connect
Your Independent HP Business Technology Community

The Connection

A Journal for the HP Business Technology Community

HP Integrity NonStop X: Welcome the NS7

Vodacom's One-Year Recovery

How Will NonStop Fit Into the Internet of Things?

May – June 2015 { Volume 36, No. 3 }

XYGATE Data Protection

All the benefits of SecureData™
No application changes

Secure Stateless Tokenization
Format Preserving Encryption
Standards Based AES
Stateless Key Management

XYPRO
Mission Critical Security

Voltage
security

Learn more at
xypro.com/XDP

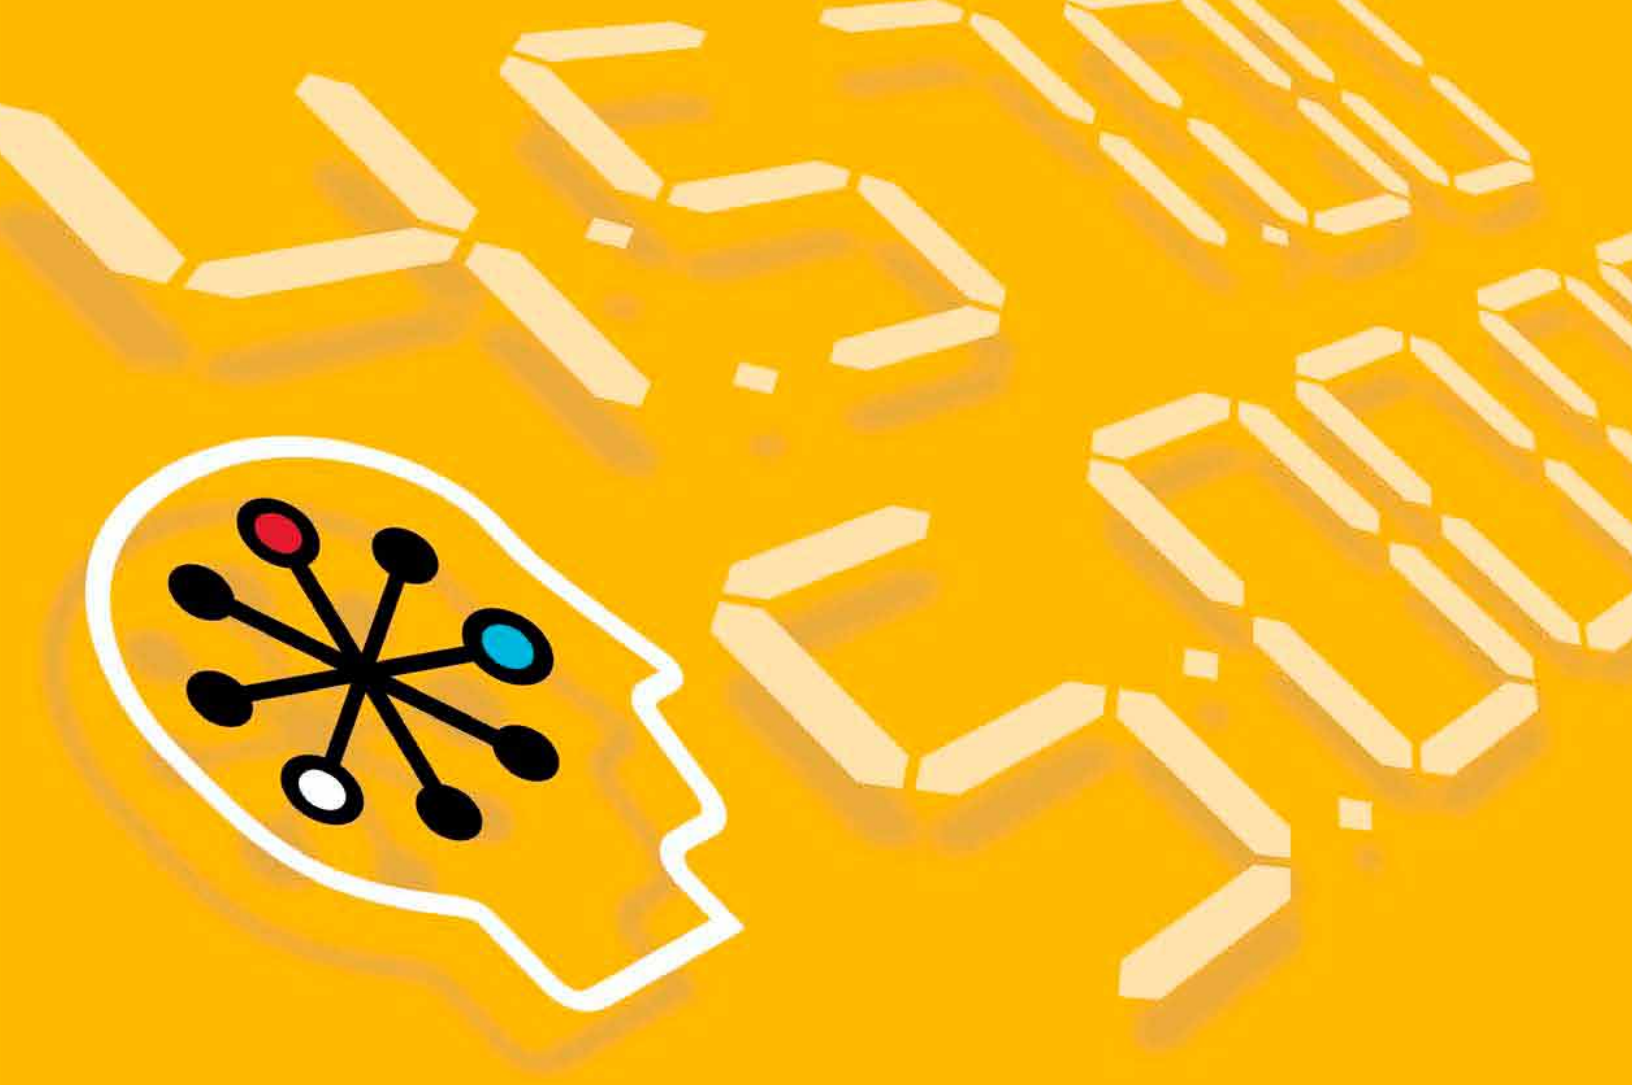

ALTERNATIVE THINKING ABOUT KEEPING TIME:

Guaranteed Time Synchronization

It's 5:00 o'clock at your workstation, but the server next door thinks it's 4:57. How accurate are your transactions?

In today's world of distributed processing and stringent audit requirements, all the systems involved in a distributed transaction must operate at the highest level of accuracy. Without synchronized computer clocks, it's impossible to correlate data from multiple system sources, much less understand transaction flow, pinpoint bottlenecks, or resolve problems.

HP NonStop Time Synchronization provides up to microsecond accuracy across all of your computers, whether you're running HP NonStop, Linux, or Windows operating environments. And, this product is fully backed and serviced by HP.

Technology for better business outcomes.

Key features and benefits

- Dramatically reduces the risk of system events or database records appearing out of sequence.
- Works across your heterogeneous environment — HP NonStop, Linux, and Windows operating systems.
- Automatically synchronizes clocks in a cluster or network; no user intervention required.
- Assures time updates occur safely on all platforms.

Contact your HP representative or partner for a FREE 60-day trial.
Visit www.hp.com/go/timesync

Replicate 100% of Your SQL/MP Database DDL Changes.

With 0% Headaches, Hassles or Human Error.

ALTERNATIVE THINKING ABOUT DATABASE REPLICATION:

Introducing HP NonStop SDR: the industry's only off-the-shelf solution for replicating changes made to the structure of your SQL/MP databases.

How many hours per month does your staff spend making tedious database structural changes? What are the costs of inevitable human error?

What if there was a way to automate database DDL replication to eliminate these headaches, hassles and costs?

Enter: HP NonStop SQL DDL Replicator (SDR) – the industry's only solution for automated SQL/MP DDL database replication.

Only NonStop SDR ensures that after performing DDL operations – creating a table, adding a column, or moving a partition – changes will automatically be replicated and implemented at the correct point in the audit stream.

Now, routine changes no longer require downtime. And best of all: Since everything's automated, the risk of human error is eliminated.

Which means you can "set it and forget it" – while you reduce downtime, remove risk, and free up your staff for much more important work. And what's not to love about that?

Technology for better business outcomes.

HP NONSTOP SDR

- Automatic replication of NonStop SQL/MP DDL operations
- Designed to work seamlessly with NonStop RDF
- Minimal setup or operator management
- Essential NonStop fault-tolerant design
- Supports DDL replication for non-TMF-audited tables

Contact your HP representative or partner for a FREE 60-day trial.
Visit www.hp.com/go/nonstopcontinuity

Table of Contents

Columns...

Special Feature Brought to You by

OmniPayments

OmniPayments' OmniCloudX Serves the Unbanked Customer

Not everyone uses a financial institution for traditional banking services. Those who don't are considered "unbanked." They do not have credit cards, bank accounts, car loans, or mortgages. Their currency is cash, and what is hidden under their virtual or physical mattresses is what makes the billions of unbanked customers worldwide so attractive to merchants eager to sell to this lucrative, untapped market.

Enter OmniCloudX, running on NonStop X. OmniCloudX hosts numerous instances of OmniPayments, the popular financial transaction switch that offers a comprehensive solution for retailers to acquire, authenticate, route, switch and authorize transactions across multiple input channels. Based on a modern Service Oriented Architecture (SOA), OmniPayments and its next-generation OmniCloudX consist of several service components, of which the critical elements are built on NonStop.

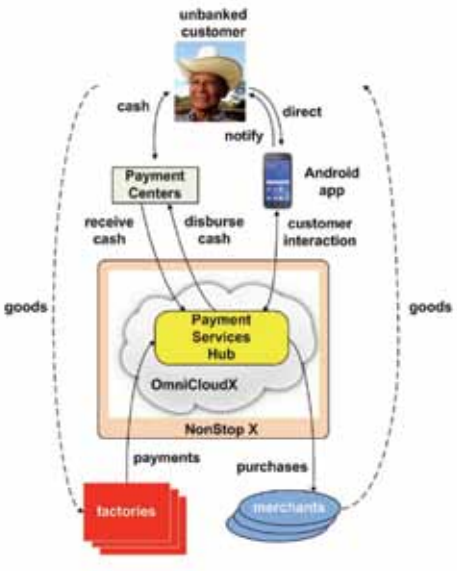

OmniCloudX now has been configured to support the needs of unbanked customers. Via their Tax IDs, Social Security numbers, National IDs, etc., unbanked customers can be paid wages by employers, can make payments to merchants, or can receive government subsidies via the use of approved OmniCloudX Payment Centers. OmniCloudX serves as the Payment Services Hub (PSH) so that unbanked customers and member merchants can interact in a vibrant, virtual marketplace.

OmniPayments, Inc.'s new deployment of the OmniCloudX Payment Services Hub is for the milk producers of a Latin American country. Over 700 factories produce milk products and sell their goods through four country-wide corporations. The milk is provided to the factories by millions of farmers, many of whom are unbanked. OmniPayments has established a network of participating stores that serve as payment centers, or cashiers, for unbanked customers. The stores are connected to the OmniCloudX Payment Services Hub. Who registers for Hub membership are the factories, the farmers, and the merchants who wish to sell goods to other members. Registered farmers download an OmniPayments-developed Android app for use as a mobile payment center. When they are notified that the factories have paid them, the farmers can receive their payments in cash at one of the payment centers or can redirect their payments to the purchase of goods from Hub-registered merchants.

The OmniCloudX Payment Services Hub keeps track of what is sold to the factories, what is owed and paid to the farmers, and what transactions are made between merchants and other Hub members. Its infrastructure provides a real-time, 24x7 centralized view of all payment activities. Fault-tolerant and cost-effective, the Hub's enormous transaction capacity comes courtesy of its NonStop X environment. OmniCloudX also provides preauthorization services via the OmniPayments Preauthorization Engine and seamlessly connects Linux/Windows-based applications to NonStop. Backup systems make any system outage immediately recoverable by switching Hub activities to another system in the cloud. Also supplied are complete security functions, including encryption-at-rest and encryption-in-flight.

OmniPayments Inc.

1566 La Pradera Drive
Campbell, CA 95008 USA
PHONE: +1 408 364 9915
sales@omnipayments.com
www.omnipayments.com

2015 Connect Board of Directors

PRESIDENT
Henk Pomper
Plusine ICT

VICE PRESIDENT
Rob Lesan
XYPRO

PAST PRESIDENT
Steve Davidek
City of Sparks

SECRETARY»TREASURER
Glen Kuykendall
Xerox

DIRECTOR
Michael Scroggins
Washington St. Community College

CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide

HP LIASON
Janice Zdankus
Enterprise Group, HP

The Connection

The Connection is the official magazine of Connect, an independent, not-for-profit, user-run organization.

Kristi Elizondo.....CEO
Stacie Neall.....Managing Editor
Kelly Luna.....Event Marketing Mgr.
Keith McLemore.....Membership Director
Janice Reeder-Highleyman.....Editor at Large
Dr. Bill Highleyman.....Technical Review Board
Karen Copeland
Thomas Burg
Bill Honaker
Justin Simonds
Joseph A. Garza.....Art Director

We welcome article submissions to the *The Connection*. We encourage writers of technical and management information articles to submit their work. To submit an article and to obtain a list of editorial guidelines email or write:

The Connection
E-mail: sneall@connect-community.org
Connect
P.O. Box 204086
Austin, TX 78720-4086 USA
Telephone: +1.800.807.7560
Fax: 1.512.592.7602

We accept advertisements in *The Connection*. For rate and size information contact:
E-mail: info@connect-community.org

To obtain Connect membership and *The Connection* subscription information, contact:

Connect Worldwide, Inc.
P.O. Box 204086
Austin, TX 78720-4086 USA
Telephone: +1.800.807.7560
Fax: +1.512.592.7602
E-mail: info@connect-community.org

Only Connect members are free to quote from *The Connection* with proper attribution. *The Connection* is not to be copied, in whole or in part, without prior written consent of the managing editor. For a fee, you can obtain additional copies of *The Connection* or parts thereof by contacting Connect Headquarters at the above address.

The Connection often runs paid advertisements and articles expressing user views of products. Articles and advertisements should not be construed as product endorsements.

The Connection (ISSN 15362221) is published bimonthly by Connect. Periodicals postage paid at Austin, TX. POSTMASTER: Send address changes to *The Connection*, Connect Worldwide, Inc., P.O. Box 204086, Austin, TX 78720-4086 USA.

