

HP Integrity NonStop Hardware and Software

VNUG, May 2010

Mittal Parekh

WW Product Manager, Multiple Product Lines
NonStop Enterprise Division



Agenda

1. HP Integrity NonStop Multi-core Hardware
2. HP Integrity NonStop Multi-core Software
3. Volume Level Encryption
4. Summary



We heard you.....and you.....



Find ways to be more efficient, so that even with less money you can still deliver the strategy

Now is the time to improve the efficiency of the IT system itself

„DO MORE WITH LESS..

³
3



We're holding off on replacements, but spending wisely where we're investing for the future...

Costs are under pressure, but we must remain competitive



NonStop customers asked for

Deliver 24 x 7 availability	Handle massive scalability
<ul style="list-style-type: none">• Minimize both planned and unplanned outages• Drive recovery time to near-zero• End-to-end availability• Instill a culture of 24 x 7 support	<ul style="list-style-type: none">• Handle the largest workloads• Scale without planned outage• Scalability of multiple dimensions—processors, database, and software
Drive to standards-based computing	Provide longevity of support
<ul style="list-style-type: none">• Lower cost hardware by leveraging “volume economics”• Modern software interfaces• Service Oriented Architectures	<ul style="list-style-type: none">• Provide product support and upgrade capability over decades• Maximize continuity and consistency
<i>“Give us bigger, better, faster...cheaper NonStop platform”</i>	

NonStop: Investing for the future

Modernize

Standardize

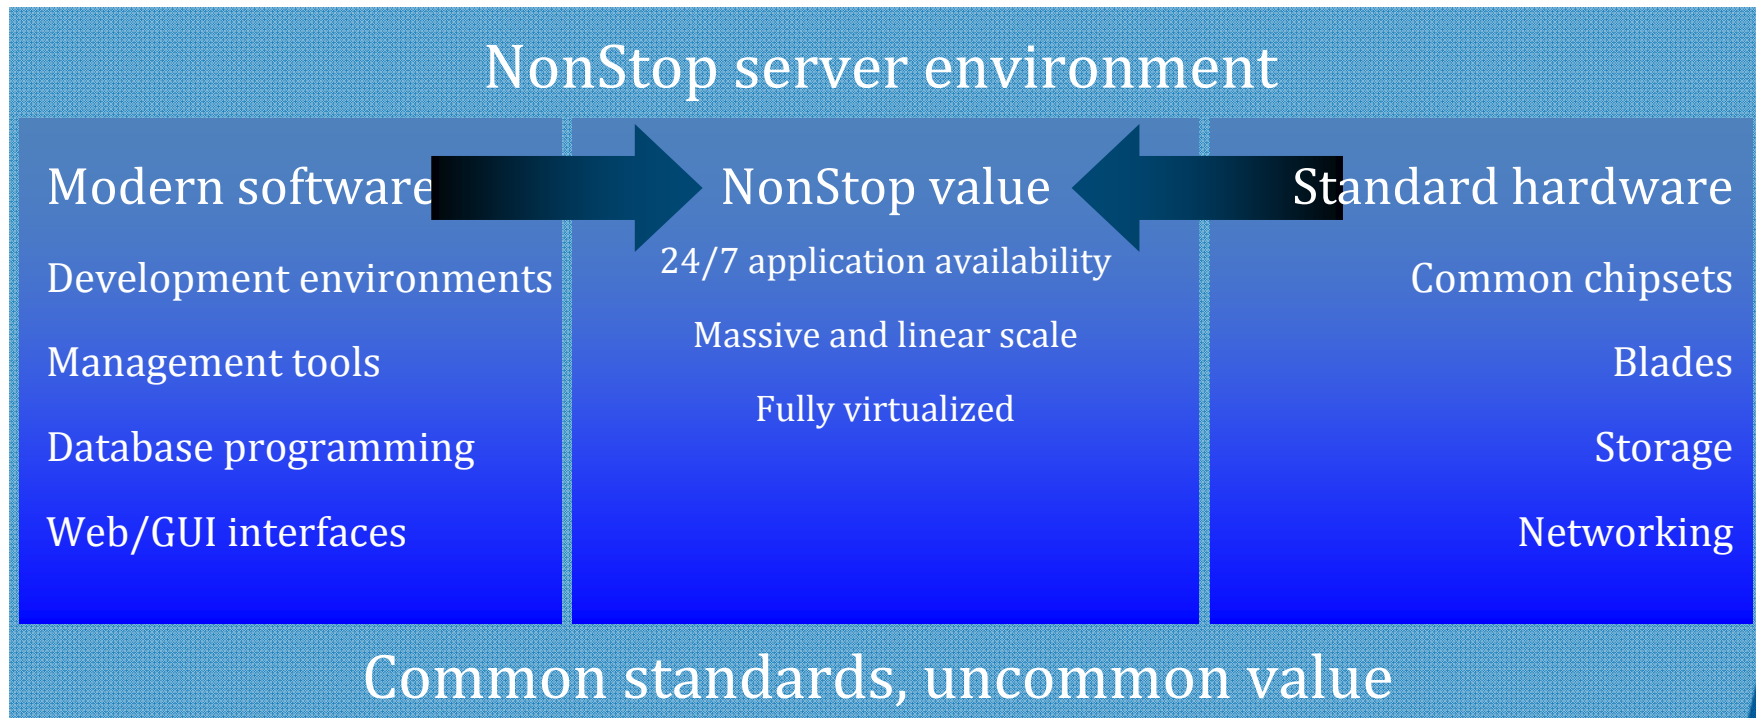


Customers tell us they want...


Modern applications
built using modern tools
running on standard platforms
with 24/7 reliability



The new NonStop... modern, standard, and 24/7



NonStop and HP's Converged Infrastructure



Virtualized • Resilient • Orchestrated • Optimized • Modular



HP Integrity NonStop BladeSystem

First-ever 24/7 mission critical
computing system built with
bladed modularity and standards

Double the performance

Half the footprint

100% NonStop



Half the footprint... Double the performance

Integrity NonStop

8 CPUs/performance = 1x



Driving efficiency via:

- Multi-core blades
- SAS storage
- Standard I/O
- Integrated ServerNet
- Integrated management

Integrity NonStop BladeSystem

8 CPUs/performance = 2x



Delivering:

Higher
performance

Higher
density

Lower cost



NS2000

The new entry-level Platform

- NSMA/J-series RVU only
- Intel's Itanium Dual-core Montvale processor
- Support new I/O Infrastructure
- Rack-mount form-factor
- Target markets
 - Development, test platform for NB50000c
 - Small stand-alone applications
 - Emerging markets



The NonStop standardization journey

An Overview



NonStop S-series

A Proprietary Design with

- Custom Rack
- Custom Power & Cooling
- Custom proprietary CPU with internally designed components
- Custom memory
- Custom IO and interconnect
- Non-Standard Disks
- ServerNet switches

Integrity NonStop

Moving to Standards with

- Standard HP Rack
- Standard Power & Cooling
- Standard BCS Server with modifications for FT
- Standard DIMMs
- Custom IO and Interconnect
- Off the Shelf Disks
- ServerNet switches

Integrity NonStop BladeSystem


Even More Standardization

- Standard ISS Chassis and Rack
- Standard Power & Cooling
- Standard Blade with unique interconnect mezzanine card
- Standard DIMMs
- Standard IO
- Off the Shelf Disks
- Only NonStop-unique HW is ServerNet



