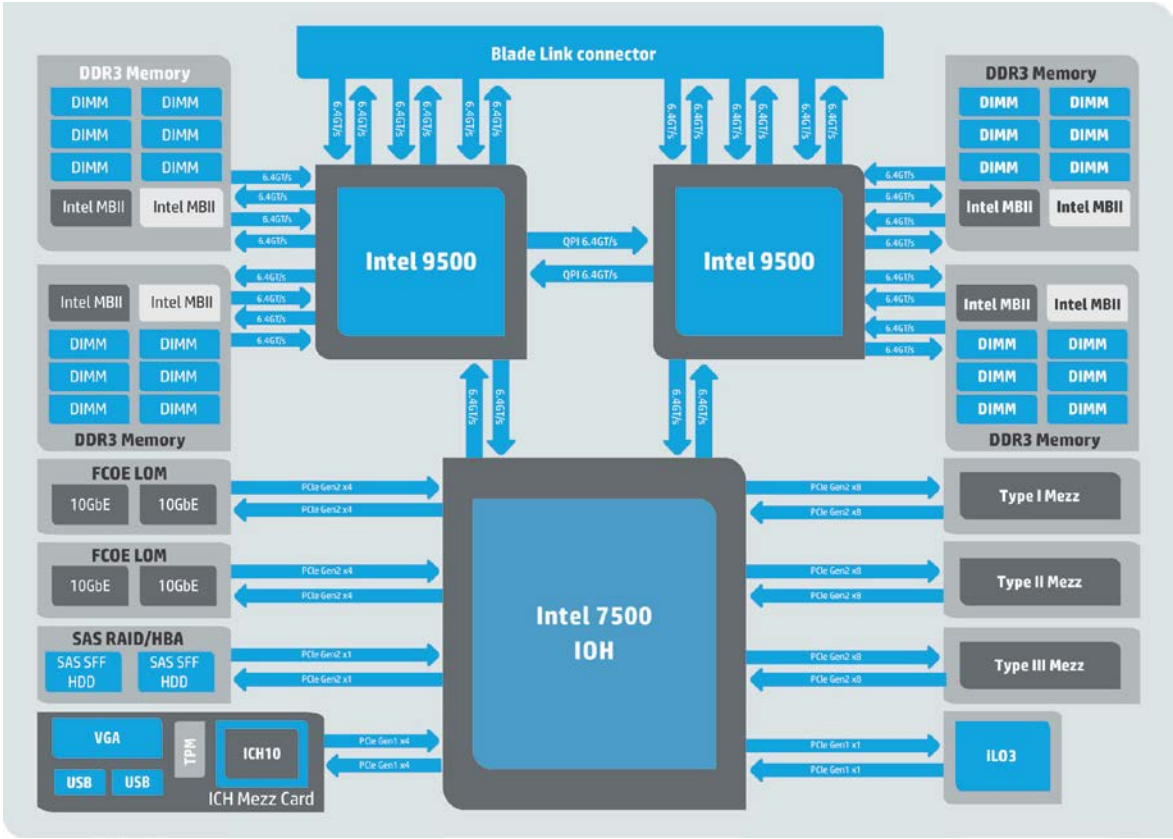
CPU architecture - Intel 9500 – “Poulson”