Ensuring continuous business for mission-critical workloads

Transform your mission-critical environment with the new HP Integrity NonStop X—the only fully-integrated, fault-tolerant Compute providing continuous availability and massive scalability on x86.

With 40 years experience running the world's most demanding transaction processing and database workloads, NonStop X is the answer when downtime is not an option.

HP Integrity NonStop X. When availability matters.

Make it matter.

For more, go to
hp.com/go/nonstop

News from HP's NonStop Enterprise Division

Continuous Business for Mission-Critical Workloads

Just over a year ago, in a ballroom not far from our Palo Alto campus, we announced a strategy to extend the 100% fault-tolerant HP Integrity NonStop platform to the Intel® x86 architecture. This past fall, at the Connect Technical Boot Camp, we gave you a preview of the new family. At HP Discover Barcelona in December, we saw NonStop take the main stage. The anticipation has been building - from the customers I've met with, to the partners who participated in the beta program in our Advanced Technology Center– all have been excited about what this means to their business.

Now I am delighted to say that the new NonStop X is here!

With the release of the HP Integrity NonStop X system, customers will experience the only flexible approach to a fault-tolerant infrastructure with the choice of Intel® Itanium® or the x86 architecture of Intel® Xeon® processors. This is a major milestone in NonStop's history, and is well aligned with the HP strategy to offer choice in mission-critical computing through our evolving Intel Itanium offerings and bringing mission-critical capabilities to the x86 architecture. Customers can leverage a proven solution for always-on business that delivers timeless value regardless of the underlying architecture.

NonStop X systems are the only fully-integrated, fault-tolerant Compute on x86. The first model in the family is shipping now, with additional models coming later in the year. For more information about the new NonStop X family, see the article ***HP Integrity NonStop X: Welcome the NS7*** by HP Distinguished Technologist Jim Smullen and HP Senior Product Manager Mark Pollans and also visit hp.com/go/nonstop.

NonStop already powers the most-demanding transaction processing and database workloads for many of the world's largest companies, and now an even larger world of possibilities is quickly emerging. As business complexity, routes to market, and availability create greater business challenges, HP Integrity NonStop systems are behind the scenes in some very interesting, and possibly unexpected, use cases:

- Connected car —NonStop enables vehicles and a wide range of devices to share information and intelligence wirelessly.
- Integrated subscriber data—NonStop synchronizes databases across geographic regions in real time.
- Roaming across countries and networks—NonStop provides seamless connectivity and authentication for mobile subscribers on different networks.
- Mobile payments—NonStop runs the infrastructure for payments processing from mobile phones to point-of-sale.
- Point-of-sale security—NonStop handles secure retail tendering and card payment switching.
- Medicaid claims processing—NonStop processes thousands of claims for concurrent online users.

For more perspective on potential use cases for NonStop, see the article ***How will NonStop fit into the Internet of Things*** by HP Master Technologist Justin Simonds.

Customers have always been clear about the business value of the NonStop platform. Our technology enables enterprises to achieve a competitive advantage and market differentiation, address rapid growth and changing customer demands, and quickly deploy new programs. NonStop also lowers their risk of downtime, security issues, and compliance delays. The NonStop X family is a breakthrough for our customers that demand the very highest level of availability, scalability, and performance that NonStop has provided for over 40 years—and want to gain the advantages of an industry-standard x86 platform.

Finally, be sure to join us at HP Discover 2015 in Las Vegas, June 2-4. We will have several sessions for mission-critical solutions and panels with customers and IDC. We'll also have a booth on the trade show floor, so please stop by to meet the experts and see a demo of the HP Integrity NonStop X and HP Integrity Superdome X systems. The full session catalog and registration information may be found at hp.com/go/discover. Discover is always a great opportunity to see the full range of solutions HP has to offer, gain insight into the latest compute trends, and meet with customers and partners. I look forward to seeing you there!

Randy Meyer
VP & GM, Mission Critical Solutions
HP Servers

ADVOCACY

HP To Separate Into Two Public Companies

Dr. Bill Highleyman >> **Chairman** >> The Connection Technical Review Committee

The Hewlett-Packard Company (HP) is currently a leader in two distinct markets – enterprise computing and personal computing. In order to provide the focus, financial resources, and flexibility to adapt quickly to market and customer dynamics while generating long-term value for shareholders, HP plans to separate itself into two distinct and independent Fortune 50 public companies:

- *Hewlett-Packard Enterprise* will provide the next generation of technology infrastructure, software, and services for leading enterprises.
- *HP Inc.* will be a predominant personal systems company specializing in PCs, printers, scanners, and fax machines.

Hewlett-Packard Enterprise will build upon its strengths in servers, storage, networking, converged systems, and software as well as its OpenStack Helion cloud platform. Meg Whitman, the current Chairwoman, President and Chief Executive Officer of HP, will become the President and Chief Executive Officer of Hewlett-Packard Enterprise. Patricia Russo will become the Chairman of the Hewlett-Packard Enterprise Board.

HP Inc. will be a leading personal systems and printing company with a strong roadmap into many new technologies such as 3D printing. Dion Weisler will serve as President and Chief Executive Officer of HP Inc., and Meg Whitman will be its Chairman of the Board. HP Inc. will retain the current HP logo.

The transaction is expected to be complete by the end of 2015. HP shareholders will own shares of both Hewlett-Packard Enterprise and HP Inc. The transaction is intended to be tax-free to HP's shareholders for federal income tax purposes.

HP is approaching the fourth year of its five-year turnaround plan initiated by Ms. Whitman shortly after she was promoted to her current HP positions in September, 2011. Interestingly, her predecessor, Leo Apotheker, had suggested selling HP's personal computer division or spinning it off into a separate business, thereby quitting the PC business altogether while continuing to sell servers and services to business customers. This approach was rejected by Ms. Whitman's board at the time. However, the strategy now has returned in a great circle and is being implemented.

"Our work during the past three years has significantly strengthened our core businesses to the point where we can more aggressively go after the opportunities created by a rapidly changing market," said Ms. Whitman. "The decision to separate into two market-leading companies underscores our commitment to the turnaround plan. It will provide each new company with the independence, focus, financial resources, and flexibility they need to adapt quickly to market and customer dynamics while generating long-term value for shareholders. In short, by transitioning now from one HP to two new companies, created out of our successful turnaround efforts, we will be in an even better position to compete in the market, support our customers and partners, and deliver maximum value to our shareholders."

Hewlett-Packard Enterprise

Hewlett-Packard Enterprise will possess a strong, multi-year innovation roadmap across technology infrastructure, software, and services to allow customers to take full advantage of the cloud, big data, security, and mobility. At one point, HP intended to compete against Amazon, Google, and Microsoft in the public cloud arena. However, HP since has realized that it is better for the company to be a supplier of equipment to the large public clouds. For smaller customers, Hewlett-Packard Enterprise plans to build private clouds that also can take advantage of the major public clouds using its Helion technology.

Hewlett-Packard Enterprise will include the USD $27 billion division that sells industrial-grade computing and networking gear and the USD $23 billion Enterprise Services business that runs the technical and IT operations for other companies. It will be traded under the ticker symbol HPE.

The company intends for HP Financial Services to continue to provide financing and business-model innovation for customers and partners.

"Over the past three years, we have reignited our innovation engine with breakthrough offerings for the enterprise like Apollo, Gen 9 and Moonshot servers, our 3PAR storage platform, our HP OneView management platform, our HP Helion Cloud, and a host of software and services offerings in security, analytics and application transformation," continued Whitman. "Hewlett-Packard Enterprise will accelerate innovation across key next-generation areas of the portfolio."

HP Inc.

HP Inc. will continue to execute against a well-defined and established strategic plan, ensuring continuity for customers and consistent value to shareholders. HP Inc. will have annual revenues exceeding USD $50 billion. It will trade under HP's current ticker symbol, HPQ.

"Since assuming responsibility for the Printing and Personal Systems Group, Dion and his leadership team have done an excellent job of building our relationships with customers and channel partners, segmenting the market and driving product innovation," added Whitman. "The creation of HP Inc. will only accelerate the progress the team has made."

"This is a defining moment in our industry, as customers are looking for innovation to enable workforces that are more mobile, connected and productive while at the same time allowing a seamless experience across work and play," said Weisler. "As the market leader in printing and personal systems, an independent HP Inc. will be extremely well-positioned to deliver that innovation across our traditional markets as well as to extend our leadership into new markets like 3-D printing and new computing experiences – inventing technology that empowers people to create, interact and inspire like never before."

The Partner Program

HP plans to support its partners through the transition by sharing the benefits of a program it calls the HP Partner Navigator. Under Navigator, HP is funding resources at top solution and distribution providers with whom HP has direct relationships. Such resources will ensure that these partners are on top of changes coming as a result of the split.

Partner Navigator also will make sure that smaller partners - those who purchase primarily from distribution providers – are not left out of the program. HP is strengthening its partner support line and is publishing a series of FAQs related to the split of the partner program. It expects to see many questions from partners closer to the split taking place.

Increasing Shareholder Value

When Ms. Whitman assumed her current leadership positions in September, 2011, HP shares were trading at USD $23 per share. By December 2014, they had almost doubled to USD $40 per share. This reflects the success of Whitman's leadership. It is expected that the corporate split will continue to improve share prices.

Learn More

For further information on the HP split, refer to HP's press release of October 6, 2014, entitled "HP To Separate Into Two New Industry-Leading Public Companies" (http://bit.ly/1Q2D0dR)

Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him a popular seminar speaker and a sought-after technical writer. Dr. Highleyman is a past chairman of ITUG, the former HP NonStop Users' Group, the holder of numerous U.S. patents, the author of Performance Analysis of Transaction Processing Systems, and the co-author of the three-volume series, Breaking the Availability Barrier.

Did Someone Say "Downtime"?

Barry Forbes >> XYPRO Technology

All I have ever really known with complete certainty in my near thirty-year relationship with NonStop has been that HP NonStop computers are mission critical servers that are truly fault tolerant and have full redundancy capabilities for a single reason: They need to be available all the time. Availability is the primary directive. Or at least it was.

Very recently, I had the opportunity to spend some time with some friends at a longtime customer. This customer is one of the top five US Banks and takes very seriously the need for NonStop and its reliability and availability. I was told that the senior executives at this bank have indicated there is a single circumstance under which they would accept, and actually prefer, downtime. That circumstance is a security breach. After all, a downtime event is recoverable. A security breach is not.

The words "downtime" in the NonStop world are sheer blasphemy. How can this be? Uptime is critical to a successful business model in the industries that rely on NonStop. Uptime ensures customers' service expectations are met, delivers financial benefits and avoids penalties for downtime. Uptime also comes with bragging rights and prestigious awards.

Figure 1. Management's concern about breach

Percentage concern level on a 10-point scale

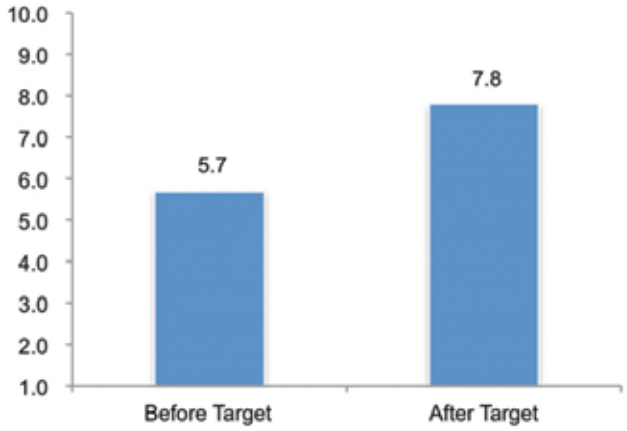

ource: Ponemon Institute, 2014: A Year of Mega Breaches, January 2015

Like so many of us in this great community, my introduction to Tandem was far too many years ago and in a very different world than we are a part of today. I was in high school.

My first introduction was not through employment, but from my father who had been working on a project to bring an ATM and Online-Teller network to the bank where he was employed. He explained to me about this very special computer system that could process transactions very quickly and had two of everything so it was really reliable. The year was 1983, I remember all of this and that my father was immensely impressed. Like a typical teenager, I didn't really care much about this. I only really began to understand a few years later when I had the opportunity to learn and work with the Tandem myself.

Working part-time in the evenings while in college, I gained some exposure to the inner workings of a bank's data center. It was a hub of activity with lots of people and with machines of all sizes. Reader-Sorters, Line printers, 9-track tape drives, massive disk packs, etc. There was also a prized area on the data center floor where the Tandems were kept. The Tandem operation also had a separate command control room where these systems were monitored. Everyone knew they were there, everyone knew they were special, not everyone knew why.

The Tandems would run all the time, literally. This was their value. In the data center, the Tandem NonStop II sat beside the gleaming new TXP. I still knew very little but I began to understand why these Tandems were special.

Later on, as we came to depend more and more on these machines, the systems in place to support their uninterrupted operation were big, important and becoming more sophisticated. As an operator, testing the UPS (Uninterrupted Power Supply) system, test-firing the diesel generator at least weekly and ensuring there was enough fuel to run for several days was a mandatory procedure. These were mission-critical computers. They had to run all the time and the Tandem systems did.

The only thing that is constant in technology is change and striving to improve and speed up the way things work. There is always something driving the need for even greater reliability and uptime. A simple fire suppression system malfunction or even worse, a fire itself, could render the system unusable. The growth of DR (Disaster Recovery) centers began in an upward direction. In the unlikely event of a disaster, the remote DR center could, and had to, be up and running in a matter of minutes. Availability was of paramount importance.

Business Continuity Planning was now the new buzzword in the Tandem community (along with remembering to call these computers NonStops following the acquisition of Compaq by HP). With natural disasters such as earthquakes and hurricanes and now very unnatural terrorist threats, the NonStop server had real-time data replication in active-active environments, spanning very large distances to ensure that these computers were operating on individual power grids and fully separate communications infrastructure that could not be affected by the loss of availability at any single site. The great Myth Busters TV show even blew up a NonStop server to prove just how quickly a failover and recovery could happen. These computers are truly mission critical and the customers who purchase and use them do so because their businesses rely on the ability to run without interruption.

