

HPE NonStop OS

Provide the availability and scalability advantage to your business at a low TCO



Contents

HPE NonStop Operating System	3
The NonStop architecture	4
The NonStop OS.....	5
Architected for linear scalability	5
Single system image.....	5
Architected for continuous availability.....	6
Architected for performance.....	7
Self-healing capabilities.....	7
Enterprise-class security	8
Data and transaction integrity in the real world	8
Database integrity.....	8
Business-critical Java support.....	8
Supports open standards and modern computing environments.....	9
Processor platform portability and investment protection.....	10
Conclusion.....	10



HPE NonStop Operating System

HPE NonStop systems have been serving the needs of the industry for mission-critical computing requirements for over four decades. They are designed for application high availability, uncompromising data integrity, and near linear scalability using a massively parallel processing (MPP) architecture. NonStop OS is the software foundation of these systems that provides a modern, open, and industry-standard operating environment for applications and helps them leverage the unique capabilities of the underlying hardware platform.

Introduction

HPE NonStop OS provides the fault-tolerant software foundation for HPE NonStop servers. These servers underpin many of the world's most mission-critical applications, including customer point of sale (POS), debit and credit card transactions, automated teller machines (ATMs), electronic funds transfer (EFT), securities trading, healthcare, manufacturing, public safety, and telecommunications. The proven openness, near-continuous availability, and virtually unlimited scalability of NonStop servers are all enabled by the NonStop OS. With its support for key industry-standard application programming interfaces (APIs) and services, including POSIX, ANSI SQL, Java Enterprise Edition (JEE), and many others HPE NonStop provides a distributed service-oriented architecture (SOA)-based solution that will meet your Mission-Critical IT needs today and tomorrow.

You have a wide choice to deploy your workloads on HPE NonStop platform. You can deploy them in a traditional way on the new HPE Virtualized Converged NonStop systems or on the bare metal HPE Integrity NonStop servers based on Intel® Xeon® processors (called NonStop X systems) or on Intel® Itanium® processors (called NonStop i servers). Alternatively you may deploy your workload in your data center built using industry standard IT gears virtualized and managed through either VMware or OpenStack based environments. In all these environments, you get the benefits of the proven mission-critical capabilities of HPE NonStop. With NonStop OS, you can take advantage of the power of parallel processing to achieve improved flexibility and application throughput, making it possible to adapt quickly to your changing business needs. For example, NonStop OS enables critical business application processing to be transparently distributed across multiple processors and multiple systems, either co-located or geographically distributed almost anywhere in the world. The NonStop OS message-based architecture efficiently connects processes running between as few as two to as many as 4,080 multi-core processors distributed over up to 255 network-connected NonStop systems.

The NonStop architecture supports online growth and reconfiguration within a single server or an entire network without disrupting application and database processing. At the same time, NonStop OS delivers fault tolerance through multiple levels of error detection, fault isolation, and workload takeover capabilities for application and database availability on both single site and multi-site, geographically separated, deployments. Due to the NonStop OS, NonStop servers are designed to detect, isolate, and recover from hardware and software failures without affecting applications and their users.



The NonStop architecture

NonStop servers are designed around a tightly integrated hardware and software architecture that enables application fault tolerance and near-continuous availability out of the box, delivering ultimate scalability and mission-critical reliability, availability, and serviceability (RAS) levels.