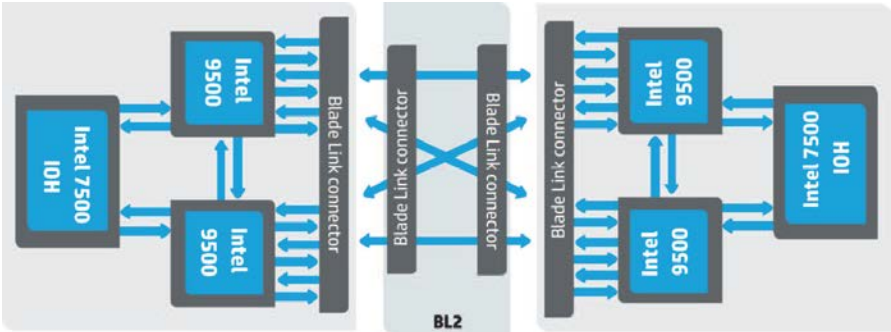
System architecture – rx2800-i4



Blade architecture – bl8x0c-i4



QPI fabric – bl870c-i4 and bl890c-i4



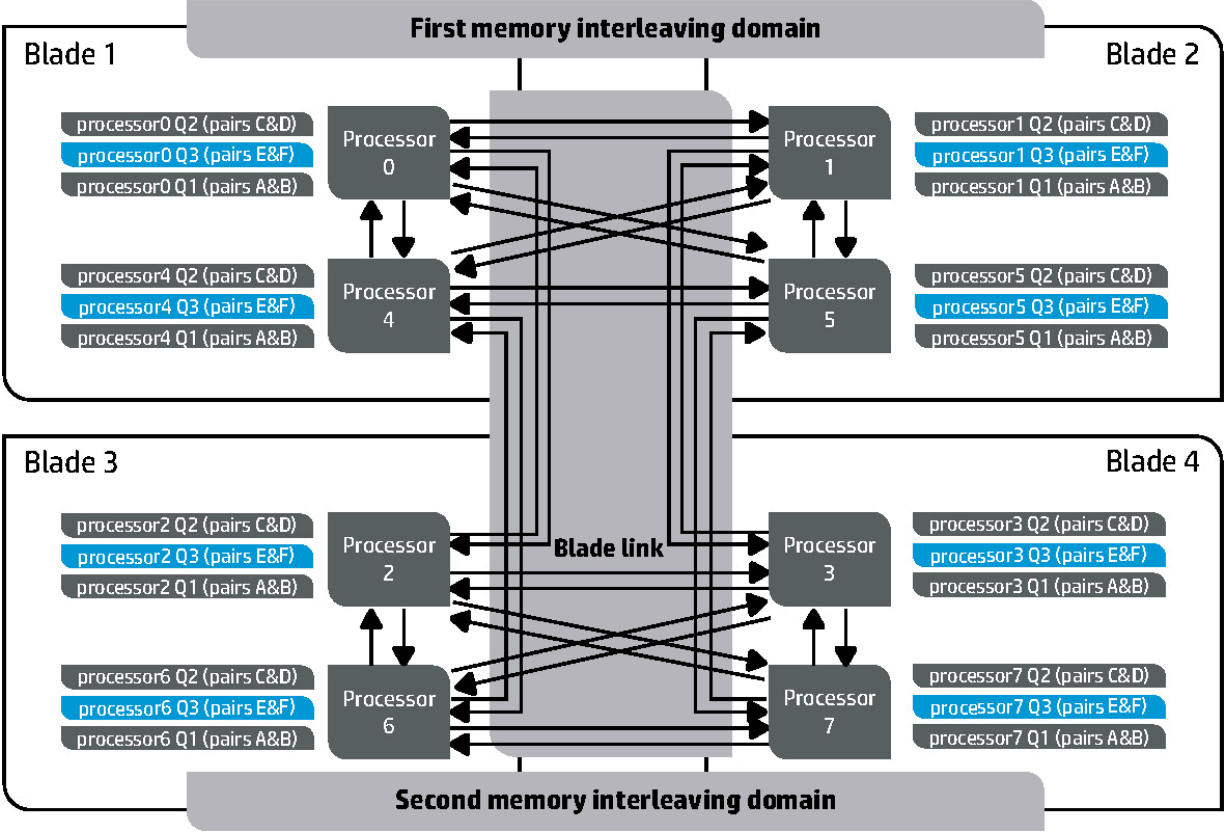
High core count

- CPU 00 is the primary CPU – try to reduce its workload
- Fastpath CPU selection – be aware of physical layout
- CPU choice for dedicated lock manager
- CPU choice for TCPIP packet processing engine
- Consider physical layout - RADs and NUMA

