

HP NonStop Is a Credible Replacement to IBM's TPF

Ian Tunnacliffe

Giga Position

The imminent demise of the TPF (Transaction Processing Facility) operating system from IBM has been confidently forecast since at least the late 1970s. It has held its leading position for more than 30 years because at the very highest levels of throughput and availability, there has simply been no other system that could match its capabilities. Major users in travel distribution and finance needed the performance that only TPF could deliver, and they were prepared to put up with the costs and inconvenience of maintaining applications in Assembler language to get it.

During the last two or three years, major players in the travel industry have been investigating the deployment of HP NonStop (formerly Tandem) systems as TPF replacements. Sabre in particular has been very aggressive in the take-up of the new platform and has successfully migrated its fare shopping system to NonStop. This is one of the most resource-hungry applications in the whole travel distribution universe, and not only has it been transferred successfully, but it is also nowhere near testing the limits of the platform's capabilities. This success and others in travel, telephony and online services convince Giga that there is no technical reason why TPF systems cannot be gracefully retired. The question of whether and when to do this is now purely a business decision based on the relative costs of continuing with TPF and those of porting or rewriting the applications — some of which have more than 30 elapsed years and hundreds of person-years development behind them.

Recommendations

Major users of TPF systems that have not already done so should evaluate the HP NonStop System as a replacement for new developments as well as existing applications. Applications that are changing rapidly in function or throughput should be the first candidates for change.

Airlines that have bought TPF-based reservations systems for in-house use should talk to their suppliers about the future availability of applications updates for the TPF platform.

Proof/Notes

Requirements for a High-Volume, High-Availability Transaction System

A typical TPF installation — or any replacement of one — must satisfy a number of onerous requirements:

- Transaction rates are high. The largest systems in current use support more than 14,000 transactions per second.
- Availability must be truly 24x7. Planned downtime has been more or less eliminated from the TPF world, and unplanned downtime is measured in minutes per month.
- Response times must be short. More than three seconds between the query being entered and the response arriving at the end user's screen is unacceptable. Given the likely transit times for the global communication networks in use in the travel industry, this means that the response time at the data center must be consistently subsecond.
- Transactions must be secure. With most Western airlines adopting electronic ticketing, the data within TPF reservations systems now carries value with it just as surely as does a banking system.

TPF Background

History

The TPF operating system has been evolving since the creation of the Sabre reservations system by American Airlines and IBM in 1960. Originally deployed on an IBM 7090 mainframe, the application and operating system were developed together. The subsequent development of an international version of the system to run on the IBM 360 by the manufacturer and a group of European airlines led to the separation of applications from the operating system, which was named Airlines Control Program (ACP). ACP was optimized for high performance at the expense of security, integrity and development time. Applications were written in 360 Assembler and were restricted to 1,055-byte segments. Applications programs addressed disk storage directly. System outages were needed to load programs and initialize certain data tables, especially those related to communications, as well as for housekeeping functions like the creation of backup tapes.

The subsequent 35 years have seen steady incremental development in the system, which was renamed TPF in 1980 when it became a chargeable software product instead of being bundled with IBM mainframe hardware. Higher-level programming languages have been added, and most new development today is done in ISO C or C++. A database manager has been introduced to insulate applications programs from the need to directly manage disk I/O. Almost all maintenance functions can be done with the system running, and the adoption of a loosely coupled architecture means that the largest systems run as N+1 complexes that can withstand the failure of a single CPU without interruption to service.

A TPF emulator has been introduced that allows TPF programs to run in a partition of an MVS system. This system — known as Airlines Control System (ALCS) — has been taken up by a number of users whose requirements do not reach the extreme performance levels available at the top of the TPF scale. For these users, the trade-off is a lower absolute level of performance in return for ease of use and access to the facilities of the underlying MVS platform. This also means that ALCS users can take advantage of new developments in IBM hardware that are invariably launched for MVS well before their eventual adoption in TPF.

Current Market

The current market for TPF and ALCS systems is approximately 50 sites worldwide. The actual number of licenses is significantly larger than this, since most installations have multiple instances of the software, for backup, system development and test, as well as because many users have split their systems along functional lines.

