



Scale Up? Scale Out? HP Scales Down NonStop for SME

Analyst: David Reine

Management Summary

After forty days and forty nights of rain, the clouds start to break and the sun begins to shine. We cannot wait for the temperature to begin its inexorable climb to the heights of where we expect summertime weather to be. Nor can we wait for the treats that summer holds in store. For example, a ride to the local ice cream stand, where we can cool off with a scoop of our favorite flavor. However, what would we do if there were only one size cone available? Instead of small, medium, and large, what would happen if extra-large was the only size available? Let's not think about calories and fat for the moment (please!), simply the cost. If you had to spend \$5.00 for a cone and you had a family of four, or a Little League team, how many servings could you afford? You could try to share, but what happens when everyone likes a different flavor? Fortunately, we do not have to be concerned. Every ice cream shop worth its sprinkles will have the three-scoop special in a waffle cone with all the mix-ins and toppings you like. They will also have a two-scoop offering of the same ice cream in a sugar cone for mom and a single-scoop portion in a cup for the kids. At Ben & Jerry's, one size definitely does not fit all!

If we can have all of these choices in something as trivial as an ice cream treat, why can't the small- or medium- sized data center — with a real-time mission to provide secure data 7x24x 365, whenever queried — have a similar choice? Not every data center has the same volume of mission-critical OLTP transactions. Not all high-availability requirements conform to a one-size-fits-all fault-tolerant (F-T) solution. Over the past few decades, we have seen a variety of mission critical, fault-tolerant solutions introduced to support the largest enterprises and government agencies, from metropolitan "911" emergency lines to Wall Street stock exchanges. All of them are capable of supporting thousands of users in environments where an outage could cost millions of dollars, or worse, a single human life. These systems, unfortunately, were costly, and not available to a city of 50,000 or a bank with 10 branches and 100 ATMs. With the advent of open system server architectures to support the day-to-day activities of the small or medium enterprise or government agency, why haven't we seen an F-T IT solution for the small-to-medium enterprise (SME)?

Apparently, HP has asked itself that same question, and answered: why not? One year after introducing a commodity-based replacement for their high-end S-Series of non-stop systems, they have introduced a downsized version based upon the same operating environment. To see if the *F-T Integrity NonStop NS1000* can fit your needs, please read on.

IN THIS ISSUE

- **Fault-Tolerant Data Center Needs** 2
- **HP NonStop Solutions**..... 2
- **Integrity NonStop NS1000 Server**..... 3
- **Conclusion** 4

Fault-Tolerant Data Center Needs

When the average data center, whether Fortune 500 or your neighborhood car dealership, installs a new server, the IT staff has certain requirements for performance and certain expectations for reliability, availability, and serviceability (RAS) functionality. These expectations normally do not include 7x24x365 availability. If the accounting server is down, you can send out the invoices an hour later. That is not a crisis. If you are the CIO of a bank, however, no matter what size, you know that your ATM network must be operational ALL of the time. If you are the community emergency services officer, regardless of municipality size, you know that your “911” emergency system has to be operational all day, every day. Lives are at stake every time the telephone rings. A system outage, whether planned or unexpected, is simply not acceptable. In fact, with the passage of time, data centers are installing more applications with higher service level agreements (SLA) than ever before. **When cost is discussed, it is the cost of not having a Fault Tolerant (F-T) system available.** CxOs and government officials around the world are becoming more concerned with mitigating risk and investing in peace of mind, and rapid response, for their customers or their constituents.

There is a variety of fault-tolerant (also called “high-availability”) systems available for these environments. Unfortunately, most if not all of them carry a very hefty price tag to ensure continuous system availability. These solutions incorporate sophisticated, and often proprietary, platforms designed for the largest Tier 1 requirements that, in fact, have no choice. If a large brokerage house on Wall St. loses its trading system, it could lose millions of dollars per hour and create anger to its best customers. No CFO would run that risk. The ROI on a multi-million dollar F-T system could be 2 hours. What would be the cost of an emergency services system crashing during a call to report a fire: 100 lives, 10 lives, 1 life? For any major city, the cost of a F-T response system is a minor factor compared to the reliability and availability of that system.

