



Hewlett Packard
Enterprise

NonStop Technical Boot Camp 2023

TBC23-TB54: What Makes NonStop non-stop; The NonStop 101 rendezvous

Vikas Kapoor, Product Marketing Manager, HPE NonStop
September 2023

Forward-looking statements

This is a rolling (up to three year) Roadmap and is subject to change without notice

This document contains forward looking statements regarding future operations, product development, product capabilities and availability dates. This information is subject to substantial uncertainties and is subject to change at any time without prior notification. Statements contained in this document concerning these matters only reflect Hewlett Packard Enterprise's predictions and / or expectations as of the date of this document and actual results and future plans of Hewlett Packard Enterprise may differ significantly as a result of, among other things, changes in product strategy resulting from technological, internal corporate, market and other changes. This is not a commitment to deliver any material, code or functionality and should not be relied upon in making purchasing decisions.



Agenda

NonStop History

The Operating System (OS) Value

Fault Tolerance

Scalability

Database, Data and Transactions

Security, Compliance & Business Continuity

Putting it all together: Value Proposition



HPE NonStop Systems – Rock Solid through the years

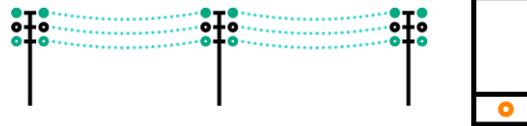
In 1974, there was no Internet

In 1984, *total* internet traffic was 15GB per month...

Tandem Computers revolutionized computing 40 years ago - and is still growing stronger today



Major Global Banks, payments processors



Over 650M+ telecom subscribers



Many of the top global Healthcare, Retail & Manufacturing Companies



One of World's Largest Settlement houses processes €800 trillion/year

Safekeeper & Custodian Holds €30 trillion in securities

Entire settlement runs on NonStop SWIFT, CREST & Bastion security

Running on NonStop systems for **30+ yrs**

“Security can be a big challenge. But with NonStop, we can be the best in class in this area.”

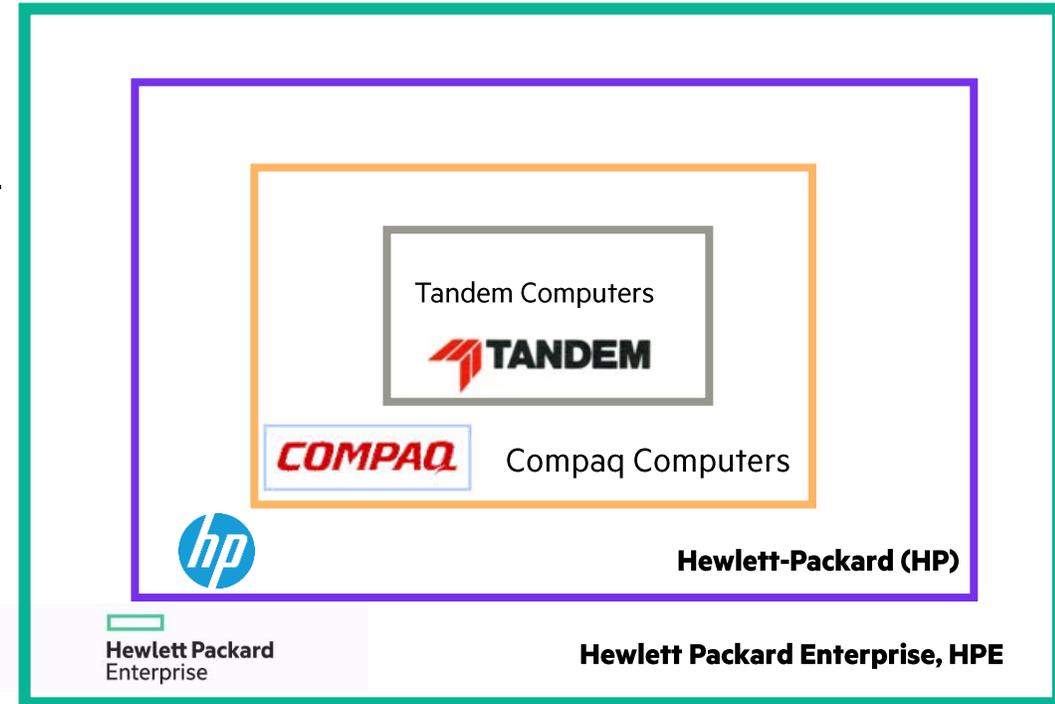
Pascal Remy, Director, Euroclear

Let's Start with some history – Tandem Computers

- Original Design goals – 1974
 - Ability of the system to survive any single hardware failure
 - Ability of a properly-coded application to survive any single failure
 - Assured data integrity
 - When in doubt, fail fast
 - System and application scalability from 2 to 16 processors per node
- What Tandem delivered in 1976 – system and software
 - Single system image across 2 to 16 NonStop CPUs
 - Message-based, client-server communication model
 - Flexible, distributed, parallel kernel architecture
 - Linear application scalability
 - Unified, virtualized view of files, processes and devices

All leading to:

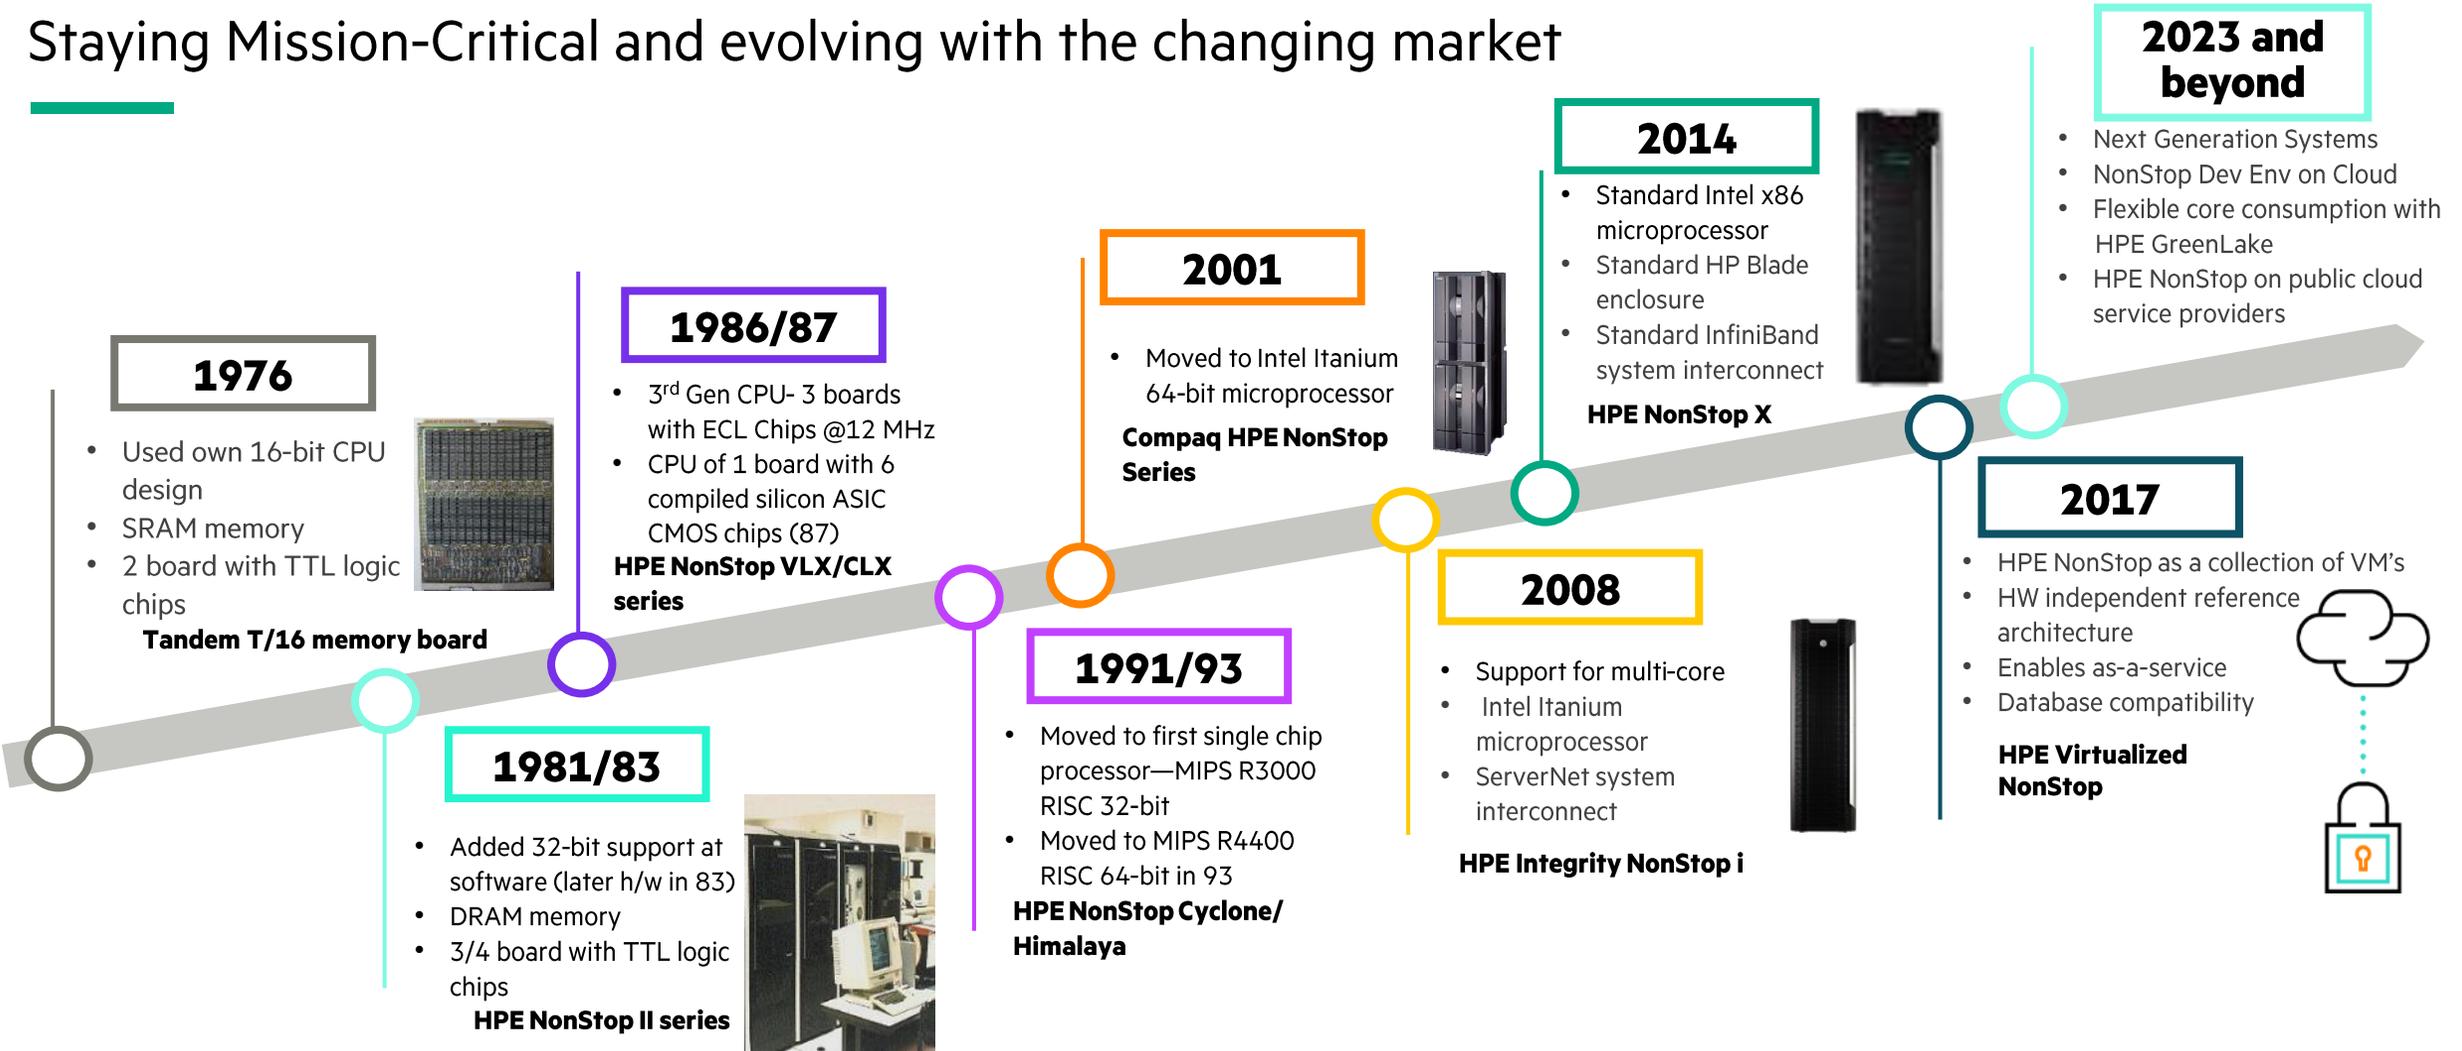
Industry-leading fault tolerance, data integrity, scalability, self-healing platform