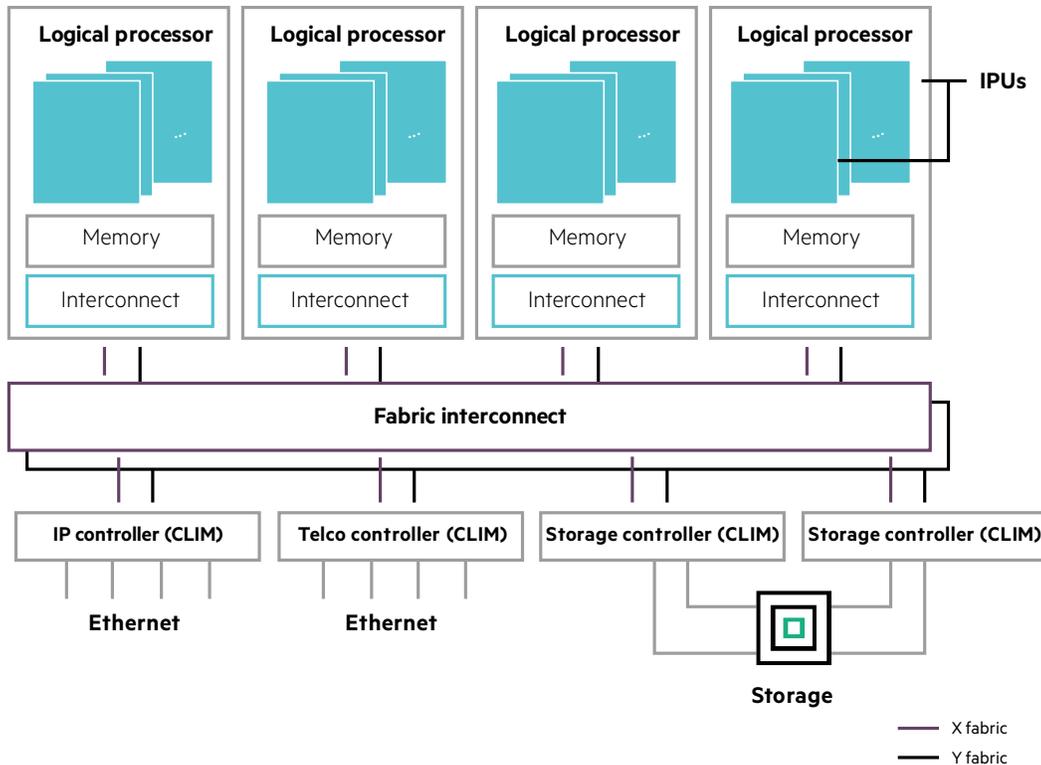


Figure 1. An HPE NonStop System Architecture

A single NonStop system is a cluster of two to 16 independent logical processors, each consisting of one or more instruction processing units (IPUs) and associated memory. A logical processor is a compute node (in HPE Integrity NonStop systems) or a Virtual Machine (VM) in a compute node (in HPE Virtualized NonStop systems). Note that throughout the rest of the document, the terms “processor” and “logical processor” are used interchangeably. Each logical processor runs its own copy of NonStop OS. The I/O peripherals (e.g., storage and communications) are connected to the processors through I/O controllers called cluster I/O modules (CLIMs) and are accessible from all processors in a system. The processors and the CLIMs are connected over a high speed switching fabric which, depending on the hardware platform, can be one of:

1. High speed Ethernet in HPE Virtualized NonStop systems
2. InfiniBand in HPE Integrity NonStop X systems
3. ServerNet (HPE proprietary interface) in HPE Integrity NonStop i systems

NonStop systems are extremely fault tolerant. All critical hardware and software components are configured so that if one component fails, another can immediately take over its work. NonStop hardware is designed to survive any single-point failure. For example, the interconnect is not one fabric, but two, referred to as “X” and “Y.” Normally, both fabrics are used for maximum bandwidth, but a failure of one will still allow communication to continue. All CLIMs are attached to the fabric via two connections—again allowing one to fail without disrupting the system. Disks are configured as RAID 1 mirrored pairs—so that there are two copies of all data—and, in HPE Integrity NonStop systems, each is connected to two CLIMs. The workloads are divided among many independent processors to reduce the fault domain in the event of a failure in hardware, system software, or application software.



In summary, this advanced parallel-processing architecture combined with the high-bandwidth and low-latency fabric enables NonStop systems to deliver exceptional performance, near-continuous availability, high data integrity, and extensive configuration flexibility.

The NonStop OS

The NonStop OS is the software foundation for the higher-level software building blocks of the NonStop system. It provides fault-tolerant, scalable services that enable the key differentiating characteristics of NonStop applications, setting them apart as the solution of choice for mission-critical workloads. The next few sections explain some of these foundational elements of the OS in greater detail.

