Tandem NonStop History

From a Hardware Maintenance Perspective

Contents

Introduction

A number of sources detail the founding of Tandem Computers Inc., its ethos, and other reasons why 'Tandem' is held in high regard by its users and former employees, but little can be found which describes the architecture from a hardware maintenance perspective. This is somewhat understandable. With any good computer system, the physical hardware should be transparent to the end user and given that fault-tolerant computers, particularly Tandem NonStop computers, are meant not to fail they should not require maintenance - right?

This paper started as a single page website overview, but rapidly developed beyond 'overview'. However, there are holes in the history; the Tandem Unix family and the entire dedicated telecommunications (Telco) sector product line for example.

This is being released as a 'work in progress' with more detail to be added. Feel free to contact the author to correct any errors or omissions.

Tandem Concepts

Before delving into the specifics of the generations of Tandem NonStop systems, the concepts of what it was that made Tandem Computers Inc. fault-tolerant system hardware unique should be covered.

From the outset Tandem's NonStop system was designed to have multiple independent processors, each with dedicated memory, dedicated power supplies, dual-ported redundant (mirrored) disk drives with dual-ported mirrored disk I/O controllers, and dual ported general I/O controllers - all providing automatic high-speed "failover" in the case of a hardware (or software) failure.

Processors connect and communicate over an Inter-Processor Bus (IPB), originally referred to as the Dynabus, with a pair of processors sharing details of process execution and the processor pair carrying out regular "I'm alive - are you?" checks. In the event of failure of one processor the alternate in the processor pair would assume execution of the failed processors processes and take control of any I/O communication. For redundancy, two inter-processor buses were designed in so that in the event of a failure of one IPB the other IPB would carry all inter-processor communications. These inter-processor buses were known formally as the IPB-X Bus and IPB-Y Bus.

A second dedicated bus was designed and integrated from the outset; the Maintenance and Diagnostic Bus (MDB). And Tandem being Tandem, two MDB buses were included. Called the MDB-X Bus and MDB-Y Bus, these carried the workload of system environmental and status monitoring without impacting on processors. They also facilitated the running of diagnostics on processors (and in later generations all system components). Additionally, the MDB Buses played a part in initial system load and processor reset/reloads.

Both IPB-X/Y and MDB-X/Y buses were intended to be readily physically extended allowing for easy expansion of a system to include more processors than could physically be housed in a single modular cabinet. And much of the success of Tandem could be down to this design feature (known as 'linear system expandability'). Customers could start off with a basic two processor system and increase application workload up to sixteen processors (and, in fact, beyond sixteen processors, but that is beyond the scope of this paper) by adding processors and cabinets to the base system.

Mirrored dual ported disk drives and dual ported mirrored disk controllers allowed for multiple paths to a mirrored disk pair with the Operating System (O/S) automatically trying all possible/available paths to a physical disk before declaring that disk as faulty. Disk data integrity was also assured as all write to disk operations were followed by a read of the written data to ensure that the data had actually been written and was not corrupted in the write process.

Failover was fundamentally built into the Tandem O/S (called Tandem Guardian Operating System and initially designated Guardian 'A-Series') - application software developers did not have to be concerned with this fundamental feature. If an application was started as a fault-tolerant application, Guardian would take care of the fault-tolerance.

With no single point of failure, ‘I’m alive’ processor messaging and automated fail-over, dual processor and maintenance buses and dual ported controllers/disk drives the Tandem system was the system of choice for customers with banking, financial trading, retail, telecommunications and defence applications where failure of that application was mission critical.

Tandem TNSI, TNSII, TXP, EXT and VLX

After just two years in development, the first Tandem NonStop system was delivered to Citibank in the USA in 1976. This system was a pioneer in the establishment of not just Tandem Computers Inc., but how the entire computer industry viewed fault tolerance. And it raised customers’ expectations of computer reliability.

The TNSI to VLX range was a gradual improvement in throughput and reliability with only the VLX introducing new hardware architecture concepts.

Tandem NonStop One (TNSI)

The first generation of cabinetry consisted of processor cabinets (which included Input/Output [I/O] controller slots), I/O only cabinets and patch panel cabinets to manage physical cable connections to peripherals and the outside world - patch panel cabinets could also optionally house a reel-to-reel tape drive. All cabinets had the same physical external appearance, cream with black doors and a smoked plexiglass cover at the front (very 1970’s!) and were bespoke to Tandem Computers Inc.

A minimum system configuration would have one processor cabinet with two processors, two disk controllers, a tape controller, a communications controller for connection to external operator terminals, and a patch panel cabinet with a tape drive installed at the front. Systems could however consist of up to four processor cabinets (sixteen processors maximum) and multiple I/O and patch panel cabinets. And, as Tandem systems could readily be expanded/upgraded from a minimum installation to a fully populated sixteen processor system, a myriad of different configurations existed in the field.

Processor cabinets comprised of a section of (up to) four processors at the top, an I/O section in the centre, four large fans below, and at the bottom processor power supplies (one dedicated power supply per processor). Mains power was supplied by a three-phase power input/filter and distribution section at the rear. The processor backplane incorporated circuitry for the dual-processor buses and dual maintenance buses with four small bus controllers at either end of the backplane. The rear of the I/O backplane was configurable using lengths of ribbon cable to determine which processor connected to which I/O slot. Each of the 24 I/O slots had its own power control card with a small switch to power off the individual controller slot. I/O cabinets were almost identical to the processor cabinets but with the processor card cage and bus circuitry removed and I/O only power supplies were installed.

The front door of the processor cabinets accommodated an array of switches, one array corresponding to each processor in the cabinet, and these were used in loading O/S to the system. LEDs above the switches gave visual indication of processor status/load and activity including a ‘heart-beat’ LED showing that the related processor was running O/S.

Each processor comprised of up to five (depending on the memory configuration) individual open integrated circuit boards with a large number of ‘plug-on’ components. The boards locating by a ‘card cage’ and connected into a backplane with finger edge connectors. Ribbon cables to connectors at the front of the boards provided additional inter-board connection.