The journey to Hybrid IT: From 1974 to now and beyond

HPE NonStop Evolution—From monolithic, converged systems to software in the cloud

Staying Mission-Critical and evolving with the changing market



HPE Partner and Customer Use Only

© 2023 Hewlett Packard Enterprise Development LP

What makes NonStop, Nonstop

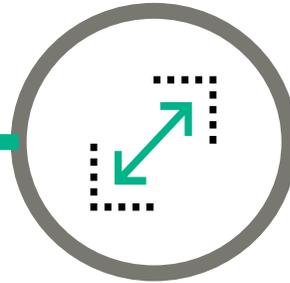
OS value

NonStop relieves the application from having to implement clustering, data integrity, load balancing



Scalability

Designed for the most linear scalability, from 4 to 24000 cores, without re-writing the application



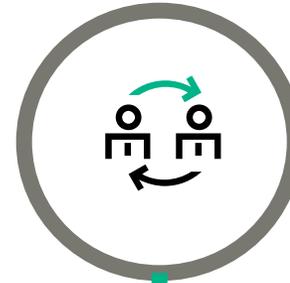
Availability

Best availability in the industry. Failures transparent to the application. No single point of failures, end to end



Trusted platform

Customer focused model:
Large partners ecosystem
Active community
High customer loyalty



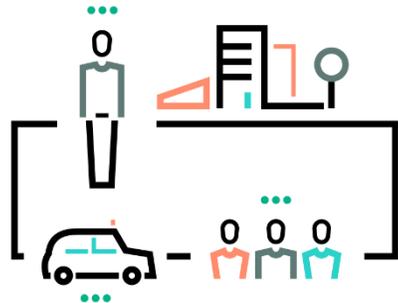
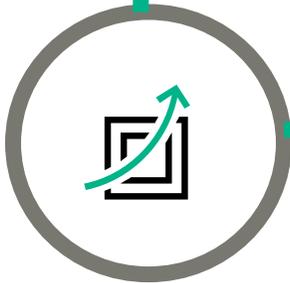
Cost and efficiency

X86 commodity hardware and cloud deployment model allowing fast and cost effective delivery



Hybrid IT

Support the business transformation by allowing a smooth transition from legacy to modern technology



The OS Value

It's about the Operating System



The heart of the system: The Virtualized NonStop OS

NonStop Operating System: Designed from day one, ground-up - for fault tolerance

- Virtualized, Process-and-message-based design for fault isolation
- Clustered, independent operating system instances running on each processor
- Critical system resources implemented as process pairs running on different processors
- Online expansion and reconfiguration of processors, disks, and controllers, transparent to applications
- Uniform access to resources, allowing a single application to run on up to 255 networked 16-processor systems

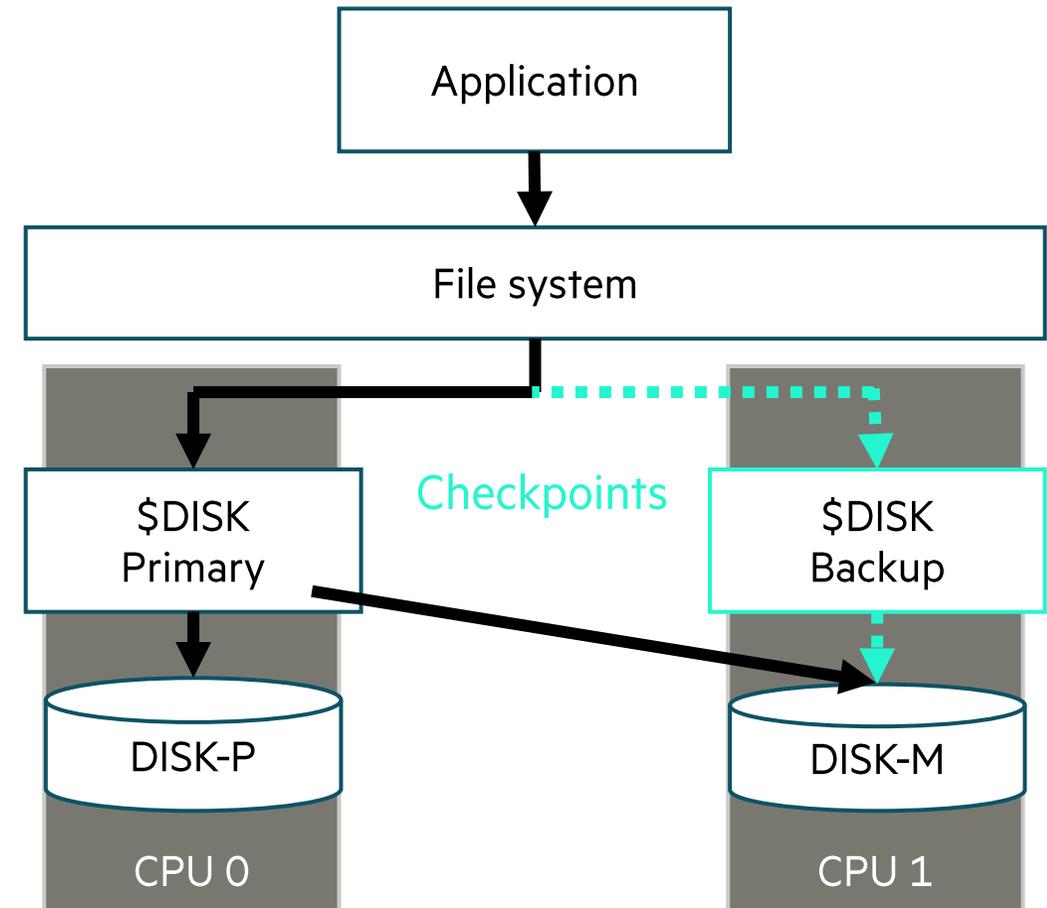
Virtualization of Resources: named Resources