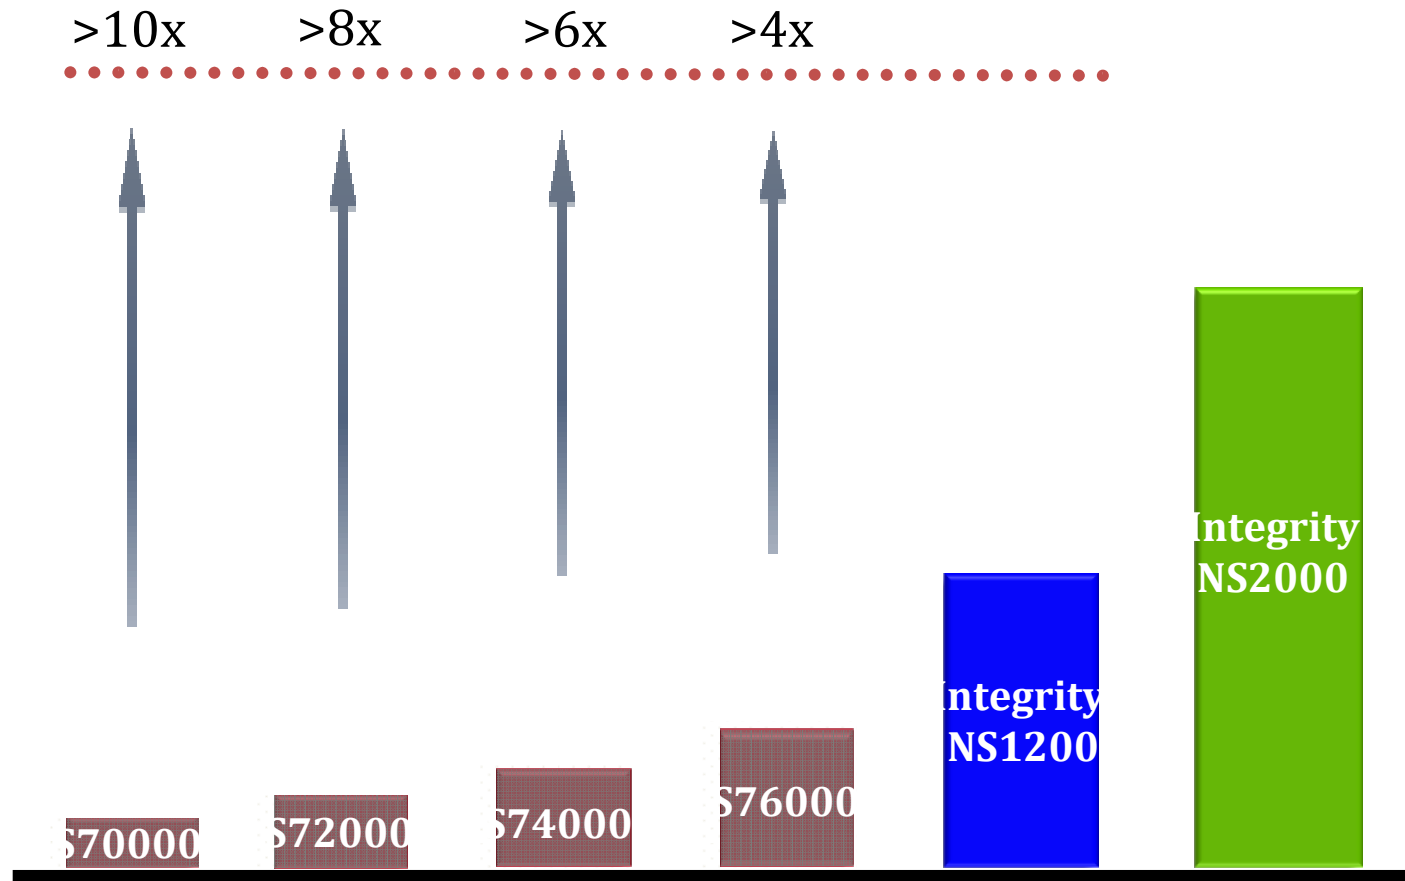
The NonStop standardization journey

Performance with RAS



The NonStop standardization journey

Modern and Affordable NonStop



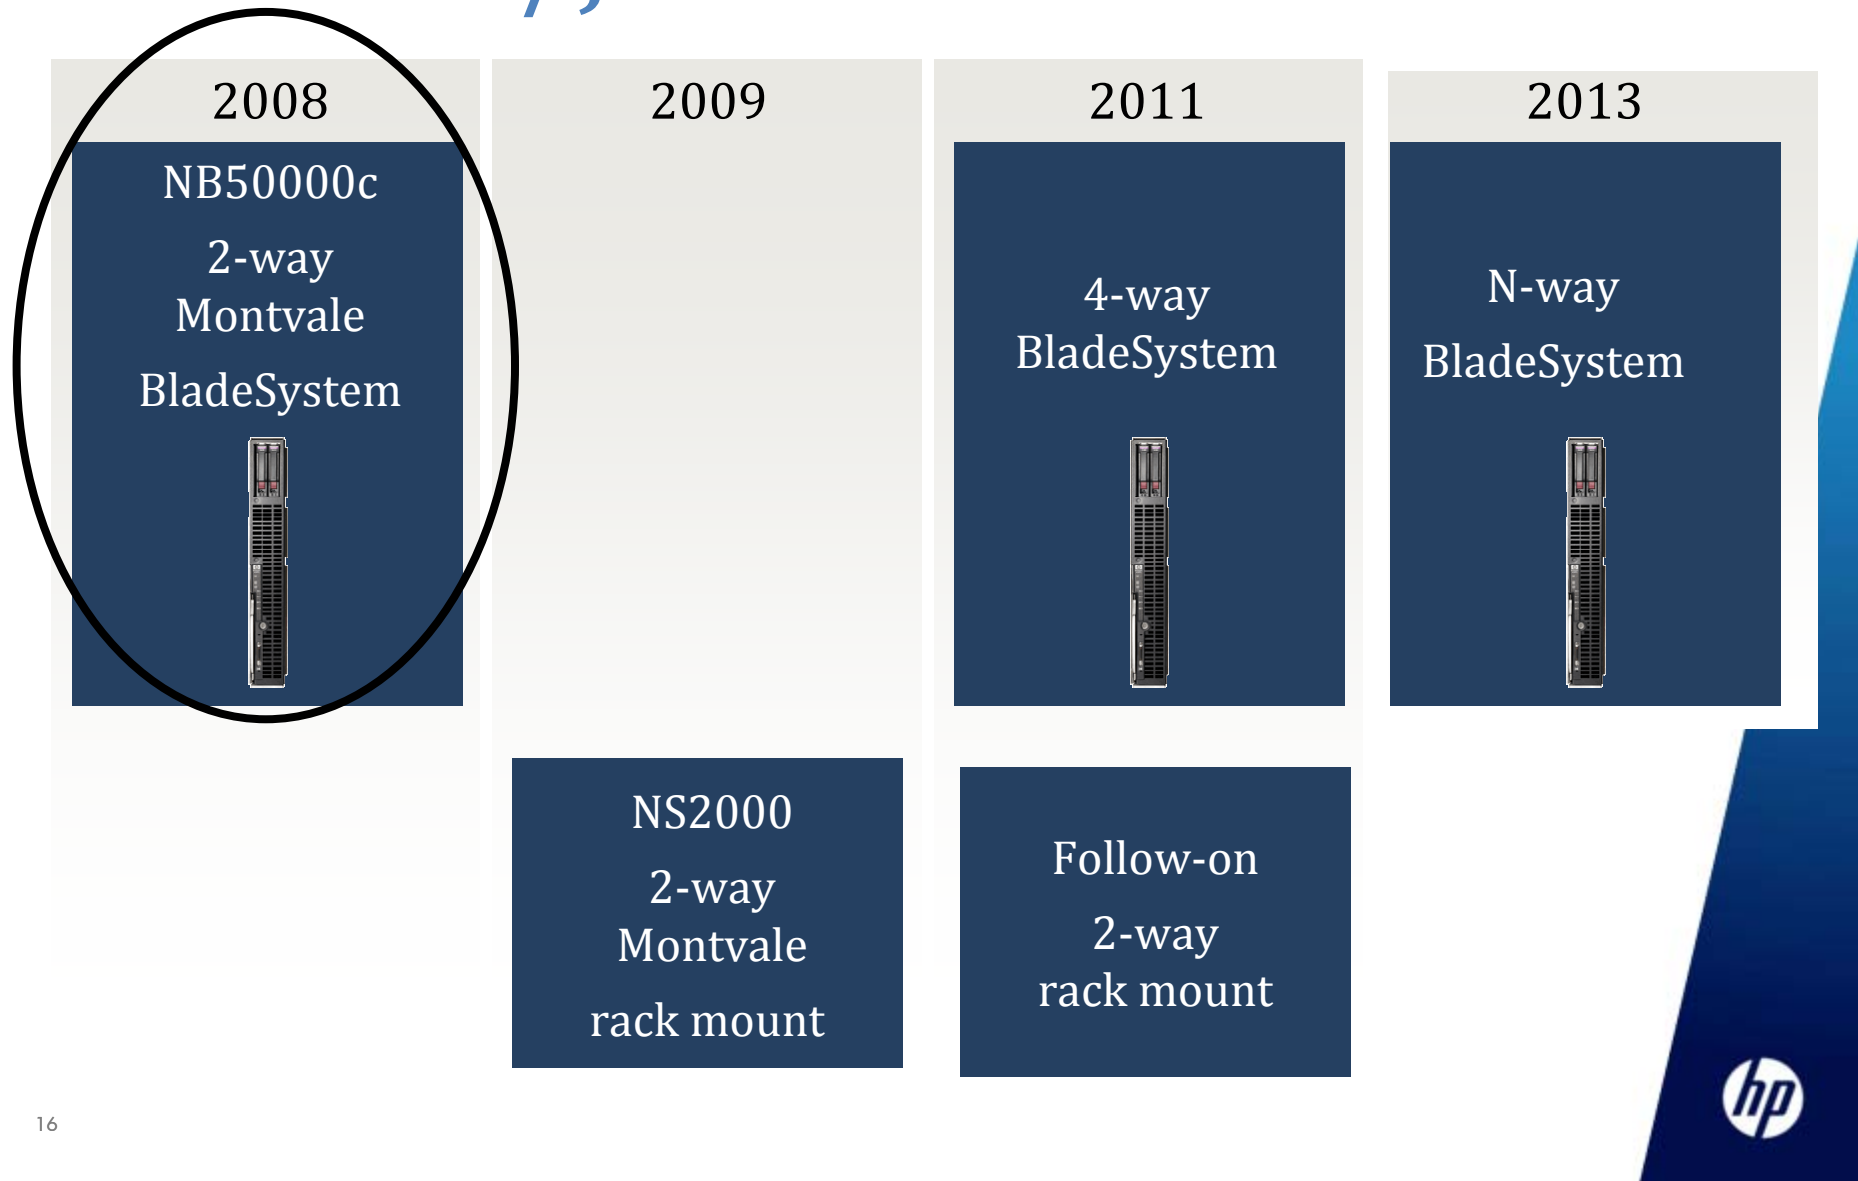
100% NonStop

- Always available
 - 24/7 continuous availability
 - Fault-tolerant NonStop OS
 - Fully-integrated fault-tolerant software stack
- Massively scalable
 - Scale-up in addition to Scale-out
 - Linear scalability
 - High-speed ServerNet clustering
- Complete investment protection
 - 100% software compatible
 - Seamless clustering with prior systems
 - Supports existing I/O infrastructure