I/O controller boards were also of an open integrated circuit board design. Dipswitches on the controller boards set the controller’s address and wire strap pins set I/O channel interrupt priority. Installing a controller with the incorrect address and/or priority set could lead to unwanted shutting down of a system - the highest embarrassment for a customer engineer and potentially very costly for the customer - so great care was needed when replacing a controller.

None of the circuit boards in this generation were hot-plug and required power, to either the individual component or group of components (such as a processor), to be switched off prior to removal/installation

Disk drives were large stand-alone units, of Winchester design (some of removal 'disk pack' type and some hybrid - half fixed platters/half removeable), with multiple cables from each drive to the patch panel cabinets and then to the disk I/O controller boards in the card cage. Removable pack Winchester drives were vulnerable to 'head crashes' due to the almost unavoidable introduction of dust particles when changing packs (the read/write heads on these drives 'flew' at 3-5 microns above the disk surface - a human hair is around 50 microns). A head crash would require thorough cleaning of the drive and replacement of read/write heads - a time consuming process needing an oscilloscope and a delicate touch to manually align the heads correctly.

Tape drives were of reel-to-reel design, initially writing at 1600bpi and later at 6250bpi, and were installed at the front of a patch panel cabinet. These were relatively trouble free provided that moving parts were maintained.

System load and initial start-up was carried out from a separate desk based (it really was a desk) micro-computer/terminal known as the Operations Service Processor (OSP). With two 5.25-inch floppy disk drives, a separate terminal with Cathode Ray Tube (CRT) display and keyboard, and often a hard copy printer for error log reports. This desk/terminal could also be used to load and run offline system diagnostics - though in practice customer engineers never found these diagnostics particularly helpful.

Configuration, preventative maintenance and replacement of faulty/suspect components on these systems required a high level of training in both the O/S and hardware. It also required a degree of manual dexterity/finesse along with physical strength in teasing boards and cables in and out of the system. Customer Engineer training required an introductory course of two weeks, explaining the concept of the Tandem system and basic O/S utilities, followed by an eight-week course covering the hardware known as the 'Level One' course.

Customer engineers (CE's) also needed an understanding of the customer application running on the system as maintenance on Tandem systems was usually carried out while the system was 'live' (Tandem systems were sold as having 99.99% availability with most systems running mission-critical applications). Customers' in-house Tandem operational/support staff would need to have faith that their customer engineer would not bring the system 'down' while working on it. And, should the system be down customers would rely on their customer engineer to get the thing back up again in the shortest possible time. Consequently, good relationships with, and responsiveness to, customers' support staff were (and remain) a key element of the make-up of a good Tandem customer engineer, as was (and remains) the ability to work under intense pressure and keep calm.

From both a maintenance and reliability perspective cabling and inter-component connection issues were the bugbear of this generation of design with multiple connections between circuit boards and backplanes, I/O Channels and external devices. Regular preventative maintenance was essential and would often show up potential problems before they became an issue. This said, reliability was remarkable and the envy of other manufacturers.

Tandem NonStop Two (TNSII)

The second generation NonStop came in 1981 and was an upgrade to the central processing chips, processor boards and memory boards only and otherwise identical to the NSI.

Tandem NonStop Extended Processor (TXP)

The TXP, introduced in 1983, was also an upgrade to the processor boards only and otherwise identical to the NSI/NSII. A new Operating Systems was required - Guardian B-Series.

Tandem EXT

We believe that the EXT was introduced in 1985-1986 but it may have been earlier *(citation needed)*. The EXT was as an entry-level (cheaper) NonStop system, essentially a repackaged TXP with a smaller footprint and limited expansion capability. However, the external appearance was completely different to its bigger sibling. All the elements were there, just smaller: a smaller processor card cage, an I/O card cage with fewer slots and, in a new direction for Tandem, some internal disk drives so that a single cabinet encompassed a complete system.

Tandem VLX

In 1986, the release of the VLX system marked a departure and modernisation of the TNSI/TNSII/TXP external design. With all black rectangular cabinets sporting a thin red horizontal accent at the waist on the front door, the VLX cut a dash in any computer room. However, opening the slick looking front doors revealed that internally the VLX was very similar to the TNSI/TNSII/TXP - with the exception of the processor card cage.

VLX processor boards, loaded with more powerful central processing chips than the TXP, were enclosed in metalwork to cover/protect the circuitry (particularly useful when installing/removing the boards) and with switch interlocks on the all processor boards to power off the processor so that the act of removing any one board in the processor set from the backplane disconnected power to all boards of that processor set. A nod to improving the serviceability of the system and a hint of what was to come in this direction.

Inter-processor bus bandwidth was also doubled which, along with processor board chip improvements, helped bring a near 50% performance increase over the TXP.

The VLX also saw the loss of a separate OSP desk/micro-computer. OSP functions were integrated into a single board in the first processor cabinet, the Remote Maintenance Interface (RMI), requiring only a dumb terminal, connected to a dedicated port, to start and stop the system, and carryout maintenance functions.

System maintenance received some welcome assistance with the introduction of a software utility which integrated hardware diagnostics, error reporting, event logging and automatic fault analysis; the Tandem Maintenance Diagnostic System (TMDS). TMDS was embedded into the operating system and allowed for status interrogation and the running of online diagnostics on most system components. Automatic fault analysis would compare errors against a known list and suggest replacement of hardware components with the fault event being logged and reported to Tandem's support service centre by modem.

The old Winchester technology removeable pack disk drives were also phased out and new modular external disk drive sub-systems, known as the V8 and XL8 disk subsystems, introduced. These were dedicated cabinets housing up to eight drives and connected to the patch panel cabinets by multiple thick copper cables. The cable patch panel in the back of these subsystems was tight so replacing a cable could easily result in the introduction of other cabling faults in the subsystem. However, the individual drives had no field maintainable parts, no requirement for head alignment, and were hot-plug requiring just the raising of an integral handle at the front of the drive to power off and release the drive from its enclosure. As with the VLX processor boards, hard drive maintenance just got a lot easier.