Bar Chart 6: Total cost per minute of an unplanned outage
Comparison of FY 2010 and FY 2013

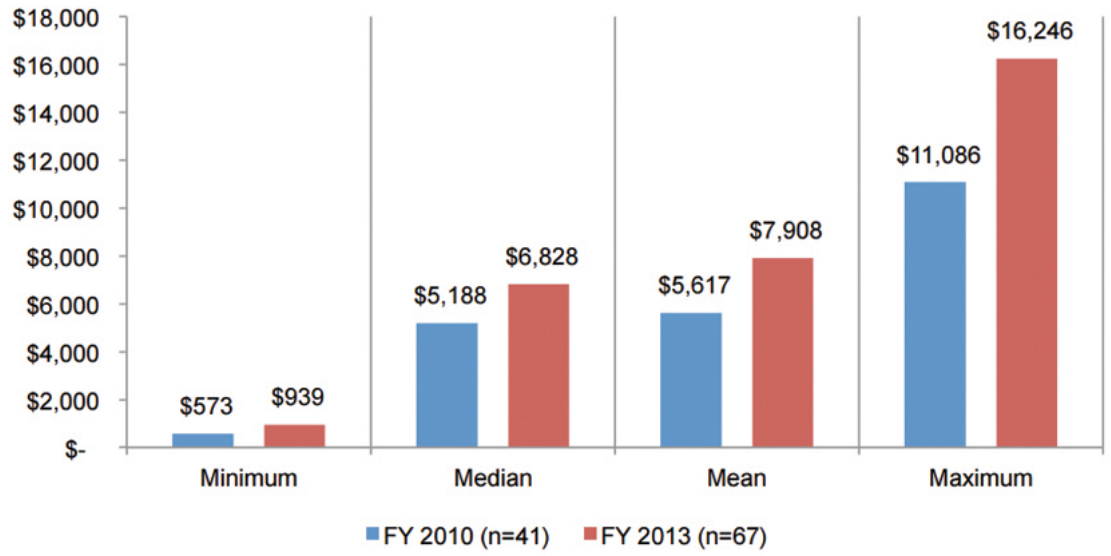

Source: Ponemon Institute, 2013 Cost of Data Center Outages, December 2013

For my thirty years on NonStop, the only direction I knew was that more uptime, and in most cases, continuous uptime, was the way to go. Never did I suspect that there would be something that was so critically important to a business they would sacrifice this near perfection. Sadly, earthquakes, hurricanes, tornadoes, and even nuclear warfare are no longer the ultimate threat to uptime. It is the cyber-criminal.

As a vendor of HP NonStop server security solutions, it's a positive thing to hear a customer say their focus on security is right up there and even ahead of availability and performance. The revelation that unscheduled downtime is more acceptable than a security breach is not only a sign of these modern times but a continental shift in priorities for the majority of companies that rely on fault tolerant, mission critical servers.

And just as the needs for more uptime drove the development of more and more sophisticated solutions to avoid possible availability catastrophes, so, too, the need to thwart the ongoing threats of cyber criminals and hackers drives the development and implementation of advanced security solutions, these days at lightning-speed.

Many of these solutions already exist in the form of strong encryption and tokenization of data, enhanced access controls, audit and analysis, continuous real-time monitoring and threat detection, security incident and event management, and more. It is a matter of time, education, commitment, investment and effort that this very present threat to downtime can be tackled. We're investing our best efforts and resources to staying ahead of the cyber criminals and hackers. It's not too difficult to imagine what will be the next phase in the evolution of the NonStop uptime story, but there is no doubt that security will always be a big part of the solution.

Barry Forbes is a Strategic Sales and Marketing Consultant and Vice President of Sales and Marketing for XYPRO Technology Corporation. With 29 years working with HP NonStop (Tandem), 21 years specializing in NonStop security and 16 years with XYPRO, a software sales and development company specializing in security solutions for Mission-Critical computing market, Barry provides insight and direction to help deliver the best security experience and technology to customers in the Mission-Critical computing marketplace.

NonStop Innovations - Spotlight

Joseph Androlowicz: The Journey from HP NonStop to HP Atalla

Gabrielle Guerrera >> NuWave Technologies

The NonStop sector has long enjoyed the knowledgeable voice of Joseph Androlowicz, a long-term player in our community. Joe has recently transitioned to pioneer product marketing within HP **ATALLA**, a security and encryption solution that safeguards essential data throughout its lifecycle on a variety of platforms. How will this new role for Joe affect the NonStop space? What can the community expect from Joe moving forward?

Atalla, founded in 1973, catered to financial service clients by focusing on security and encryption for companies that value continuous data protection from the moment of creation. Featuring a robust portfolio that began with payment and ATM security, Atalla joined HP (along with the rest of Compaq) and contributed a diverse collection of security platforms. Now HP Atalla, the division, provides protection for sensitive data throughout its lifecycle--whether it is on-site, mobile, or in the cloud. Atalla is one of four pillars of enterprise security at HP; the other three are TippingPoint, Fortify, and ArcSight.

Steven Wierenga of HP Atalla shed some light on the history of the division in a previous Connection article:

"HP Atalla's focus for much of the past 35 years has been protecting the financial services, banking, and retail industries and their customers from fraud and attacks using strong cryptography and key management. Atalla, founded as a startup in 1972, quickly focused on the emerging card payments industry and the need for strong security controls. Atalla became part of Tandem Computers in 1987 and HP's NonStop Division in 2002. In 2012 Atalla joined HP's new Enterprise Security Products group under HP Software."

Gabrielle: Joe, what does your history in the NonStop space look like? Could you bring us on that journey?

Joe: So a brief history of me in the NonStop division: I was there for eighteen years starting with Tandem, then we transitioned to Compaq, where we eventually merged with HP. I spent the first ten years in NonStop education and training, where I was a training manager. One of the bigger programs I worked on involved designing and developing the NonStop certification programs. I managed a team of program managers who interfaced throughout the organization, I sat in on R&D core teams, and I helped bring new products to market. We were responsible for transferring the education information out from the development core teams to the field for "field readiness", which is essentially making sure the support and sales teams were ready whenever something new was being released or updated.

The training department got smaller over the years, and five years ago I moved over to the NonStop product management group. There I did a lot of the technical and outbound communications along with various marketing programs for the division. In this position I usually owned anything media related: Video, social media, photography, and so on. I was also tasked with sourcing, coordinating and sometimes writing many of the HP articles that would go into The Connection magazine and Connect Converge. I often recruited writers or would help people develop content, while sometimes ghost writing pieces that were aimed at the themes for each issue. Over the years, I also managed to drive elements of some of the events like ITUG, ETUG, HP Discover, and Connect NonStop Technical Boot Camp. I would usually be the person coordinating the messaging, the marketing, and the expo booth. I would also coordinate and pull together NonStop hardware and system demos for the event booth, hands-on labs for half- and full-day sessions, and set up and run certification testing centers on-site.

Gabrielle: Could you describe your new role in HP Atalla?

Joe: My new role is in Atalla product management, but rather than owning a single product, I work on all of the Atalla offerings. My role is about 60% product management and 40% marketing. On the marketing side I work with Sheryl Wharf, who is our marketing manager, to pull a lot of stuff together. We have our content in a lot of different places, so my working title is "Knowledge Management Manager", and I'm pulling a lot of our content together to make it easier to find in the field. We are trying to make it simpler for people in the enterprise sales and field enablement organizations to find, learn, and sell Atalla products. I'm also working on identifying, consolidating and updating a lot of our internal content sites and looking for ways to improve internal and external communications. Of course my day is filled with lots of tactical activities that include updating and editing data sheets, use case documentation and system integration content. I'm glad I also still get to be involved when we're producing short videos or product photography.

Another thing that I'm doing is looking into developing new Atalla product certification, which would be similar to what I did with NonStop. Atalla doesn't have any certification programs for end users, customers, or partners, so I'll be working on some of that this year. I'm a bit all over the map for a little while as I start to settle into my new position, but the NonStop audiences will be hearing from me moving forward as I work on helping to push Atalla attachments with NonStop solutions. I'm also working on a new fact guide for the products and helping to communicate

things like demos that people can download and then run a virtual version of our Enterprise Secure Key Manager product. This way they can get a feel for it and see if that works before they buy it. I'm not directly involved in the development, but more so in the communication around it and making sure people can access it and they understand what it is. It's exciting--we've got a lot of things going on!

Gabrielle: How would you describe the relationship between NonStop and Atalla?

Joe: Atalla has traditionally been close to the NonStop product. They were actually both part of, and owned by, Tandem many years ago. When Compaq came on board, the Atalla guys sat next to us, then at some point they moved over to the HP networking organization and subsequently ended up in the HP software organization, and are now one of the four main pillars of Enterprise Security products for HP. To this day, NonStop field support folks are intimately familiar with Atalla because Atalla security products go hand in hand with a lot of the NonStop systems that are sold. Outside of our direct sales engineers in Atalla, it's the NonStop field and sales teams that know the Atalla products best. The Atalla team is also supported by other folks in the enterprise space, including NonStop employees and other enterprise security sales reps, who help support and sell the Atalla products.

Gabrielle: Is there anything new happening at Atalla that we should know about?

Joe: One exciting development is that HP and Atalla recently announced the acquisition of Voltage Security (now HP Security Voltage). The Voltage team is joining with the Atalla group and is now part of the same organizational pillar, which greatly increases the size of our team. Some of my goals surrounding this acquisition are to sort through the changes that will be happening and help the Voltage team integrate with current HP processes. We are also looking for ways to incorporate some of their processes, activities, and methods within Atalla, such as the Agile and Scrum software development processes. We are hoping to combine the best of both worlds.

Gabrielle: Joe, how are you staying connected to the NonStop community?

Joe: I think I'll be staying in touch with the NonStop folks for sure because a lot of them do keep in touch personally and others are still actively selling Atalla products. I still hear from many people in the NonStop world regularly; sometimes due to the transition people think, "Joe used to do that, call Joe - he's still here! How do we do this, Joe?" So in that respect it's nice to be missed. Other times it is people looking for the most up-to-date Atalla information and asking questions about how new things will be integrated, or they could be working with us in leveraging processes and tools to improve communication between the hardware and software security worlds in HP.

Gabrielle: Are you still involved with NonStop, other than through informal communication?

Joe: Not directly, but when it comes to providing product and sales certification around Atalla, the NonStop folks are interested in assisting as we roll out the programs and exams this year. I'll probably have a few of them involved in helping us with the development because they're intimately familiar with the product as well as what it takes to create a solid and bulletproof certification exam.

From the Connect organization standpoint, there's a Connect Advocacy Board within HP, and I started attending meetings with them when they got a new director, who serves as the main liaison to Connect from HP. I attended several meetings before I switched over to Atalla, but I'm staying on the board because of my history and connections with the Connect team. So in this regard I will still be able to help influence things on the HP side based on my continued interaction with Connect, and I'll help them move forward in areas HP is looking to explore, such as finding ways to grow our user group audiences. I'll also have greater input around things that we know are needed from the security side of the business, but also from the NonStop side because the NonStop folks will always have my ear. If anything is going on or needed specifically around the NonStop audiences, I will be one of their strongest advocates. The NonStop team maintains direct connections in Connect as well, but my seat on the advisory board can help push conversations to a higher level when necessary, and I'm always happy to help.

Gabrielle Guerrera is the author of the NonStop Innovations blog, which is located at www.nuwavetech.com/hp-nonstop-innovations. The blog focuses on the latest products, services, partnerships, and other news in the HP NonStop space. Some of her most recent topics were: how NonStop could help to curb fraud in the healthcare industry, how to manage your multi-platform application portfolio, and the future of NonStop X.

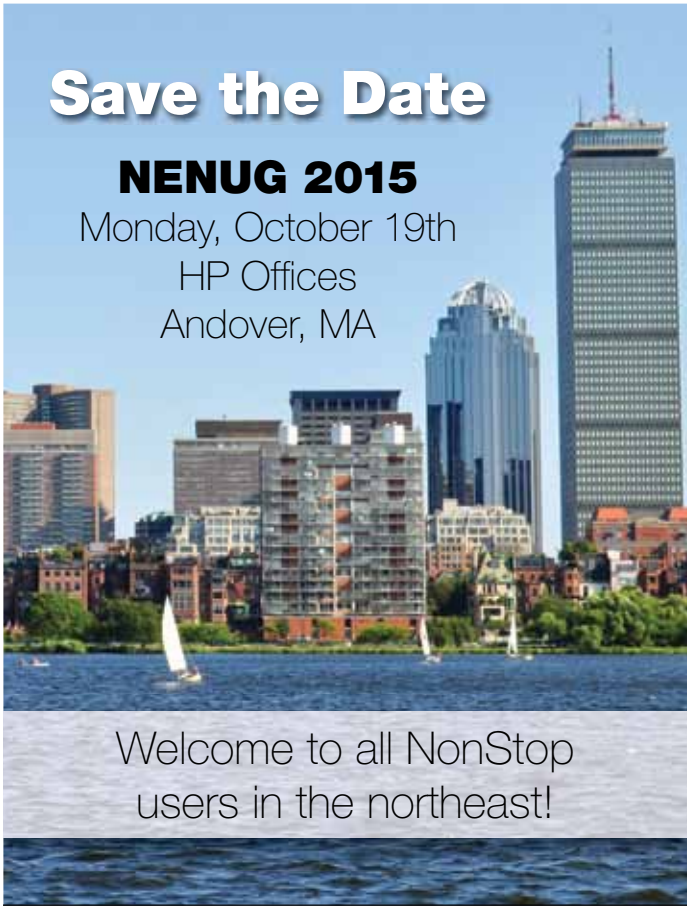

When You Make Sausages, Do You Really Want to Know What Is Inside?

Randall S. Becker >> **President** >> Nexbridge Inc.

Version control technologies have evolved significantly since the 1970s, particularly as a result of workstation-based development. Processes and policies built around trying to manage how software flows to and from workstations is challenging. With flexible process enablers, like git, different approaches to this problem are available, which has unfortunately led to conflicting development cultures. In the git sub-culture, there is a core question that boils down to whether you really want to know what is in your sausages (change history) or not. In concrete terms, there are process questions of how much interaction developers have with central repositories, including when and why, and then what library managers do with what is delivered by developers. This article describes why you need to look at this question and how to choose which approach to use to manage your software. It also will discuss some roles and responsibilities you need to consider that may change when using git.

Preamble

Bob and Steve are good friends and colleagues. Every Thursday, the two of them go out to the food truck at the corner. Bob always orders a sausage with the works. Steve likes his simple: just brown mustard. While Bob joyfully eats his in large bites, juices rolling off his chin, Steve savers every morsel and flavour that comes from the tasty indulgence. Bob often looks at Steve with a bit of disdain, thinking that he overthinks his lunch, while Steve wonders how Bob does not get indigestion.
Back at their desks, side by side, they continue working on their shared project. Just as he ate his sausage, Bob prefers to check-in his code only in big bites, once everything is tested and perfect. Steve, on the other hand, saves every change he makes, recording lengthy dissertations about the purpose of each change, even if it does not work yet. Bob thinks Steve is a bit obsessive about his change history, while Steve thinks Bob is being obsessive about not wanting to show any work that does not test cleanly. To each, this is a matter of pride. Their philosophical disagreements have even become legendary, to the point that the library manager, Jan, has had to break up some rather heated arguments over who is working the right way.
Jan has become an expert in eye-rolling at the two colleagues. At the end of the day, when it comes time to release code, both developers deliver the same high quality delicious sausage. The code moves cleanly through the deployment process and rarely, if ever, is there an issue, so why the arguments? This question of how to work seems to mean so much to both Bob and Steve, but when Jan gets the code, it all seems so nonsensical.