HP Integrity NonStop roadmap

Multi-core/J-Series



NonStop BladeSystem

System configuration overview

Blade chassis

- c-Class enclosure
- ServerNet double-wide switch modules
- Ethernet single-wide switch modules (maintenance connections)

Logical processors/blades

- Two to eight blades per chassis, each with:
 - One 1.66 GHz dual core Montvale processor (one logical CPU)
 - ServerNet Mezzanine card
 - 8, 16, 24, 32, 48 GB main memory per logical CPU


Input/output

- Networking CLIM
 - Five GBit Ethernet ports (five copper or three copper/two fibre)
- Storage CLIM
 - 2 SAS HBA default; 2 additional: SAS/FC HBA choices
 - SAS enclosure: Hosts up to 25 disks
SAS 146 GB @ 15K drives, 300 GB @ 10K drives
 - XP connection option; FC tape option
- IOAME is supported, S-series I/O for traditional TDM based SS7 only



NonStop BladeSystem enclosure


- NonStop ServerNet switches
- The NONSTOP midplane
- Pre-integrated, connected
- Pre-populated with 10 fans and six power supplies (2250 watts each)
- Facilitates power and cooling for all the server blades
- Drastic reduction in cabling
- Two varieties based on power type
 - North America/Japan
 - International
- One enclosure max per 42U rack



10U high enclosure holds up to eight logical processors

NonStop BladeSystem BL50000c configuration — 8 processor


Server blades



Front


 NonStop value add

Switch modules



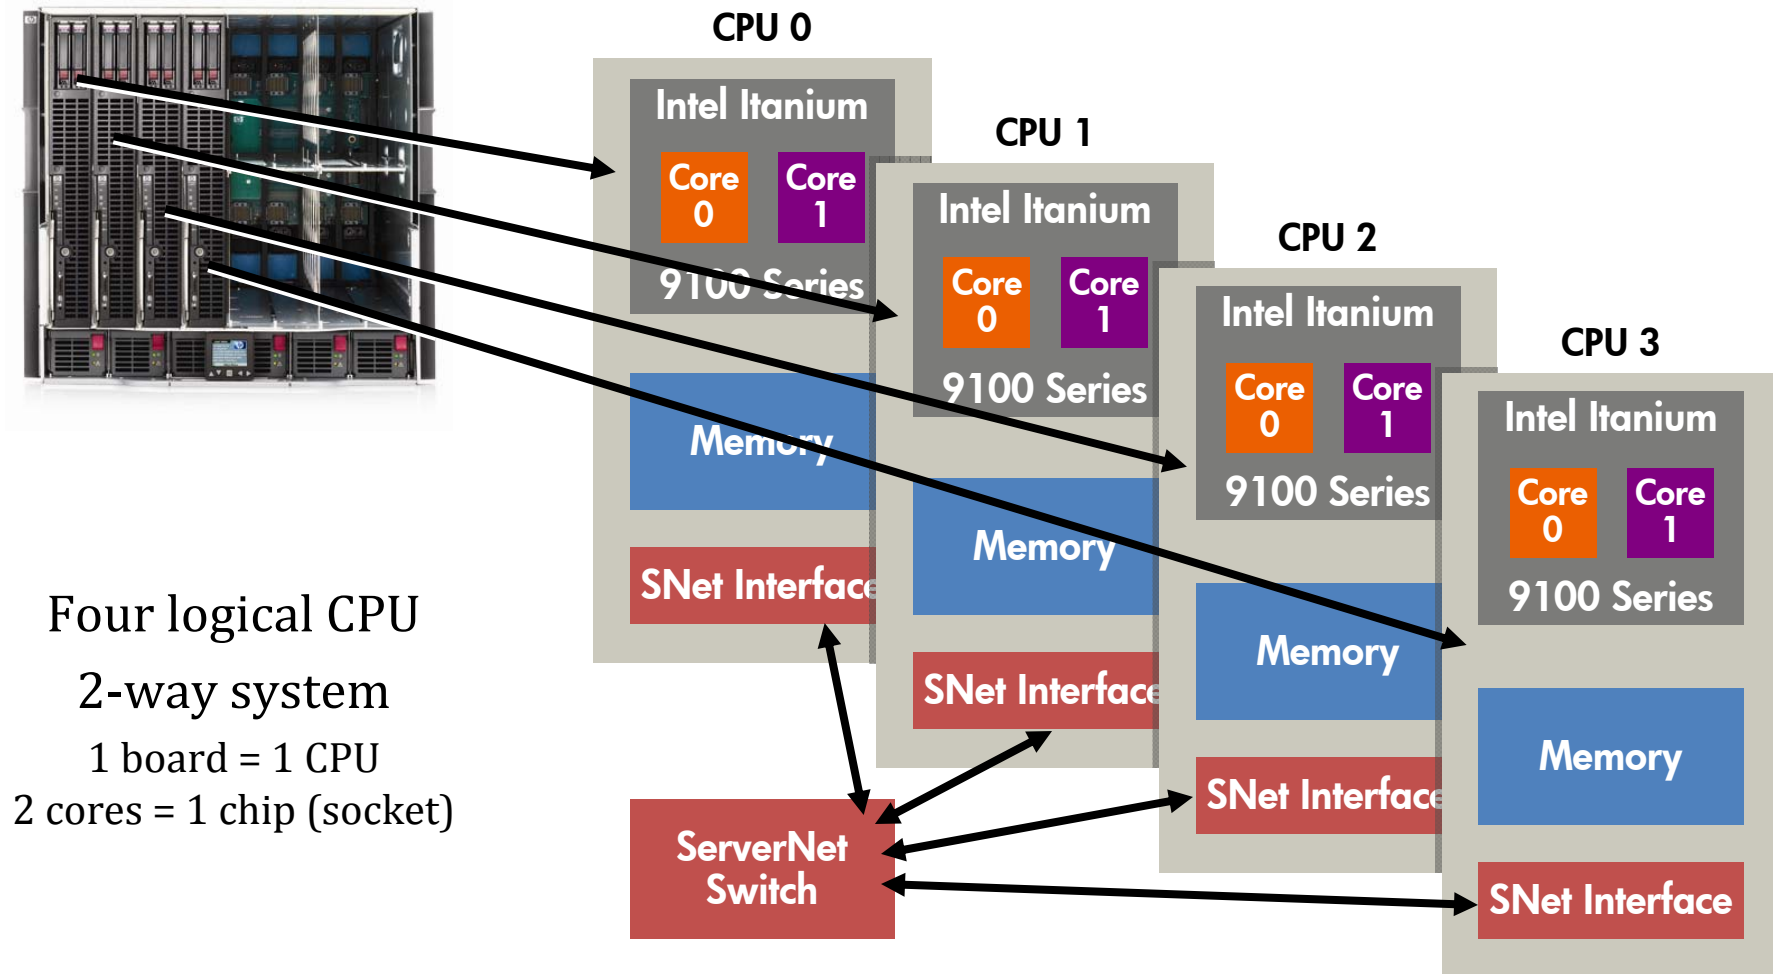
Rear

Why NonStop deploys all 10 fans?



More fans, fans required to spin slower to move the same volume of air, so each fan uses less power. Eight fans are almost always more power-efficient than four fans. As the air flow rate increases, ten fans are even more efficient (Figure 8). Slower spinning fans also create less noise.