- Most system resources are instantiated as named processes
 - This includes devices, the OS, the database and application environment
 - The actual location of a named process is transparent to processes that wish to access it
- The system maintains the name/location mapping in a per-processor (per NonStop CPU), replicated table
- Virtualization enables:
 - Transparent distribution of processes across multiple processors or Load-balancing
 - Single-image management of the cluster, the db, and application services



Process Pairs - Fault tolerant design

- Process pairs are used to provide services that must be available at all times with no loss of context
- **Normal operation**
 - Primary process sends backup small amounts of critical state for use during recovery (restart context for request), typically just prior to transaction commit
- **When primary process fails**
 - Backup process takes over
 - OS hides the redirection of application request to the backup process
- They are used primarily to provide **system services** that allow customers to write recoverable single processes



Message system

- The Inter-Process Communication (IPC) system over InfiniBand
 - Low latency, high bandwidth, reliable message delivery
- Used directly by system processes (internal), not applications
- Fault tolerance features include:
 - Heartbeat support (“I’m alive” messages)
 - System quorum algorithm (“Regroup”)
 - State Synchronization among all the processors in a system (“Global Update”)
- Expand – is a logical extension of the Message System that provides aspects of a single system view of files and processes
 - Up to 255 systems (4080 processors) are supported within an Expand network



Fault Tolerance

It's about the Availability



IDC Availability Spectrum - HPE NonStop is AL4

Availability Level	Characterization	Impact of Component Failure	System Protection Factor
Availability level 4 (AL4)*	Fault-tolerant server	Switch to alternate resources is not perceptible to end users	100% component & Functional resiliency
Availability level 3 (AL3)	Clustered server	Short outage is needed for failover to take place	User workload fails over to alternate resources
Availability level 2 (AL2)	Workload balancing	Balancing may not be perceptible to end users because of retry	User request is redirected to alternate Resources
Availability level 1 (AL1)	Not shipped as highly available	Need to switch to redundant resources before processing resumes	No special protection for availability

Source: An IDC document (#US46640020), July 2020

The HPE NonStop Integrated Stack and its value

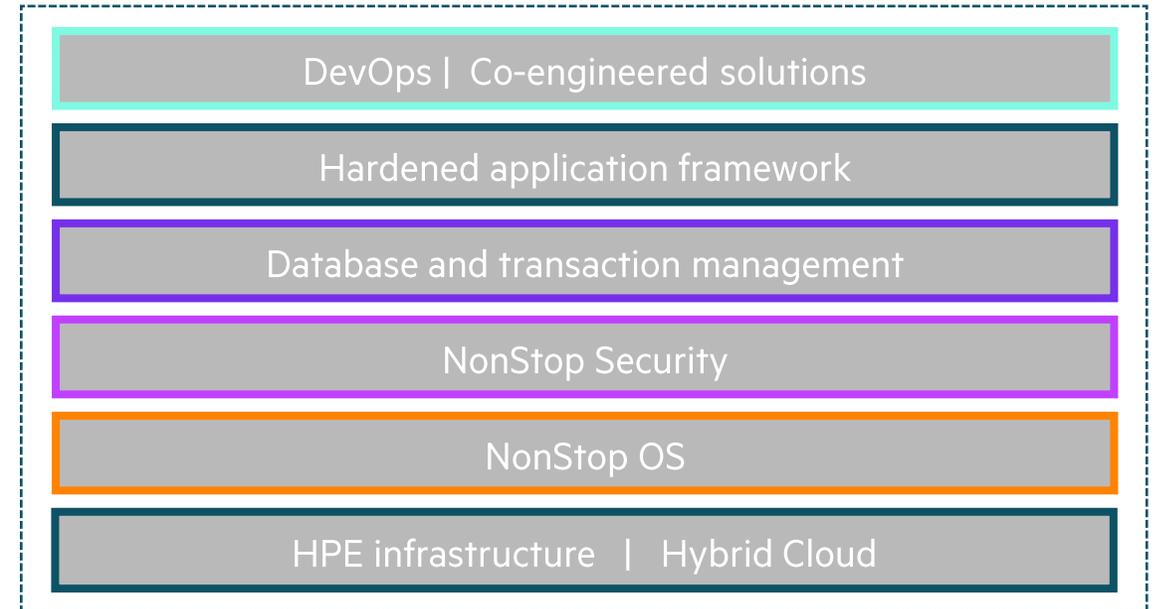
- **Full & Integrated s/w stack**

- Single vendor certification, sale, commitment and support
- End-to-end security
- Cohesive SLAs, security and compliance
- Co-engineered architecture and performance optimizations
- Make the most of the platform potential

- **HPE global excellence, reach, support and services at scale for the full stack**

- Certified and highest security infrastructure
- Global skills set with a Mission Critical culture
 - Solutions Architects
 - Advanced Technology Center
 - Support 24x7 and follow-the-sun
 - Professional Services
- HPE Managed Services
- HPE GreenLake consumption and co-location
- Long term protected investment
- 50 years of experience in mission critical IT

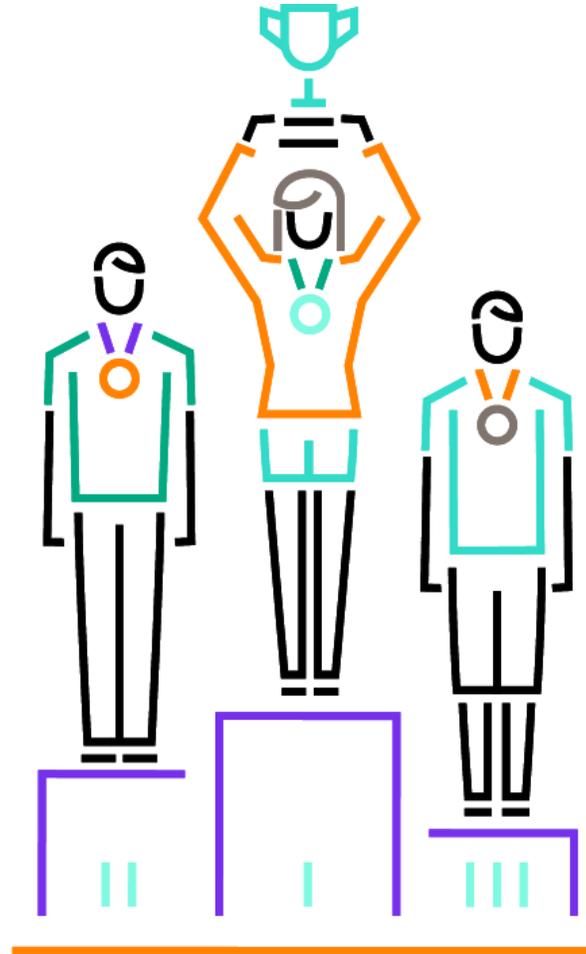
The peace of mind



No integration effort for the customer
Reduced risk of a component lacking support or security

No single point of failure

NonStop feature	Scope
Online operations	Online data backups, database alterations,...
Application availability	Automatic restart and load re-distribution of application processes
Middleware availability	Ensure application servers or web services are always on
Process pair	Implements OS services with context always on
Message System	Reliable, fault-tolerant IPC mechanism
Redundancy of all hardware resources	Disks, networking, processors, fabric



NonStop Middleware offers



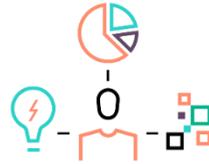
Interoperability

Use standard communication paradigms with simple APIs



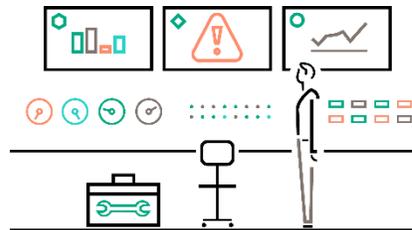
Concurrency control

Manage simultaneous activity to ensure transactional data integrity



Simplified Programming

Provides common services via high-level APIs



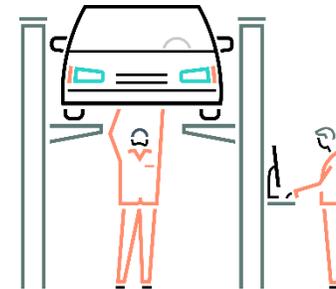
Resource management

Operate large pools of application resources, with security controls, status monitoring and error reporting



Continuous Availability

Access NonStop's ability to automatically detect failures and take action to recover