Tandem CLX, CYCLONE and Cyclone-R Range

Introduced with the CLX in 1988, the CYCLONE in 1989 and Cyclone-R in 1991, this range was a complete inside and out redesign and departure from the original 1976 system but kept all the key fault-tolerant hardware concepts of: no single point of failure, dual processor/maintenance buses and dual ported controllers/disk drives. No components were carried over from the TNSI/TNSII/TXP/VLX. Operating System for this range was the C-Series.

Tandem CLX600, CLX700 and CLX800

Intended as a value entry to NonStop computing system and designed for the ‘copier room’ environment, the CLX shared much of the design innovations with the CYCLONE but on a very much smaller scale. A single cabinet (with styling cues from the Design Award winning Tandem EXT system) was of a similar size to a full height filing cabinet and yet encompassed a complete fault-tolerant NonStop system: two processors (ASIC CMOS CPUs), six I/O controllers (two of these being multi-functional [disk/tape/communications]), six SCSI disk drives, dual power supplies, dual fans and all the processor/maintenance bus logic. System start/stop and initial load was carried out from a small operator panel on the front door.

All components were hot-plug with metalwork protecting the circuits and all cabinet components, including processor boards and I/O channel bus controllers, were now enclosed in metalwork with power interlocking so that physical removal of a component automatically disconnected power to that component. Whereas power interlocking on the VLX processor boards was through the use of micro-switches, on the CLX this was achieved very simply by manufacturing the power carrying pins of the connectors to be shorter than the ground and data carrying pins - genius!

Additional processor cabinets could be added for more processing power up to a maximum of sixteen processors by extending the IPB/MDB buses. Also, an I/O only cabinet, externally physically identical to the processor cabinet and known as the MC-8 (it had eight I/O controller slots), could be added (one per processor cabinet) to increase connectivity and came with six additional SCSI disk slots.

A major change heralded by the CLX/MC-8 design was the abandonment of patch panels. I/O controllers now consisted of two parts; a Logic Board (LB) at the front of the MC-8 cabinet, and the second part in the corresponding slot in the rear of the cabinet, the Backplane Interconnect Card (BIC). External devices, disk drives, tape drives and communications cables connected directly to the BIC. This minimised component inter-connection leading to a further improvement in reliability.

Another significant change was that I/O controller addressing and I/O channel priority (rank) issues were eradicated by the hardwiring of slot addresses to the backplane and the fixed ranking of controllers according to their capability. Furthermore, all boards in the CLX processor cabinet and MC-8 I/O cabinet where physically disabled from being put in the wrong slot through physical keying.

Internal disk drives were standard 5.25-inch dual port SCSI drives housed in a bespoke case with its own power supply. Tandem’s ‘standard’ drives were slightly altered at the manufacture stage to ensure that only ‘Tandem OEM’ standard drives could be used. As with the V8/XL8, the CLX internal disk drives were hot-plug requiring just the raising of the integral handle to power off and release the drive from the cabinet. External disk sub-systems could be added via dedicated I/O controllers but seldom were on CLXs.

Tape drives could be connected to a SCSI port on the rear of the multi-function controllers as well as via dedicated I/O controllers. Many flavours of tape drive were available for connection to the CLXs but most popular was 4mm DDS DAT drives with both single cartridge and automatic cartridge loader versions available. The initial release of CLX even came with the option of an internal cassette style tape drive which took up the top left disk drive slot. All older systems of the CLX generic style have a flap in the front door to accommodate changing the cassette. However, the storage capacity of the cassette was tiny and disk drive slots too valuable, so most customers did not take up this option or soon abandoned it.

These systems were designed from the outset to be customer maintainable (‘customer self-repair’ in modern Hewlett Packard Enterprise [HPE] parlance) and the marketing releases for the CLX touted that the only tools required to maintain a CLX system was a hand. And this was pretty much a fact. All field replaceable items could be removed/installed without tools and with system status and power/environment monitoring built into the O/S customer engineers no longer needed to visit site armed with a tool case stuffed full of tools and a multi-meter. And, with TMDS embedded, and extended to all major components of the system, diagnostic tools were built in too.

Having mentioned the customer self-repair concept it has to be said that this was not a popular option with customers. Most customers preferred to pay for a hardware maintenance support contract and have their Tandem customer engineer diagnose faults and replace parts. This, however, would change in later years with many third-party maintenance companies taking advantage of the ease of maintenance and availability on the market of used systems for use as replacement spares.

The initial release system was the CLX600, but this was quickly superseded by the more powerful processor version, the CLX700 and again the same with the CLX800. All cabinetry and other components were identical with just processor and memory replaced. Upgrading was simple; a matter of stopping the system, replacing the processors and memory and loading the new O/S (Tandem were always protective of their product and would often lock the delivered O/S to a specific processor type) - a thirty-minute operation.

The CLX design was, from a customer engineer's perspective, an absolute dream. It was easy to work on and extremely reliable. And the design proved seriously robust lasting, in another guise (K1000/K2000 - see below) until the mid-2000's.

Tandem CYCLONE

Looking to compete with IBM mainframe levels of processing capability and connectivity, the 1989 CYCLONE was an up-scale and complete re-working of the cabinetry, processor and I/O elements to facilitate significantly more powerful processors and more I/O connectivity, and with multiple I/O channels per processor.

CYCLONEs were physically large and imposing systems. Continuing the black cabinet with a thin red accent design from the VLX, they stood out in a computer room and looked almost menacing. And being air-cooled they sounded like they meant business too.

Dedicated processor cabinets housed two processors cooled by three massive fans with power provided by three equally bulky power supplies. As with the CLX, all processor cabinet components, including processor boards and I/O channel bus controllers, were enclosed in metalwork with power interlocking so that physical removal of a component automatically disconnected power to that component.

I/O cabinetry came in the shape of the all new Multi-Channel 32-Slot (MC-32) I/O cabinet. The same physical size and shape of the Cyclone processor cabinet, the MC-32 I/O Cabinet had capacity for up to 32 I/O controllers split over four sections known as quadrants. Multiple MC-32s could be connected to each processor cabinet and as each Cyclone processor had (up to) four I/O channels, and with multiple options for connection to the quadrants, configuration of these beasts was an extremely complex task.