Hypertexting

- Hypertexting is extremely workload dependent
- In general the OpenVMS scheduler does a better job
- Enable / disable hypertexts and reboot
- “CPU” count will appear to double when enabled
 - V8.4-1H1 supports 32 CPUs (or “scheduling units”).
 - V8.4-2 will support 64 CPUs

Memory architecture – bl890c-i4




NUMA (non-uniform memory access)

- OpenVMS uses large shared memory regions:
 - XFC (50% available memory by default)
 - RMS global buffers
 - Global sections (especially database caches)
 - Memory disc driver (MD devices)
- Useful starting point for OpenVMS is “mostly UMA”

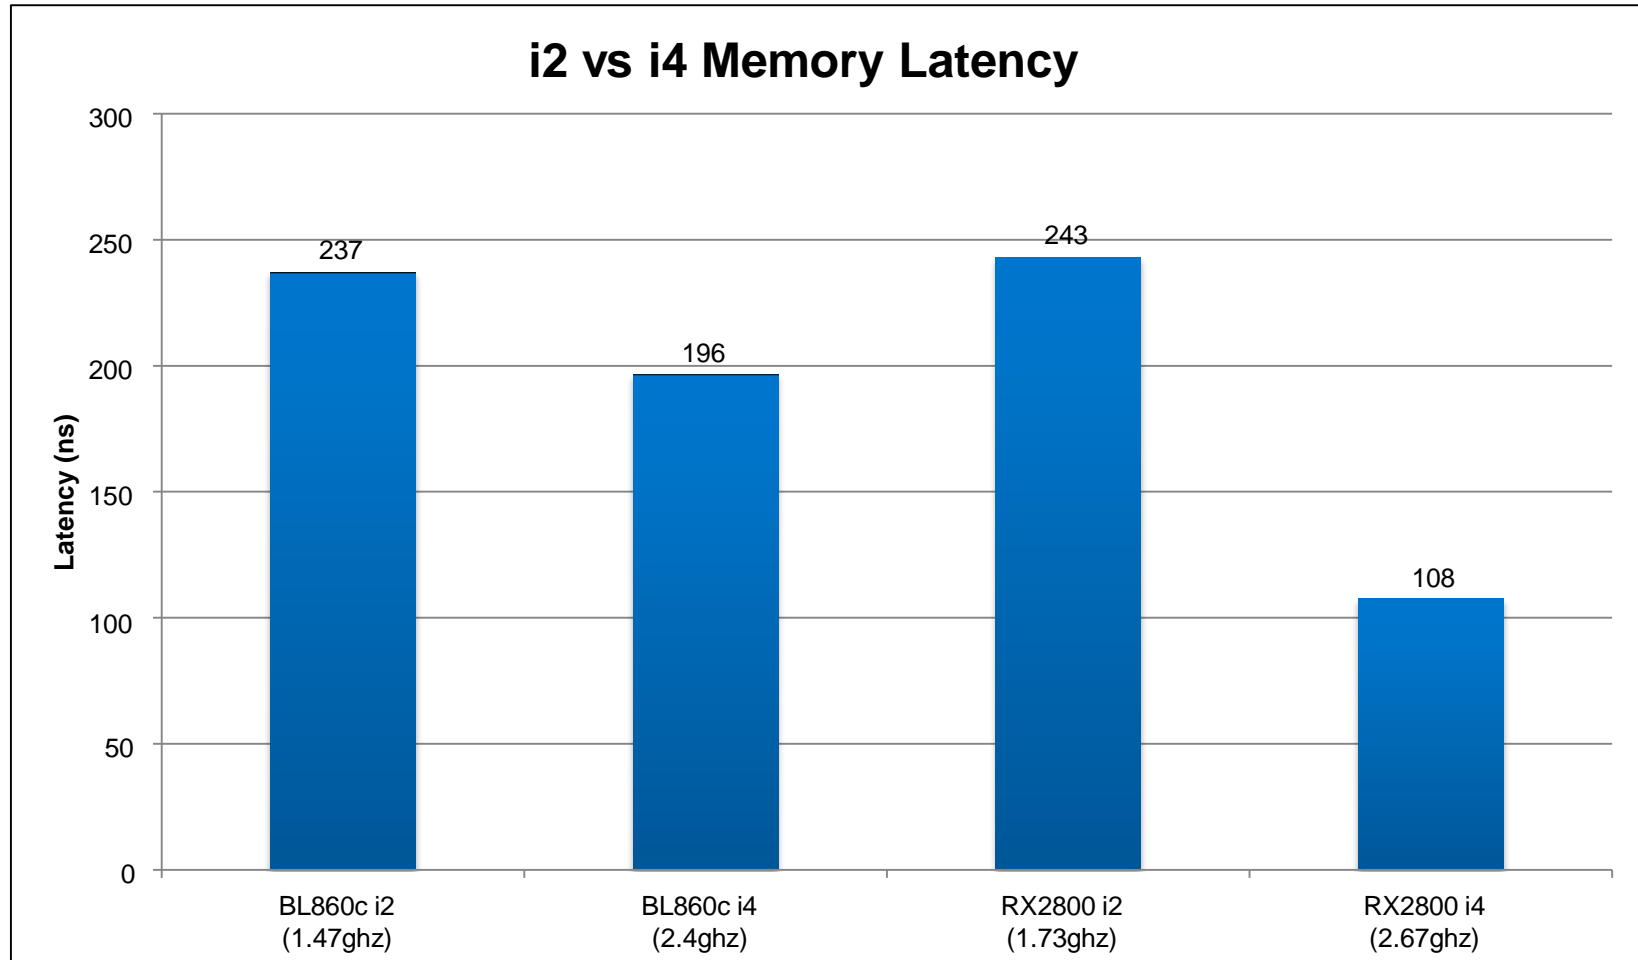
Preliminary Performance Results: Integrity -i2 vs. -i4