Architected for linear scalability

As mentioned earlier, a NonStop server node can have between two and 16 single or multi-core processors connected over dual fabrics. A copy of the NonStop OS resides in every processor and provides the infrastructure for software fault tolerance. It monitors and synchronizes all processors, coordinates resource sharing, and interfaces to connected devices. Each processor has its own memory and uses high-speed fabrics (including direct memory access with address bounds checking) to access I/O devices, other processors, and other NonStop servers. HPE Expand software, which is a logical extension of the NonStop message system infrastructure, can be used to interconnect up to 255 local or geographically distributed NonStop server nodes over IPv6 or IPv4 networks transparently and efficiently. The clustering solution provides high-bandwidth, low-latency clustering for up to 24 nodes. The number of processors supported is only one part of the scalability picture. The other part is the added contribution from each additional processor. Symmetric multiprocessing (SMP) systems, for example, begin with good single-processor performance, but with each additional processor, there is a noticeable loss of contribution as a result of the inherent bottlenecks of the SMP architecture. This creates a fixed throughput limit and a decline in financial return on investment (ROI) as each processor is added to the system. Massively parallel processing (MPP) systems do not face the same barriers.

The multi-core logical processors within the MPP-architecture of NonStop systems do possess SMP characteristics, as they represent a single instance of the OS sharing multiple IPUs of a processor. However, this relatively small-scale level of SMP scale-up (up to 4 cores per processor in NonStop i systems and up to 6 cores per processor in NonStop X and Virtualized NonStop systems) is effective in improving system performance because of the limited impact of SMP bottlenecks.

The individual instances of the NonStop Operating System (OS) that run independently on every logical processor share no computing resources with the instances running on any other processor. The small amount of state that needs to be shared by all processors in a system is synchronized through an efficient global update protocol. The processes running on different processors of the system communicate with one another using a lightweight messaging system, over the high-speed interconnect fabric, forming the inter-process communication (IPC) mechanism within the system. As a result, addition of a new processor to the system adds almost linearly to the processing throughput of that system.

In summary, NonStop systems are based on a “shared-nothing” MPP architecture with an added SMP element and demonstrate near-linear scalability and flexibility matched by very few in the industry.

Single system image

The NonStop OS presents a single system image of virtualized resources that enables applications to interact with the services within the entire network of interconnected NonStop systems using the highly efficient message passing mechanism as if it were a single system. No matter how large the configuration or how geographically dispersed the system resources are, the single system image enables all users and applications to have efficient access to files, databases, processes, communication networks, and hardware resources anywhere within the network as long as they have appropriate authorization.

The system virtualizes resources by representing them as logical names. A process running in a NonStop system can be assigned a unique logical name that abstracts out its addressing information such as its process ID (PID) and processor number. Such “named” processes are reachable from the rest of the processes in the system through their unique names, regardless of their actual location. This provides a virtualized environment within the system whereby one process doesn’t need to know the actual location of another, and the primary (request-handling) member of a process pair (explained in the next section) can change from one process to another without requiring the client application processes to be notified of the change. Therefore, services and applications can run efficiently on any processor or group of processors within the NonStop system. Even when NonStop systems are expanded to include thousands of processors on a global scale, processes are globally addressable from anywhere within an Expand network by simply including the name of the process or process-pair system as part of its logical name.



NonStop OS and NonStop compilers automatically create an efficient re-entrant code execution environment that makes replication and parallel processing highly efficient. This means that a system manager can distribute processing functions and data according to the requirements of the business, taking into consideration such things as local staffing needs and national boundaries. The single system image provided by the NonStop OS also supports system management either from a single location or from multiple locations. This capability provides flexibility in locating the operations staff; assigning cooperative, distributed, or centralized management responsibilities; and establishing remote, lights-out environments.

Architected for continuous availability

NonStop Operating System is designed and built for the highest service levels of business application availability. NonStop OS supports the process-pair technology used to create pairs of software fault-tolerant processes that run on different processors to perform critical system tasks such as disk volume management. Figure 2 shows primary processes “A” and “B.” Each one processes incoming requests and sends regular updates with respect to its state to its backup process (“Backup A” and “Backup B,” respectively) through a method called “checkpointing.” Upon detecting a failure of the primary process (due to either hardware or software fault), the backup process takes over as the primary and starts processing using the context it received through checkpointing. The backup processes require only minimal processing to consume checkpointing data. The primary benefit is that the backup process has the context, and upon abrupt termination of the primary process, can continue processing using that context.