The I/O controllers themselves were the same as used in the CLX, with logic board and backplane interconnect board, so no patch panels.

Just as with the CLX, I/O controller addressing issues were eradicated by the hardwiring of slot addresses to the cabinet backplane and all boards in a CLX system where physically disabled (through keying) from being put in the wrong slot.

The V8 and XL8 disk subsystems introduced with the VLX received an upgrade for connection to the Cyclone with fibre-optic cables replacing the thick copper cables of the V8/XL8 connecting directly to the disk controller BIC in the back of the MC-32 I/O cabinet. The fibre-optic connect disk subsystems were renamed the V80 and XL80.

Like the VLX, CYCLONE OSP functions were integrated into a single board RMI in the first processor cabinet requiring only a dumb terminal, connected to a dedicated port, to start and stop the system, and carryout maintenance.

As with previous systems, configuration could be complex, and the power situation with three power supplies providing power to two processors was a tad quirky, giving many engineers (including this one) a bad day, but in comparison CYCLONEs were a breeze to maintain (every pun intended given the size of those fans).

Tandem Cyclone-R

Rather than, as the name would imply, being an upgrade to the CYCLONE, the Cyclone-R had more in common with the CLX sharing many components and cabinetry. The big difference over CLX system was the installation of significantly faster MIPS 'Reduced Instruction Set Computing' (RISC) R3000 CPUs and an upgraded Guardian O/S, the D-Series, to make use of the RISC abilities of the Cyclone-R processors.

Confusion over the Cyclone-R name led to a re-think of marketing strategy, and along with the introduction of a completely re-designed CYCLONE, possibly because the CYCLONE mainframe's uptake was lacklustre, the new CYCLONE range became the Tandem NonStop Himalaya K-Series family with the Cyclone-R being renamed the K1000.

Tandem NonStop Himalaya K-Series

Good engineering design had always been part of the Tandem ethos and this included the external appearance. A decision was made to unify the entire Tandem range with common external design features and a common name - the K-Series.

No equipment to wear the Tandem/NonStop badge before or since is or was as well designed and constructed as the K-Series. The K-Series were made of good quality bespoke components and this included the metals that make up the internal component covers and chassis. High quality manufacturing processes and stringent quality control processes, added to the superbly engineered design, gave for systems that were, and still are, a joy to maintain.

Tandem NonStop Himalaya K1000

As described above, the K1000 was a rebadged Cyclone-R using CLX components and cabinetry. As part of the re-branding, however, the external cosmetics where changed to the by now Tandem synonymous black cabinet with a thin red line.

Also introduced with the K1000 was an 'enhanced availability kit' consisting of a dual/split power distribution unit (PDU) and split IPB/MDB bus components. The dual/split PDU allowed for a K1000 cabinet to be supplied from two separate mains AC power sources and the split IPB/MDB bus allowed for replacement of faulty bus components without the need to bring the system down. Most customers took up the enhanced availability kit option.

As with the CLX range, processors cabinets and MC-8 I/O cabinets could be added by extending the IPB/MDB buses with a maximum of eight processor cabinets (sixteen processors) and eight I/O cabinets. External disk could also be added including a new range of small modular SCSI disks mounted in a 4500 disk sub-system (six disks per sub-system).

Like it's older cousin the CLX, the K1000 was easy to work on and very robust. We have seen systems heavily damaged in transit and from building disaster (including an explosion) that have started right back up when power was re-instated.

Tandem NonStop Himalaya K10000

The behemoth that was the CYCLONE was quickly replaced by a MIPS RISC R4400 processor system under the Himalaya K-Series family. With significantly more processing power and I/O capability than the K1000 and named the K10000. Its processor cabinets had the same physical dimensions as the CLX/Cyclone-R/K1000 cabinets and the same black cabinet with a thin red accent external appearance, but internally shared nothing with their junior family members.

The K10000 processor cabinets contained just two processors, IPB/MDB bus controllers, dual power supplies and two fans. System OSP functions were integrated into a single board in the first processor cabinet requiring only a dumb terminal, connected to a dedicated port, to start and stop the system, and carryout maintenance.

I/O connectivity came from either MC-32 I/O cabinets or MC-8 I/O cabinets as described above. Like the CYCLONE, the K10000 had up to four I/O channels per processor enabling connection to a serious number of I/O devices.

Tandem NonStop Himalaya K100

The K100 was a K1000 system housed in a shorter and slightly wider cabinet with just two optional I/O controller slots and no dedicated memory slots. Expansion was limited one add-on processor cabinet (I/O only cabinets could not be connected to the K100).

This was a cut-down, cut price, entry level K1000 system. It was popular though and just as robust and easy to work on as the rest of the K-Series.

Tandem NonStop Himalaya K2000

The K2000 was an upgrade of the K1000 using all the same components except for more powerful processors with more memory, either 128MB or 256MB (with all memory integrated on the processor board).

These systems were last manufactured in the 1990s and yet a number are still giving excellent service to this day.

Tandem NonStop Himalaya K200

The K200, like the K100, is a K2000 in a shorter and slightly wider cabinet with just two optional I/O controller slots. Again, expansion is limited to one add-on processor cabinet.

Tandem NonStop Himalaya K20000

The K20000 is a processor upgrade to the K10000 using all the same components except for more powerful processors.

Tandem NonStop Himalaya S-Series (ServerNet)

Tandem experienced a downturn in fortunes during the early to mid-1990s. This is mentioned here as, unfortunately, during this time it was developing the replacement for the K-Series. This replacement was to become known as the S-Series. Why unfortunately? Because value engineering was engaged at the cost of good design and quality construction.

The S-Series comprised of two ranges; the Snn00 (two trailing zeros) range and the Snn000 (three trailing zeros) range. With the additional zero came more powerful processors, generally more memory, improved expansion capability and a bigger price tag, both for the hardware and the new O/S – the G-Series.

The S-Series introduced ServerNet technology. ServerNet was the name given to a very efficient bus message routing system using ‘worm-hole routing’. The ServerNet busses incorporated the functions of the IPB, MDB buses and I/O Channel, and being Tandem two ServerNet busses (known as the X-Fabric and the Y-Fabric) were provided. Most components in an S-Series system have a ServerNet router on board with some having both X-Fabric and Y-Fabric routers on board.