What you have just read was mostly fiction, based on true events, with the names changed because our legal department told us we had to – you know who you are. The question of Sausage Making, as the git community calls it, is a core philosophical disagreement in many organizations, and you will probably encounter it. The root cause of this disagreement calls into question how people fundamentally view their jobs, the type of communication they have with others, and the pride in their own work. For the author's part, the key advice here is not to ignore or minimize the impact of this question. In order to understand it, in a NonStop context, we need to look back in time, into our own history, and look at how code has been managed historically.

A Brief History of Culture

In the beginning, we had EDIT and TAL, but no tools. Our code sat in groups of Guardian sub-volumes that we fixed in place. As time passed, and we patched code by compiling it in production, we realized we needed to keep official copies and work in progress. Deployment became people typing and later scripting "FUP DUP". In the early 1980s, someone had this idea to try to port some standard SCM tools. Then came CONTROL and later RMS, which both did things differently, but we finally had version numbers, change logs, and central repositories. At the same time, the outside world was evolving through SCCS, RCS, PVCS, CVS, and then Subversion. Still, version numbering and central repositories reigned. The cultural impact was not entirely subtle, but important: quality == good; clutter == bad.

In the beginning, we were motivated only to record finished products.

In order to reduce the massive numbers of versions that would result from normal developer activities, we, as a culture, kept our interactions with the central repository to a minimum. It was natural. That and no one likes SCM systems anyway, right?

Then came the Enterprise Toolkit (ETK) and everything went into the proverbial loo. The idea of checking out code, making a change, and checking it back in became impractical for the NonStop-based repositories. Tools soon emerged to deal with that but not before customers had moved their code to off-platform repositories like CVS and Clear Case™. Suddenly, NonStop was isolated. Interactions with the NonStop repositories became even more infrequent (and painful). Some customers stopped storing code on the NonStop and just pushed objects.

But by 2007, the World as a whole had moved on and left NonStop behind, which is unfortunately where we still are today. In the Wider World, people were playing with Linux-style approaches to code management that had theoretical roots back in the late 1980s, but would not see production for some time – the Distributed Version Control System (DVCS), like git and Mercurial. These systems greatly improved workstation-based development by allowing history to be shared and changes to be added without constant contact with the central repository. This was actually

revolutionary – not technologically, systems like ClearCase™ had some of that, but because of its cultural impact. The DVCS allowed two developers to commit their code in isolation, and then combine the code later. Branches became lightweight, easy to create, deliver, and destroy when no longer needed. Culturally, this change allowed developers who were previously frustrated by the need to always be connected to the central repository, to keep their own private change histories – some of us call those research notes. These private histories became a trace or footprint of your work.

For security managers and intellectual property policy writers, workstation-based development was the stuff of nightmares that continues to this day – but that is another lengthy topic to be discussed in the next article.

In git, you have control of your own footprint.

For managers, being able to see all of the research notes was a boon to track what people were doing and whether there were potential quality issues slipping in. As you might guess, this was not received well by some developers. This negative reception was actually unnecessary but by the time git got to it, it was too late, culturally. In git, this footprint is only visible if the developer chooses to share it. The impact pushed further down the path of minimal interactions with the repository; but in a DVCS, this is rather like not saving your work in the editor. Resist committing at your own risk, at least according to Steve. Bob was in the other camp, wanting only to commit work that was perfect.

The DVCS Divide

So here we are, in the present, facing a question of sausage making. It is your development group who really needs to decide, do you want to know what is in the sausage - do you want to see the developer footprints, or do you just want to see the final product. In git, you can do either, or both. This is mostly a matter of culture and process. In my previous article, we saw how you can use merge-squash to make the final deliveries immutable – the git community calls it atomicity if you are Googling this. The same technique applies to changes being pulled into your integration branch. Let's take a look at how the sausage makers work.

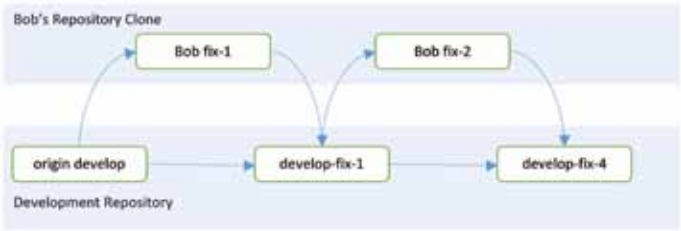

Figure 1 Bob's Ideal View of the World

From Bob's perspective, the repository where he works contains the development branch and his own topic branches. Topic branches are the light-weight branches where you make changes for a specific unit of work, from a bug-fix to a project. Bob's topic branches are very small, containing only one commit for the work he has completed. Once he is ready, Bob pushes his topic branches to the development repository where Jan merges them into the main development branch. It is actually up to Jan how to sequence Bob's changes into the integrated development branch. Bob can continue to work on the second fix either off the same origin point as his original change, or build on top of it. That is actually dependent on the nature of the fix.

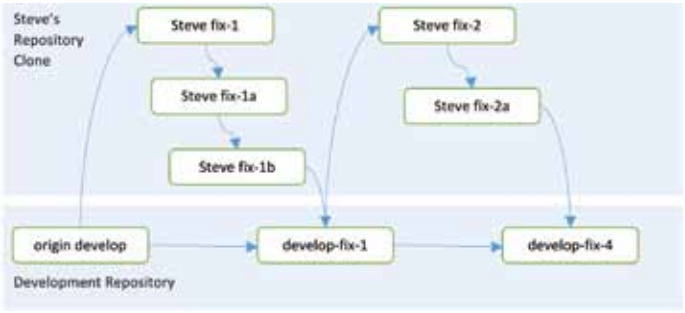

Figure 2 Steve's Ideal View of the World

As with Bob, from Steve's own perspective, the repository where he works contains the development branch and his own topic branches. Steve's topic branches are longer than Bob's because they contain his intermediate changes though the final commit, at which point, he pushes his topic branches to the development repository where Jan merges them into the main development branch. Jan can choose to perform a merge-squash to treat Steve's work as a single commit, or can take the entire history. Like Bob, Steve also can continue to work on the second fix either off the same origin point as his original change, or build on top of it. If you are a small shop, this is probably the extent of the process and decisions you need to look at initially.

Working Together on a Project

However, if Steve and Bob are part of a larger team, you will need to add a layer to the picture when the two work on a project together, because things become a bit more complicated as they share code. In fact, the roles and responsibilities are pretty much the same, but unless Jan wants to stay in the middle of the arguments between Steve and Bob, the two of them are going to have to learn to put on Jan's proverbial hat – that also means that the team maintains proficiency and backup skills for the merge process, something that is essential in an environment with many branches.

Working together means having your own integration branch.

When you decide to have two members of your team work on a change together, and the change takes more than a few days – like an ATM key management enhancement – they are probably going to have to have an integration branch for that change. There is really no difference compared to the main development branch managed by Jan, except it only contains changes for the specific fix that Bob and Steve are working on. Bob can keep his sausage making to himself, while Steve can continue to write a lot of history. When it comes time to sharing changes with each other, Bob can merge his changes onto their **ATM_Key_Enhancement_1234_Branch**, while Steve rebases to pick those up, or Steve can merge his changes and Bob can rebase, or both. That is how they keep in sync with each other on their ever-growing common branch. Pretty cool? Actually, mostly for Steve. Bob probably might think he won't like this approach much because it does mean that their common branch may contain more than one commit. However, if they agree, collectively, that only working changes must be shared, then Bob actually can be the happy one, while Steve will have to wait to rebase or merge until he is at a stable point. It is still pretty cool, because the history is kept intact for both of them. If Bob insists

SQLXPress

Not just another pretty face

An integrated SQL Database Manager for HP NonStop.

Single solution providing database management, visual query planner, query advisor, SQL whiteboard, performance monitoring, MXCS management, execution plan management, data import and export, data browsing, and more...

With full support for both SQL/MP and SQL/MX

XYPRO
Mission Critical Security

Learn more at
xypro.com/SQLXPress

that he does not want to see Steve's history, then Steve can use a merge-squash to – remember what Jan can do? – publish his changes, so their joined history is also clean. By the time Jan gets it, all conflicts are gone – both in code and interpersonal. And better still, the development branch still only contains functionally consistent commits. And that makes the quality of delivered code higher.

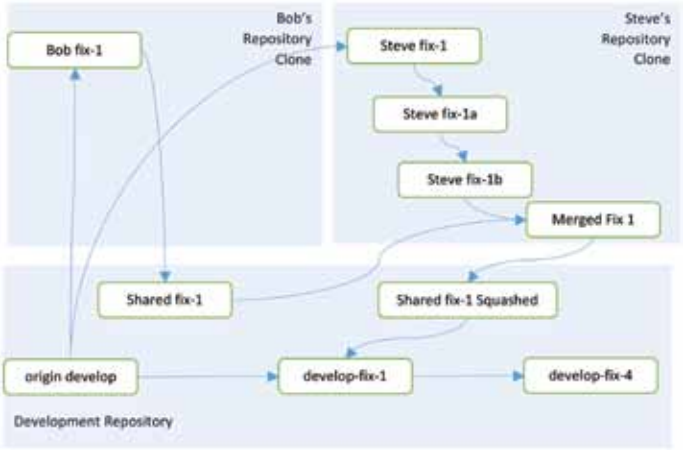

Figure 3 Three Repositories Working Together

In the above diagram, Bob and Steve are collaborating on a fix. Either Bob or Steve could have done the rebase merge and delivery, but in this example, Steve was chosen. Bob's repository contains an abridged image of the history, much to Bob's delight, although if he chooses to fetch all branches, he can see Steve's work. Technically, the **Shared fix-1** commit is actually the same as Bob's initial fix, and the squashed commit could be done either in Steve's repository or on the server, depending on who is performing that operation. Having a good process design and cookbook of work instructions is important so that everyone knows and agrees on what they are supposed to do and when. The branch pointer, **ATM_Key_Enhancement_1234_Branch**, would first be on **Shared fix 1**, and then move to **Shared fix 1 Squashed**. Jan would use that branch pointer, to merge the change into the **develop** branch.

The team can even choose to drop Steve's and Bob's branches as part of regular maintenance, which will orphan the commits along their branches and cause them to be marked for removal during a git garbage collection. But Steve still has the option to retain his own history in his repository for his own research notes.

Production should never see the insides of a sausage.

But none of the details of these changes should ever make it to production. As the previous article discussed, when the final code was delivered, the detailed footprints disappear into the atomic immutable commit that represents the release itself.

Conclusion

The details of sausage making are of concern when thinking about how developers work and keeping them productive. Often, the repository manager is the arbiter of the philosophical disputes discussed here, but their primary role is to ensure that quality code deliveries make it to production. These footprint discussions are valuable because, ideally, they improve how developers work together and how both options can co-exist. Consider carefully whether footprints are essential to your business, particularly if you are conducting audited research. But mostly, just make sure that the sausages taste good by the time they are finally served to your customers.

Randall S. Becker is a speaker, author, and consultant on Policy and Process that delivers continuous availability. He is an expert in Software Configuration and Change Management since 1989 and has spoken at many NonStop and community events. Randall can be contacted at: +1.416.984.9826 or rsbecker@nexbridge.com.

Vodacom's One-Year Recovery

Brett Dismore >> Principal Engineer >> Business Connexion

Paul J. Holenstein >> Executive Vice President >> Gravic, Inc.

Vodacom (Pty) Ltd and Prepaid Front End (PPFE)

Vodacom (Pty) Ltd is one of the largest cellular telephone service providers in Africa. From its roots in South Africa, Vodacom has expanded to include cellular telephone networks in Tanzania, Lesotho, and the Democratic Republic of the Congo. It currently offers voice and messaging services to over 55 million customers. Vodacom is owned by Vodafone, the world's second largest mobile phone company (behind China Mobile).

Prepaid calling cards, the fastest growing cellular option in Africa, are a major Vodacom service. If this service is unavailable, much of Africa's cellular capability comes to a halt as subscribers exhaust their cellular minutes. Therefore, Vodacom uses HP NonStop server pairs to provide prepaid calling card services via its Prepaid Front End (PPFE). The PPFE, primarily a NonStop OSS application, is used by customers to recharge their accounts. If the PPFE fails, subscribers cannot add money to their cellular accounts; if subscribers exhaust their accounts, they no longer have cell phone service. To minimize PPFE outages, the HP NonStop server pairs implement the Shadowbase bi-directional Sizzling-Hot-Takeover (SZT) data replication architecture to ensure multi-second recovery times.

As within its other geographical areas of service, Vodacom provides a PPFE HP NonStop SZT pair in Tanzania. However, in August, 2013, a battery explosion downed Tanzania's production PPFE NonStop server. Vodacom was able to switch PPFE operations to its backup system in just a few minutes; however, problems with the Online Charging System/Intelligent Networking (OCS/IN) platform restricted full operational capability of the remote (backup system) until the next day. Even worse, a year later, the backup capabilities for this installation were still being brought back into service. This article investigates the causes of the lengthy failure/recovery cycle as a sober lesson for other mission-critical businesses that are implementing business continuity architectures.

The Tanzanian PPFE Configuration

Call Handling

A simplified overview of how cell calls are handled is shown in Figure 1. In this figure, cell phone subscribers connect to the Vodacom cellular network via Intelligent Network (IN) systems. Each subscriber is assigned an IN based on the first digits of his telephone number. Each cell phone tower connects to an IN. Various information is collected about a cell phone user (such as the number of minutes he has left in his account, and the geographical location of the subscriber).

When a subscriber turns on his phone, his mobile signal is picked up by the closest cell phone tower. The IN to which the cell phone tower is connected determines the subscriber's assigned IN from his mobile- phone number and notifies the assigned IN of the subscriber's location. The subscriber is connected to an IN, which also obtains the subscriber's account data from his assigned IN.

When a subscriber places a call, it is handled by the connected IN, which contacts the assigned IN of the called party to determine the location of that party and then routes the call to the connected IN of the called party. The call has now been established.