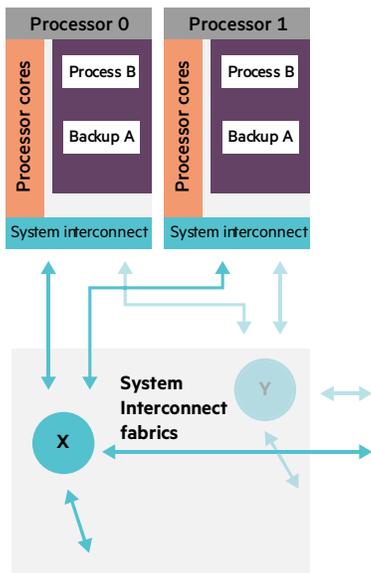


Figure 2. NonStop process-pair technology

Writing process pairs does require some specialized coding; the architectural model is that they are intended primarily for fault-tolerant delivery of certain key system services. Application-level availability is supported through the Transaction Services Massively Parallel (TS/MP) middleware, which removes the need for any specialized application coding.

When a processor does fail, the NonStop OS instances running on the surviving NonStop processors automatically detect the failure and remove the faulty processor from the configuration. The NonStop OS then notifies higher-level processes and applications so that they may take appropriate actions. This could mean, for example, backup processes taking over tasks from primary processes that no longer exist or TS/MP instantiating additional (server class) processes on other processors in the system to handle the additional load.

For applications and services that can tolerate a brief outage and, possibly, even a loss of context, the NonStop OS includes a built-in Persistence Monitor that will automatically re-launch a fresh copy of the process if the original one fails.



The following important features of the NonStop OS further improve system availability:

Automatic system monitoring: On Integrity NonStop systems, the management subsystem provides automatic monitoring and diagnosis of processors, interconnect fabric, I/O devices, power and cooling subsystems, and enclosure environments. It can also automatically initiate diagnosis without disrupting application availability, take corrective actions as needed, and optionally contact the HPE Support Center to report potential problems and initiate on-site service calls for critical repairs.

Standards-based distributed management facilities: NonStop OS includes a systems management infrastructure that is based on the Distributed Management Task Force's Common Information Model (DMTF/CIM). This provides the foundation for system and network management applications. Using distributed systems management, NonStop servers can participate in open, standards-based operations management environments using open standards such as the Simple Network Management Protocol (SNMP).

Architected for performance

The NonStop OS and NonStop system architecture are designed to increase the total throughput performance by both scaling out and scaling up to accommodate the smallest to the largest and the most demanding customer workloads. NonStop servers scale from two to 4,080 multi-core processors combined with dedicated storage and networking hardware to deliver unparalleled business application throughput.

As the software foundation of a vertically integrated hardware and software platform, the NonStop OS works closely with other system components to achieve high efficiency for complex computing operations. For example, database operations that include real-time extract, transform, and load (ETL); large-scale online transaction processing (OLTP); enterprise application integration (EAI) tasks; real-time analysis and reporting; batch extracts; complex direct querying; and online database maintenance all can occur simultaneously. You can perform most of the common maintenance tasks on the database online while maintaining write access to all data.

The integration between the HPE NonStop SQL database management system and the NonStop OS enables tight control over the concurrent mixed-workload environment, making the NonStop SQL database unique in effectively handling all types of workloads executing concurrently across a potentially very large cluster. The NonStop system software plays a critical role in allocating processor, disk, and I/O resources in complex environments where competing priorities have to be dynamically monitored and processes deemed to be of the highest priority need to be given precedence without starving other work. Additionally, many key NonStop SQL database operations take place within the disk access method (DAM) process, which in turn, leads to performance efficiencies that have a positive impact on very large database (VLDB) and real-time database performance.

Self-healing capabilities

Apart from processor level redundancy achieved through the MPP design, a NonStop system has hardware redundancy designed to survive any single-point failure. As illustrated by Figure 1 on page 4, connectivity is not through one fabric, but two, called "X" and "Y." Both are used for bandwidth purposes, but a failure of one will still allow communication to continue. All CLIMs are attached to the fabric via two connections, again allowing one to fail. Disks are configured as RAID 1 mirrored pairs—so that there are two copies of all data—and are actually connected to two CLIMs, and so on.