All S-Series systems share the same cabinets, now called enclosures, and these enclosures are essentially a bespoke 19-inch, 36U high rack. Indeed, a number of customers removed the internals from the Tandem enclosures for installation in their standard data centre racks. The black with thin red accent design theme was maintained to a degree (the enclosure front doors) on early systems but all main chassis parts were painted with a hammered grey finish.

Cosmetic changes to the appearance were made over the production span but nothing could take away from the perception that the systems looked like a collection of bits from a parts bin. The enclosure chassis metalwork was also thin, it would bend or dent easily and did not fit together well. Much the same could be said of the internals where further economies were made. Metals would discolour after just a couple of years and in the worse cases corrosion would set in. As the systems aged it also became apparent that quality of connectors on the backplane, and the soldering of those connectors was not great.

A processor enclosure consisted of two processors, space for up to sixteen SCSI disks, two fans, a couple of environment/power monitoring boards, two batteries and four optional I/O slots. The four I/O slots would be reduced to two slots if the system was expanded beyond one enclosure as two I/O slots would be taken up with ServerNet Extender Boards (SEBs).

S-Series processors combined the function of CPU, ServerNet fabric controller, SCSI disk/tape controller and DC power supply interleaved to half the system (processor boards were known as the Processor Multi-Function [PMF] board). This was not good design - the failure of the DC power supply would result in the loss of the fabric controller associated with that processor, the SCSI controllers and the processor itself. And, a fault with a non-mission-critical element such as the tape controller element would require replacement of the entire PMF.

In the first processor enclosure the processors also handled communications for the OSP functions. The OSP, now called the Tandem Support and Management (TSM) Console, was a Windows PC running bespoke software, directly connect to dedicated PMF TSM LAN ports.

For I/O expansion, I/O only enclosures were available and were identical to the processor enclosure except for replacement of the processors with I/O Multi-Function boards (IOMF) - essentially a PMF with the processor deleted.

I/O controllers were specific to the S-Series and consisted of a range of single boards combining both the logic circuitry and physical interface (a combined LB and BIC as it were) with interfaces appropriate for the required function. These were all pretty robust electronically but suffered dreadfully from the metal discolouration and corrosion issues mentioned above with controllers even becoming seized in the card cage. Later revisions of controllers came without much of the metalwork from earlier versions again hinting at value engineering.

Disks were also specific to the S-Series, though older disk sub-systems could be connected to an S-Series through optional I/O SCSI disk controllers. Most customers plumped for the internal S-Series disk satisfied that the sixteen slots available per enclosure would give sufficient disk capacity. The internal drives were standard dual port 5.25-inch SCSI drives mounted on a small metal frame, to locate and secure the drive within the enclosure's drive cage, with its own integral power supply. Storage capacity ranged from 4GB to 144GB. Compaq (see below) continued Tandem's practice to slightly alter 'standard' disks at the manufacture stage to ensure that only 'Tandem OEM' standard drives could be used. These drives gave little trouble, though reliability of some models was not sparkling.

Tape drives connected either directly to the SCSI port on the front of a PMF (copper cable) or through an optional SCSI I/O controller (copper or fibre-optic). Many different flavours of tape drive were made available for connection to the S-Series from stand-alone desktop DDS2 4mm DAT drives through to StorageTek Silos.

With the purchase of Tandem Computers Inc. to first Compaq and then HP the future of Tandem looked in doubt as HP seemingly had no idea what they wanted to do with the Tandem product line - other than standardise components with off the shelf HP Proliant product. The die was cast. Anyhow, S-Series sold well and fundamentally changed little over its production span though work on a replacement for the S-Series using HP Proliant product started soon after HP took control.

Compaq/HP ServerNet S7000 through S7800

The initial release of S-Series was the S7000. The S7000 differed from all later S-Series family PMFs in that AC power connected directly to each processor. Again, this was not good design - the loss of one mains AC mains supply would result in the loss of the fabric controller associated with that processor, the SCSI controllers on that PMF and the processor itself.

Up to four processor enclosures could be connected making an eight-processor system but further expansion was limited to one I/O only enclosure per processor enclosure.

Following the S7000 came the S7400 with a more powerful processor and a power shelf installed *(see S70000-S88000 below)* but otherwise was identical. A line of more powerful range of processors followed until the pinnacle if this range the S7800 all just PMF upgrades with more powerful processors and usually more memory.

S7000	*Two MIPS RISC R4400 CPUs, 256KB cache, 128MB-512MB memory*
S7400	*Two MIPS RISC R10000 CPUs, 1MB cache, 512MB Memory*
S7600	*Two MIPS RISC R12000 CPUs, 2MB cache, 1GB or 4GB Memory*
S7800	*Two MIPS RISC R14000 CPUs, 2MB cache, 2GB or 4GB Memory*

Compaq/HP ServerNet S70000 through S88000

The second release of S-Series was the more powerful S70000. The S70000 (and the whole S70000-S88000 range) differed generally from the S7000-S7800 range in that the Snn000s had more powerful processors, could generally be configured with more memory, and every processor enclosure could be connected to more than one I/O only enclosure.

S70000	*Two MIPS RISC R4400 CPUs, 1MB cache, 128MB-1GB Memory*
S72000	*Two MIPS RISC R10000 CPUs,1MB cache, 256MB-2GB Memory*
S74000	*Two MIPS RISC R12000 CPUs, 4MB cache, 512MB-4GB Memory*
S76000	*Two MIPS RISC R14000 CPUs, 4MB cache, 1GB-16GB Memory*
S78000	*Two MIPS RISC R14000 CPUs, 4MB cache, 2GB-4GB Memory*
S86000	*Two MIPS RISC R14000 CPUs, 8MB cache, 1GB-16GB memory*
S88000	*Two MIPS RISC R16000 CPUs, 8MB cache, 2GB-16GB Memory*

Starting with the S70000 and adopted across the entire S-Series range thereafter, a separate AC 'power shelf' was fitted to the enclosure with two power supplies converting AC mains to high current DC which was delivered to the PMFs by two short cables. The AC power shelf was interleaved so that the loss of one AC PSU or one AC input would not mean the loss of a PMF, but again, this was not good design as it was not possible to shut off power to one PMF so the DC power cables always remained live. Compounded by a cheap connector at the PMF end of the short cable, replacement of a PMF in a live system always ran the risk of misconnecting the DC cable or even shorting the high current DC against the PMF metalwork - que a big spark, a big bang, a raft of damaged system components and an unhappy customer.