When the call has been completed, both the IN handling the calling party and the IN handling the called party update their respective subscribers' account data (for instance, deducting minutes from the subscribers' accounts of prepaid minutes). The updated data is then sent to the subscribers' assigned INs to update their databases.[3]

Adding Minutes

In order to add airtime to a prepaid cellphone account, a subscriber has a number of choices:

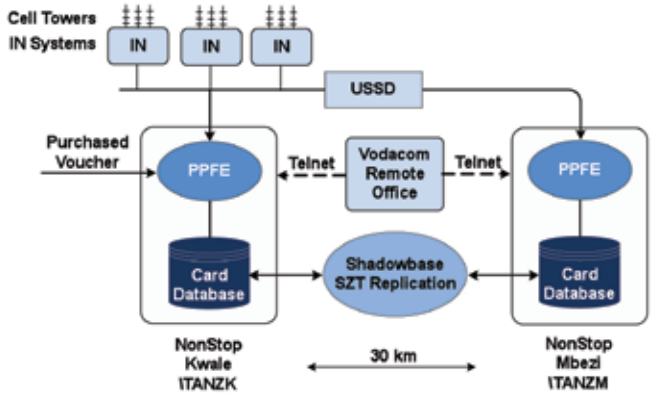

Figure 1 – PPFE Architecture

- *Purchase a voucher at a store*
  To activate this airtime, the subscriber has to redeem the voucher by sending the unique number printed on the voucher to the Vodacom PPFE system using Unstructured Supplementary Service Data (USSD).[4] The PPFE system then performs a lookup in its database to get the amount of airtime associated with this unique number from the voucher and applies it to the subscriber's IN.
- *Purchase airtime from a Bank ATM*
  This transaction is sent to the Vodacom PPFE system via

[3] The actual details of how calls are handled are much more complex than described here, which is just a simple overview for the purposes of this case study.
[4] USSD is a protocol used by cell phones to communicate with the service provider's computers.

a standard financial format message (called an ISO8583 message). The PPFE system verifies the request by using data in its database and applies the correct airtime to the subscriber's IN account.

- *Purchase airtime from a street vendor*
  Street vendors have Vodacom accounts (similar to bank accounts) that are stored on the Vodacom PPFE system. The vendor uses USSD to interact with the PPFE system and the PPFE applies the airtime to the subscribers IN account when the vendor completes the transaction.

Continuous Availability for the PPFE

If the PPFE experiences an outage, subscribers will not be able to add minutes to their accounts. Cell phone service will become unavailable to subscribers once their balance is exhausted. Therefore, the PPFE is implemented as a dual HP NonStop server pair using a bi-directional Shadowbase SZT data replication configuration, where one system is the production system and handles all transactions (see Figure 1). However, the backup system is actively running the PPFE application and is ready to take over in seconds if the production system fails. Its database is kept synchronized with the production system by the Shadowbase data replication engine from Gravic, Inc. (www.gravic.com/shadowbase).[5]

Shadowbase data replication is accomplished in sub-second times so that the backup database is synchronized with the production database. If the production system fails, then rerouting transactions to the backup system is all that needs to be done. By configuring bi-directional replication, any changes made on the promoted system queue are delivered to the original system to resynchronize its database once it is recovered. The backup system can be assessed continually with test transactions against test accounts to ensure that it is functional and working end-to-end. This assessment is quite useful since validating the backup system's application processing can be done at any time without requiring a production application outage. This testing helps to ensure the backup system will be able to take over processing without any problems if the production system fails.

On behalf of the customer, Vodacom installed one PPFE NonStop system in Oyster Bay, Kwale Street, Tanzania, and the other in Mbezi, Tanzania (Figure 2). The two systems are about 30 kilometers apart and are situated on high ground to address the primary cause of disasters in Tanzania – tsunamis. The systems are managed for the customer by Vodacom personnel located in a remote office connected to the systems via Telnet links.

The node name for the NonStop system in Kwale is \TANZK, and the node name for the NonStop system in Mbezi is \TANZM. Node \TANZK was the production node and the Shadowbase architecture replicated its database changes to its SZT backup node, \TANZM.

The IN systems at the backup site are typically down, and require a restart sequence to bring them up in the event of a failover. This means that during normal processing at Kwale, all IN systems are connected to the production node in Kwale. IN access to the backup site in Mbezi is provided by an Unstructured Supplementary Service Data (USSD) gateway.

The Production PPFE System is Downed by an Explosion

Disaster Strikes

On Friday afternoon, August 16, 2013, disaster struck. A UPS battery exploded and caught fire in the Kwale datacenter. The explosion, along with the fire-suppression system, damaged the \TANZK PPFE production node, taking it out of service. The sudden outage of \TANZK caused significant database corruption, as updates in progress could not be completed.

Even worse, the IN systems were in the same data center as the \TANZK PPFE node, and their local database was corrupted by a hard down when the failure occurred. Their communication links with the backup PPFE server in Mbezi were also damaged.

Failover to the Mbezi IN systems failed, with the Mbezi IN systems being unable to complete the DB startup sequence and allow access to the IN customer database application. Later, it was learned that this was caused by a massive number of simultaneous

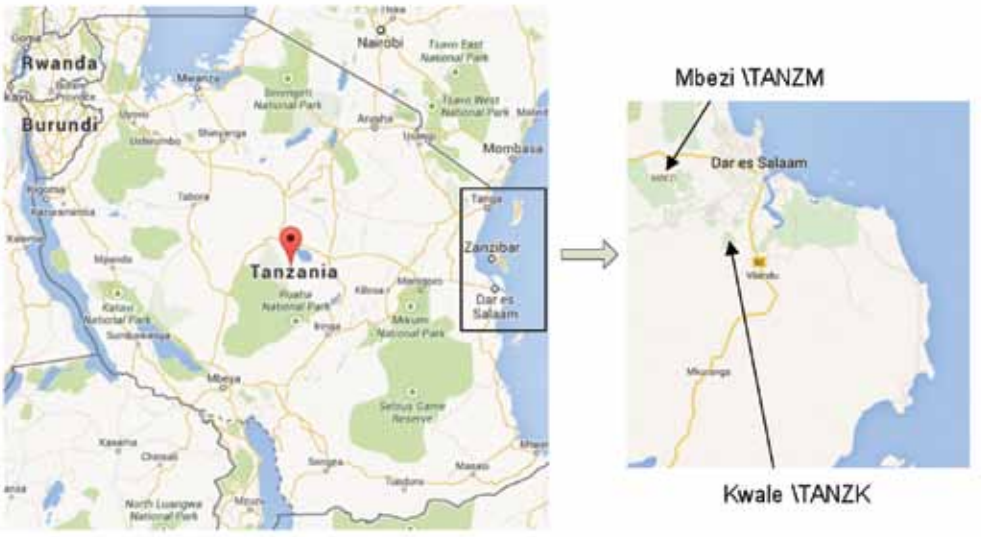

Figure 2 – Location of HP NonStop PPFE Servers

[5] For more information about Shadowbase and the Sizzling-Hot-Takeover business continuity architecture, please read the Gravic white paper: *Choosing a Business Continuity Solution to Match Your Business Availability Requirements*.

HP SHADOWBASE

NOW AVAILABLE

FOR HP INTEGRITY NONSTOP X

CALL +1.610.647.6250 • **EMAIL** Shadowbase@gravic.com
VISIT www.gravic.com/shadowbase

Business Partner

GRAVIC® Shadowbase

farm node requests to connect to the Mbezi IN database coming in and overloading the environment, causing timeouts and retries.

If the IN system failover to the backup systems installed in the Mbezi datacenter had been successful, subscriber calls to redeem prepaid vouchers could have continued with little, if any, interruption. Instead, cell phone requests to redeem prepaid vouchers for additional minutes could not be serviced by the backup PPFE system in Mbezi. Subscribers running out of minutes lost their cell phone service.

Lesson Learned: Avoid single points of failure. All network, system, application, and data components necessary to provide service must be redundant, must exist in geographically seperate locations, and need to be tested for failover periodically.

The Initial Recovery

By the next morning, the \TANZK node was once again operational from a hardware viewpoint. However, none of its software had yet been restored. There were no startup scripts, no subsystems, and no applications running. An attempt was made to start TMF, but it immediately ran into problems. Since the database had been left in an inconsistent state from the sudden node outage, TMF could not recover the database from the Audit Trail. The database had to be recovered from its backup, provided by ETI-NET's virtual tape facility, BACKBOX.[6]

By noontime Saturday, communication between the IN systems in the Kwale datacenter was established with the PPFE backup in the Mbezi datacenter via the USSD gateway system installed in Kwale. The USSD gateway ran on a Windows server that had not been damaged by the explosion. By this time, the Tanzanian prepaid voucher redemption service had been unavailable for a day.

The Long Database Recovery

The database backup strategy is that a full backup (via a TMF online dump) is done over the weekend with incremental backups (via dumping the generated audit trail files) during the week. Therefore, the backup system's corrupted database had to be restored, followed by rolling forward through nearly a full week's worth of Audit Trail change data. This recovery was accomplished via a TMF Recover Files operation for the corrupt database partitions, and it had to be completed before the database could be made available for application access. The database comprised SQL/MP tables with SQL/MX aliases. The database recovery from the online dump was completed by Saturday afternoon. However, it took an additional two weeks to complete the Audit Trail roll forward recovery.

During this time, even though the PPFE service was restored to subscribers, the production node now in Mbezi (the original backup node) had no backup, and a failure of that system would have again rendered the PPFE service unavailable. When database recovery was finally completed, the Tanzanian PPFE system was restored to an active/passive SZT configuration with production accomplished by the \TANZM node and backup services provided by the recovered \TANZK node.

[6] See http://www.etinet.com/ for more information about the BACKBOX product line.

Lesson Learned: 1) Back up more than once per week – perhaps every few hours if you are using virtual tape. Otherwise, to recover the database, you may need to restore considerable amounts of data as you roll the database forward from the last backup point, which can take a long time, extending the service outage. 2) Parallelize the backup and restore operations as much as possible (and test the recovery). 3) Consider whether a single backup system/site is sufficient, since an extended outage of the primary system/site or the backup system/site will leave you vulnerable to another prolonged outage if either system/site also fails.

The Erratic Backup System

However, even though the PPFE was now running in an active/passive SZT mode, the backup node in Kwale proved to be somewhat unstable. Vodacom made many service calls to HP. HP was very helpful, but the service calls were for a backup system and not for a production system and hence were not of the highest priority Service Level Agreement (SLA). The service contract allowed for longer service intervals for non-production system issues. However, no matter what the time of day or night, HP was very helpful and consistently went above and beyond its contractual obligations under its support agreement.

Lesson Learned: Carefully consider the support SLA's required for backup systems for mission critical applications.

Vodacom finally asked HP to visit onsite and verify that the Kwale node hardware was in good condition. HP performed exhaustive tests on the hardware and determined that the server was unusable due to smoke, water, and other fire-related damage.

Replacing the Backup System

Vodacom turned to its insurance company to obtain funding for purchasing a new backup node. The insurance company at first balked because the damaged node was still functioning, albeit erratically. However, after an extended negotiation, Vodacom convinced the insurance company to finance a new node.

Lesson Learned: If you do not have a total loss, be prepared to fight with the insurance company to obtain coverage.

The new node funded by insurance was installed in a different building than the original \TANZK node. It was given the name \TANZKW. At the same time, Vodacom replaced its IN network with the IN Advantage system from Siemens. The multiple INs in the original configuration were replaced with just two of the more powerful Siemens' INs, freeing considerable space in the new datacenter.

Bringing Up the New Backup System

As shown in Figure 3, bringing the \TANZKW node into service was the next task. The first problem was the shipping

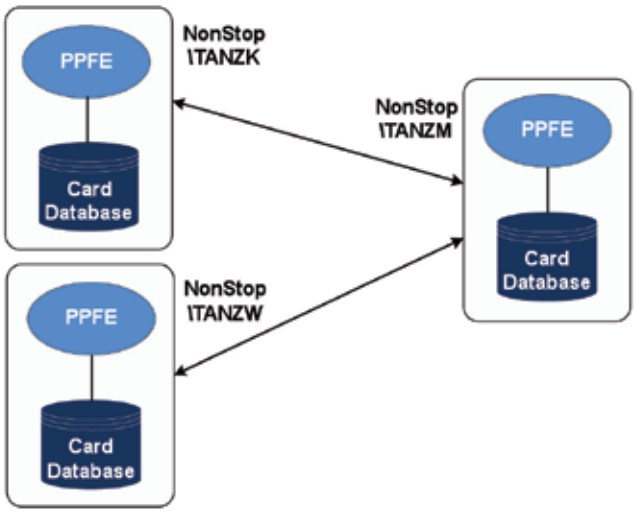

Figure 3 – Recovering \TANZKW as the New-Hot-Standby

of the replacement NonStop to the Kwale data center. When the new hardware was shipped, the cabinets had to be stripped and reassembled on site because the cabinets could not be accommodated standing up on the airplane. Every single device was unracked/de-installed from the cabinets in the factory, and the cabinets were shipped in the horizontal position. Devices – servers, switches, disk shelves – were all separately packed and shipped on pallets. All of this shipping resulted in a much longer time to get the new node hardware rebuilt and up and running.

Once the \TANZKW hardware was operational, the node's subsystems and communications were installed. Vodacom personnel were onsite for six days bringing up the \TANZKW node. At this point, further restoration procedures for the node could be accomplished from the Vodacom remote office via Telnet.

Next, the support team began the migration of the PPFE application to the new node. They started by loading the \TANZKW databases from the \TANZM production system. Several challenges slowed this procedure to a crawl. The first problem was that the software versions on the older \TANZM node and the newer \TANZKW node were different. The original \TANZM and \TANZK nodes used the J06.12 version of the NonStop operating system, and the \TANZKW node used J06.16. Furthermore, the original nodes used SQL/MX version 2.3.4, and the new node used SQL/MX version 3.2. The result was that the standard NonStop export/import utilities could not copy the database directly from the source environment to the target environment.

Lesson Learned: Understand the limitations of different hardware and software versions on systems that must interoperate, and be prepared to address the complexities.

The support team switched to NonStop's UNPAK2 and PAK2 utilities to perform block moves via FTP for the SQL/MX tables. Macros were created to generate SQL LOAD jobs to run in parallel for the large SQL/MP tables. Because the production system, \TANZM, had to be quiesced to PAK the SQL/MX tables, the PPFE application was quiesced at midnight when it was handling a minimum load and these tables were copied during overnight periods. Each morning, before the application was restarted, Shadowbase replication was initiated on the newly copied tables to keep them synchronized as the application updated them.

The Shadowbase SOLV Online Loader was employed to copy the smaller SQL/MP tables since it could replicate these tables while the application was running. It also kept these tables synchronized with the production system as the tables were being created and loaded. Thus, small-table migration proceeded continuously throughout the day.

