

HP NonStop Servers Migrated to x86

May 2015

Fault-tolerant HP NonStop servers have come a long way from their origin as Tandem Computers. Along the way, NonStop servers have moved from proprietary hardware and software to become more and more open, with a Unix-like operating system and support for modern programming languages such as Oracle Java. HP has now taken the final step in moving NonStop servers to commodity technology by migrating these servers to an x86 architecture. The new servers are known as HP NonStop X servers.



Tandem Computers – The Origin of HP NonStop Servers

Forty years ago, Tandem Computers introduced the first fully fault-tolerant computer with no single point of failure. Based on proprietary hardware and software (the Guardian operating system and the TAL programming language), Tandem focused on providing continuous availability mainly for large financial and telco transaction-processing applications. At one time, it became one of the fastest growing publicly traded companies.

Tandem Computers was ultimately acquired by Compaq Computers in 1997, and Compaq in turn was acquired by HP in 2001. Tandem computers are now known as HP NonStop servers.

Tandem's Fault-Tolerant Architecture

The original Tandem fault-tolerant architecture is still in use today in the HP NonStop servers. It is based on a shared-nothing architecture. Up to 16 processors can be configured in a single node. The processors are interconnected by a dual high-speed backbone fabric called ServerNet. Applications communicate via a messaging system – they never use shared memory or any other kind of interprocess communications.

If a processor fails, all of the applications are moved to surviving processors. There are two methods for moving a process:

- **Check-Pointed Processes** – each critical process is backed up by a copy running in another processor. The backup process is kept in the same state as the primary process via the checkpointing of state changes. Should the primary process fail, the backup process immediately takes over production processing.
- **Process Monitor** – For less critical processes, the Pathway monitor can detect a process failure and restart the process in another processor. More generally, Pathway manages “server classes.” A server class is a set of like servers running in multiple processors. Pathway distributes requests to the various servers in a server class. If a processor fails, subsequent requests are distributed only to the surviving processes. Pathway can expand or reduce a server-class pool as required.

In all cases, recovery from a failure of any kind is measured in subseconds to seconds.

Tandem supported two database systems – an hierarchical file system, Enscribe, and a relational SQL database, SQL/MP (MP for massively parallel). To guarantee no data loss, all disks in a Tandem system are mirrored. Transactions are applied to the disk mirrors by Tandem's Transaction Management Facility (TMF). TMF uses the two-phase commit protocol to ensure that data changes are made to both disks of the mirror, or else the transaction is aborted and changes are made to neither disk.

Pathway, TMF, and the disk processes are examples of critical processes that are protected by checkpointing.

The Tandem architecture is massively scalable because of its shared-nothing approach. Up to 255 nodes, each containing 16 processors, can be connected in a network and can provide the full capacity of all 4,080 processors.

The Tandem architecture was awarded the coveted classification of AL4 under the IDC (International Data Corporation) availability classifications. According to this classifying method, the following levels are defined:

- AL4 – a fault-tolerant server in which switchover to alternate resources is not perceptible to end users.
- AL3 – a clustered server in which a short outage is required to fail over to alternate resources.
- AL2 – workload balancing in which a user request is redirected to an alternate resource following an outage.
- AL1 – the need to switch to redundant resources following an outage before processing is resumed.

According to IDC, “[Tandem Computers (now NonStop servers)] provide true fault tolerance, enabling continuous data processing, even in the event of the failure of one hardware or software component. The end user experiences no perceived interruption based on the use of fault-tolerant servers. In this level, the combination of multiple hardware and software components allows a near-simultaneous failover to alternate resources so that business processing continues as before without the interruption.”

HP NonStop Servers

Today's HP NonStop servers follow the same architecture as the original Tandem systems. However, there have been some significant improvements:

- The original Guardian operating system has been separated into a NonStop kernel on top of which is a Guardian personality. A Unix-like personality (OSS – Open System Services) also runs on top of the NonStop kernel. This allows NonStop systems to support applications written in Unix as well as being able to run open-source software.
- Modern languages are now supported, including C, C++, and Oracle Java.
- A modern Eclipse development environment is supported.
- A new SQL database, SQL/MX, was developed which more closely complies with the SQL specifications.
- NonStop servers were ported from their original proprietary hardware to HP's Itanium blades so that common hardware is used for all of HP's systems – HP-UX, NonStop, OpenVMS, Linux, and

Windows. The only difference between a NonStop Itanium blade and a standard HP Itanium blade is an interface to the ServerNet backbone fabric.