NonStop BladeSystem- logical view



Integrity NonStop BladeSystem is MPP, not SMP

NonStop server blade

- Leverages Integrity server blade – two socket full-height server blade featuring the latest Intel® Itanium® 9100 series dual-core processors
- Includes NonStop ServerNet controller
- Includes 1.66 GHz, dual-core Itanium 9100 series processor
- Offered in 8, 16, 24, 32, or 48 GB memory configurations (bundled)
 - Memory upgrades offered in multiples of 8 GB (e.g., two 4 GB DIMMs)
- Shared memory within a logical processor




New I/O infrastructure

- Three Cluster IO Module (CLIM) products
 - **IP CLIM** for networking protocols and Ethernet connectivity
 - **Storage CLIM** for attaching Serial Attached SCSI (SAS) disks, Storage XP Array family, and fibre channel tape
 - **Telco CLIM** for SS7 over IP and other
- Can co-exist with existing I/O Infra such as
 - IOAME
 - SS7 over T1/E1 via S-series I/O



NonStop BladeSystem I/O infrastructure

Cluster I/O Module (CLIM)



CLIM

Overview



- Industry standard server running Debian Linux
 - Features ServerNet controller
 - Uses NonStop enhanced Linux to transfer data
 - Does not allow any customer code or applications
 - Linux “personality” is hidden
 - Fully integrated with NonStop manageability subsystems
- 2U high rack-mount server
- Makes innovation faster and easier
- Improves TCO



IP CLIM




- New NonStop Networking Subsystem can replace IOAME and G4SA
- Supports
 - Gigabit Ethernet connectivity: Copper/Fibre ports
 - **IPSec**
 - TCP/IP v4
 - TCP/IP v6
 - SCTP (Telco streaming protocol)
- Offloads some of TCP/IP load from NonStop CPU to CLIM CPU

IOAME does not and will not offer IPSec



Next generation of NonStop Networking

- S-series I/O based
 - E4SA (10Mb/s)
 - FESA(100Mb/s)
 - GESA (1Gb/s)ServerNet I and ServerNet II
- IOAME based
 - FCSA (100Mb/1Gb/s)ServerNet III
- IP CLIM
 - Gigabit Ethernet
 - ServerNet III
 - Offloading of TCP/IP, IPSec, SCTP



IP CLIM footprint comparison



$$\begin{array}{r} \text{IOAME} \\ + \text{Two I/O Switches} \\ + \text{Two G4SAs} \\ \hline = 11\text{U} \end{array}$$




Two IP
CLIMs

$$= 4\text{U}$$

**REDUCED
FOOTPRINT**



IP CLIM stack



IP CLIM Software relieves
NonStop BladeSystem of
TCP processing while
requiring
no application changes

Storage CLIM



- New NonStop Storage Subsystem
 - Co-exists with/Replaces the combo
 - IOAME / FCSA
- Runs on same industry standard servers as IP CLIM
 - Based on a HP ProLiant server with PCIe cards for ServerNet and IO
 - Replaces function provided by IOAME/VIO and FCSA
- Supports SAS disks, FC tapes, XP Array, VTS
- CPUs can boot from either SAS disks (via Storage CLIM) or FC disks (via IOAME)
- Advanced Caching technology on SAS storage
- Integrated Volume Level Encryption



Why choose SAS?

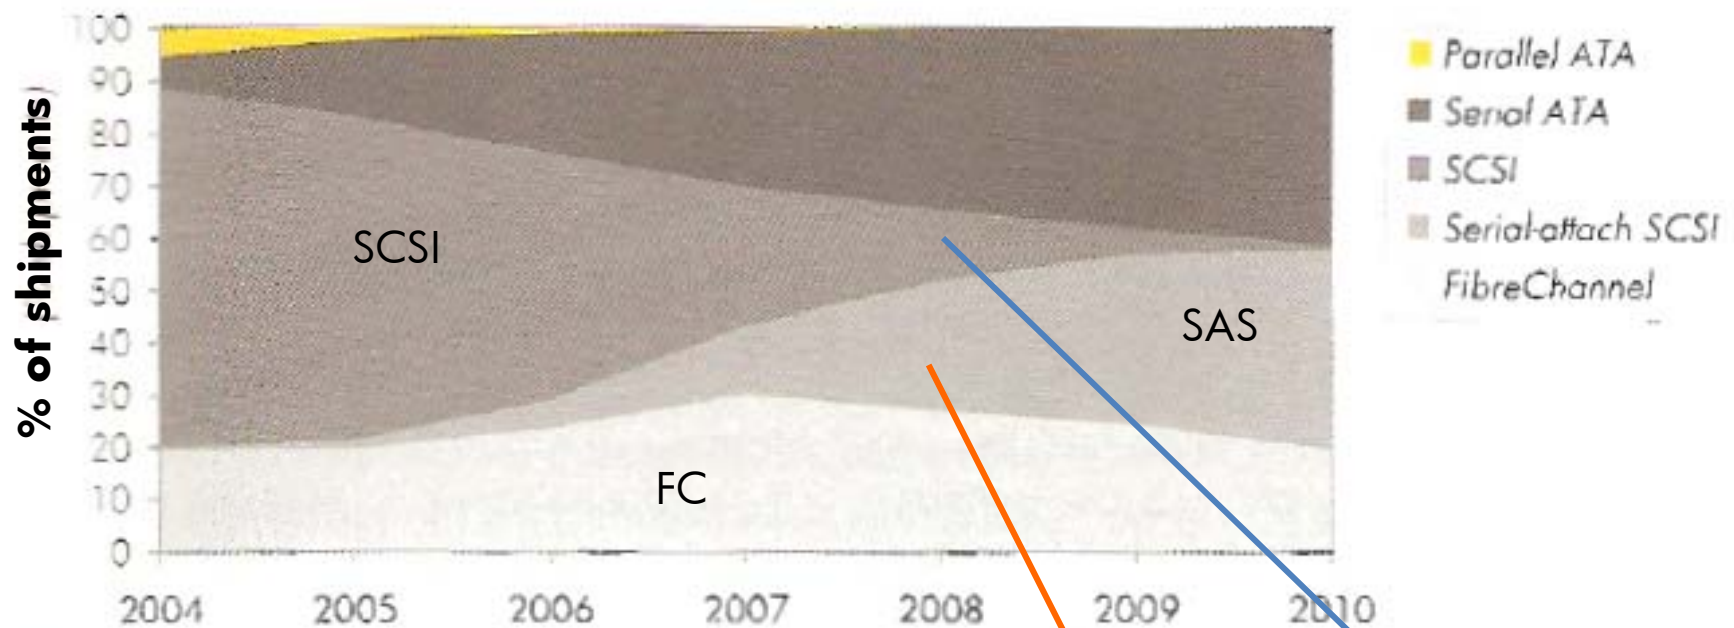


Figure 2.5: enterprise disk drive shipments by interface type


Source: IDC, Worldwide hard disk drive 2006-2010 forecast and analysis: record-breaking years may lie ahead, May 2006, IDC report #201478)