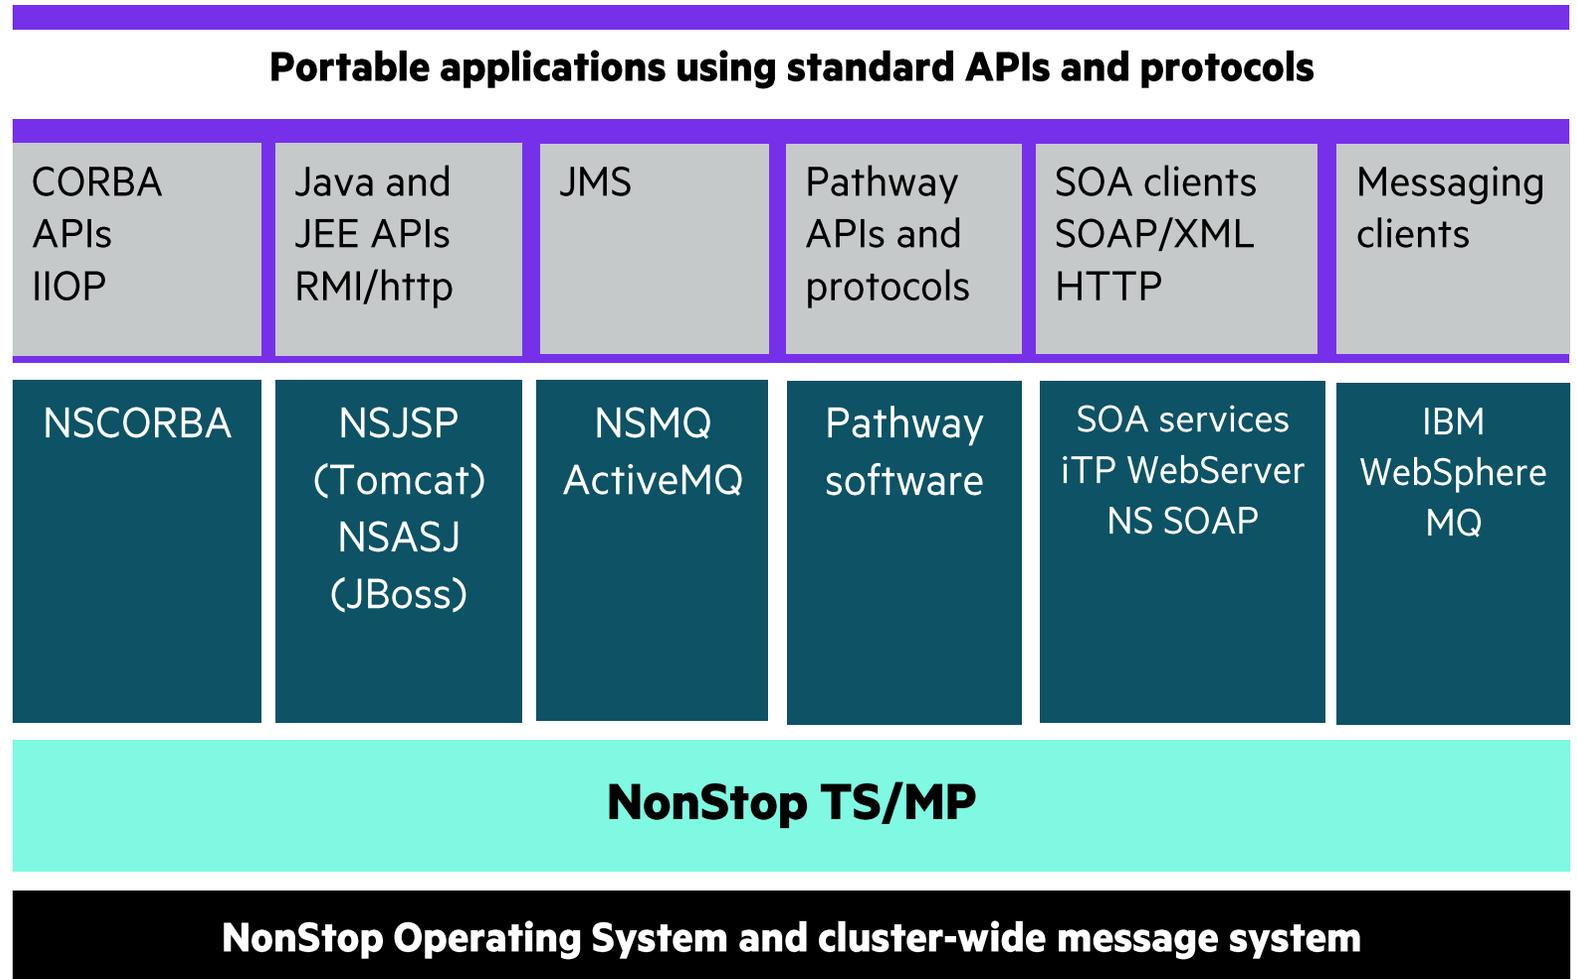


Workload Management

Adjust or scale resources for workload to protect application performance

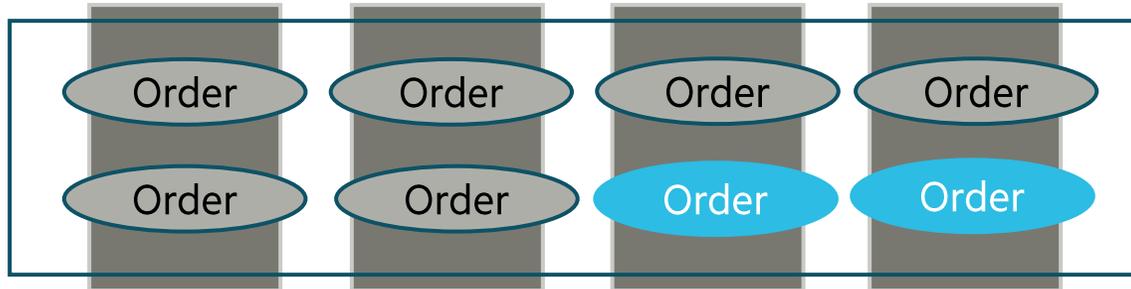
Transaction Services Massively Parallel - TS/MP

- TS/MP provides a run-time environment or application server functionality to simplify management of highly available and scalable OLTP applications
- Used internally by many NonStop products
 - NSASJ, NSJSP, iTP WebServer, SOAP 4, NSCORBA, NSJI.....
- Supports applications written in C/C++, COBOL, Java and TAL

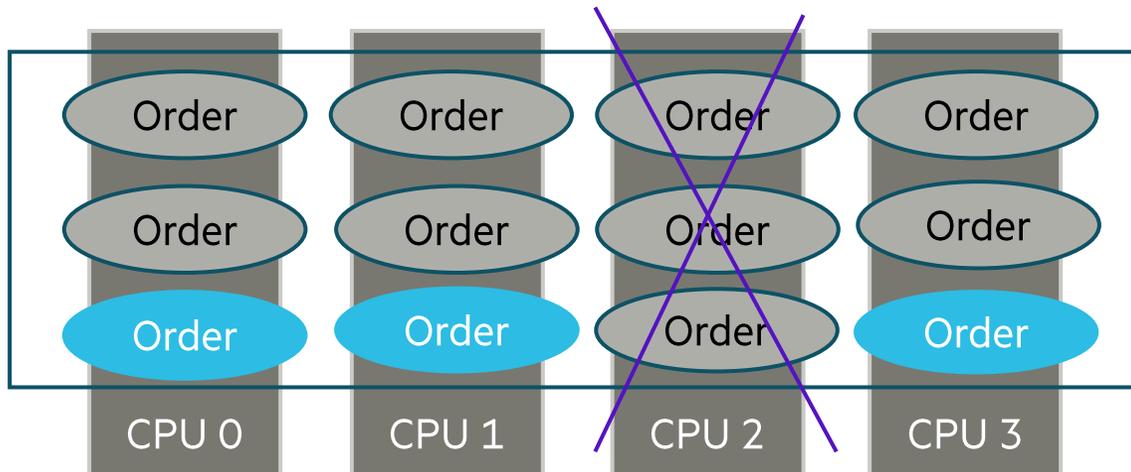


TS/MP makes your application Mission Critical

Propagates scalability, availability and load balancing to the application.



Auto scale
When the service load increases, TS/MP creates more processes to handle the load and stops them when load decreases



Auto reconfiguration
If a node fails, TS/MP restarts additional servers in the remaining processors