Once the SQL/MP tables were loaded, their indices had to be built, which took an extraordinary amount of time. Because of the size of the tables and the amount of system resources available (disk and cpu), the indexing had to be undertaken one table at a time. Twelve terabytes of data had to be indexed, which took a total of 55 days. The support team worked long hours through March and into April to accomplish this task. Progress slowed considerably in April, which is vacation month in Tanzania. Progress was further gated by the lack of overtime authorization to complete the restoration task. All work had to be accomplished during normal office hours while the employees were still tasked with their regular support and systems administration jobs.

Lesson Learned: Allocate additional budget for recovery operations after a failure to allow for the additional effort. Such factors should also be considered when developing or updating a business continuity plan, to ensure sufficient resource availability when it's needed, rather than after the fact when it only further slows the recovery process.

After the database was fully loaded onto \TANZKW with all of the indices created and Shadowbase replication keeping the database synchronized with the production database on \TANZM, the support team brought up the applications. To their horror, the disks immediately disconnected. It turned out that the logical unit locations for the various volumes were different from the ones that had previously been used on the older system's XP storage area network. It took another two weeks to fix these issues.

Once the database and disk storage was correctly configured, application environment issues arose with the OSS environment that affected the PPFE application, and further challenges with SQL/MP and the SQL/MX aliases came to light. These issues were finally resolved, and testing of the new \TANZKW node began in June, 2014.

Finally, on July 10, 2014, nearly a year after the original explosion, \TANZKW was brought into service as the backup system for \TANZM, and \TANZK was shut down. During all this time, the damaged system \TANZK had continued in an erratic fashion as the backup PPFE system. It even survived another fire that erupted in its data center in December, 2013. The fire caused further damage to the \TANZK system, but the support team was able to return it to service.

Lesson Learned: It ain't over 'til it's over! In other words, even though you think the end is in sight, it often isn't...

As of this writing, the systems have now been switched so that \TANZKW is the production system, and \TANZM is the backup

system. The XP storage array used by \TANZM is currently being rebalanced to eliminate OSS hot spots that impose heavy loads on some disks in the array, which is necessary so that \TANZM will have the capacity to handle peak PPFE loads if a failover occurs.

Summary

Vodacom's PPFE is a critical system for its mobile services in Africa. The PPFE allows subscribers to top off their cell phones with purchased vouchers. Without the PPFE, cellular service will gradually grind to a halt as subscribers use up their minutes and cannot add additional minutes. To ensure continuous availability, the PPFE systems are duplexed across HP NonStop servers using a Shadowbase bi-directional SZT configuration that provides data access failover times measured in seconds.

An explosion in August, 2013, nearly destroyed one of Vodacom's Tanzanian PPFE systems. It took Vodacom nearly a year to rebuild the damaged system using new hardware. Challenges included hardware problems due to smoke damage, software version issues, lengthy database loads, and limited staff hours. During this time, the PPFE system limped along as an active/passive SZT system with the damaged node providing erratic backup to the production node. Fortunately during this period, the fault-tolerant HP NonStop production node experienced no problems that required failover to the erratic backup system.

Though prepaid voucher services to Vodacom subscribers should have been restored in seconds by a rapid SZT failover following the explosion, PPFE services were lost for a day because the IN communication channels linking the production site and the backup site were damaged. The IN environment did not have a backup system; probably the single most important lesson of this entire saga is to eliminate single points of failure.

Lesson Learned: If you want high availability, make sure that you can always failover to a backup environment quickly. Active/Passive environments can provide this level of avalability. If you want continuous availability, the "backup" environment must be up and running and in a known working state at all times. Sizzling-Hot Takeover and fully Active/Active environments can attain this level of availability. Pick your architecture to meet your business objectives.

A series of unfortunate events caused a mission-critical production environment to fail, with loss of services for over a day. The recovery of the failed environment took considerably longer than even the most pessimistic business continuity planner could have conceived. Failures happen. It is not a matter of if; it is only a matter of when. The Vodacom experience provides valuable insight into how to plan to avoid having a failure turn into a catastrophe. Consider whether any of the Lessons Learned might apply to your IT systems, and address them before disaster strikes.

The Shadowbase Data Replication Product Suite

The Shadowbase solution suite comprises several products addressing business continuity, data replication, data and application integration, zero downtime migration, and other utilities to deliver a true 24x7 "nonstop" enterprise. Shadowbase sales and support are now directly available globally from your HP NonStop account team, Business Connexion (Pty) Ltd in Africa, or contact Gravic, Inc. for more information for local resellers in your region.

Brett Dismore began his IT career in 1982 in the South African banking operations environment, moved through the RACF, TSO and JCL systems programmer ranks, and then worked as a DBA on the IBM mainframe. In 1995 he was offered the opportunity to implement and support the Tandem K Series platforms, which he did for four years, before moving to a NonStop third-party presales support career. Mr. Dismore has spent the last 13 years working as a principal engineer at Business Connexion (Pty) Ltd, providing support on the NonStop platforms for Vodacom SA, Vodacom Mozambique, Vodacom Tanzania and Vodacom DRC.

Paul J. Holenstein has direct responsibility for the Gravic, Inc. Shadowbase Products Group and is a Senior Fellow at Gravic Labs, the company's intellectual property group. He has previously held various positions in technology consulting companies, from software engineer through technical management to business development, beginning his career as a Tandem (HP NonStop) developer in 1980. His technical areas of expertise include high availability designs and architectures, data replication technologies, heterogeneous application and data integration, and communications and performance analysis. Mr. Holenstein holds many patents in the field of data replication and synchronization, writes extensively on high and continuous availability topics, and co-authored Breaking the Availability Barrier, a three-volume book series. He received his BSCE from Bucknell University, a MSCS from Villanova University, and is an HP Master Accredited Systems Engineer (MASE).

Data ain't Data – a case for mission-critical data on NonStop

Richard Buckle >> CEO >> Pyalla Technologies, LLC.

Margo Holen >> Cofounder >> Pyalla Technologies, LLC.

In a very successful advertisement campaign Castrol Oil created the punchline, "Oils ain't Oils!" In doing so it effectively separated itself from every other producer of oil and lubricant products. And as Castrol's only touch point with consumers was its line of motor oils and lubricants, it projected an image for Castrol that resonated with consumers everywhere. Today, several decades later, Castrol has stepped up its game suggesting with its latest punchline, "It's more than just oil. It's liquid engineering."

It's along these lines that industry analysts are starting to think that Data ain't Data – not all data holds the same level of importance. The significance of newly arrived data depends on the perception of the business or agency acquiring the data as well as on timing or "freshness" of the data with respect to all that is in play surrounding the data – in other words, one man's gold can be another man's dross. The more we have talked with the NonStop community, the more we understand that interrogating the right data at the right time can be every bit as mission critical as some of the processes running on the NonStop systems. Mission-critical transactions and applications, and indeed mission-critical systems, are well understood but when it comes to data there's clearly a case to be made for mission-critical data and it needs to be handled as diligently as any mission-critical transactions.

Any discussion on data tends to gravitate towards the latest developments involving Big Data and indeed, Big Data Analytics. There's just so much new technology and new products appearing almost daily in support of Big Data and often the mere reference to a technology, product or service triggers substantial debate within the IT industry, but the reality is that it's still all about the data and the value that can be derived from the data. On the Forbes magazine web site, the post of April 19, 2015, For Marketers Every Tech Trend Hinges On Big Data And Analytics, contributing blogger, Daniel Newman, makes the crucial observation that, "gadgets and trends are changing the way we connect, communicate, market products, and engage with our employees and customers. But as we talk incessantly about new technologies, we may overlook the fact that it isn't really technology that's driving this revolution, but the data and the insights these technologies create."

As for gadgets and trends, Newman quickly focuses on "IoT and wearable technology are the hottest trends in the enterprise, IT, and communications landscapes. We want smarter gadgets, smarter homes, and smarter offices — where every object 'knows' us, and exchanges that knowledge with smart sensors and devices to create an ecosystem of hyper-connectedness, or the IoT."Perhaps even more important, from the perspective of Pyalla Technologies, is a later observation that, "Social media is one of the biggest contributors to the ever-increasing mass of big data. Essentially unstructured, the data generated by millions of social media interactions grows by the minute."

When we listen to, watch, or read the news nearly every week there is a story of the authorities finding warning signs and often, alarming information, on social media that leads to crime prevention and arrests. The other side of this coin,, the infamous YouTube posts that show police harassment as well as simply people acting badly, all captured on mobile phones, help fuel the growth in data about all of us. As reported by Newman, "Social media data is a veritable goldmine of critical business information: Customer behavior, purchase patterns, customer satisfaction, and even an indicator of upcoming trends."

There we have it. The worlds of Big Data, Social Media, the Internet of Things (IoT) and Mission Critical are all connected and share a common theme. When we want or need something and it's important to us, then we expect to get it. No questions asked. No misleading information. And most important of all, a response in a timely manner! "Time may change me but I can't change time", so wrote singer David Bowie back in the 1980s as he lamented lost opportunities and events that no longer can be revisited Even as Bowie has been known to sing these words, the actual written lyrics are a little different – "but I can't trace time" which is probably a more powerful statement in that it's not just the impossibility of revisiting the past but rather, the sense of loss that comes with the passage of time – time is the enemy of relationships, our beliefs and even our processes. Could time also be detrimental to data as well?

"Not all data will hold the same significance for business even as not all data will be considered mission critical, but when a company identifies data as being mission critical, having it supported by NonStop will become a clear choice that CIOs will be reluctant to ignore."

In the post of April 14, 2015, ***Time to talk about 'Things'*** to the ATMmarketplace blog, reference is made to earlier conversations with T.C. Janes, a managing consultant for enterprise solutions and architecture at Hewlett-Packard. What came from this exchange ended up in the August 28, 2014, post to Real Time View, ***A time to put the hammer down! NonStop accelerating adoption of Clouds and Big Data.*** "From our perspective, we should be talking about 'time value' of information. Time is the enemy of data," said Janes. "Certain types of data have enormous value at its moment of creation but may have less value an hour from now and perhaps no value tomorrow." This led to the observation that even as we often emphasize two of the "V"s associated with big data - volume and velocity - timing is of the utmost importance when integrating the world of big data with transaction processing.

In a recent white paper, ***WebAction ensures IoT is never too late***, big data proponent, WebAction, highlights the ultimate risk to IoT companies of being unable to process, visualize and act upon such high-volume, high-velocity data, such that the data may simply be

"too late." Too late to make a marketing offer at an emotional inflexion point! Too late to correlate multiple events and realize there's been a security breach on a wearable device; too late to pick up on identity compromise as an ATM transaction commences. "Of course there is a big play for WebAction in the IoT world," said WebAction cofounder and EVP, Sami Akbay. "Transactional systems capture only events of interest. IoT systems capture everything and you need to filter out the unwanted noise in order to capture any event of interest. WebAction allows you to make real-time and complex decisions to identify those events of interest from IoT sources."

As we travelled to different conferences and events it has given us the opportunity to catch up with good friends and former colleagues. As was reported in the post of March 12, 2015, to Real Time View, ***My financial advisor talks IoT?*** "Our friend, who had a pacemaker installed, was called by the hospital while she was going home – they asked her to get right back (to the hospital) as her pacemaker was sending disturbing messages. WTH! She had no idea she had a talkative device planted inside her!" The IoT embraces a multitude of sensors and who could argue that the sensor keeping our heart pumping isn't mission critical or the data coming from the sensor extremely material to our well being?

Even more recently, as we enter the hospital where our just-born, premature, identical twin grandchildren are fighting to survive we observe multiple devices, sensors, and graphical displays that are attempting to make sense out of all the data available. Some of this data being so critical that it is being monitored 24 X 7, with alarms preset so the changes that are undesirable can't be missed. Yes, it's hard to watch at times, but it is a dramatic illustration too of the value of timely data making it into the right hands.

No, data ain't data! More importantly, some data is more mission critical than other data and needs to be treated quite differently. In the April 16, 2015, post to the WebAction blog, ***Data Is Mission Critical Too – the Case for Industry Standard NonStop!*** The question is asked, "Would we want to see data important to saving lives left to some of the better known commodity Big Data implementations, or would you really want to see it on NonStop?" Before ending the post, there was a critical observation made, "Not all data will hold the same significance for business even as not all data will be considered mission critical, but when a company identifies data as being mission critical, having it supported by NonStop will become a clear choice that CIOs will be reluctant to ignore."

From the time the first Tandem Computers began shipping usage of terms like online, front-end processing, and transactional processing has resonated with the communities that came to depend on the fault tolerant properties of what eventually evolved into the NonStop systems we have today. Along the way, it was easy to separate mission critical transactions from less important transactions. Looking up a name and address is one thing, but to change the address that subsequently changes the destination for a purchased product, that's something entirely different. Within the NonStop community this is well understood business fundamental and we don't even need to be reminded of the NonStop users who have separated the "look" from the "book" transactions. However, as we head to a data-driven world where changes in behavior, the emergence of patterns, not to mention social media "trending" data too has become every bit as mission critical as the transactions themselves.

In recent presentations by HP Master Technologist, Justin Simonds, it would be hard to miss the focus on Big Data, the Internet of Things (IoT), and in particular, his recent presentations featuring the Internet of Mission-Critical Things (IoMCT). A highlight of the post of March 17, 2015, to the blog Real Time View, Can you believe it? "Pragmatism trumps spectacle"! Simonds said, (when it comes to NonStop systems, everyone involved knows of the) "very clear mission – to build a computer that won't fail. All our development is in sync with the mission and every developer hardware and software has to answer the question 'what happens when this fails?' Not if, but when." Yes, all those years ago, Simonds notes, "Tandem was the first to have no single point of failure. That has been copied. Tandem was the first to have online repair. That has been copied. But NonStop after 40 years is still rated by IDC as an Availability Level 4 (AL4), a full level higher than any clustered system."

Ultimately, NonStop is the platform for data that we can't afford to loose or to see its importance diminish over time. Returning to the reference made to our friend's interaction with her hospital following surgery to implant a heart pacemaker, as addressed in the post ***My financial advisor talks IoT?*** question was asked, "who is going to process, store and act upon the information all the 'things' are providing. Most of the applications that my financial advisor mentioned, and most that I've encountered so far, are critical in nature, and it is unthinkable that the data could get lost, action delayed, or even worse, be compromised. Almost every hospital went ahead and added a power generator so as not to be caught by power outages in the middle of the medical procedure; have they all got a live backup system for the medical data they are collecting?" Again, data simply ain't data and for the data you absolutely rely upon at the most critical of times, it's simply got to be on NonStop!