Though we have never done this, nor seen it done, we have seen damaged PMF boards with what looks like an arc-weld on the DC connector - enough to make the hand of the most resilient customer engineer tremble when plugging the DC connector in to a PMF on a live system. The other flaw with this design oversight is that plugging in of the DC cable has to be done in a rush as from the point of pushing the PMF into the enclosure backplane the DC cable has to be fully connected to the PMF within 20 seconds otherwise the PMF will not be recognised by the system and the action has to be repeated.

HP Integrity NonStop NS-Series

The replacement for the S-Series, the Integrity NonStop NS-Series lined up with HP's Integrity family and shared Intel's Itanium II CPUS across the range as well as many of the standard off-the-shelf HP peripherals. As with the S-Series, the NS-Series has two ranges; the NSnn00 (known as the NonStop Value Architecture [NSVA]) range and the NSnn000 (known as the NonStop Advanced Availability [NSAA]) range. And, similar to the S-Series, the additional zero brought more powerful processors, more memory, improved expansion and a bigger price tag, both for the hardware and the O/S. However, the two ranges differ completely in their processors. H-Series O/S was used across the range.

Both ranges were housed in a standard HP 19" rack, either 42U or 36U - the later is very rare - with an HP ProLiant 1U server running Microsoft Windows Server O/S acting as the management interface to the system, now known as the Operations Service Manager (OSM), connected to a 1U video/keyboard/trackpad device. The systems are generally configured with a UPS at the bottom, the processor section above, then the OSM/KVM, then the I/O section and at the top the disks. At the rear of the rack are the PDUs and maintenance LAN switch(es) are at the very top. There is some fluidity to this layout depending on the number of processors and/or disks installed.

Inter-connections between elements of the system, and there are lots of inter-connections with NS-Series systems, are either by red ethernet cables or orange fibre-optic cables. This lack of variation of colour in the cabling is one of the main bugbears with this generation for customer engineers - tracing a cable fault is near impossible and replacing a faulty component requires care not to get the cables mixed up.

Additionally, routing management of the inter-connection cables was problematic as it was reliant on the original installer allowing sufficient slack in cables and bindings to permit the cables to move when system elements are pulled forward out of the rack, but tight enough not to be snagged on a neighbouring element should the neighbouring element be pulled.

The range and diversity of the NS-Series and the rapidity with which components of these NonStop systems were made obsolete brought further maintenance issues - tracking what generation of product a particular system had originally installed and keeping spares holding for all the different flavours. By example, one customer had three separate systems delivered within two months of one another, yet each system came with a different model of HP ProLiant OSM console.

The spares holding/supply problem deepened once SAS disks were introduced as the sales life (End of Sale Life [EoSL]) for a particular SAS disk could be measured in days rather than years.

HP Integrity NonStop NS1000

The first generation of the NSnn00 range, introduced in 2006, and the first NonStop to use HP's off-the-shelf processors in the form of an HP Integrity rx2620 with one Itanium 2 9100 CPU (1.3 GHz with 3 MB cache) was the NS1000. Two memory options were available, 1GB or 4GB with the former readily upgraded. The rx2620 was fitted with a single bespoke PCI ServerNet card with dual fibre-optic connections, acting as the X/Y-Fabric inter-processor bus - a clear single point of failure. Standard HP 'Integrated Lights-Out' (iLO) LAN connection was used for maintenance reporting/management. Each processor is known as a 'blade element' and this naming convention is continued through the NS-Series and later NB-Series. A maximum of eight blade elements were permitted in an NS1000 system.

The blade element processors had no hard drive installed. At power-on the processors boot from a CD which had to be left permanently in the DVD tray - we have no idea why a hard drive could not be used. Once the CD loaded the processor was ready to receive O/S through the PCI ServerNet connection.

Processor connection, X/Y-Fabric ServerNet control, I/O and OSM interfaces were all combined in an 'I/O Adapter Module Enclosure' (IOAME - or IOAM), a bespoke 19-inch device fixed in the rack. The IOAM was split between X-Fabric and Y-Fabric with multiple power supplies and fans so that fault-tolerance integrity was maintained. Slots for optional I/O expansion were available but were limited to ethernet (copper) or fibre-channel disc/tape - other flavours of interface connection had to be through external convertor devices. As with the S-Series, cheap materials were used in the build of the IOAM and corrosion of the chassis leading to seized components was a problem.

Later NS1000s were available with the IOAM replaced by two 'Versatile I/O' (VIO) Modules. These were effectively an IOAM split in two to separate the X/Y-Fabrics into two separate 4U high chassis modules. Though the chassis was based on an Integrity rx4400, the VIO main circuit board and some PCIE I/O adapter cards which plugged into the main board were bespoke to NonStop. The VIO has dual hot-swap PSUs and dual hot-swap fans and is mounted on rails with articulated cable management arms for ease of maintenance, but cables were frequently installed (during initial system build) so tightly that pulling the VIO out from the rack was a lottery.

Disks came in the form of (at least) two HP Fibre-channel fourteen slot sub-systems with 72GB, 146GB and 300GB disks installed. One sub-system acted as the primary sub-system and the second sub-system as the mirror.

Tape provision was either LTO or DAT over fibre channel using proprietary interface though Virtual Tape (VT) solutions were also popular.

Reliability was reasonable with just the occasional HP server PSU failure and ServerNet card failure in the rx2620 (usually through working itself loose from the I/O daughter board). As systems got older, the fibre-channel disks started to fail, but not alarmingly.