- The following slides contain preliminary data on performance differences between selected i2 and i4 servers running OpenVMS E8.4-1H1.
- The data was generated from VSI-written programs used to measure certain aspects of system performance.
- The results shown here should not be used as a general characterization of overall system performance or as an indication of how any specific application may perform.

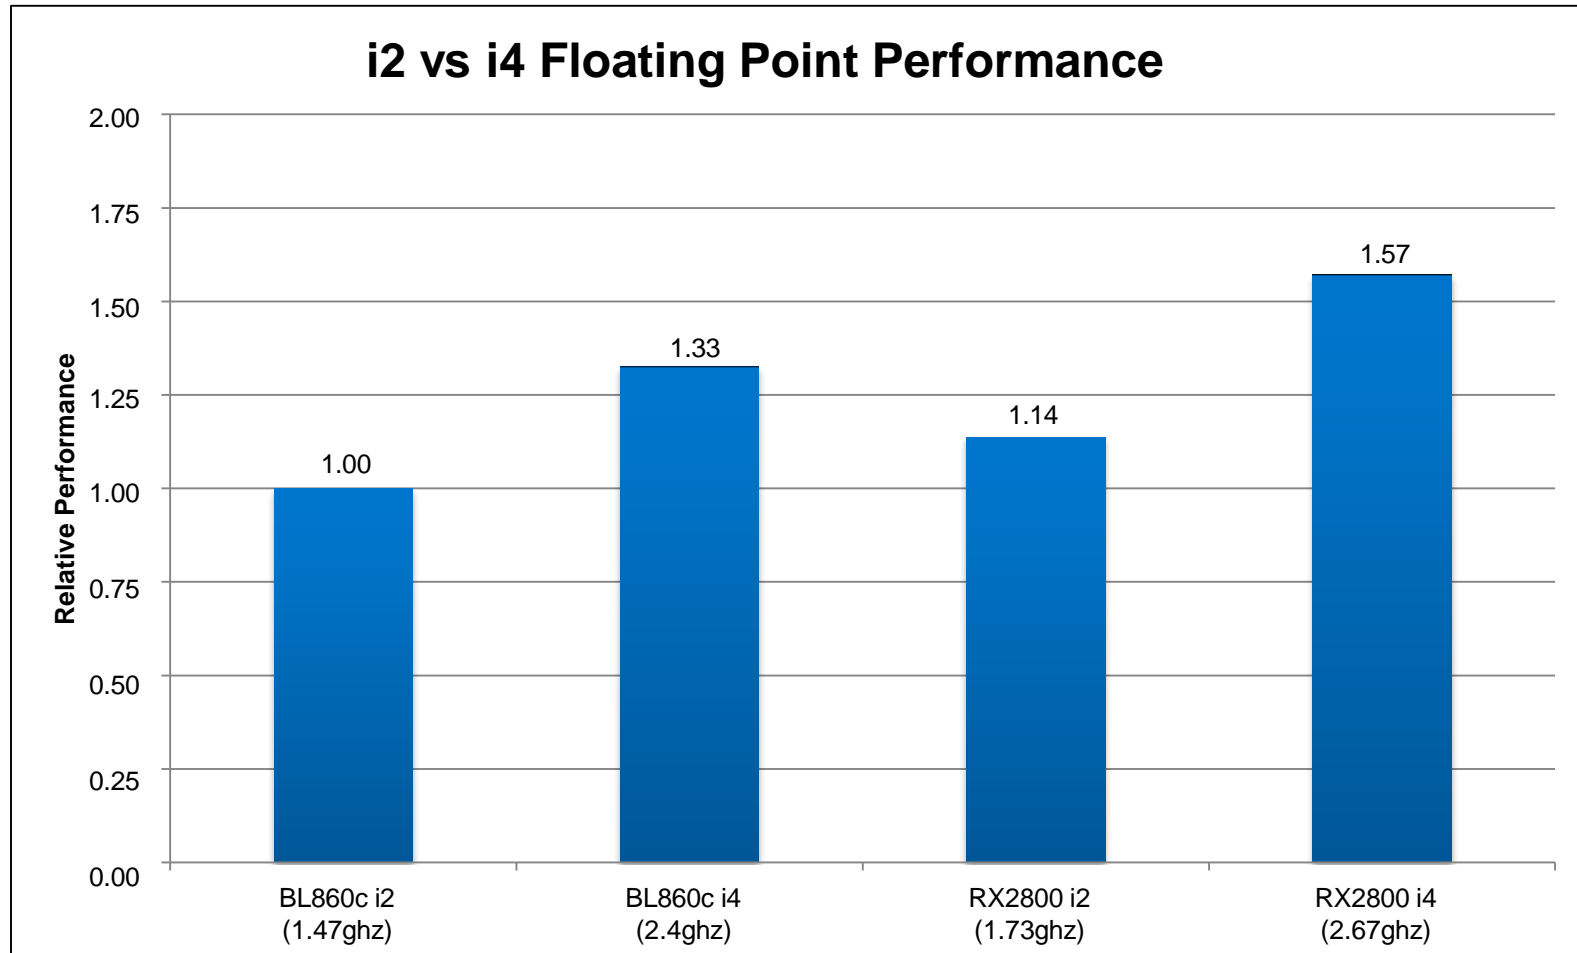
i2 vs. i4 Memory Bandwidth



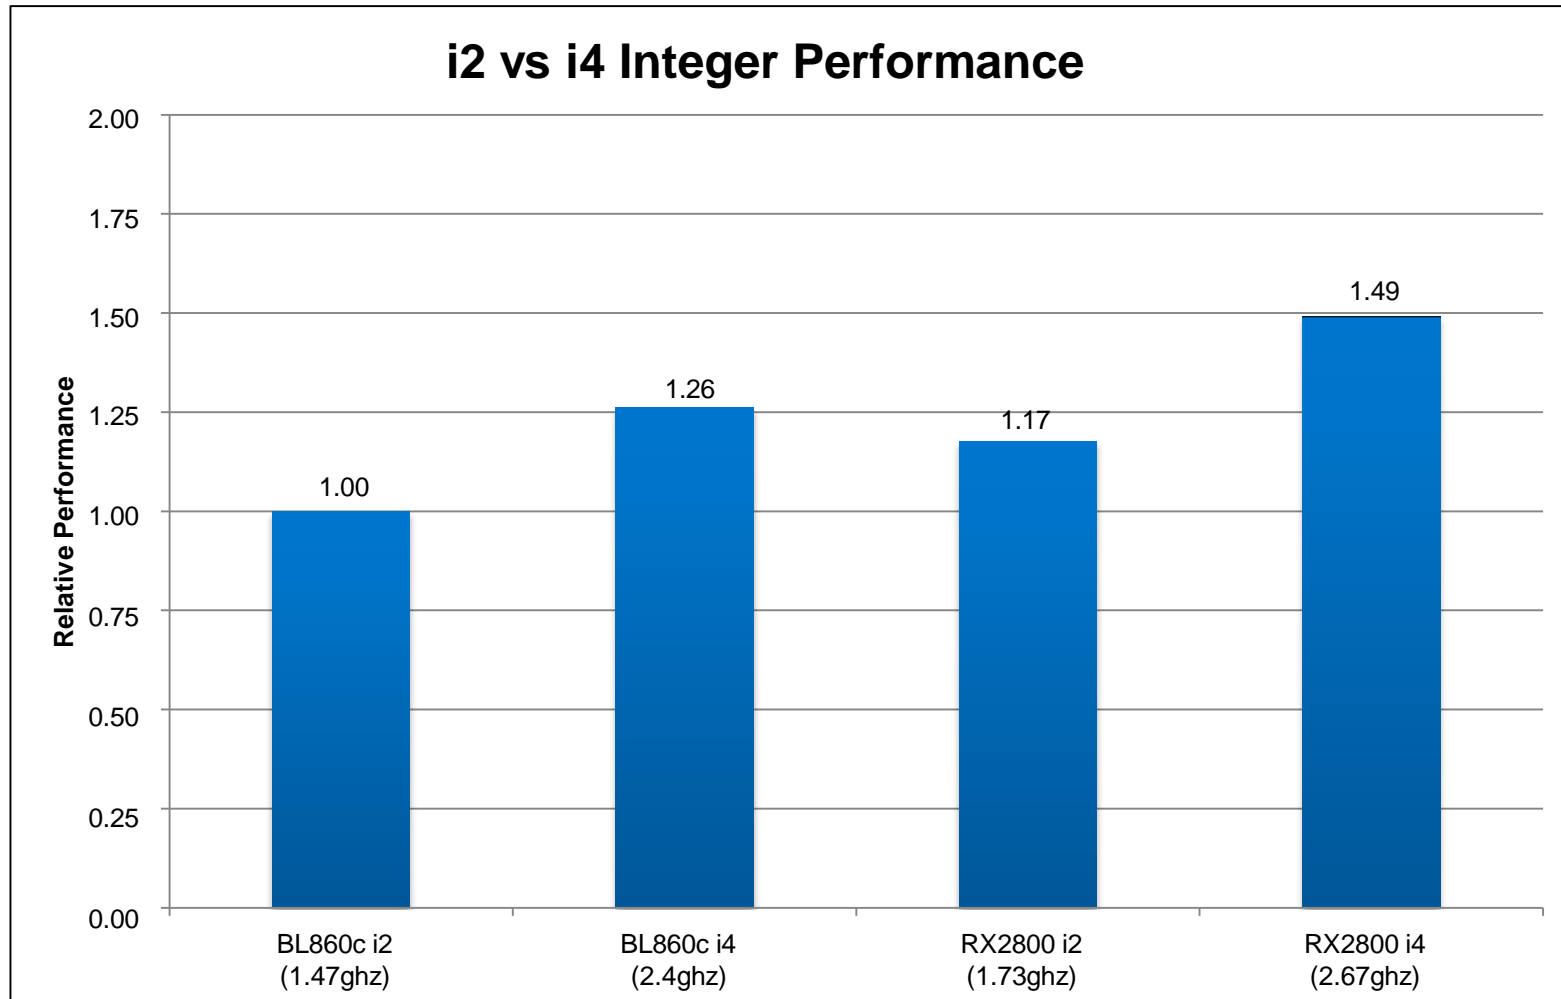
i2 vs. i4 Memory Latency



i2 vs. i4 Floating Point Performance



i2 vs. i4 Integer Performance



Performance engineering – use T4

- Avoid guesswork - run T4 all the time
- Without good data you cannot do good performance work
- A faster machine just waits more quickly
- Don't make it go faster, stop it going slower
- The fastest IO is the IO you don't do
- The fastest code is the code you don't execute

Summary - VMS V8.4-1H1 on -i4 servers

- Disable devices you don't use
SYSMAN IO SET EXCLUDE=(EWC,EWD,...)
- Experiment with memory interleave setting