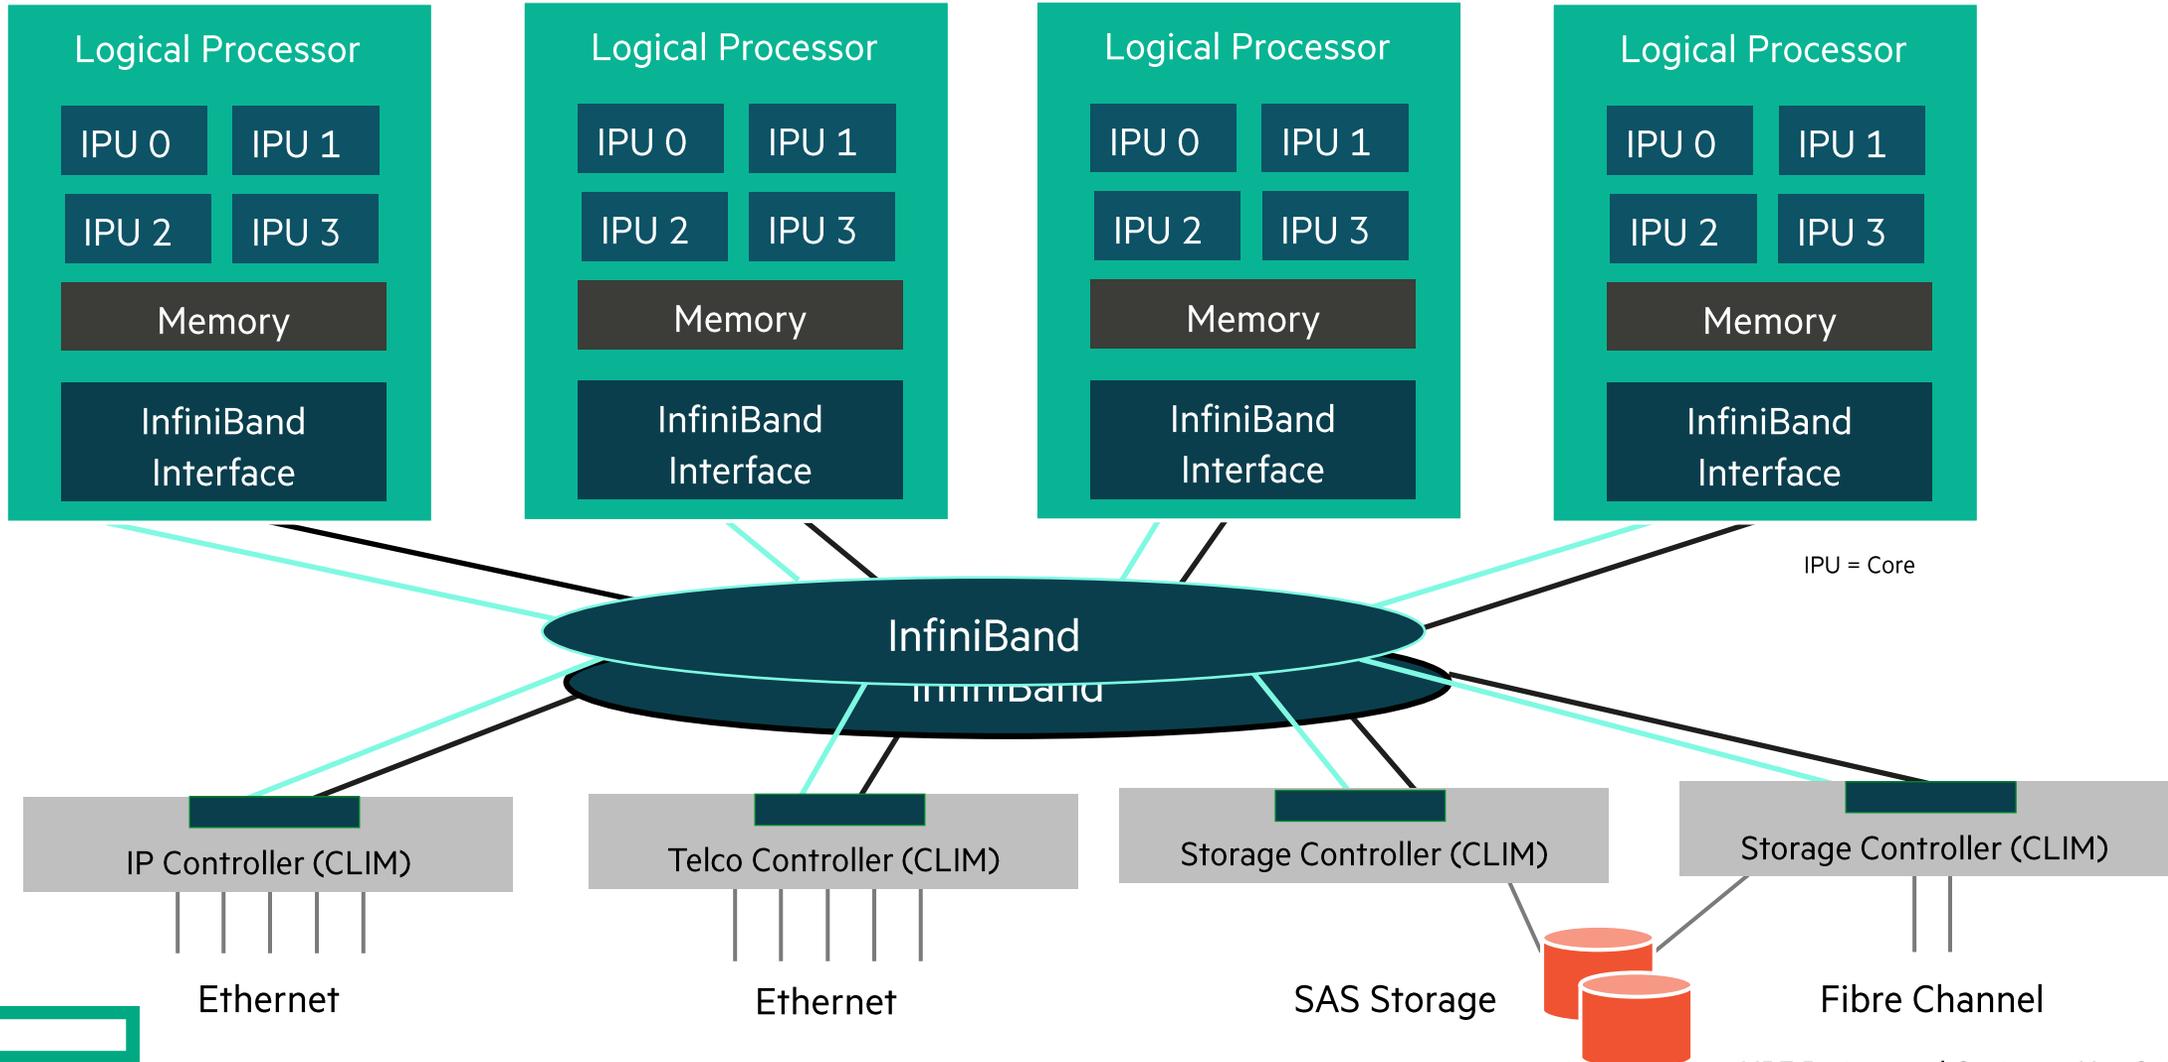
Start Small – keep Growing

It's about the Scalability

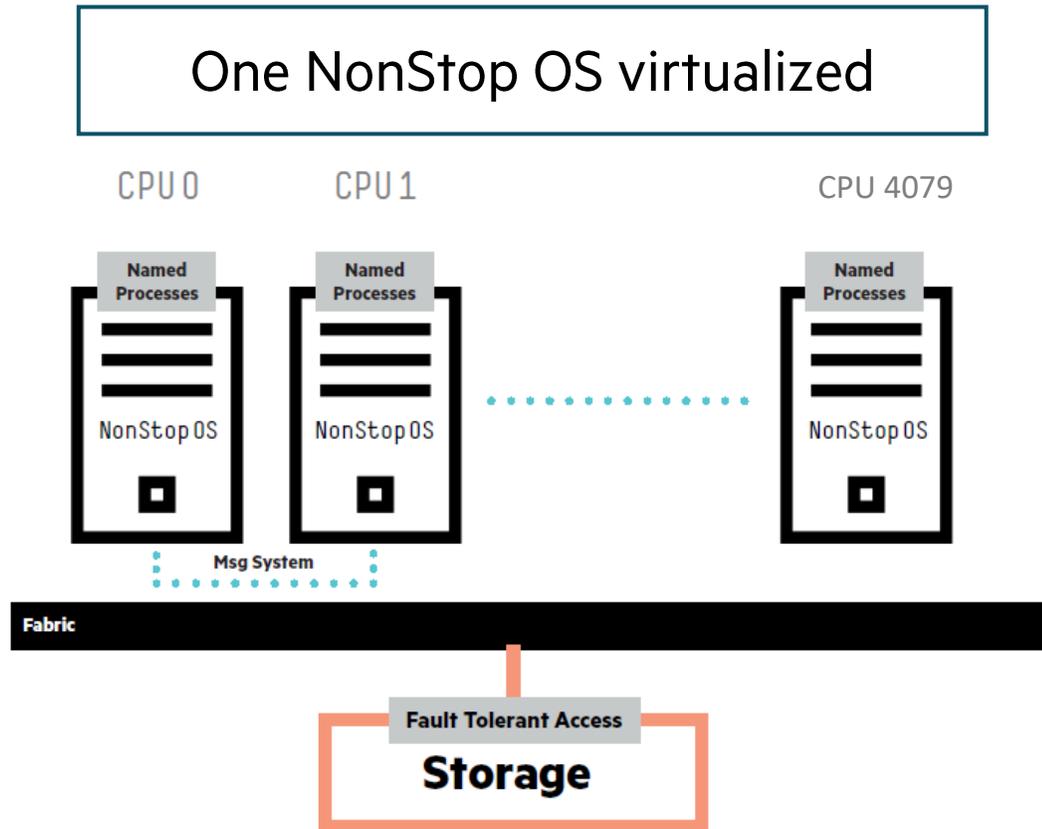


System architecture, a physical view

Each logical processor (NonStop CPU) runs a copy of the NonStop OS



Combining virtualized resources and message system



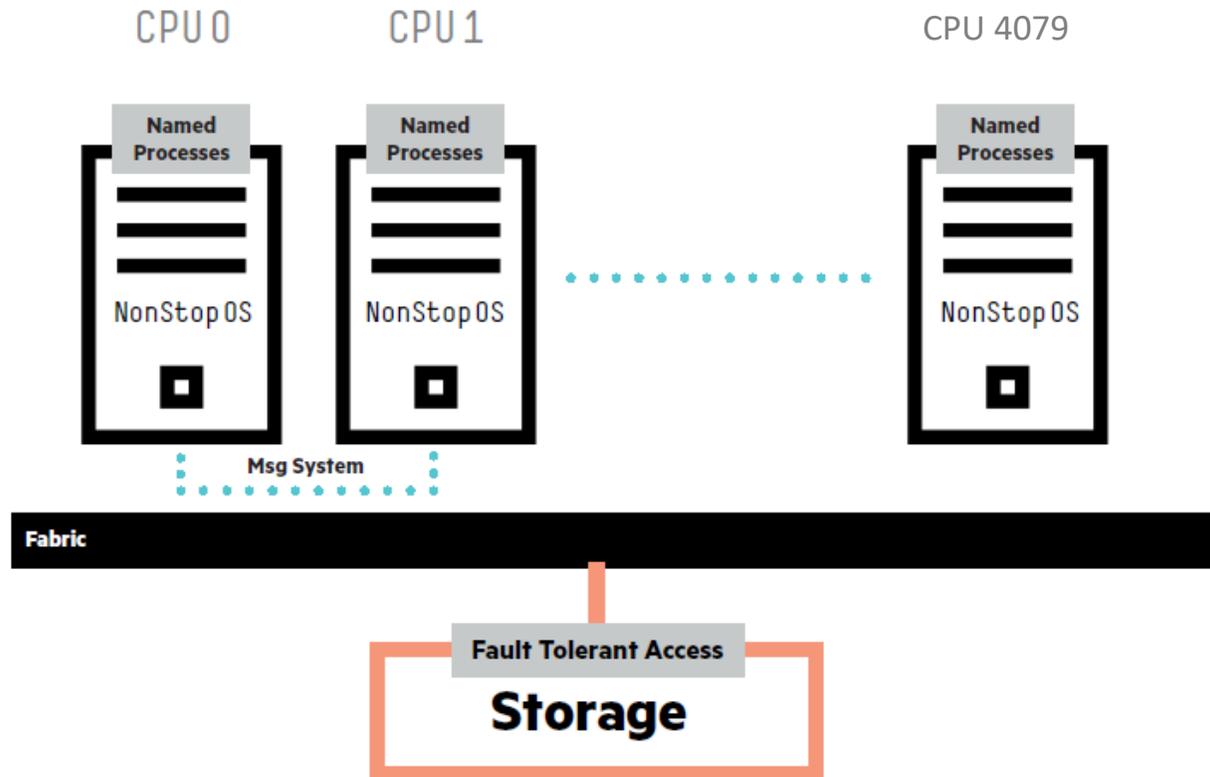
The OS is presented to the application as a single system image

“Simplicity is prerequisite for reliability”

- Edsger W. Dijkstra (1970)).

Same architecture from 2 to 4080 Processors

Application can scale without being re-architected



- Each NonStop CPU has its own copy of OS, memory and IO Interconnect
- No computing resources are shared
- NonStop message system (Expand), interconnects upto 255 nodes or CPUs
- NonStop server node can have between 2 and 16 single or multi-core processors
- $16 * 255 = 4080$ processors
- The NonStop OS presents a single system image of virtualized resources
- The single system image enables easy access to resources, with minimal loss in latency – applications can scale linearly
- MPP Vs SMP

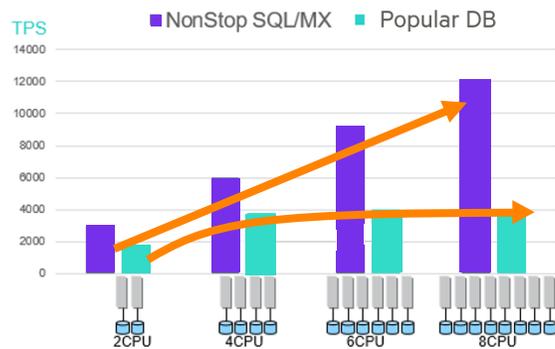
* RoCE = RDMA over Converged Ethernet

Scalability and resilience using multi-dimensional linear scale