HP Integrity NonStop NS14000 and NS14200

Also introduced in 2006 alongside the NS1000, the NS14000 and the NS16000 had an entirely different processor setup using either two or three HP Integrity rx4640 chassis with up to four separate CPUs installed in each chassis. With two rx4640 chassis, a CPU in the first chassis independently executes the same instructions as the CPU in the second chassis and the results of the two executions are checked for data integrity giving Dual Modular Redundancy (DMR), also known as Duplex operation. With three rx4640 chassis three CPUs in separate chassis execute the same instructions independently giving Triple Modular Redundancy (TMR), also known as Triplex operation, for even greater data integrity. With NS14000 system, duplex/triplex rx4640 blade elements each had dual inter-connections and dual connections to a Logical Synchronisation Unit (LSU) which manages lockstep, data integrity and redundancy across the installed CPUs so that the application only sees the logical processor regardless of the fact that two or three physical CPUs make up that processor. The LSU has dual connections to the ServerNet connections for each logical processor to the ServerNet ports on the IOAM.

NS14000 systems can be expanded to a maximum of eight logical processors (corresponding to sixteen or twenty-four physical CPUs with each CPU being an Intel 2 Itanium 9100 CPU, 1.5 GHz/4 MB Cache and either 4GB or 8GB dedicated [partitioned] memory). Upgraded CPUs were available from 2008, the NS14200, offering 20% performance improvement over NS14000. These could co-exist with NS14000 provided that they were installed as matched pairs.

I/O capability came initially from the same IOAME as used with the NS1000 and was limited to on the NS14000/NS14200 systems two one IOAM with up to six I/O adapters. Later systems were offered with dual VIOs in place of the IOAM.

As with the NS1000, disks came in the form of two to eight HP Fibre-channel fourteen slot sub-systems with 72GB, 146GB and 300GB disks installed.

We have no direct experience of maintaining an NS14000/NS14200 system so cannot comment on whether they have/had any vices. From the architecture, the prospect of replacing a physical CPU looks daunting.

HP Integrity NonStop NS16000 and NS16200

The NS16000 shares the duplex/triplex blade element setup but differs in that bespoke P-SWITCH module sits between the LSU and the IOAM to manage connection to multiple LSUs. This increases the expansion options to a maximum of sixteen logical processors and a maximum of up to ten IOAMs. Each physical CPU are is Itanium 2 9100 CPU, 1.6 GHz/6 MB cache, so slightly more powerful than the NS14000, and available with 4GB, 8GB or 16GB or dedicated [partitioned] memory.

Upgraded CPUs were available from 2008, the NS16200, offering 20% performance improvement over NS16000. These, unlike the NS14000/NS142000 these could co-exist, with NS16000.

Unlike the NS14000/NS14200, the NS1600/NS16200 systems were only available with IOAMs installed.

Disk were fibre-channel, as NS1000/NS14000/NS14200 but a maximum two hundred and thirty-two FDCM sub-systems could be connected.

Tape provision was either LTO or DAT over fibre channel using proprietary interface though Virtual Tape (VT) solutions were also popular.

HP Integrity NonStop NS1200

Introduced in late 2007, the second generation of the NSnn00 (NSVA) range used HP Integrity rx2660 servers as processor blades with one Itanium 9100 CPU (1.6GHz/12MB cache) so more powerful than the NS1000, and either 4GB or 8GB of memory installed.

I/O was standardised on the Versatile I/O (VIO) module though NS1000 systems could readily be upgraded by replacement with these newer processor blades so NS1200 systems could be found with an IOAM.

Disk and tape were the same as for the NS1000.

HP Integrity NonStop NS2000

Introduced in 2008 as an entry model to the NVSA range, the NS2000 used the same rx2660 server as the NS1200 with one Intel Itanium 9100 CPU running at 1.42GB/12MB cache and the option of either 8GB or 16GB memory installed. System size was a maximum of four processors.

With the NS2000 came the introduction of Cluster I/O Modules (CLIMs) for storage and I/O expansion offering more expansion capability than was physically possible with VIOs alone. CLIMs are separate individual HP ProLiant servers (initially DL385-G5 and later DL380-G6) running an adapted version of Debian Linux from hard disk drive and connect to the two VIO I/O Core modules through a dual port ServerNet PCIE card (connecting to the X-Fabric and Y-Fabric). A standard ProLiant server Ethernet port and the iLO port are used for maintenance bus connection.

CLIMs were offered in two I/O expansion types with distinct I/O interfaces:

'Storage CLIMs' - for connection to fibre-channel disk sub-systems and fibre-channel tape drives using HP standard Fibre-Channel Host Bus Adapter (HBA) PCIE controllers and/or connection to the new HP MSA70 twenty-five slot SAS disk sub-system using SAS HBA Host Bus Adapter PCIE controllers, then

'IP CLIMs' - for connection to customer ethernet networks and wide area network communication adapters using HP standard copper or fibre-optic ethernet PCIE controllers.

The DL385-G5s CLIMs only had a single hard drive but from the next CLIM generation on, the DL360-G6, all would have two hard drives in a RAID array improving fault tolerance of the CLIM. Storage CLIMs are configured in pairs as dual disk controllers. IP CLIMs are configured singularly but with internal interfacing bonding and software failover to an alternate IP CLIM in the event of the loss of a complete IP CLIM.

HP Integrity NonStop NB50000c

In 2009 came the first NonStop system to use the HP C7000 enclosure. The C7000 is capable of housing up to eight full height blade processors, and HP chose the BL860c blade each with one Intel Itanium 9100 CPU, running at 1.66GHz/18MB cache, and from 8GB to 32GB memory installed. Every processor has a small dual port ServerNet mezzanine card installed for inter-processor bus communications with internal enclosure (backplane) connection to two separate ServerNet controller/routers in the rear of the C7000 enclosure. Systems could have a maximum of sixteen processors housed in two C7000 enclosures. The rear of the enclosure also included dual maintenance bus routers and Dual Onboard Administrator Modules (OAM) wrap up the package.

I/O was provided though either IOAM enclosures (maximum of 60) or CLIMs (maximum 48) connecting to the ServerNet controller/routers in the rear of the C7000 enclosure.

Disk and tape options were the same as for the NS-Series.