- Use memory reservations
- Fastpath settings for device types
- Dedicated CPU for TCPIP + LCKMGR
- Experiment with hyperthreading

Benefits - VMS V8.4-1H1 on -i4 servers

- Significant step up from AlphaServer GS1280, with modern storage and network infrastructure
- Significant step up from HP Integrity -i2 blade servers
- rx2800-i4 delivers more than 2x rx2800-i2 in same space
- New VSI OpenVMS releases will bring new features and prepare the way for transition to OpenVMS on x86

OpenVMS Rolling Roadmap

February 2016

2016 / 2017

2018

OpenVMS V8.4-2

HP Integrity System Support

- 64 Cores (CPU hyperthreads off)
- UEFI 2.3
- 3PAR 8200 Flash Array

VSI TCP/IP 1.0

- Installation Option
- **HP TCP/IP will not be supported in a future release**

Performance

- More alignment faults eliminated
- Improved network throughput with VSI TCP/IP

Software Component Updates

- ACME-enabled LOGIN
- Enterprise Directory
- HSM
- Apache 2.4.12
- LDDRIVER 9.7

OpenVMS V8.x,.....

HP Integrity System Support

- 16Gb Fibre Channel
- HP Kittson-based servers
- Additional device support, depending on customer feedback

Software

- JAVA 1.8
- Large Disks (64b LBNs)
- GNV & CRTL (Application Portability)
- OMNI / OSAP
- gSOAP, Samba, PHP, git, OpenSSL

New File System

- In addition to ODS-2 and ODS-5
- Eliminate 2TB volume size limit
- Faster create/delete
- Directory scalability
- 4 billion files per volume
- Fast crash recovery

OpenVMS V9.0

Itanium & x86-64

- Network stack performance
- Containers
- OpenJDK

Itanium

Additional servers & I/O, depending on customer feedback

x86-64

- Selected HP Servers (Intel and AMD)
- OpenVMS as a virtual machine guest
- Dynamic Binary Translator (Alpha/Itanium)
- Updated Language Standards
 - C
 - C++
 - FORTRAN

Multiple releases may occur between V8.4-2 and V9.0. The order in which various improvements are added to these releases will be determined by readiness, hardware availability, and customer feedback.

VSI OpenVMS on -i4 servers

Why you should upgrade to VSI OpenVMS on
the HP Integrity -i4 servers

Colin Butcher
XDelta Limited

www.xdelta.co.uk
+44 117 904 8209