• Linear scalability

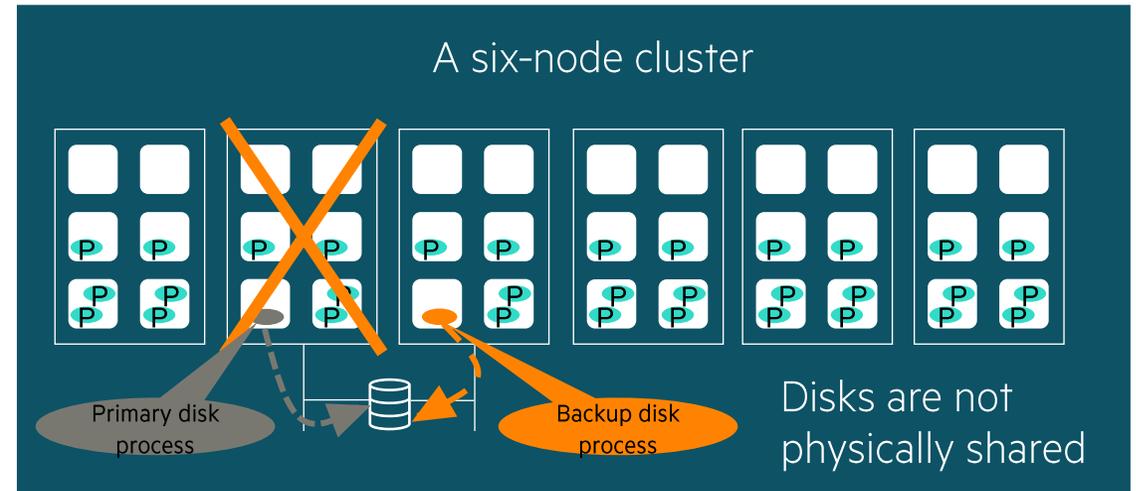
- Unlimited expandability and capacity using “shared nothing”
- Combine vertical (cores) and horizontal scale (nodes) for optimal performance
- Grow from prototype to mission critical without re-architecting or rewriting the application
- No architectural bottleneck when adding capacity
- No increase in latency for high availability purposes
- Add cores or nodes without stopping the application

Adding nodes translates to 98% increase in throughput



98.3%
Linear
Scalability

Better design drives better outcomes



Only one node owns the disk at a time (shared nothing), other nodes ship the I/O to this node via message passing. If the node fails, the backup process becomes the owner without visible impact to the application. Combining “shared nothing” and “process pair” for storage access relieves the NonStop architecture to use replicas or complex distributed locking.

Using replicas as it is generally done in the rest of the industry raises many challenges and limitations such as split brain, CAP (*) theorem restrictions, need for consensus algorithm, double writes or ghost I/O and overall solution limited scalability