Welcome to the NonStop X

At the NonStop Technical Boot Camp held in November, 2013, HP announced its plans to port NonStop to Intel's x86 Xeon processor. With this move, NonStop would now run on commodity hardware. The new NonStop system was named HP Integrity NonStop X (the Itanium versions are now known as HP Integrity NonStop i).

Several improvements came with the move to the x86 architecture. Like the NonStop Itanium processors, the NonStop X processors are configured as standard HP blades. However, the NonStop X blades are only half-height, allowing 16 processors to fit within a standard HP c7000 chassis (only eight NonStop i blades can be put into a c7000 chassis). The new blades use the Intel Xeon e5-2600v2 series of Intel microprocessors with four cores active. Available memory sizes range from 64 GB to 192 GB per CPU, giving a per-system total of over three TB of main memory. Tests with an Order Entry benchmark show that the new x86 NonStop processors are 50% faster than their Itanium counterparts.



As in the earlier NonStop systems, disk and communication I/O are provided by CLIMs (CLuster I/O Modules). Up to 56 CLIMs can be configured in one system. A communication CLIM can handle four 10GbE Ethernet channels and one 1GbE Ethernet channel. A storage CLIM can support 100 SAS hard disk drives or solid-state disk drives. It can also support Fibre Channel connectivity for SAN attached storage (such as HP XP7 and HP XP P9500 storage subsystems), LT06 tape drives, and virtual tape.

Perhaps the most dramatic improvement for the NonStop X was the replacement of the dual ServerNet backbone fabric with a dual industry-standard InfiniBand fabric running at 56 Gbps for inter-processor communications. InfiniBand has 25 times the capacity of ServerNet. This greatly reduces interprocessor message and direct-attached I/O latency times. Even more important for the future, connecting NonStop servers with InfiniBand means that they also potentially can communicate with HP-UX, Linux, Windows, and OpenVMS systems. This opens the door, for instance, to using Linux servers to front-end NonStop transaction processing systems.

The new operating system is the L-series. The first operating-system version is L15.02. The naming convention is L<year><month>. Thus, L15.02 is the operating-system version that was released in February, 2015. The new operating system has been optimized to take advantage of the very high bandwidth and low latency of InfiniBand. A compiler is provided to recompile programs to run on the NonStop X.

A major hurdle for the NonStop X x86 operating system was that the NonStop software supports the big-endian architecture but the x86 is little-endian. To accommodate this, NonStop X compilers store all program variables in big-endian format. Therefore, messages sent between NonStop X and NonStop i systems require no transformation.

A partner Beta program was introduced in early 2014. Virtually all major middleware and solutions vendors ported their products to NonStop X to have them certified. It was gratifying that there were a minimum of problems in the product ports. NonStop X proved itself to be ready.

The NonStop X NS7 X1 (NS7) is Formally Announced

On March 30, 2015, HP announced the general availability of the HP Integrity NonStop X NS7 X1. This is the first system in the new NonStop X NS7 product line based on the x86 architecture.

It is expected that more systems in this family will be announced. This should include smaller systems running on two-core x86 microprocessors as well as larger systems running on six-core x86 microprocessors. However, all systems will retain the NonStop fundamentals of fault tolerance, availability, and scalability that NonStop customers have come to expect – over the last forty years.

Summary

In the early Tandem days, it was said that Tandem computers were the best-kept secret in the IT industry. Almost no one outside of big financial and telco firms had ever heard of Tandem or were even aware that fault-tolerant systems even existed. With HP's firm commitment to the NonStop product line, hopefully this perception will become a thing of the past. The NonStop X and NonStop i family of processors provide customers with the flexibility and focus of an unparalleled portfolio of fault-tolerant systems for mission-critical applications.