Castrol may have started with the message, "oils ain't oils" but their introduction of the new message, "It's more than just oil. It's liquid engineering" should resonate strongly with the NonStop community. When it comes to mission-critical data and the mission-critical applications that depend upon it, we know that "data ain't data" but perhaps it's not a stretch to now consider, "It's more than just data. It's timely mission-critical information."

Richard Buckle is the founder and CEO of Pyalla Technologies, LLC. He has enjoyed a long association with the IT industry as a user, vendor, and more recently, as an industry commentator. Richard has over 25 years of research experience with HP's NonStop platform, including eight years working at Tandem Computers, followed by just as many years at InSession Inc. and ACI Worldwide.

Well known to the user communities of HP and IBM, Richard served as a Director of ITUG (2000- 2006), as its Chairman (2004-2005), and as the Director of Marketing of the IBM user group, SHARE, (2007-2008). Richard provides industry commentary and opinions through his community blog and you can follow him at www.itug-connection.blogspot.com, as well as through his industry association and vendor blogs, web publications and eNewsletters.

The quotes come from some of Richard's clients including HP, Integrated Research, comForte, DataExpress, WebAction, Inc., InfraSoft, and OmniPayments, Inc.

Margo Holen is a co-founder of Pyalla Technologies, LLC and serves on the Board of Directors of InfraSoft, Pty. Previously Holen was Vice President of Engineering and Technical Support at XYPRO Technology, a Senior Vice President at ACI Worldwide and VP Engineering at InSession Technologies. Holen also spent 14 years at Tandem Computers, the last years as Vice President of Operations and Chief Operating Officer of the NonStop Software Business Unit.

Holen served the NonStop user community first as the ITUG Vice Chairman and then as the first Vice President of CONNECT, the HP user community combining ITUG, Encompass, and HP-Interex EMEA. Holen was first elected to the ITUG Board in 2006 and served as Finance Director before becoming Vice Chairman of the Board in 2007.

Got 30 minutes? Then start MODERNIZING.

SOAPam® can be up and running in less than 30 minutes and is so intuitive that you'll be exchanging data in no time.

SOAPam Increases Efficiency Throughout Your Business

- Access information on your NonStop, securely, from other platforms
- Receive orders placed by customers and business partners, directly on your NonStop
- Access external information from your Pathway applications
- Prepare for a zero-downtime migration

"We purchased SOAP/AM® Server and SOAP/AM Web Service Client... They have met our expectations and then some."

"SOAP/AM is easy to install, easy to use, and it is a pleasure. Night and day between it and NonStop SOAP."

WHY SOAPam?

- ✓ *GREAT VALUE*
- ✓ *QUICK IMPLEMENTATION*
- ✓ *INTUITIVE*
- ✓ *SUPERIOR TECHNICAL SUPPORT*
- ✓ *RUNS IN GUARDIAN*

Download a trial and get started now:
http://hub.am/1ngyjOJ

Scan to learn more about SOAPam.

NuWave TECHNOLOGIES

HP Integrity NonStop X: Welcome the NS7

Jim Smullen >> HP Distinguished Technologist **Mark Pollans** >> WW Sr. Product Manager >> HP

Introduction

On 30 March, 2015, Hewlett-Packard announced the availability of the HP Integrity *NonStop X NS7 X1* (NS7), the first system in the new *NonStop X* product line based upon the Intel® x86 architecture. Bringing the NonStop environment to x86 represents a major milestone in the NonStop platform's evolution. The NS7 product announcement makes our NonStop on x86 strategy a reality as we redefine continuous availability and scalability for x86. Introducing *NonStop X* products based on the Intel Xeon® processors alongside the current HP Integrity *NonStop i*[1] systems based on the Intel Itanium® architecture, provides customers with the flexibility and choice of an unparalleled portfolio of HP Integrity NonStop fault-tolerant systems for high-value business workloads and customer-facing applications—each with the same NonStop fundamentals.

HP is delivering the software for the NS7 through a new L-series software product line. The first release delivers the NS7 with all the NonStop fundamentals for fault tolerance, availability, and scalability that our NonStop customers have come to expect and depend upon from HP NonStop.

Inside the rack

The NS7 is based upon the HP BL460c server blade using Intel Xeon® E5-2600 v2 series microprocessors contained in a standard HP c7000 enclosure; the infrastructure is similar to that used on the *NonStop i* systems, but with a different configuration to support half height blades. As is the case with other high-end NonStop systems, the NS7 supports up to 16 NonStop CPUs per node. With the smaller blades we can fit all 16 NonStop CPUs in a single c7000 enclosure. Available main memory sizes range from 64GB to 192GB per NonStop CPU for a per system maximum greater than 3TB!

The NS7 leverages the CLuster I/O Module (CLIM) architecture from the Itanium based NonStop systems, increasing the number of CLIMs supported per system up to a maximum of 56. The IP CLIM Ethernet ports, in combination with the InfiniBand based system interconnect, provide support for 10GbE networking. The Storage CLIM supports both Solid State Drives (SSDs) and Hard Disc Drives (HDDs) for internal storage and Fibre Channel connectivity for SAN attached (e.g. HP XP7 Storage and HP XP P9500 Storage) and tape storage.

A system interconnect based upon industry standard InfiniBand (IB) 4X FDR running at 56Gbps is central to the NS7 system. This provides tremendous interconnect capacity among NonStop CPUs and direct attached I/O while also considerably reducing the system interconnect latency.

The initial release of the NS7 is licensed for 4-core software

Figure 1: System Diagram

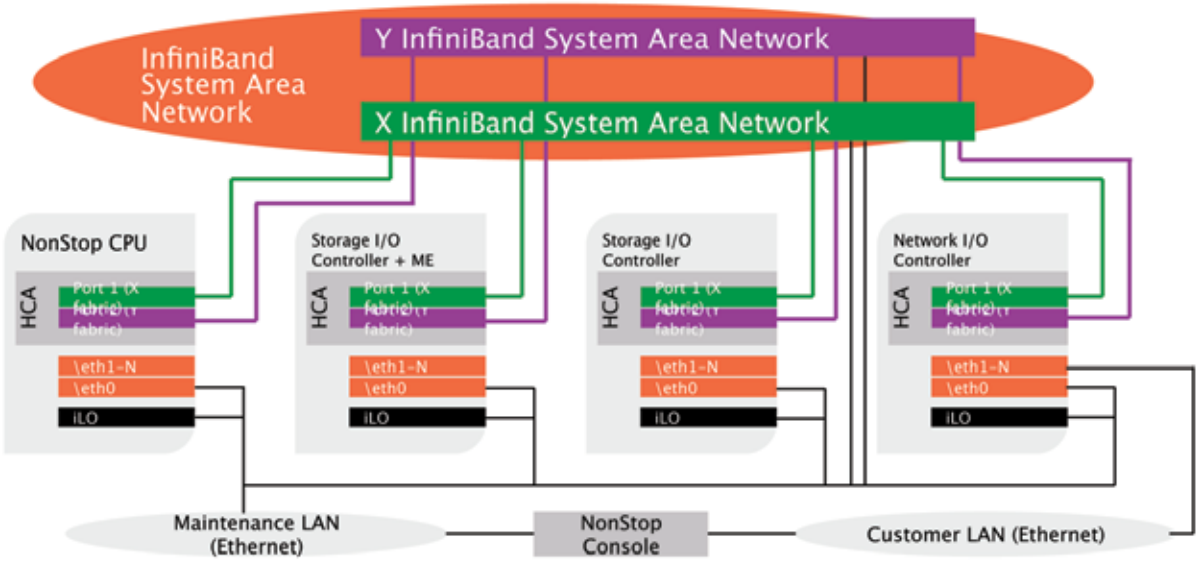

[1] The Itanium based systems are now the HP Integrity NonStop i family of systems.

and runs 4-cores per NonStop CPU. As additional Core Licensing options could become available, the same hardware should support other options.

The Software

The Release Version of NonStop Software that runs on the *NonStop X* family is the L-series, with the first Release Version Update (RVU) being L15.02[2]. The functionality of this RVU is very similar to that of the J-series. The NonStop OS has been updated to take advantage of the very high bandwidth and low latency of the InfiniBand system interconnect.

The system runs exclusively in the x86-64 mode of operation, which means the NonStop system always runs in 64-bit mode, even when executing 32-bit applications.

The system architecture for *NonStop X* is called TNS/X (the Itanium based system architecture is called TNS/E).[3]

How is the *NonStop X* the same as *NonStop i*?

In many fundamental ways, the *NonStop X* is the same as the *NonStop i*. A logical processor (NonStop CPU) is still defined as a set of cores sharing the same memory and system interconnect interfaces. The processor blades (NonStop CPUs) are contained in a c7000 enclosure. The I/O subsystems are connected via a system interconnect using CLIMs. The 32-bit NonStop environment is implemented by sign extending 32-bit addresses into 64-bit addresses (as done on the *NonStop i*). The system is still big-endian (see section "Under the Covers"). TNS (CISC) programs are fully supported on the platform and a new accelerator is available for maximum performance

How is the *NonStop X* different from *NonStop i*?

Since the processor blades are half-height blades, a 16-processor system can be contained in a single c7000 enclosure whereas Itanium blades are full-height and thus require two c7000 enclosures for a 16 processor system.

Native programs must be recompiled to run on the *NonStop X* and the development tools (compilers and linkers) have the same look and feel as the ones on J-series. Virtually all non-priv programs should require no source level changes. Enterprise Toolkit (ETK) and Visual Inspect (VI) are functionally replaced by NSDEE (an Eclipse-based development environment) on the L-series. NSDEE is also available on the J-series.

The NS7 performs 50% better for an Order Entry benchmark[4] compared to a 4-core NB56000c (apples-to-apples comparison) and for CPU bound tasks, performance is even better. Interprocessor message latency is also significantly reduced. As always, your specific application performance will vary depending on the workload characteristics.

The Maintenance Entity (the software that configures the fabric) is hosted on a pair of storage CLIMs (the ones that host $SYSTEM) for a *NonStop X* (IB based) system. For *NonStop i* (ServerNet based) systems, it is hosted on the ServerNet switches.

The NS7 system supports a new diagnostic feature, where in the unlikely event of a halt on both CPUs 0 and 1, their memory is dumped to a CLIM (where it is stored encrypted) so that full diagnostic data can be obtained. On *NonStop i* systems, only one of the CPUs' memory could be captured, so it is possible that adequate diagnostic information might not be available.

Supported Features

The most current versions of the NonStop software stack are supported by *NonStop X* systems. This includes:

- SQL/MP
- SQL/MX 3.3 (which also will be available on J06.19)
- TS/MP 2.5
- JAVA™[5] version 7 in both 32b and 64b flavors
- LTO6 tape drives (which can read and write LTO5 tapes and read LTO4 tapes)
- NSDEE (AKA Eclipse Development Environment) 5.0, this is an enhanced version of the product which is already available on J-series.
- VTC (Virtual Tape Controller) for virtual tape applications
- CLIM based I/O

Under the Covers

Big-endian

The x86 Instruction Set Architecture (ISA) is a little-endian architecture, so how does *NonStop X* support a big-endian environment? The *NonStop X* system's compilers store all program variables (i.e. the named objects in a program, such as A from the c source "int A") as if they were stored on a big-endian machine. So if you viewed that variable in memory using its address and did a HEX display, it would be big-endian.

There are many non-source data items in a process, for example a procedure return address that was pushed onto the stack. That data is maintained in little-endian format. The system debuggers display them properly because those tools and the NonStop OS understand what is big and what is little-endian. However, displaying the memory that stores the execution stack reveals a mix of data types and a mix of endian status, so non-symbolic debugging is more difficult than on a *NonStop i* system.

The major benefit of all this design is that a binary message sent from a system to a *NonStop i* system is big-endian and requires NO transformation on the receiving end, nor does a message sent from a *NonStop i* system to a *NonStop X* system. This also means that data files (including database files and tables) can be copied directly from a *NonStop i* system to a *NonStop X* system without need for translation.

TOSVERSION

The L-series RVUs naming (e.g. L15.02) is now not directly related to the TOSVERSION of the system. The TOSVERSION for all L-series RVU is "L06" (or "V06" in its unconverted format), just like "J06" was the TOSVERSION for all J-series RVUs.

Improved Scale-up

To better exploit IB, there is now a message system interrupt

[2] The L-series software starts a new naming convention for NonStop RVUs: Lyy.mm where yy is the last two digits of the RVU's release year and mm is the month number for the RVU, so L15.02 is February 2015.

[3] The acronym TNS which did stand for Tandem NonStop, is now simply defined as TNS.

[4] Order Entry is a variation on TPC-C where the drivers are on the system under test instead of external to it.

[5] Java is a registered trademark of Oracle and/or its affiliates.

process per core within each CPU. This allows improved and smoother scale-up as the number of cores per CPU increases.

Coming Events for *NonStop X*

The next L-series release is expected to provide additional functionality to the NS7; in particular, clustering over IB which will allow *NonStop X* systems to cluster with each other for very low latency operation.

We also expect the release to include Core Licensing functionality. Core Licensing permits the system to run as 2, 4, or 6-core NonStop CPUs. With this feature, an NS7 system (licensed for 4-cores) that was purchased prior to the introduction of the Core Licensing feature could be upgraded without hardware changes to a 6-core SW licensed system. NS7 Core Licensing allows upgrades from 2 to 4, 2 to 6, or 4 to 6-cores as an online operation.

Summary

The introduction of the HP Integrity NonStop X NS7 X1 product brings to reality HP's multi-year strategy of extending the NonStop environment to the x86 architecture. The NS7 is an exciting new product with impressive specs as the first member of the new *NonStop X* product family.

Jim is an Operating System Architect with responsibility for new platforms. He has been with HP/Compaq/Tandem since 1987. Jim was a lead architect for the move from MIPS to Itanium resulting in the development of the NS16000. He was also a lead architect for the multiprocessor NB5x000c NonStop systems, and for the x86 based NonStop X (NS7) system.
Jim has Masters and Bachelor degrees from the University of Texas at Austin and the University of Florida respectively. Jim lives in northern California.

Mark is HP's Worldwide Senior Product Manager responsible for the HP NonStop systems portfolio, including the NonStop X (Xeon® based) and the NonStop i (Itanium® based) systems, disk storage and new solid state drive technologies.
Most recently Mark introduced the NonStop NS7 system. Prior to this, he orchestrated the release of the second and third generations of the HP Integrity NonStop BladeSystem and four generations of HP Integrity NonStop entry-class systems. Earlier, he introduced the NS16200 along with various NonStop platforms and storage solutions.
Mark has several years of experience at HP, largely in enterprise computing and networking. During his tenure with HP, he has held various management and engineering positions in R&D and marketing for hardware and software projects.

How Will NonStop Fit Into the Internet of Things?