SAS is increasing and SCSI is tapering

HP white paper on SAS over SCSI

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00302340/c00302340.pdf>

Storage CLIM footprint comparison



IOAME
+ Two I/O Switches
+ Two FCSAs
+ Two FCDMs

= 28 disks/17U

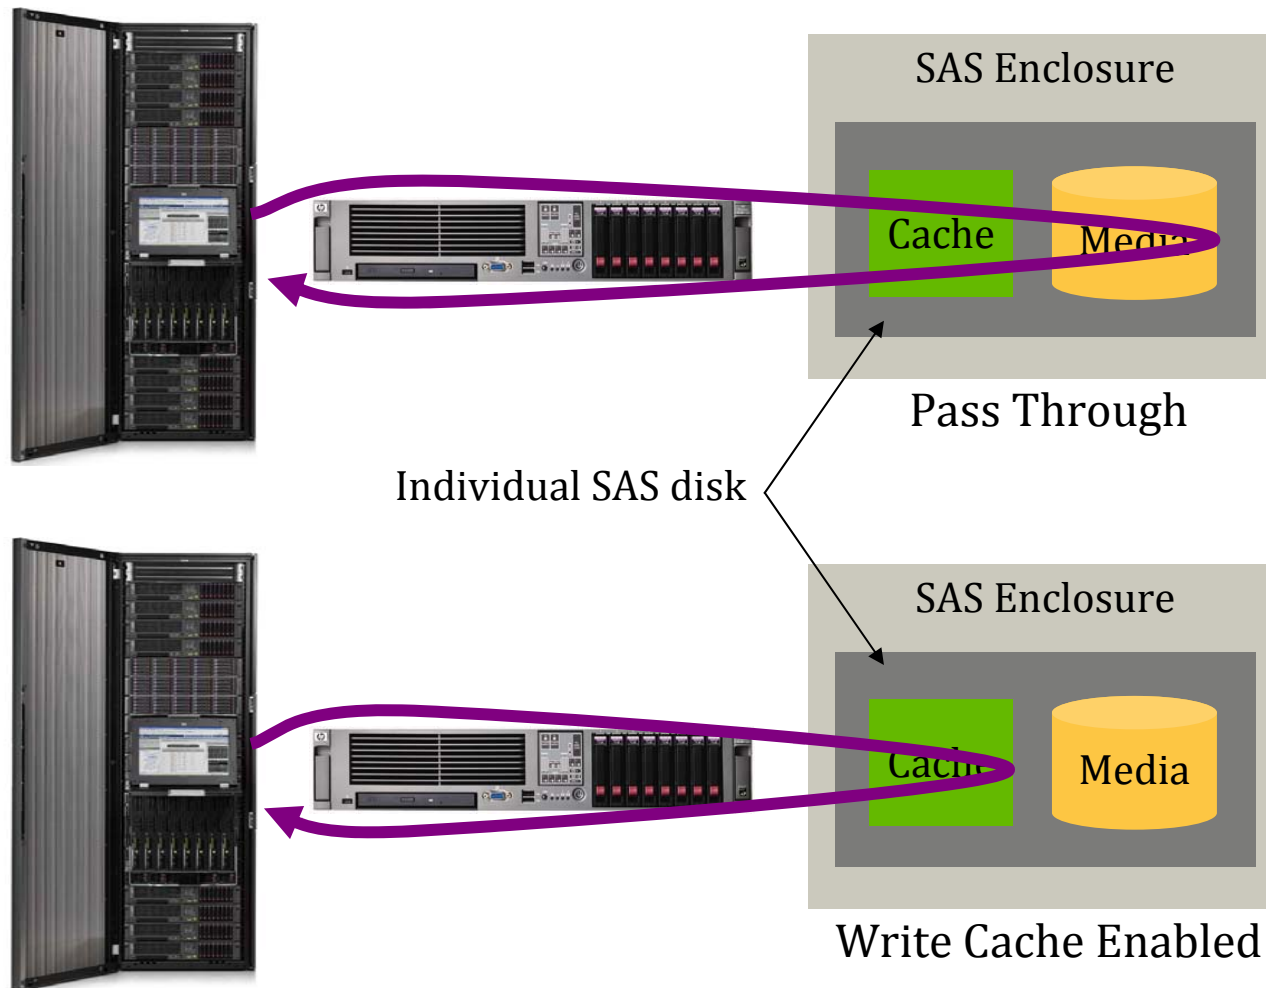


Two Storage CLIMs
+ Two MSA70s

= 50 disks/8U

**REDUCED
FOOTPRINT**

Storage CLIM with write cache enable




Lab Test

Random writes are
30% faster with
WCE compared to
Pass Through

Sequential writes
are up to **6 times**
faster

Your mileage will
vary...

Next generation of NonStop storage



Each SAS HBA has 1 SAS port. Each SAS port has 4 physical links. Each link is 3 Gb/s

Each FC HBA has 2 FC ports. Each FC port has one FC connector. Each FC connector is 4 Gb/s

Each FCSA has 2 FC ports. Each FC port has one FC connector. Each FC connector is 2 Gb/s

Telco CLIM



- New NonStop Networking Subsystem that delivers M3UA (SS7 over IP) and SIP (Session Initiation Protocol)
- Offers 5 physical ports per CLIM with up to 16 associations per port allows 80 configurable links per controller
- High Performance: 32K MSU/sec per CLIM
 - Approximately ~25X performance of M3UA over SS7TE3
- Supports SCTP Multihoming to other ports on the CLIM



ServerNet backbone comparison

NS-Series Rack Mount System



Processor
Switch
Y ServerNet

Processor
Switch
X ServerNet

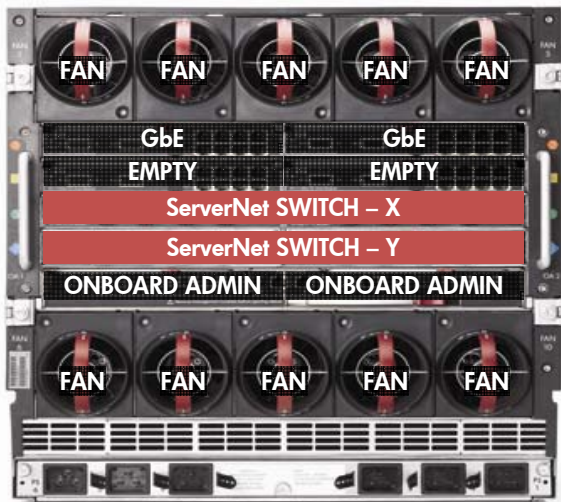
P-Switch 3U x 2 = 6U

One type of P-Switch

Total of 24 I/O links per fabric

One P-Switch per Fabric per system
(total of two P-Switches per system)

NonStop BladeSystem



No additional rack height consumed

ServerNet switches are embedded inside c-Class enclosure (at the back)

REDUCED
FOOTPRINT

Two types of ServerNet Switches

1. Std I/O: Total of 12 I/O links per fabric
2. High I/O: Total of 24 I/O links per fabric

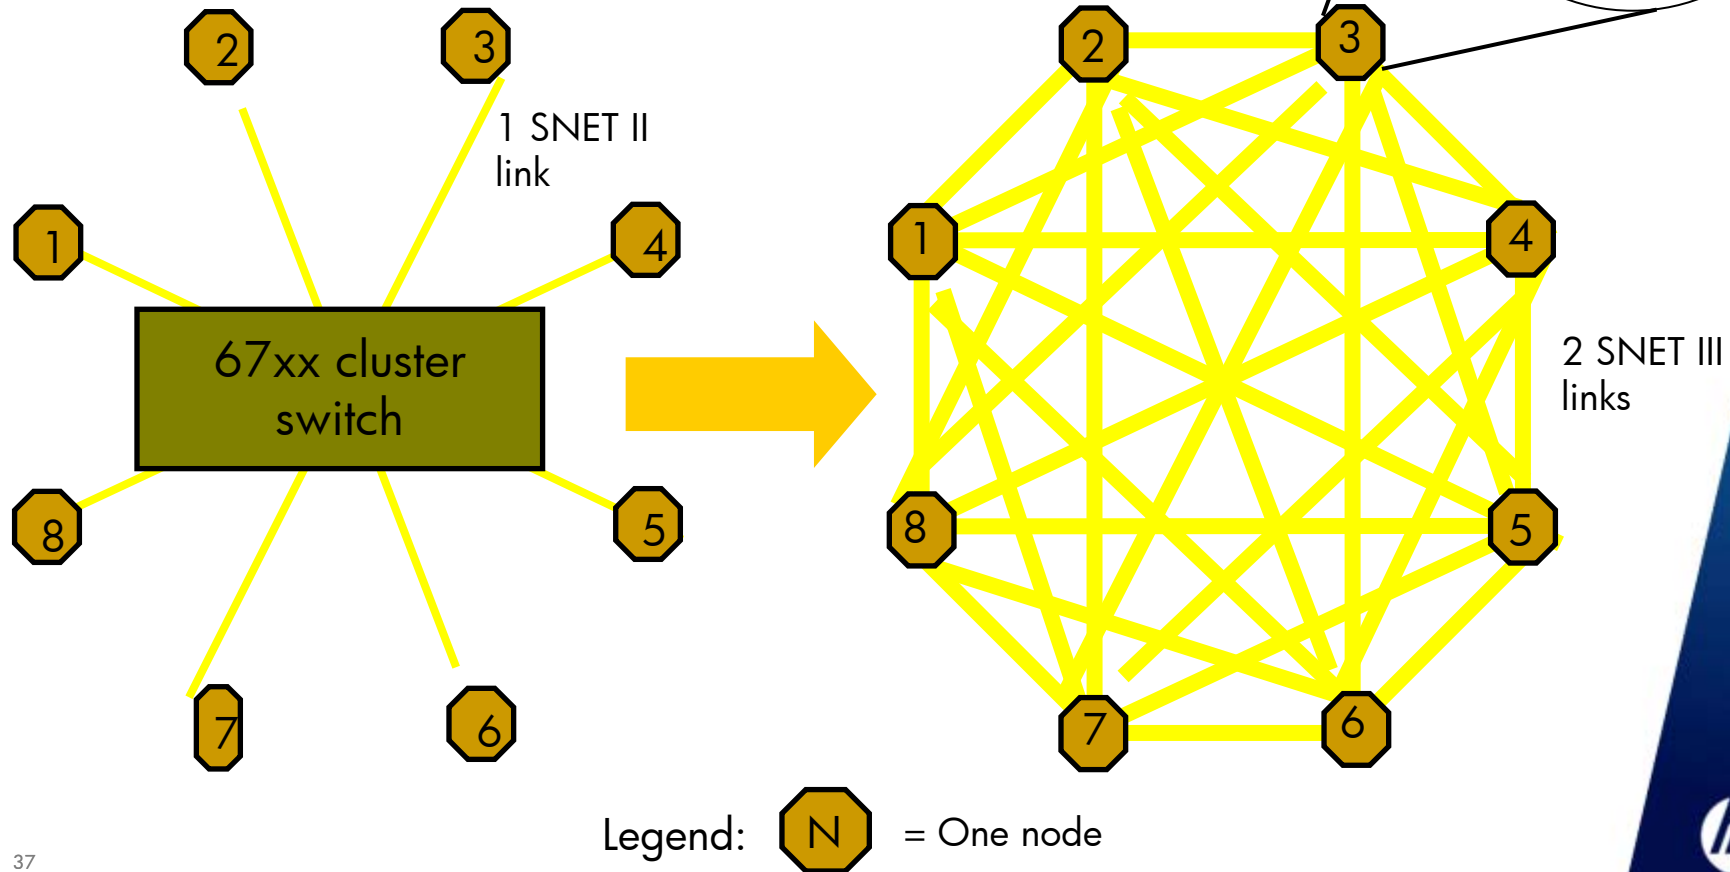
One ServerNet Switch per fabric per c-Class enclosure
(total of two ServerNet Switches per c-Class enclosure)

16p system = two c-class enclosure
= four ServerNet Switch


ServerNet Switches are FRUs

Next generation of NonStop Clustering: BladeCluster

BladeSystem




Next Generation of Clustering: BladeCluster Breaking performance barrier



BladeSystem characteristics

- Dynamic Smart Cooling
 - Easy to retrofit or spec for new construction applications
 - Reduces cooling energy costs by 25-40%
 - Increases available cooling capacity for additional IT loads
- Reduction in power consumed per transaction from 16000 DMR to BladeSystem
- Dynamic Power Savings
 - Fans: All ON
 - Power supplies N+N



"Dynamic Smart Cooling is the most remarkable development for data center critical support systems."