The largest implementations are the four Global Distribution Systems (GDSs) that supply booking services to travel agents, airline offices and Web sites. These are the systems that are pushing hardest at the limits of the technology, with Sabre having recorded a peak hour at more than 14,000 transactions per second. The original use of TPF in airline reservations systems is still well represented with many airlines' in-house systems running on the platform as well as the multi-host systems SHARES and Atraxis, which are both owned by EDS. A few TPF systems provide the reservations environment for other travel companies such as railways and hotels. There is a handful of financial services companies running TPF systems to manage part of their banking and credit card transaction handling requirements. Finally, there is a small number of systems used by public services in the United States.

In the past few years, the trend has been for a reduction in the number of sites running TPF but a large growth in the traffic being handled by each. This reflects a trend for airlines in particular to decide that running a major IT system is not part of their core competence and to outsource reservations. This move is almost complete in North America with no major airline now running its reservations system in-house. In Europe, several significant carriers have begun the move to outsourcing, especially those associated with the Amadeus GDS, while in Asia-Pacific and Latin America, in-house systems are still the norm for larger airlines.

HP NonStop

Background

The current generation of NonStop systems from HP brings the history of the platform full circle. Tandem Computers was founded by a group of former HP employees in 1973 and delivered its first fault-tolerant computer system in 1976. The product has grown and evolved during the following 27 years, despite two changes of ownership — acquisition by Compaq in 1997 and then the merger of that company into HP in 2002.

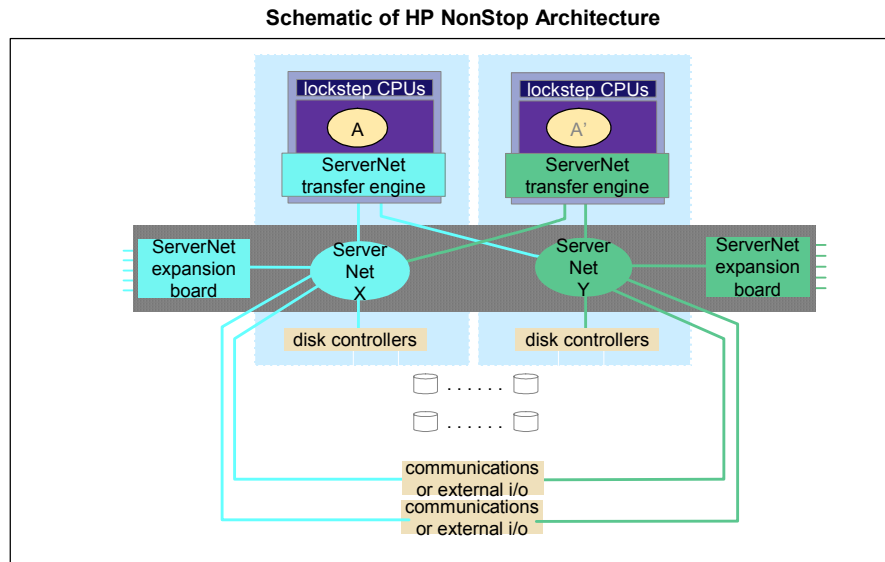
Since its launch, the NonStop system has found a ready market among businesses that need to ensure that critical systems maintain a true 24x7 operation. These include banks, stock markets and telecommunications companies. While the same level of fault tolerance, high availability and scalability can be achieved using other technologies, the unique proposition of the NonStop servers is that they are all built into the base product and work straight out of the box. This more than anything has ensured that the system has maintained and enhanced its position in the marketplace and continues to do so.

General Features

The architecture of the NonStop server is referred to as “loosely-coupled shared-nothing” in the company’s literature. Each logical complex is linked by HP’s ServerNet architecture, which is a high-performance messaging system providing interconnections between processors and I/O devices (see figure below).

The basic processing component is the Processor Multifunction (PMF) card. This consists of a pair of processors running in lock-step along with local memory, two ServerNet ports, SCSI and Ethernet interfaces. Output from the two processors is continuously compared by the operating systems, and in the event of a disagreement, the PMF is shut down and its workload shifted to another PMF containing the backup process. Thus, the minimum system configuration contains two PMFs and four processors in all. Each PMF runs its own copy of the NonStop Kernel operating system. PMFs may be added in increments of two to allow system growth up to 4,080 PMFs, organized in nodes of 16 pairs, in a complex connected by local or wide area networking.

Error checking and redundancy are built into each layer of hardware and software. Before a transaction alters the database, the affected row is written to a log. This allows the restoration of the database to any specific point in time and the backing out of failed or erroneous transactions.