For the smaller community with a 150-bed

hospital, or a brokerage house branch, however, the lack of resources to acquire a full-scale, mission-critical, F-T system is a major concern. An affordable enterprise class system is a necessity. These environments still have the same 7x24x365 availability expectations, the same requirements for secure and immediate response provided by a Tier 1 solution; however, they do not require the scalability inherent in them. They simply do not have the same high volume OLTP requirement. In some instances, such as the brokerage house, the data center is looking for a solution that can be sized to the environment – for Tier 1 environments on Wall Street and for SME deployments in branches around the country.

In order to downsize a F-T system effectively and to make them more affordable, vendors, such as HP, have transitioned their proprietary lines of no-single-point-of-failure systems to commodity platforms, taking advantage of the economies of scale available in the x86 world. HP has evolved the most recent version of their *S-Series* line of *NonStop* servers to a complete line of F-T systems based upon their *Integrity* series of *Itanium-2* servers.

HP NonStop Solutions

HP has been a leader in F-T solutions for over thirty years, via their acquisition of Tandem (via Compaq) in 1997. Their first F-T system came into being in 1974 using a proprietary architecture, since the PC with Intel chips would not appear for another 10 years. Using this experience, along with years of work with Intel in developing the Itanium-2 processor, HP introduced, in 2005, a new commodity F-T solution for implementing processor and data integrity - the HP *Integrity NonStop NS16000 Server*¹. Using the *NonStop Advanced Architecture (NSAA)*, the *Integrity NonStop* provided a new high-end solution for the *S-Series*, protecting the investment that enterprises had made in HP systems, and opening a new arena for enterprises looking for an open systems F-T solution,

¹ See The Clipper Group Navigator dated June 15, 2005, entitled *HP Gives Green Light to NonStop – Adapts Integrity for Failsafe Environments*, available at <http://www.clipper.com/research/TCG2005036.pdf>.

using a virtualized application environment and standard Ethernet and Fibre Channel interfaces.

The NS16000 has excellent linear scalability. It is capable of growing from a single node of two Itanium-2 processors to 255 nodes, up to 16 processors in each, for up to a total of 4080 CPUs. In order to control costs, HP uses the *Integrity rx4640* as the primary building block, modifying the unit with the logic required to implement *Dual Modular Redundancy (DMR)* and *Triple Modular Redundancy (TMR)* for the highest levels of availability and reliability. While costing less than a comparable S-Series solution, a typical eight-processor NS16000 still carries a significant price tag at \$1.8M. In order to make the NonStop technology more affordable to a wider range of customers, earlier this year HP introduced the *Integrity NonStop NS14000 Server*. Using the same NSAA technology, this system has 50% of the scalability of its big brother, with the capability to expand to 255 nodes, with up to eight processors in each, for a maximum of 2040 processors. While the NS16000 uses Itanium 2 processors at 1.6GHz with 6MB of cache, the NS14000 runs at 1.5GHZ with 4MB of cache. This provides a performance level of 75% of the NS16000; however, the NS14000 retains the same DMR and TMR reliability characteristics as the NS16000, with a TMR configuration rated at seven nines availability (99.99999%), although it does not support the S-Series I/O supported in the NS16000². With a hardware price at just under \$300K and a software price at just over \$200K for a DMR 2-processor system, a fully configured 8-processor solution has a TCO of just under \$1.5M, still not exactly within the budget of most enterprises.

Recognizing the opportunity that the commodity NonStop technology offered, HP has now decided to really scale down the architecture, making a more affordable, simplified version available to SMEs and smaller government jurisdictions, alike – the *HP Integrity NonStop NS1000 Server*.

Integrity NonStop NS1000 Server

In order to make the NonStop architecture available to smaller enterprises and for departmental functionality, HP reduced the scalability and went to off-the-shelf Integrity components. Using a standard *Integrity rx2620* as a building block, the HP Integrity NonStop NS1000 can scale to eight mono-processors using an Itanium 2 processor at 1.3GHz, each with 3MB of cache. An NS1000 will perform at 65% of the processing power of the NS16000 with an equal number of processors. At an average list price of \$85K per processor, the NS1000 doubles up on the price/performance of its two big brothers and outdistances the S-Series by a factor of four. When you compute a total cost of ownership (TCO) for three years, the results are more revealing. A two-processor NS1000 configuration has a TCO of \$225K for hardware and software. This compares quite favorably to competitive solutions from IBM (AIX) estimated at \$350K and Sun (*Solaris*) at \$250K. The four-processor results are equally revealing with the NS1000 at a TCO of \$400K vs. IBM estimated at over \$900K and Sun coming in at \$500K³. This pricing puts five-9s non-stop reliability and availability within the budget of enterprises with reduced scalability requirements and on the agenda of smaller municipalities around the globe.