Peter Gross
CEO and CTO
EYP Mission Critical Facilities Inc.



Migration to NonStop BladeSystem

Complete investment protection

- 100% software compatible
- Seamless clustering with prior systems
- Supports existing infrastructure




Ease of upgrade

NonStop BladeSystem



NB50000c
Blade



4-way Blade

No Hardware change anywhere else

HP Integrity NonStop Multi-core Software



NonStop software investments

In a nutshell

- Support **industry standard technologies** for application development
- Differentiate these standard applications by deploying them into the **most scalable and available platform infrastructure** (without change)
- Make this infrastructure easily accessible, open, highly secure, and simple to manage

Common standards, uncommon advantages

The same application runs better on NonStop



NonStop Operating System

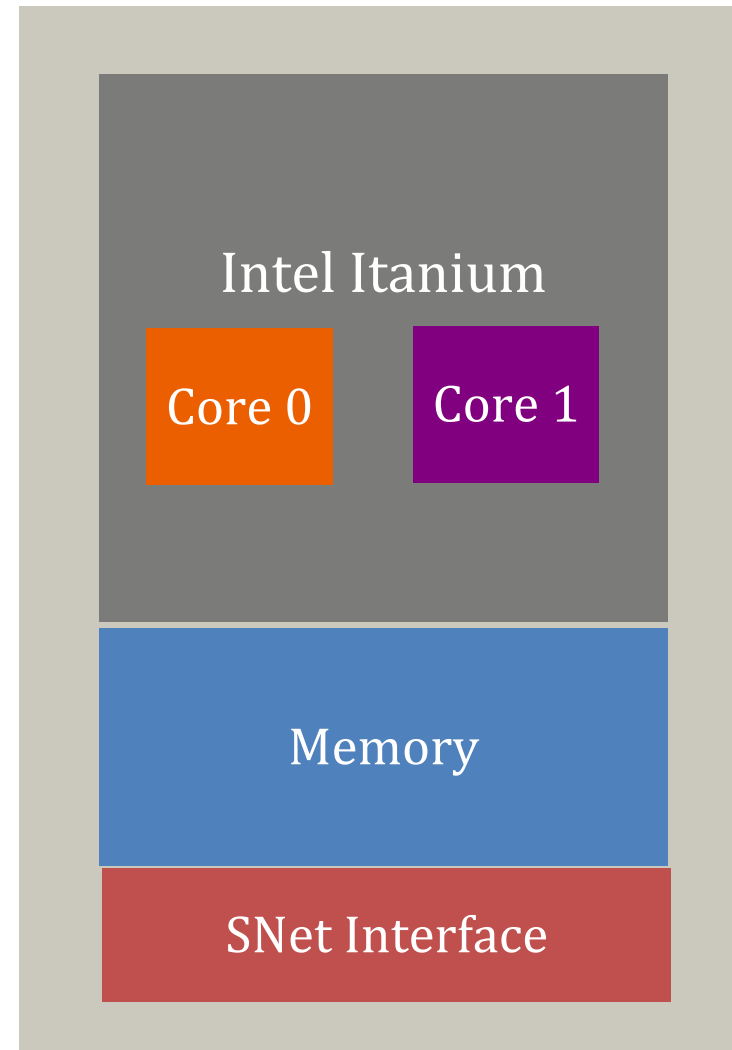
NonStop Multicore Architecture(NSMA)

- IPU – “Instruction Processing Unit”, a core
- Monarch (IPU) – the initial IPU upon power on (Intel calls it the boot processor)
- CPU – logical processor, a set of cores sharing the same memory. The traditional NS Logical CPU extended to be a multiprocessor
 - One X and Y SNet interface per CPU
 - All IPU's in CPU share same memory map (except small per-IPU anchor area for low-level software)
- n-Way – traditional indication of IPU's in a multiprocessor: 4-way means 4 IPU's per CPU
- Process Scheduler: the new NSK subsystem that distributes and redistributes processes to IPU's

NonStop Operating System

NonStop Multicore Architecture(NSMA)

- IPUs (individual cores) share...
 - NonStop OS system image
 - Memory (e.g. system globals)
 - Locks and synchronization
- IPUs have own...
 - IPU-specific data area (128Kb)
 - Register set
 - Pipeline
 - Data and instruction caches
 - Ready list of processes



Where are NSMA IPU's user visible

- IPU's are system resources and just make more CPU cycles available
- No knobs or levers
- Only a few places where individual IPU's are externalized
 - Measure's CPU entity shows IPU level busy/idle
 - PEEK shows the number of IPU's in a CPU
 - PROCESSOR_GETINFOLIST_ has new attributes on the number of IPU's in a given CPU
 - Multiprocessor CPU model numbers are distinct from uniprocessor model numbers, they denote the number of IPU's in the CPU

NSMA Attributes shared across IPU's



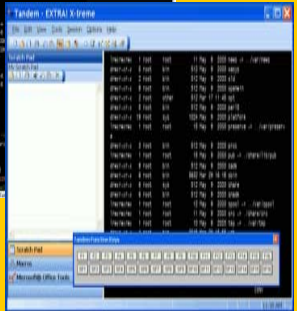
- Main Memory
- Mapping of Region 7 (upper part of the system address space including SG)
- Software Synchronizers
 - MUTEX
 - NSK Locks
- Software Global Data Structures
- Application's Shared Data Segments

NSMA Attributes per IPU

- Per IPU Area (128KB)
 - Carved out of SPAD
- Ready List
 - The Process Scheduler rearranges the lists to balance the CPU
- CPCB / CMTCB
 - Note that CPCB is not in SG: there is more than one "current" process for the CPU.
- Live Register Set
- Pipeline
- Data and Instruction Cache (Levels 1, 2, and 3)

NonStop software investments

We are with you... all the way

Flat Files

id	name	team
1	Amy	Blues
2	Bob	Reds
3	Chuck	Blues
4	Dick	Blues
5	Ethel	Blues
6	Gilly	Blues
7	Hank	Reds



SOA Services



Flat Files

id	name	team
1	Amy	Blues
2	Bob	Reds
3	Chuck	Blues
4	Dick	Blues
5	Ethel	Blues
6	Gilly	Blues
7	Hank	Reds