When a failure does occur, transaction processing continues using the surviving resources. For hardware, the service management subsystem can initiate low-level problem diagnostic procedures on the affected parts to isolate the problem further, repair it if possible, and optionally notify one of the HPE global support centers. At the software level, the surviving members of process pairs continue to provide service and management processes, such as TS/MP Pathmon, and the OS persistence manager recreate processes as needed to restore capabilities. Once the failed resource has been restored to health, it is reintegrated into the system automatically.

This design contributes to the self-healing capabilities of the overall system thereby acting as a key enabler in providing continuous uptime for solutions hosted on NonStop systems. Users and other servers are rarely even aware of system problems.



Enterprise-class security

NonStop provides the enterprise-class security infrastructure and products necessary to deploy solutions in environments with stringent security requirements. The Safeguard subsystem, which is integrated into the NonStop OS, forms the foundation for user administration and other security management services. Safeguard allows you to define and implement a strong and fine-grained security policy for your NonStop system through features such as user account management, fine-grained file access controls, and auditing. You can further customize rules for user authentication, password quality, or access control, using a Security Event Exit Process (SEEP). You also can use XYGATE User Authentication to support two-factor authentication and to integrate NonStop systems as clients into your corporate wide authentication, authorization, and accounting (AAA) solution such as LDAP.

The NonStop OS provides secure shell (SSH) and Transport Layer Security (TLS) for protecting data in motion. HPE NonStop Volume Level Encryption (VLE) is an optional product that provides a fully integrated encryption solution for storage devices connected to NonStop systems.

You can centrally monitor your NonStop system security audit trails and event logs using XYGATE Merged Audit (XMA). XMA also supports real time streaming of event data to Security Information and Event Management (SIEM) enterprise solutions, such as Micro Focus ArcSight, which can store and analyze NonStop event logs and security audit trails in conjunction with other data center activity.

A number of third-party solutions are available that provide additional options for tailoring NonStop security configurations to specific environments. NonStop security products from both HPE and partners enable you to secure your computing infrastructure, defend against ever increasing security threats, and help you comply with industry-specific laws and regulations such as the Payment Card Industry Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), and EU's General Data Protection Regulation (GDPR).

Data and transaction integrity in the real world

NonStop OS is the product of decades of expertise in implementing highly distributed computing solutions that provide unmatched data and transaction integrity. NonStop Integrity servers are built on Intel Xeon or Intel Itanium processors and other hardware building blocks that provide superior RAS features.

NonStop OS tightly integrates with the hardware error-checking foundation and implements additional software data and transaction integrity. When an error is detected, advanced error-containment techniques have been designed to prevent data corruption from being propagated to application databases, other processors, or external computer systems.

NonStop OS also integrates with the NonStop servers' built-in environmental sensing and management processors, providing automatic system restart and recovery features that can protect application databases against such events as power outages or potential errors that might be caused by extreme temperature variations. Additionally, the NonStop OS extends the fault-tolerant software concepts such as the process-pair and persistent process techniques explained in the earlier sections, to enable the integrity and continuity of business applications.

Database integrity

NonStop OS provides the file system and the infrastructure for NonStop SQL/MX and the NonStop SQL/MP relational databases. To ensure database integrity, the NonStop OS includes the NonStop Transaction Management Facility (NonStop TMF) software, which provides distributed two-phase commit protection for local and distributed database changes globally across all NonStop server nodes participating in the transaction and integrates with resource managers on other platforms to support heterogeneous transactions.

NonStop TMF software is designed to help protect the Enscribe file system and SQL databases from damage due to failures. With NonStop TMF software, a failure in an application, a system, or a network component does not result in a corrupted database because of a partially completed database update. Before a transactional update changes the database, an "image" of every affected record or row is captured in memory and written to an audit log. If any part of an update fails or is aborted, NonStop TMF software automatically backs out the change in its entirety, returning the database to its state just prior to the start of that transaction.

Business-critical Java support

NonStop OS integrates with and supports a full range of Java-oriented services, allowing Java components and applications to benefit from the scalability, availability, and manageability advantages of NonStop systems. NonStop provides a high-performance, standards-compliant Java Virtual Machine (JVM), made scalable for enterprise environments. NonStop also supports widely adopted open source Java frameworks, such as Spring, Hibernate, Apache Axis2, and Apache MyFaces.