Source and Copyright: HP 2003

High Availability

The NonStop architecture is built for high availability from the ground up. An important point to consider is that it is the application that is able to run at high availability rather than the operating system. Many Unix clustering solutions are able to offer high availability of the O/S but may require substantial time to restart applications after a component failure.

The NonStop Kernel operating system detects, isolates and recovers from component failures without applications being affected. The fail-fast design philosophy stresses takeover rather than failover. Transactions in progress in the affected component at the time of failure are backed out and redirected to another processor. End users will typically see a delay of two or three seconds if their transactions are redirected this way.

Scalability

The shared-nothing aspect of the architecture removes bottlenecks from the system and allows for effectively linear scalability up to a maximum of 4,080 processors. By contrast, a symmetric multi-processor (SMP) solution may become markedly non-linear as it scales up beyond as few as eight processors.

Within a large NonStop complex, all of the processors are doing useful work. There is no requirement to have processors waiting idly in case of a failure.

Current Deployments

HP NonStop is currently deployed in some of the largest financial, telecommunications, retail and online services environments. Key customers include Barclaycard, Citicorp, New York Stock Exchange, Sprint and Tandy/RadioShack. In these segments it is well established and a clear leader. It is in the area of travel and transport that the true test of whether this is a viable TPF replacement will be played out. Established customers for NonStop in this sector include Outrigger Hotels and the Scandinavian Leisure Group, which has long operated a complete commercial system based on the NonStop platform. However, the biggest test for the system so far has come with its adoption by Sabre to take on the full burden of its fare shopping application.

Sabre

As the first user of the system that became TPF and the operator of the largest TPF complex in the world, Sabre is the organization that more than any other understands the requirements of large-scale transaction processing in the travel industry. During the last five years, Sabre has seen an exponential increase in the traffic through its fare shopping application, primarily due to the rapid increase in online sales of travel. Not only has the traffic increased, but also the rate of change of data and functionality has increased. Consumers shopping on the Internet have a much lower threshold of boredom than the professional travel agents and airline staff who were Sabre's main users only five years ago.

With these factors in mind, Sabre had been examining various technology alternatives to TPF and in 2000 implemented a proof of concept with the then Compaq Himalaya technology. The system was built by five full-time and five part-time developers with support from Compaq between May and September. It consists of 350,000 lines of C++ code and the database has around 180 tables. The system went live in the Sabre complex in August 2001 and in the first 60 days of operation it handled the full load of transactions with zero downtime. Since then, Sabre has further developed the application and upgraded the hardware to the current S86000 servers — all without application downtime. The current application is handling the entire operational load with a total of 17 nodes of 16 processors each — just 6.6 percent of the maximum processor capability of the platform.

Sabre is so convinced of the success of the NonStop platform that it is already moving a number of other applications, including ticketing, to the NonStop platform.

Futures

The HP NonStop platform is undergoing constant development, with a detailed road map set out for the next three years. The most significant change in that time is a planned replacement of the current MIPS processors with standard Intel Itanium chips from 2004 onward. Tuxedo and WebLogic are already available for NonStop. This enables common applications code to be deployed across NonStop servers for mission-critical applications and commodity Unix/Linux boxes for background functions. This commonality also provides reassurance for customers who might be nervous about a single supplier replacing the incumbent single supplier of TPF. With common applications code, the progress that is being made on scalability and availability by suppliers of Unix clustering solutions may be leveraged to prevent HP holding a complete stranglehold on the market.

Disadvantages of HP NonStop

Although it is true that the current users of HP NonStop report very positive experiences, there is no such thing as a perfect system, and this one like others has its disadvantages. These include:

- High initial cost: Compared to TPF, this is not an issue, but HP NonStop is significantly more expensive to acquire than standard Unix or Linux servers.
- Lack of productivity tools: Travel industry users have more than 30 years' experience with TPF, and in that time they have built a formidable array of programmer tools. The equivalent toolset has not yet been built for the NonStop environment.
- Single supplier: There is only one supplier of this technology, and many users will be nervous about being tied to a single source. This concern is alleviated by the availability of standard middleware environments as noted above.

Conclusion

HP NonStop is without doubt capable of taking over the load currently carried by TPF systems in the world of travel, financial and public services. One of the world's biggest travel booking systems has migrated its

most resource-intensive application to NonStop, with more than 90 percent headroom for future expansion and a net improvement in system availability. The real questions to be answered now are concerned with the business case.