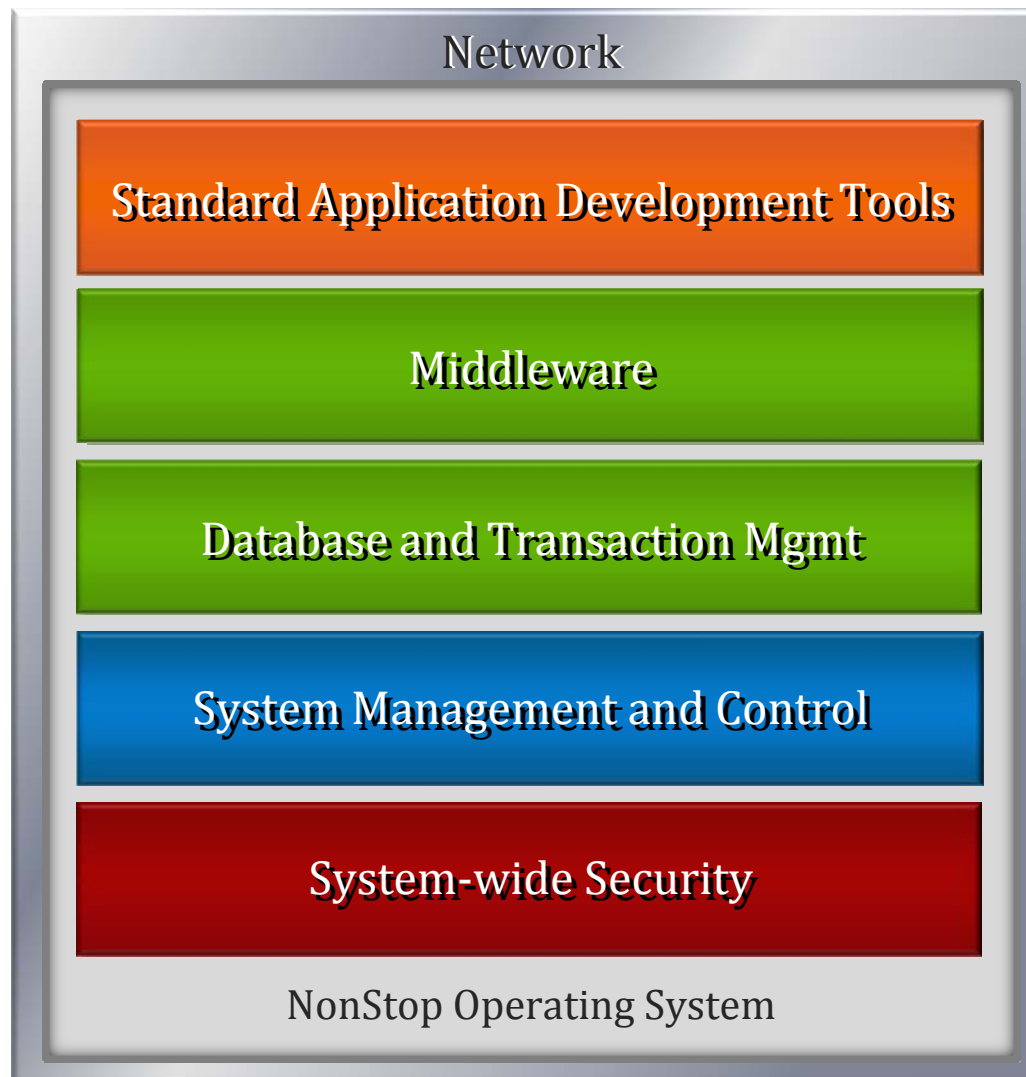







NonStop J-series Software

Investments across the stack



NonStop Operating System Infrastructure

- Scheduler, Locks and Synchronization
 - 64-bit address space for OSS user processes - **FUTURE**
 - *Think Terabyte of virtual address space*
 - *Think In-memory database*
 - OSS file open per CPU, disk file open per CPU
 - OSS open sockets per CPU
 - Improved concurrency for Pthreads-based application
 - OSS System I/O calls to regular files are thread-aware (i.e. do not block the process)
 - Guardian Binary Semaphores per process
 - Increased from 64 to 24K
- ... And many more such innovations/developments..



Operating system infrastructure – plans

Adherence to industry standards, preparing for Quad-Core

February 2009

- **Guardian Binary Semaphore – Limits Relief**
 - Increase the number of binary semaphores per process from 64 to 24K

May 2009

- **OSS File Open – Limits Relief**
 - Increase OSS file opens (per CPU including sockets, terminals, disk, ...) from 12K to 64K
 - Increase OSS disk file opens (per CPU) from 12K to 48K
 - Increase OSS open sockets (per CPU) from 4K to 16k

2H 2010

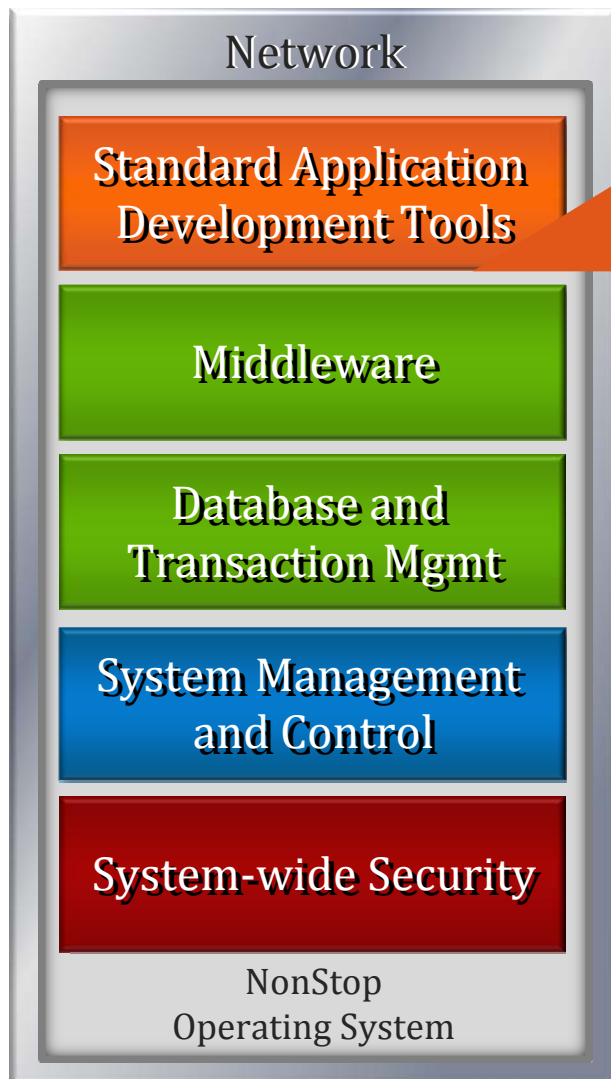
- **Standard Library Support for Non-blocking IO**
 - Non-blocking IO for threaded applications using standard C libraries
- **System Limits Relief**
 - Increase OSS PIDs (per 16P system) to 128K
 - Increase number of Guardian processes to 10K
 - Increase OSS file opens (per CPU) to 128K
 - Increase OSS disk file opens (per CPU) to 96K
 - Increase OSS open sockets (per CPU) to 32K

Time



NonStop Software

Application development

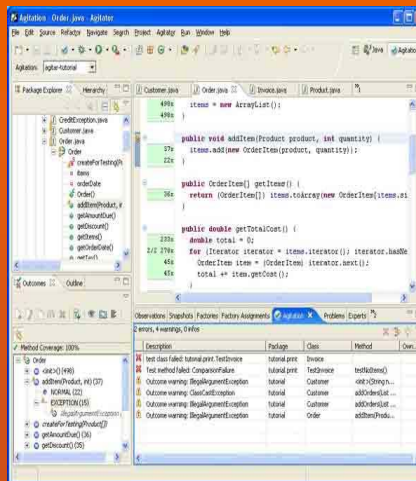


Java Support

Support latest Java releases
Support latest Tomcat releases
Open-source frameworks for ease of appl development

Application Development

Increase developer productivity
Make apps easier to port
New optimization and debugging capabilities



NonStop Application Development Integrated Development Environment

NSDEE – NonStop Dev Env for Eclipse

- Integrated Debugging
- Improved NonStop connectivity
 - (Not just project files but others too)
- Launch programs on NonStop
- Remote build error correction
 - (wherever compile error is; file is presented to you in IDE)
- Improved Eclipse integration
 - (Eclipse update manager)



NonStop IDEs Quick Comparison

			2010	
	ETK (Visual Studio)	EPE (Eclipse)	NSDEE (Eclipse)	
Languages	C/C++ Java COBOL pTAL	C/C++ Java COBOL pTAL	C/C++ Java COBOL pTAL	Visual Inspect: 3DES Encryption for Login credentials
Local build	Yes	Yes	Yes	
NonStop build	No	Yes	Yes	ETK: Support for Visual Studio 2008
NonStop launcher	No	No	Yes	
Integrated debugging	No	No	Yes (optional)	NSDEE: Successor to EPE
Connectivity	Telnet / FTP	Telnet / FTP	SSH Telnet / FTP	

Develop for NonStop using state of the art tools without knowing about the target platform... more or less



Standards-based app development with Java

Java platform with open-source frameworks



Apache MyFaces

Component based web UI framework



Apache Axis2

Web services framework



Spring

Framework for developing apps using
POJO components



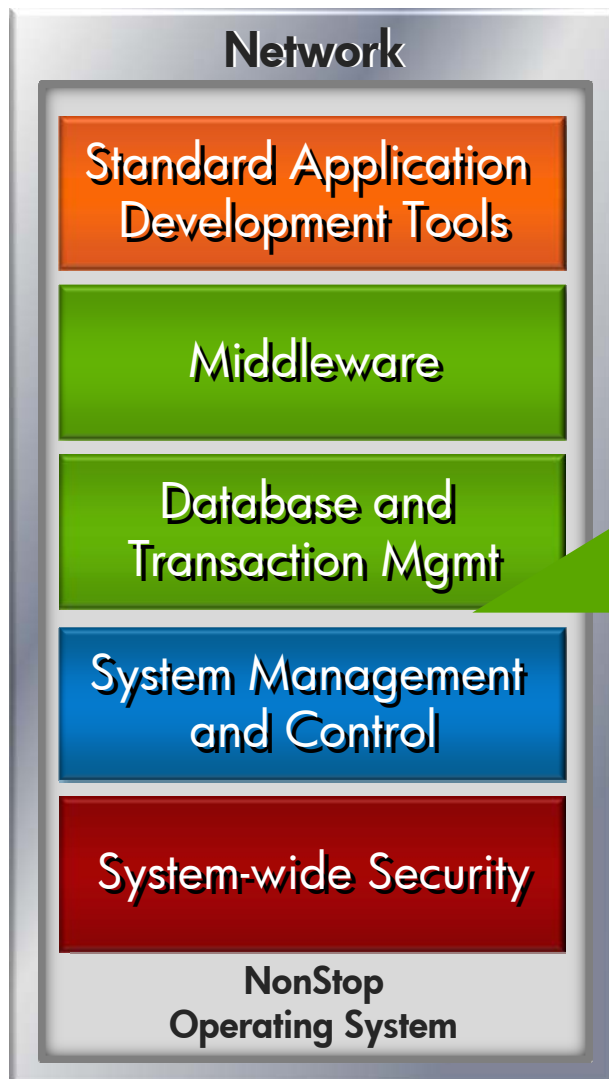
Hibernate

Object Relational Mapping framework

Integrated and tested by NonStop development
End user customer support by HP Global Support