* CAP stands for consistency, availability, and partition tolerance

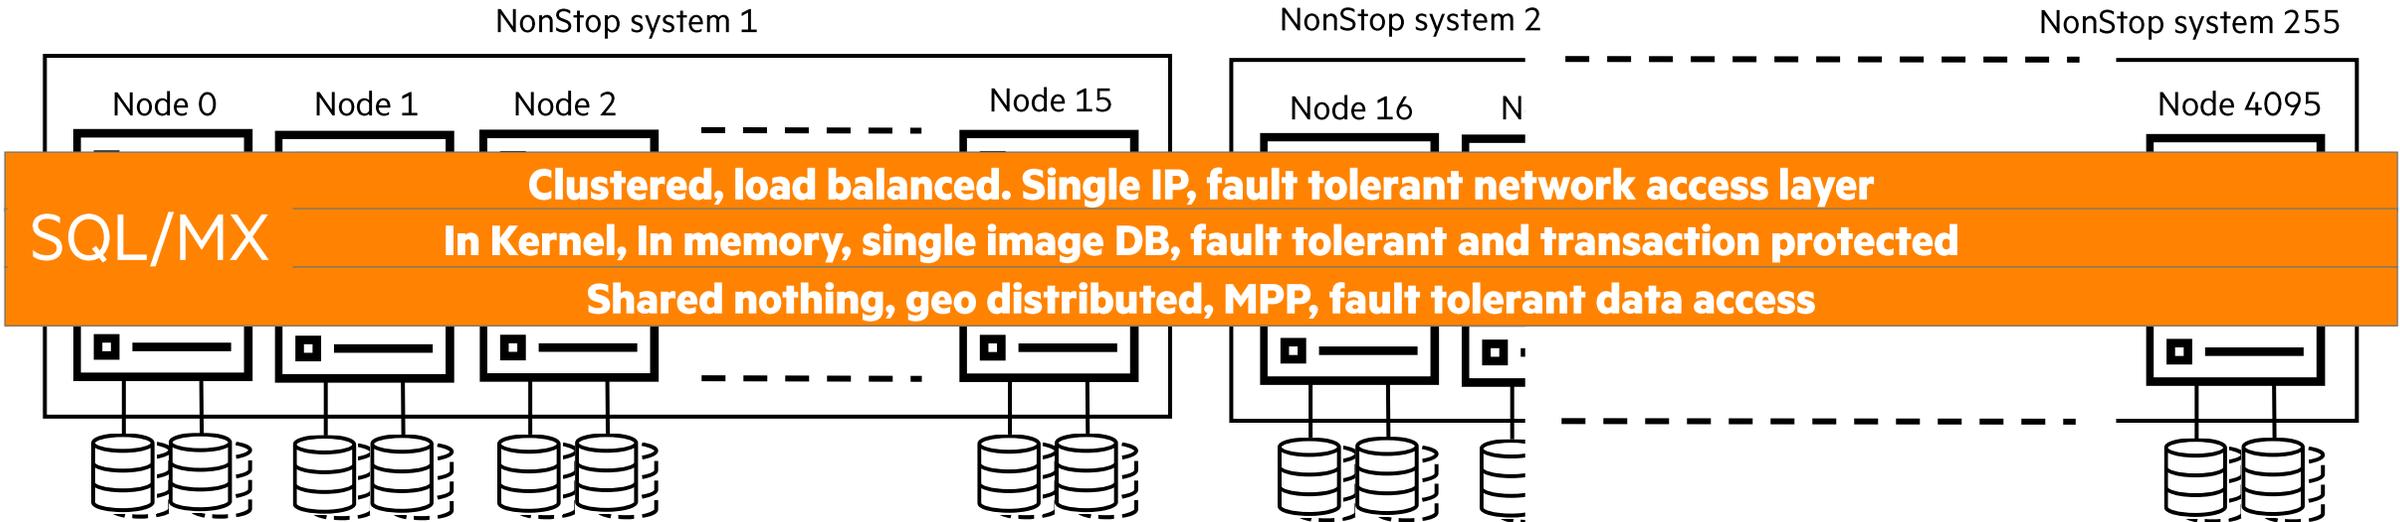
Database and Transaction Integrity

It's about the DATA



The SQL/MX unique architecture

An RDBMS that was not designed to run only on a single computer



- Each node runs its own copy of the OS but that is not visible to the end user who sees a single system
- Each NonStop system can have 16 nodes, each with 2 to 6 cores
- The SQL engine is automatically available as a shared library within each node when starting the node
- Each data access is fault tolerant with processes based mirrored I/Os via backup nodes
- Each node is active, participating in massively parallel execution of the queries

HPE NonStop SQL/MX database advantages



Massive scalability

Single database image across more than 24,000 cores



Parallel processing

Leverages NonStop MPP architecture



Absolute data integrity

Checksums, ACID, business continuity



Availability

Online manageability, NonStop fundamentals, AL-4



Mixed workload support

OLTP, batch and OLAP query workloads



Standards based relational database

ANSI compliant, JDBC, ODBC



Virtualized data access



Maintain data integrity in all scenarios

TMF: Transaction Management Facility

- TMF is a fault-tolerant disk and transaction manager that provides transaction atomicity and database integrity
- All database modifications are captured in the NonStop **TMF audit trail**
 - Before and after images efficiently captured
 - Guaranteed log of database changes
 - Audit trail can be used to replicate and rebuild the database
 - Single logical audit trail across an entire system
- The audit trail which enables:
 - Online database backup
 - Online database rebuild
 - Point in time recovery
 - Database replication
 - Change Data Capture
- If any part of the transaction fails, the entire transaction fails

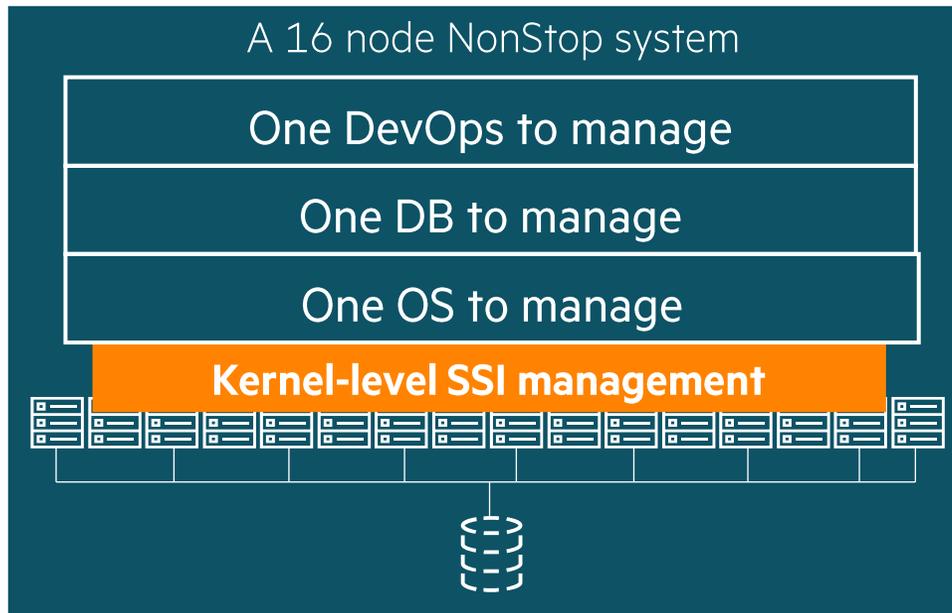
Because TMF is fault-tolerant,
it ensures zero data loss
regardless of any type of
failures

NonStop Database and
Middleware take advantage
of TMF to provide their own
data integrity

NonStop Platform and its DB manage complexity like no other

• Full-system cluster & Single System Image

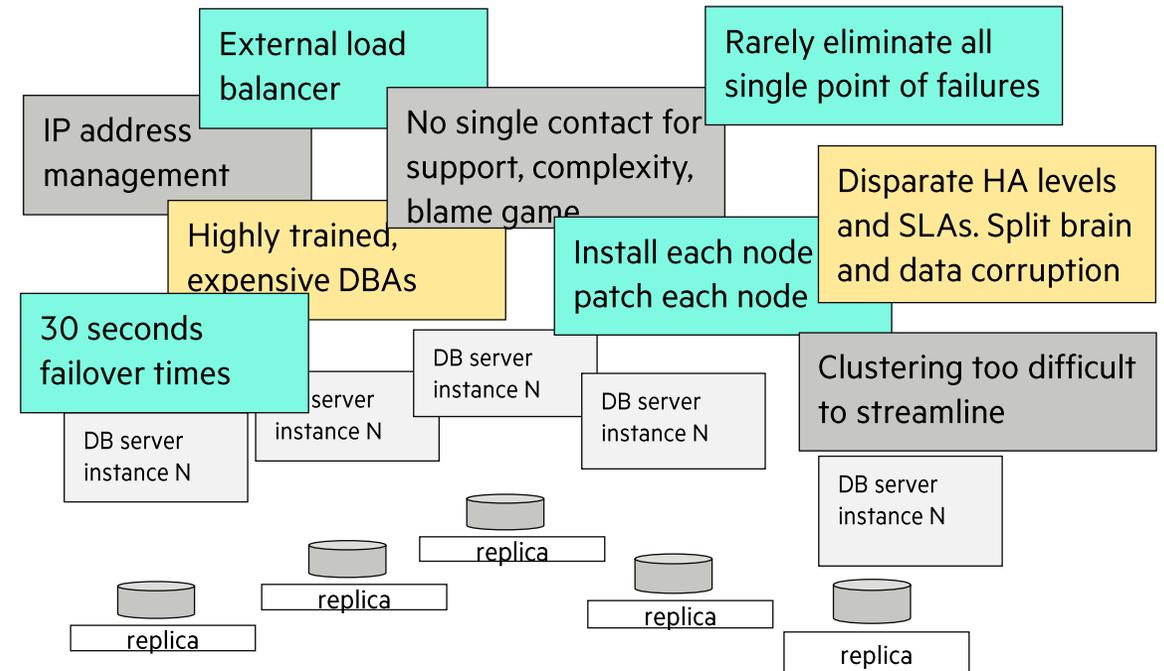
- Simpler for the application, relieved of clustering efforts
- Simpler for administration to manage a single system
- Simpler to secure a single system
- Only a small team required to manage NonStop



Kernel-level SSI management is the most desirable solution to manage clusters (*)

Avoid unnecessary complexity

Other “assembled” clusters



Handling Cyber attacks, System and Datacenter Failures

It's about Security, Compliance & Business Continuity



HPE NonStop Business Continuity & strategy

Provide high performance, reliable capabilities

Business Continuity

- Enable continuous availability (lowest RTO)
- Enable zero data loss where required (lowest RPO)
- Enable sophisticated architectures, from active/passive to active/active
- Facilitate system migration and upgrade without application downtime

Data Integration and Application Integration

- Offer sophisticated heterogeneous data replication and integration in real-time
- HPE NonStop technology makes localized fault-tolerance a reality
- HPE Partner products and NonStop RDF extend this technology to geographic fault-tolerance