The decision to use standard rx2620 nodes with a modified non-stop architecture appears to be a key factor in making non-stop processing affordable to thousands of new enterprises. Unlike the NS16000 and the NS14000, the NS1000 implements the *Non-Stop Value Architecture (NSVA)*, an N+1 environment with software fault tolerance and hardware redundancy. The NS1000 does not support the DMR/TMR architecture of the more scalable models, but it does support the latest release of the *HP NonStop Operating System, Mission Critical Edition*, optimizing the relationship between server hardware components, NonStop interconnect technology, and the NonStop OS and database environment. A fully configured NS1000 consists of eight mono-processor rx2620s. They are not paired. If any node fails, any of the other rx2620 nodes can take on the workload of the

² S-Series I/O cabinets that house internal SCSI storage cannot be directly connected to the NS1000. However, the NS1000 can connect to systems with S-Series I/O cabinets via Expand over TCP/IP or ATM.

³ Competitive pricing provided by HP.

failed node. It is a simple and efficient architecture to deliver the high availability required by many emerging SME markets. (See Exhibit 1, at the right.)

The NS1000 is an ideal alternative to the complex high-availability solutions available from a variety of UNIX vendors. The typical UNIX H/A environment relies upon tailored scripts, configured as a result of expensive professional services engagements, and on-going maintenance updates to maintain an evolving IT environment. The NS1000 also offers application virtualization based upon specialized characteristics of the NonStop environment, such as the message-based architecture of the NonStop OS and the ServerNet interconnect.

Instead of simply announcing the end-of-life for the S-Series, HP has created a growth environment based upon commodity architecture. With dual-core Montecito scheduled to be qualified during 2H06 and a defined roadmap for Itanium 2 thru 2015, HP has re-energized an already stable fault-tolerant community.

Conclusion

While enterprises around the globe are fixated upon the debate between *scale up* versus *scale out* technologies, HP has expanded their real-time mission by **scaling down** their mission-critical NonStop environment with the NS1000. By re-engineering the fundamental architecture, HP has improved the performance and reliability while reducing the price, significantly.

In IT, as in all aspects of life, perception becomes reality. If you believe that all fault-tolerant systems have to be complex, proprietary engines, then in your mind, they are! HP has changed that paradigm with the introduction of the Integrity NonStop series, in general, and the NS1000 in particular. By starting with an off-the-shelf, commodity node, HP has managed to reduce the TCO of fault-tolerant data centers. With the introduction of the NS1000, HP now provides the data center with choice, the same choices that any consumer gets outside the IT world.

As with the ice cream cone, the IT staff has their choice of small, medium, or large,

Exhibit 1 – Emerging SME NonStop Markets

- **Finance** –POS/EFT environments and real-time financial services;
- **Telecommunications** – Infrastructure for developing countries and a platform for real-time services;
- **Healthcare** –Real-time patient care in smaller hospitals;
- **Government** – Extends the opportunity for digital empowerment of state and local constituencies; and
- **Cross Industry** – Provides ideal platform for RTE opportunities worldwide.

depending upon their scalability requirements. Most significantly for the CIO responsible for an environment of F-T servers that includes both data centers and branch offices, a solution for each based upon their needs can be deployed, knowing that the operating system and application software, like the ice cream in the cone, remains the same, regardless of size. This common operating environment helps to reduce the distribution and maintenance costs, reducing the overall TCO, and improving the bottom line. This is especially true for data centers already using HP's S-Series F-T platforms.

As a result, over 250 partners have already ported mission-critical applications to the NonStop environment, supporting a network of dozens of enterprise-class data centers. With this software foundation already available, HP has simplified the task of converting an SME-sized IT environment from an environment where a "blue screen" indicates a coffee break for the staff and a loss of revenue. If you are one of many with increasing SLA responsibilities, take a look at the NS1000; it may be less expensive than you thought.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group *Captain's Log*, The Clipper Group Voyager, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.