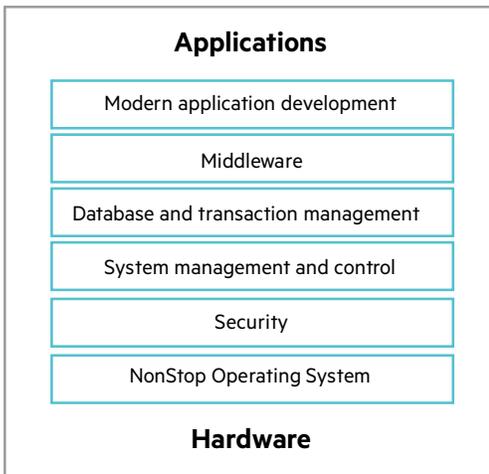


Figure 3. The HPE NonStop integrated stack

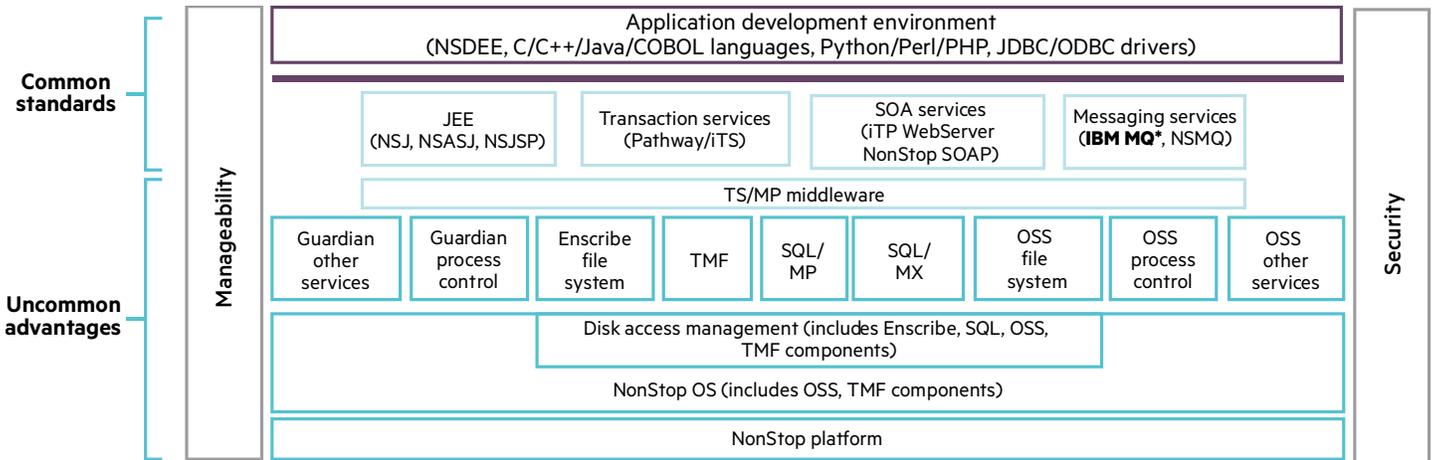
As the foundation for web technology, HPE provides middleware products such as the iTP WebServer and NonStop SOAP 4 (based on Apache Axis2/C). In the domain of Java Enterprise Edition (JEE), NonStop offers middleware containers such as the NonStop Servlets for JavaServer Pages (NSJSP) based on Apache Tomcat and the NonStop Application Server for Java (NSASJ) based on the JBoss Application Server. All these are “deep-ported” middleware solutions that leverage the underlying NonStop platform features of reliability, fault-tolerance, scalability, and manageability out of the box and enable development of scalable and available SOA-based solutions on NonStop. Other Java-based services include the NonStop Message Queue (NSMQ) based on Apache ActiveMQ and the robust Java Database Connectivity (JDBC) APIs to enable application access to the NonStop SQL/MX and SQL/MP databases.