Although the initial acquisition costs of HP NonStop are higher than for standard Unix systems, research by The Standish Group shows that on a total cost of ownership (TCO) basis, the picture can change quite dramatically, especially if the cost of downtime is factored into the TCO calculations. On that basis, the HP NonStop system is among the most cost-effective platforms in the market.

Undoubtedly, there are large savings to be made by retiring TPF systems, but these must be set against the cost of replacing the applications that have been built over several decades. Clearly, Sabre has decided that the business case is proven for a range of its applications, but it is important to note that the company has no current intention to completely replace its TPF systems. As is often the case, the key is to use the most appropriate technology for the task at hand. For applications that are relatively stable and well within the capacity of the platform, there is not yet a compelling case for moving away from TPF. Where an application is evolving rapidly in either functionality or scale, then HP NonStop is now a technology that must be considered.

For more information, see the table titled [Comparison of TPF, HP NonStop and Unix for Large-Scale Transaction Processing Applications](#), located at the end of this document.

Footnote: As previously discussed by Giga, **Unisys** is developing an all-new Passenger Services System (PSS) called AirCore based on J2EE and WebLogic. Given that the HP NonStop platform now supports WebLogic applications and that it has all the advantages of availability, scalability and integrity that airlines need in their PSS, it seems to us that a combination of the AirCore application on a NonStop platform would be a very powerful offering in the marketplace. Whether this could ever come about within the current landscape of industry politics and allegiances is another matter, but it is one that Unisys, SITA and their customers should consider carefully.

Alternative View

Despite the advantages offered by HP NonStop and other newer technologies, TPF is still the incumbent system in the very largest transaction processing systems. As a fully mature platform with a wealth of existing applications, the costs of moving away from TPF are high. IBM could make the barriers to exit even higher by dramatic reductions in the license costs of TPF. The current generation of IBM mainframes needs no more environmental or operator support than the equivalent NonStop hardware, so a substantial reduction in the cost of licensing the operating system could slow down the move to other platforms by several years.

References

Related Giga Research

IdeaBytes

[SITA and Unisys Partnership Creates a Credible Force in Airline Hosting](#), Ian Tunnacliffe

Relevant Links and Other Sources

“Dollars to Cents: TCO in the Trenches 2002,” The Standish Group,
<http://h71033.www7.hp.com/object/TCORPT.html>

Comparison of TPF, HP NonStop and Unix for Large-Scale Transaction Processing Applications

TPF	HP NonStop	Unix
-----	------------	------

Cost of Acquisition	High: IBM mainframe hardware and TPF licenses	Medium	Low: Commodity servers and Linux make this option the cheapest to acquire.
Cost of Development	High: Dedicated staff required with specialist skills. Code is not reusable on any other platform.	Low: Uses standard development environments and tools. C++ and Java with CORBA, SOAP and other tools.	Low: Mainstream development environment. Greatest supply of development staff and productivity tools.
Speed of Development	Medium/Low: New developments are mostly written in C and C++. Maintenance of existing code requires expertise in Assembler language.	High: C++ and Java are both available on the platform. Standard middleware products such as Tuxedo and WebLogic are supported.	High: Mainstream of IT development and the focus of most open source projects.
Future Development Path	IBM has indicated that it is developing a new release of TPF that will use a 64-bit architecture and provide open toolkits for C and C++ developments.	Detailed development path set out, including move to Itanium processors.	Mainstream of IT development. Able to access new hardware capabilities as they are released.
Database	Non-relational database: Some applications address data records directly. More modern programs use TPF/DF — a limited database management utility	NonStop SQL built into the platform.	Must be acquired separately from commercial vendor. Open source databases are also available.
Availability	Loosely coupled N+1 complexes have reduced downtime to minutes per year. System recovery times after any outage are very short ~ 2 minutes.	100% availability is a realistic target for any but the very smallest installations.	High availability can be achieved by building clustering solutions on top of the basic O/S.
Scalability	Proven ability to scale to more than 14,000 transactions per second in a travel industry application.	Able to scale linearly up to 4,086 processor pairs.	SMP clusters are subject to a law of diminishing returns as processors are added.

Source: Giga Research, a wholly owned subsidiary of Forrester Research, Inc.

[Return to text](#)