Remote Database Facility - RDF

Extending NonStop's Fault Tolerance to Disaster Recovery

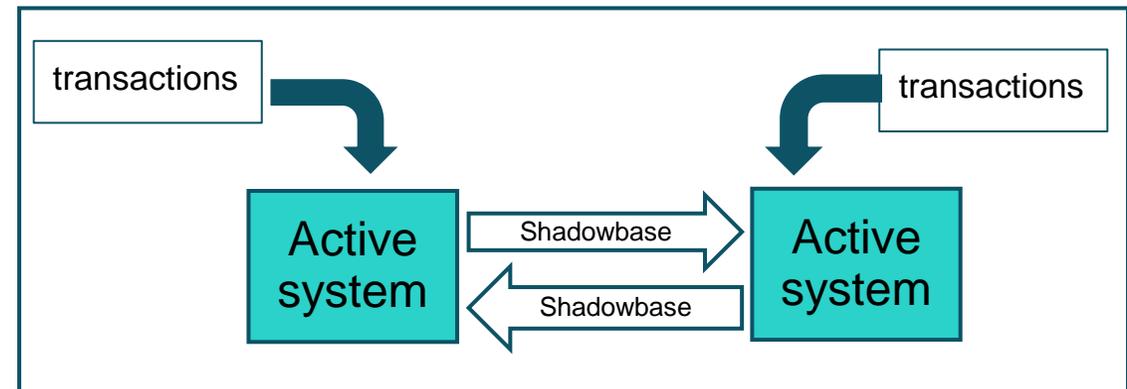
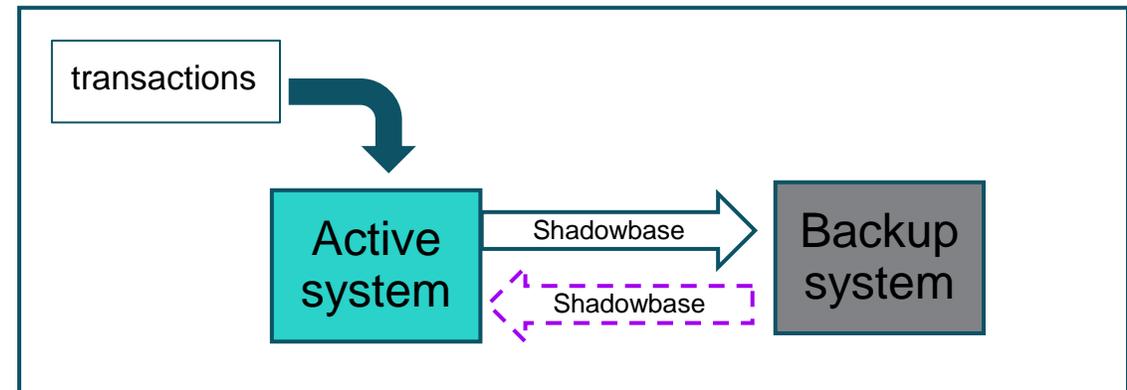
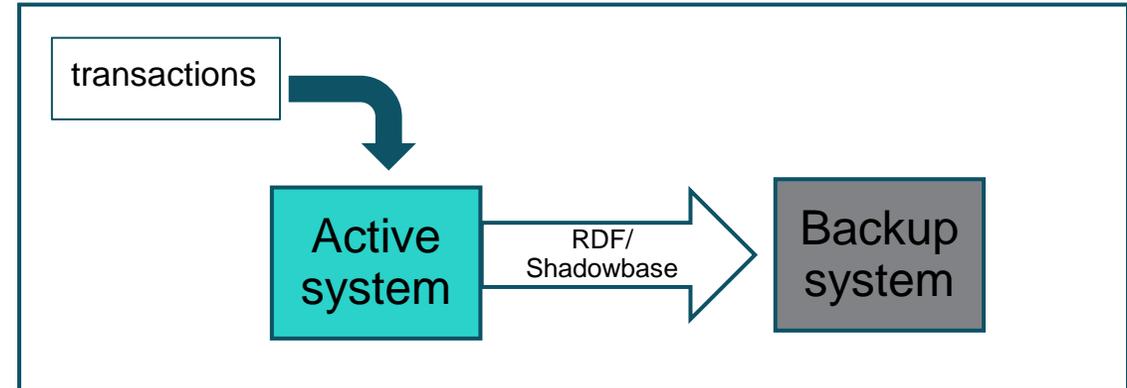
- Remote Database Facility, RDF replicates audited db changes to identical target db on one or more target system
- Provides active/passive (A/P) architecture solution for disaster recovery
- Fast, scalable, asynchronous NonStop-to-NonStop data replication, with low CPU overhead
- Can replicate data over unlimited distances
- RDF Uses –
 - audit trail generated by NonStop TMF
 - operating-system level services
- Backs out any incomplete transaction on the target system (when it takes over processing)
 - If this is not acceptable – there is Zero Lost Transaction (ZLT) which ensures no data loss



- TMF commits all updates to the primary and remote mirror audit trail disks, therefore NO transactions (data) are ever “lost”

HPE NonStop Continuous availability

- Active/Passive
Also referred to as uni-directional replication
- Active/almost-Active
Also referred to as Sizzling-Hot-Takeover (SZT) with bi-directional replication
- Active/Active
Also referred to as bi-directional replication



Security and Compliance

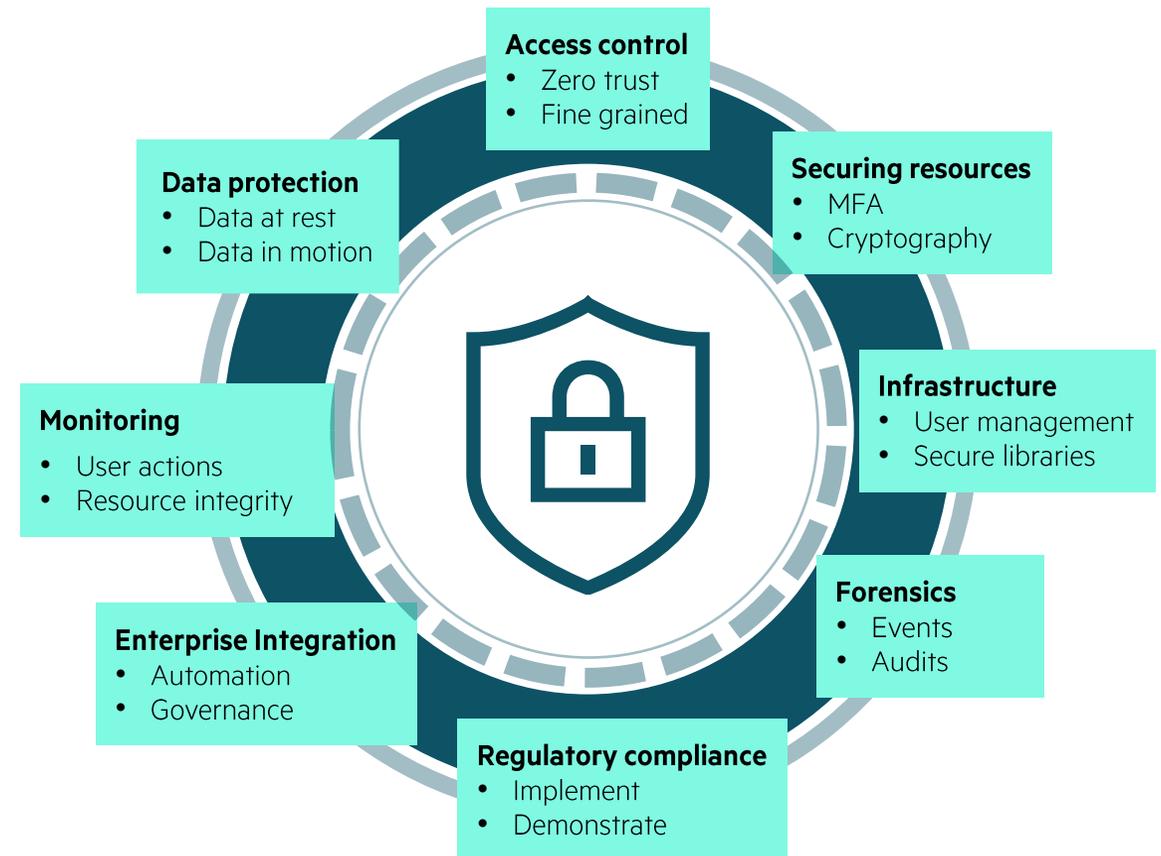
• Security

- Comprehensive security portfolio for the full s/w stack
- Unique OS design that offers superior security to the SW
- Out of the box security infrastructure for customers and applications to protect and secure the resources
- Implement digital resilience using Zero Trust security design
- Advanced forensics to trace user actions
- AI/ML based monitoring to detect the “low and slow” attacks
- Integrate easily with your enterprise’s authentication, monitoring and other security infrastructure

• Compliance

- Easily implement industry standard compliance standards (PCI, HIPAA, GDPR, DORA,...)
- Demonstrate compliance through reports and tools

Increased and comprehensive security



Putting this all together

Explaining the NonStop Value Proposition – to your Management, LoB leaders, Decision-Makers



HPE NonStop Systems: Always-On and Always-Adapting

High-performance, distributed data transaction-processing software stack

Fault tolerance
Ranked as AL4 by IDC*

Massive Linear Scalability,
Robust Security

Best-in-class, Scalable
Database; Unmatched
Data-Integrity

Integrated Software Stack,
HPE Support & Services

High Performance, True
Business Continuity



Cloud Experience
HPE GreenLake, Hybrid Cloud exp.

Reduced Risk, Costs and
Efforts

Vast Eco-system of
solution/system partners

Open-Systems, Modern
apps, DevOps support

As-a-Service, Flexible
consumption models

NonStop Partnership– It’s a Beautiful Thing!



HPE Partner and Customer Use Only

© 2022 Hewlett Packard Enterprise Development LP

Thank you for attending this talk

TBC23-TB54: What Makes NonStop non-stop;

The NonStop 101 rendezvous

Vikas.kapoor2@hpe.com

Marketing Manager for HPE NonStop



HPE Slides and Materials Usage

This content is protected

This presentation is the property of Hewlett Packard Enterprise and protected by copyright laws of the United States. The material in this presentation is provided to attendees of the NonStop Technical Boot Camp 2023 as part of their registration and attendance at the event. Attendees are free to use this material and share it with others within their own company.

This material may not be quoted, copied, communicated or shared with third parties or mutual customers without permission from HPE. To request permission to share material in this presentation outside of your company, send an email to mark.pollans@hpe.com explaining the usage you are intending and your request will be considered.