Supports open standards and modern computing environments

NonStop OS provides industry-standard POSIX-compliant APIs and services that form the foundation for application portability, interoperability, manageability, and web-centric computing in a multi-vendor environment. NonStop provides a development and runtime environment, called Open System Services (OSS) that complies with X/Open® Common Application Environment (CAE) specifications, incorporating the POSIX.1, POSIX.2, and some XPG4 standards. NonStop also provides a rich suite of GNU utilities such as the bash command shell, Vim editor, make tool etc., to enable developers with Linux® backgrounds to work on the NonStop platform using familiar tools and utilities. NonStop supports industry-standard programming languages such as Java, C/C++, and COBOL (in addition to the NonStop proprietary pTAL programming language) and scripting languages such as PHP, Perl, and Python. Business applications for NonStop operating environments, including ones based on the original, proprietary Guardian environment as well as the more modern OSS environment can be developed using the NonStop Development Environment for Eclipse (NSDEE) product that can be integrated into the Eclipse™ Integrated Development Environment (IDE) as a plug-in.

The NonStop OS and system software supports multiple industry-standard middleware services (see Figure 3) that provide the infrastructure for the service-oriented architecture (SOA)-based solutions on NonStop. These solutions automatically benefit from the unique advantages of availability, scalability, reliability, and performance of NonStop servers. To summarize, with NonStop, you not only are able to use platform services, development environment, database, and application environment based on popular software industry standards but also greatly benefit from the uncommon advantages that NonStop offers due to its unique architecture that has evolved and matured over several decades of innovation.





* Refers to IBM WebSphere MQ-Series Server, a product sold by IBM

Figure 4. NonStop OS supports industry-standard APIs used by modern business applications

Processor platform portability and investment protection

NonStop OS has a proven history of hardware architecture portability. Over the last 2+ decades, NonStop OS has been migrated from CISC-based platform to MIPS (in the 1990s), to Itanium (in 2005), and to x86-based (2015) platform architecture, keeping its fundamental architecture intact through this evolution. NonStop OS is designed with a modular structure to isolate the hardware-specific layers from both customer applications and higher-level system software. This approach protects your investment in software applications critical to your business by enabling easy migration of these applications to the latest generation of NonStop servers with little or no modification required. HPE Virtualized NonStop offers more flexibility by enabling applications built for the NonStop X architecture to be deployed in a virtualized and cloud environment managed by OpenStack or VMware. By doing this, you can now deploy NonStop workloads in an IT environment built using an enterprise’s preferred IT hardware and leverage operational benefits offered by standardization of IT.

The standards-based architecture of HPE NonStop servers and the NonStop OS provide exceptional system performance, investment protection, and low total cost of ownership (TCO) by leveraging standard hardware and software components. The fault-tolerant design, near-linear scalability, manageability, security, and openness provided by NonStop OS are the pivotal enablers of its industry-leading TCO advantage. The NonStop OS has time and again proven its ability to leverage technological evolution taking place in IT and offer long-term investment protection for HPE NonStop customers.

Conclusion

NonStop systems are based on a shared-nothing, massively parallel processing architecture designed to provide continuous application availability, linear scalability, and data integrity. NonStop OS provides modern, secure, industry-standard operating environment and developer tools to enable easy development and migration of software applications and solutions on the platform. Higher-level HPE NonStop products, such as SQL/MX database, TS/MP application management software, and Java-based middleware, leverage the capabilities of NonStop OS to deploy highly available and scalable service-oriented architecture-based solutions on NonStop systems. A rich suite of software solutions are offered on the NonStop platform by independent software vendors (ISVs) for banking, financial services, communication, defense, manufacturing, and several other industry verticals.

NonStop systems have been the mainstay of mission critical computing for the industry for over four decades and have consistently kept pace with the evolving technological trends. NonStop OS has been the key differentiator that enabled NonStop systems to provide unique availability advantages to its long standing customers.



Learn more at
hpe.com/info/nonstop



Make the right purchase decision. Click here to chat with our presales specialists.



Sign up for updates

© Copyright 2010, 2015, 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel Itanium and Intel Xeon are trademarks of Intel Corporation in the U.S. and other countries. Java and Java Compatible logo are registered trademarks of Oracle and/or its affiliates. X/Open is a registered trademark and the X device is a trademark of X/Open Company Ltd. in the UK and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. All other third-party trademark(s) is/are property of their respective owner(s).