The O/S used was H-Series, like the NS-Series, to now anyhow.

HP Integrity NonStop NS2100 and NS2200

Using HP Integrity rx2800-i2 servers as processors, with one Intel Itanium 9300 CPU installed, and introduced in 2010, the NS2100/NS2200 were a performance upgrade to the older NS1200 but announced multi-core CPU capability. Prior Itanium CPUs installations over the previous NS-Series and NB-Series range had CPUs chips with more than one core onboard but did not make use of the additional/potential processing power. The NS2200, however, came with two cores enabled.

From here on specific details of CPUs are hard to get a handle on. Product releases failed to mention speed and cache, potentially this was a result of multiple cores now processing resulted in a significant performance increase so that a few percent MHz performance uplift was irrelevant.

A new generation of O/S software/license, J-Series, to enable multi-core processing was required. Even with J-Series O/S the NS2100 remained hobbled, by the O/S, to processing with one core - keeping to the ‘entry level’ NonStop System marketing ethos. Both NS2100 and NS2200 systems were available with, or expandable up to, eight processors.

VIO modules and CLIMs were used for I/O with MSA70 SAS Disk sub-systems.

HP Integrity NonStop NB54000c

Following the NS2200, in 2012, the NB54000 was a processor upgrade to the NB50000. The same C7000 enclosure, but fitted with BL860c-i2 blades, each with one Intel Itanium 9300 CPU, and the NB54000 runs Guardian J-Series Operating System with multi-core processing capability. A choice of either two or four cores could be enabled.

With new orders I/O was now standardised on Storage/IP CLIMs, using Gen6 servers, and the IOAM dropped, but existing NB50000 systems could readily be upgraded with the new BL860c-i2 blades so NB56000 systems with IOAMs could be found.

Upgraded systems would include fibre-channel and MSA70 SAS disk, but the new order NB54000s heralded the introduction of the HP D2700 SAS Disk Sub-system, much the same as the MSA70, but with 6.0Gb/s bandwidth (3.0Gb/s bandwidth on the MSA70).

HPE Integrity NonStop i NS2300 and NS2400

Re-branded as members of the HPE Integrity NonStop i Series (with the ‘i’ denoting Itanium CPU to differentiate from the NonStop X range with the ‘X’ denoting Xeon CPU) the NS2300/NS2400 are based on rx2800-i4 servers with a single Intel Itanium 9500 CPU. The NS2300 was hobbled to one core enabled and the NS2400 two cores. Memory options were 16GB, 32GB or 48GB.

I/O was through VIO (now denoted as VIO-II) and storage/IP CLIMs with D2700 SAS Disk sub-systems as the default (D3700 sub-system on later systems).

HPE Integrity NonStop i NB56000c

Under the NonStop i Series, the NB56000c was another processor upgrade, with BL860c-i4 (one Intel Itanium 9500 CPU) this time, with either two or four cores and up to 96GB memory.

HPE Integrity NonStop X Series

The Itanium CPU range had run its development course, so HPE had to rethink. Intel Xeon was the future for NonStop.

ServerNet was also due an upgrade. ServerNet concepts had, for some time been shared with Intel and others, developed, and performance improved. Infiniband was the product of that development and was adopted as the inter-processor and I/O channel communications highway for the new NonStop X Series. Being a Tandem system, multiple Infiniband buses were used.

Initially this new range was based on the C7000 Platinum Enclosure, now using half-height blades so that up to eight processors could be housed in a single enclosure, with I/O shared over 4xFDR (Fourteen Data Rate) InfiniBand double-wide switches with 56Gbp/s bidirectional bandwidth.

From the X4 series the C7000 enclosure and associated blade processors are to be dropped in favour of DL360-G10 Servers.

L-Series Guardian Operating System is a required to run the new Xeon CPUs.

HPE Integrity NonStop X NS3 X1

The entry level NonStop X system, the NS3 uses BL460c Server Blades with Intel Xeon E5-2600 v2 series CPUs, supports from 32GB to 64GB of memory per CPU, and has either one or two cores enabled. A maximum of four processors per system are permitted. Using InfiniBand 4 x FDR (fourteen times data rate) switches with 56Gbps bi-directional bandwidth for X/Y-Fabric and I/O.

I/O is through GEN8 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NS7 X1

The full-fat NonStop X system uses the same BL460c Server Blades with Intel Xeon E5-2600 v2 series CPUs, supporting from 64GB to 192GB of memory per CPU, but with two, four or six cores enabled. A maximum of sixteen processors per system permitted. Using InfiniBand 4 x FDR (fourteen times data rate) switches with 56Gbps bi-directional bandwidth for X/Y-Fabric.

Again, I/O is through GEN8 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NS3 X2

Same as the NS3 X1 but uses BL460c Server Blades with Intel Xeon E5-2600 v2 series CPUs, supports from 32GB to 64GB of memory per CPU, and has either one or two cores enabled. A maximum of four processors per system are permitted.

I/O is through GEN9 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NS7 X2

Same as the NS7 X1 but uses BL460c Server Blades with Intel Xeon E5-2600 v2 series CPUs, supporting from 64GB to 192GB of memory per CPU, but with two, four or six cores enabled. A maximum of sixteen processors per system permitted.

Again, I/O is through GEN9 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NS3 X3

Same as the NS3 X1 but uses BL460c Server Blades with Intel Xeon Silver 4100-Series CPUs, supports from 32GB to 64GB of memory per CPU, and has either one or two cores enabled. A maximum of four processors per system are permitted.

I/O is through GEN10 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NS7 X3

Same as the NS7 X1 but uses BL460c Server Blades with Intel Xeon Silver 4100-Series CPUs, supporting from 64GB to 256GB of memory per CPU, but with two, four or six cores enabled. A maximum of sixteen processors per system permitted.

Again, I/O is through GEN10 IP/Storage CLIMs with HP D3700 subsystem SAS Disk and LTO-6 Tape or virtual tape.

HPE Integrity NonStop X NSn X4

Believed to use HPE Proliant DL360-G10 Servers rather than C7000 Enclosure and Blades. InfiniBand increased to 100Gbps. No firm details available as yet.

INDEX