NonStop Software

Middleware, Database & Transaction Mgmt



Middleware

Reduce planned downtime
Increase OLTP capacity
Support SOA standards/protocols



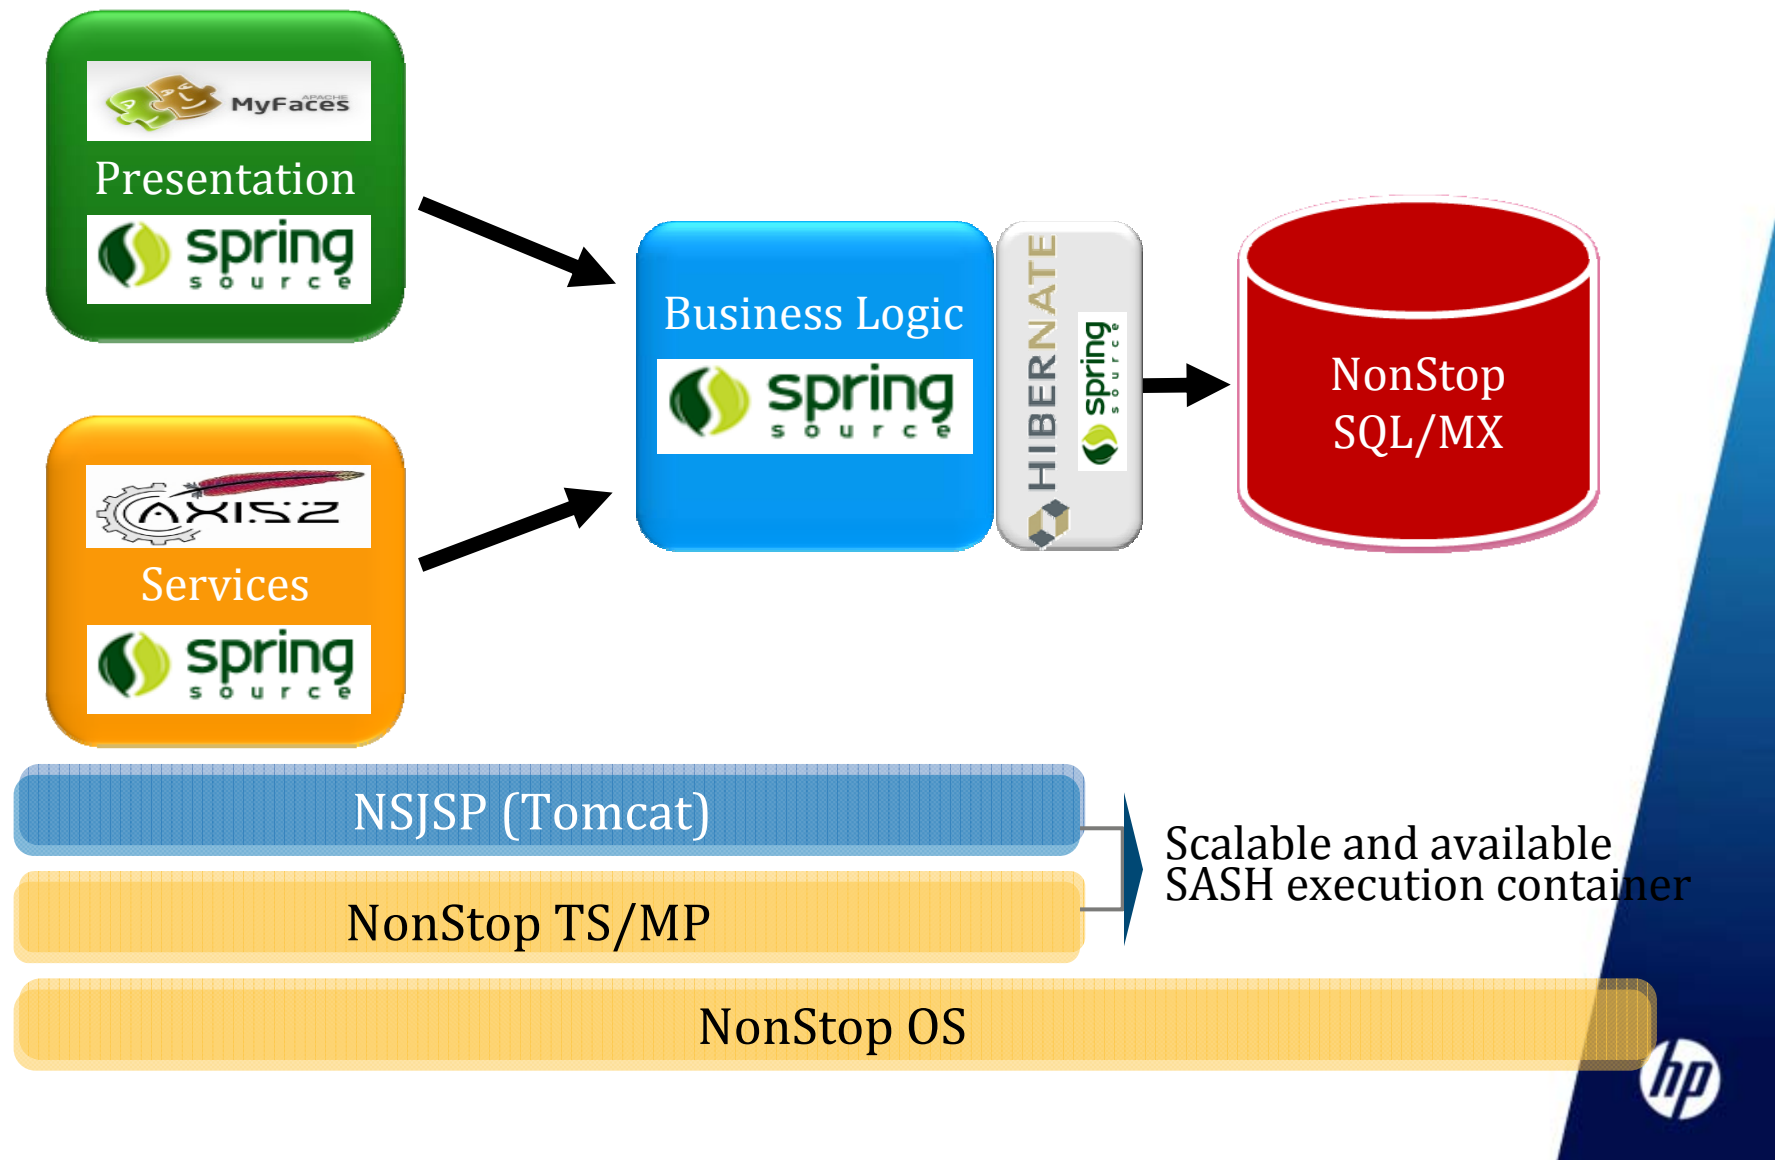
Database & Transaction Mgmt

Improved performance
Enhanced ease of use
Richer coding capabilities




Platform for enterprise Java development

Deploy on standard Tomcat, optimized for scale




Standards-based application integration

SOA/Web-Services technology



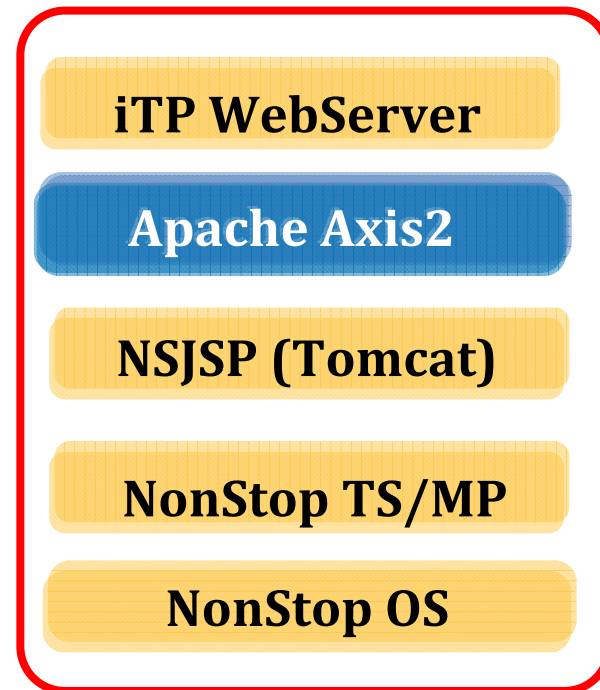
SOA enables legacy apps to retain their legacy interfaces while providing new ones

NonStop provides 2 Web Service stacks



NonStop SOAP

- C-based
- No coding
- Included in J-series OS



Apache Axis2

- Java-based
- Some coding
- Open source, free

NonStop Software in a nutshell

Modern environment based on NonStop fundamentals

Develop Application programming models	ECLIPSE	Open Source Java Frameworks			
		Apache Tomcat			
		Certified Java SE Platform (JDK and JRE)			
		SOA Infrastructure (SOAP, XML, HTTP, WSDL)			
Deploy Application infrastructure	NonStop TS/MP				
	NonStop OS				
Differentiate	Network access	SOA infrastructu re	Open source Java frameworks	Business logic	Database
Transparent Scalability	✓	✓	✓	✓	✓
Transparent Fault Tolerance	✓	✓	✓	✓	✓

Delivering Uncommon advantages by leveraging Common Standards





THANK YOU