Justin Simonds >> Master Technologist >> Americas

Dean Malone >> NonStop Architect >> Caleb Enterprises Ltd.

PART I – Foundational Premises

There has been quite a bit of excitement about the Internet of Things (IoT) AKA Machine to Machine (M2M), Industrial Internet of Things (IIoT), Web of Things (WoT) and I'm sure there will be more identifiers that will emerge. It is the idea of machines or sensors providing information of interest to people, processes and other technology. This will usher in the 3rd wave of the World Wide Web.

The 1st wave involved the linking of information – documents to documents and search capabilities. This could be described as human to information. It was based on the ability to do quick research. Web 2.0 added many more human-to-human interactions with such web-based programs as YouTube, Skype, Twitter, Facebook, LinkedIn and the list keeps growing; the underlying element is person-to-person or person-to-group interaction.

This new and growing market under Web 3.0 uses the Internet as a large World Wide Service bus for passing machine information to interested end-points; be that other machines, humans or services/applications. How much information are we speaking about? IDC is forecasting the Digital Universe to be 44 Zetabytes by 2020 (Tera, Peta, Exa, Zeta – that's 10 followed by 21 zeros times 44). Consultants seem to be agreeing that 40% of that Universe will be M2M in 2020. Yikes! Even if they're off by a large factor, yikes!

Historically, the capabilities we are talking about for Web 3.0 have been monitored on a smaller scale by SCADA (Supervisory Control and Data Acquisition) systems that fundamentally monitor two types of telemetry points – control and analog. Control points are basically binary – on/off, open/closed, yes/no, etc. Analog points typically have integer or floating-point values and set-points against them such as low-critical, low-warning, normal, high-warning and high-critical. Tied to those telemetry points are rules that react to them. Gartner has been talking of the integration of IT with OT (Operational Technology) like the SCADA systems just discussed. Businesses that have internal systems such as these - and most do - might consider integrating them as a first step toward IoT integration. With IoT we're talking about making expanded telemetry data available on the web and aggregating sets of them against particular entities such as:

- A member of a hierarchy such as a department, within a group within an organization.
- A person
- A geographical address
- A GPS coordinate

HP has its own particular vision of global SCADA. Please visit http://www8.hp.com/us/en/hp-information/environment/cense.html#.U_uOB_ldWSo for a brief overview of HP Labs CeNSE project (Central Nervous System for the Earth). The vision is to distribute nanoscale sensors around the globe to collect all sorts of information. To take an excerpt from the site:

> "By providing real-time information on the physical environment, the networks are intended to improve the way governments, businesses, and society respond to and manage environmental, biological, and physical/structural changes. Examples of potential CeNSE uses include roads, buildings, bridges, and other infrastructure; machines such as those used in airplanes and manufacturing plants; and organizations that work on health and safety issues, such as the contamination of food and water, disease control, and patient monitoring."

These are all inspiring ways of using the new technology and we are already seeing the rollout; one example being connected cars. OnStar generates monthly reports on covered cars: engine overview, transmission, oil life and even the tire pressure. This data is collected, transmitted and prepared as a report for its customers and this is just the beginning. Google already has a self-driving car with over 300,000 miles on it. Cars will communicate with each other; maintaining safe distances and drastically reducing - if not eliminating – accidents but the beginnings are already here.

Most people are aware of smart meters; the idea of providing feedback between water and power use within a home and trying to distribute requirements within a community. Let's not all wash our clothes at the same time, for example. Smart metering causes a great increase in volume and velocity. Back in the day (not so long ago) when a meter reader used to come around every month, a meter reading would generate 12 data points per year. To make things easy, let's assume an energy company has a million customers. That's 12 million readings a year to collect, generate bills for, report on and

derive statistics and trends against. With smart meters installed, the company can now get a reading every 10 minutes. That's 6 per hour, 24 hours a day and 365 days a year; that is 52,560 readings per year. This is a three-orders-of-magnitude increase from the old 12 readings per year! With the same million customers the company now has 52.6 billion readings per year. That's a lot more volume, but it's also quite a bit more velocity because it is expected to be delivered in near-real time. What if the utility company decides it wants to read meters every minute? Every second?

Here are a few examples of how the information can be used:

- For each time interval, the consumption will be compared to the average mean of consumption to determine trend usage by customer, sub-station, grid sector and network.
- Customers who are consuming inordinate amounts of power that are outside the mean can be investigated for opportunities to conserve or move their consumption to a time when the network is under less load – particularly during peak power usage.
- Incentives can be provided to those particular customers who may have inefficient energy consumption appliances.
- Customers can be billed a higher rate for peak usage if it exceeds the norm by a particular percentage.
- Unusual occurrences can be flagged for exceptions. For example, a building's power going to ground or the sensor going offline might be indicative of the house being on fire so the fire department could be dispatched to investigate.
- Not all of this information needs to be stored in a database; just the data that is outside the normal range. Since the vast majority of the data will fall into the norm, this can radically reduce the amount of secondary storage needed to hold the trending data.

What makes this so exciting from a NonStop perspective is that most of this generated machine information will be in a very structured block of information. Sensors will not be generating snide comments, will not employ sarcasm, and won't use emoticons or any of the many things that makes textual, unstructured human information so hard to process for machines. It will be reliable. It will be structured. The length will be fixed. It will come in fast - possibly randomly and in bursts - and it will in many instances look very much like OLTP transactions, for which the NonStop system was designed.

At a recent talk Dr. Michael Stonebraker - a professor from MIT and a co-founder of Vertica, VoltDB and Postgres among others - said that 75% of all data stream applications will be OLTP-like. The rest will need Complex Event Processing (CEP) models. The IoT will create many if not most of the new data stream applications Dr. Stonebraker

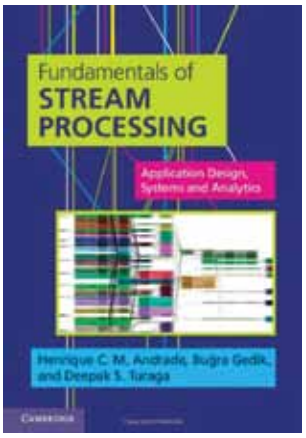

was describing. The NonStop operating system has always favored an OLTP/message-switch type processing load so this burgeoning market hits NonStop in its sweet spot.

In the book "*Fundamentals of Stream Processing*" by Andrade, Gedik and Turaga, the first chapter contains a section titled "Towards Continuous Data Processing: The Requirements." We will be paraphrasing but encourage anyone interested to invest in the book. The chapter describes the challenges of ingesting, processing and analyzing information as it is continuously produced. The authors make the point that processing must keep up with ingestion rates; the velocity portion of Big Data or as we might state, be able to scale. Further on, they discuss parallel and distributed systems; clearly showing that a parallel architecture is ultimately the only way to achieve the scalability required.

Problems (or ingestion, processing and analysis) must be broken into sub-problems that can then be worked on and solved simultaneously. They must be broken into critical regions. The authors also state that the continuous nature of stream processing requires fault tolerance; although they also point out that different segments of stream processing applications may require different levels of reliability. As they phrased it, "For instance, tolerance to sporadic data loss is, in many cases, acceptable for certain parts of an application, as long as the amount of error can be bounded and the accuracy of the results can be properly assessed. Other application segments however cannot tolerate any failures as they may contain a critical persistent state that must survive even catastrophic failures." (ibid p. 6). The requirements for high velocity, critical segments of the stream processing environment are for a massively parallel, highly scalable, fault tolerant system. Sound like any system you know? From a requirements standpoint, NonStop starts looking like a pretty good contender.

The new processing requirements of IoT are largely concerned with velocity and volume. There will be a lot of data coming in very quickly. To use an overused analogy, it will be a fire hose – a data deluge. What architectural enhancements might allow NonStop to thrive in this new environment? InfiniBand is part of that new thinking. In Part II of this series we will discuss some of the technical underpinning which will enable NonStop to uniquely target the new IoT market.

Justin is a Master Technologist for the Americans Enterprise Solutions and Architecture group (ESA), a member of the HP IT Transformation SWAT team, and a member of the Mainframe Modernization SWAT team. His focus is on real-time, event-driven architectures, business intelligence for major accounts and business development. Most recently he has been involved with modernization efforts, Data Center management and a real-time hub/Data Warehouse system for advanced customer analytics. He is currently involved with HP Labs on several pilot projects. He is currently working on cloud initiatives and integration architectures for improving the reliability of cloud offerings. He has written articles and whitepapers for internal publication on adaptive enterprise, TCO/ROI, availability, business intelligence, and the Converged Infrastructure. He is a featured speaker at HP's Technology Forum and at HP's Executive Briefing Center. Justin joined HP in 1982 and has been in the IT industry over 34 years.

Dean is one of the pioneers of Message Oriented Middleware (MOM), having chaired three panels on MOM in '93, '94 and '95 at COMDEX. He developed the world's first fault-tolerant shared memory (XIPC on NonStop in 1995) deployed that product as the first customer implementation of active NonStop process pair (four programs implemented) and also ported Seer HPS/NetEssential 4GL-middleware to the NonStop. His biggest middleware achievement was the porting of IBM MQ-Series to NonStop as Chief Architect in 1998. He was the infrastructure architect for the Province of Ontario responsible for implementing the world's first wireless WAN-based mobile workstations for OPP, regional police, carrier enforcement and ambulance services. His customers include banks, brokerages, retail, EFT/POS switches, funds wire, vendor products, airlines, reservation systems, industrial automation and more. He has built systems on NonStop, VMS, Stratus, Unix and PDP-11 and has played roles as architect, technical lead and hands-on technical problem solver as a consultant for over 30 years. He is presently completing an RDMA Middleware product that will implement distributed shared memory, semaphores and queue-based messaging between NonStop, Linux and Windows servers over InfiniBand.

Your journey into the Cloud, begins with a single step. Now you can take that step from your home or office.

Now available in a self paced version, this Introduction to OpenStack® course assists administrators and users in operation and architecture of the various OpenStack® components, called Projects. This 8 hour Web Based architectural overview ensures understanding of various OpenStack® projects and their functions.

If you are a developer new to OpenStack®, including engineers who will develop code using OpenStack® APIs and those who will submit code to OpenStack®, then this course is the perfect first step. Architects, Solution Designers, Technical Consultants and equivalent technical roles may also benefit from this course.

After completing this course, you will be able to describe the purpose and features, list the major components and describe the steps to sign up for OpenStack® development. You will learn HP's involvement with OpenStack® and be able to identify the interfaces for engaging OpenStack® services, describe the architecture and provide an overview of OpenStack® code development and testing methods.

Why HP

- 35+ years of experience
- lead the industry in IT training, certification training, and education consulting services
- focused on end user acceptance during a technology change
- award winning IT training covers all elements of HP cloud computing and converged infrastructure from services to storage to networking and security.

Let HP and HP Helion Education Services help you with your cloud implementation.

Learn more at
hp.com/learn/cloud

Available Here

2015 - The Year of the Leap Second

Dr. Bill Highleyman >> Managing Editor >> Availability Digest

There is a reason that no space launches will take place on June 30 or July 1, 2015. Scientists do not want to risk a computer malfunction due to a leap second being added at the stroke of midnight.

What is a Leap Second?

Our day is one rotation of the Earth about its axis. We call this a *solar day*. We break the day up into 24 hours, with 60 minutes per hour and 60 seconds per minute. Therefore, 86,400 seconds comprise a day (24 x 60 x 60). Each second is timed very accurately with an atomic clock, which counts 9,192,631,770 periods of radiation emitted by a cesium-133 atom to time one second (see Wikipedia – Atomic Clock). This timing is used to determine Coordinated Universal Time (UTC).

However, the rate of the Earth's rotation is slowing down. The solar day is getting longer. To account for this, one second is added on occasion to the UTC time to synchronize it with the solar day. This is the leap second.

At the current rate of deceleration, the Earth's rotation is slowing about 2 milliseconds per day or about one minute every 60 to 90 years. In the past 40 years or so, 25 leap seconds have been added to our clocks to keep them synchronized with the solar day.

The major cause of the Earth's deceleration is believed to be the result of tidal action. The rotation also is guided by a fundamental law of physics – the conservation of angular momentum. Just like a spinning skater who brings in her arms to spin faster, the Earth will rotate faster if its mass becomes more concentrated toward its center. Likewise, it will spin slower if its mass becomes more distributed away from its center. Earthquakes tend to push tectonic plates upwards or downwards, thus slowing or speeding the Earth's rotation. However, melting ice caps move water off the ice caps into the oceans and lead to a more rapid rotation.

Leap seconds were introduced in 1972, but they do not occur with regular frequency and cannot be predicted because of irregularities in the Earth's rotation. Nine leap seconds were added in the eight years from 1972 to 1979. No leap seconds were added in the seven years from 1999 to 2005. The last leap second was added on June 30, 2012. The next leap second will be added at midnight on June 30, 2015.

Leap seconds are generally announced six months in advance to give everyone time to prepare. They always occur at midnight UTC time (Greenwich Mean Time) on the last day of a month. Since 1972, these months have been June or December. Also since 1972, there have been no subtraction of leap seconds. Only additions. The Earth's rotation seems to be slowing down continuously.

The Form of the Time Change

Not all systems handle the time change in the same way. For a positive leap second, some clocks will add an additional second and will count as follows:

23:59:59 23:59:60 00:00:00

Others will repeat the last second (Unix does this):

23:59:59 23:59:59 00:00:00

Still others will freeze the last second (NTP, the Network Time Protocol, does this):

23:59:59 ..... 00:00:00

If a negative leap second is required, the last second is deleted:

23:59:58 00:00:00

The Last Leap Second Caused Major IT Problems

From 1999 to 2005, there were no leap seconds. Cloud services and multiprocessors came into existence during this time. Many facilities were implemented by software programmers who didn't even know that leap seconds existed. The programmers were incapable of allowing for the fact that the addition of a leap second makes the time appear to go backward.

The leap second is problematic in many cases in computing:

- The exact time between two UTC past dates cannot be accurately computed without taking into account leap seconds.
- The exact time between two UTC dates more than six months in the future cannot be accurately computed because it is unknown how many leap seconds, either positive or negative, will be inserted during this time period.
- Compute actions driven by a timer may get confused if they are triggered twice in two seconds, or they may not be triggered at all if a negative leap second occurs.

The 2012 leap second caused a myriad of problems, mainly with systems locking up and needing rebooting. Most sites running Linux had problems, as did those running Solaris. Web sites suffering Linux problems included LinkedIn, Reddit, Mozilla, Yelp, Gawker, and StumbleDown.

Java also had its challenges. Reddit suffered an outage with its database, Apache Cassandra. Mozilla had a problem with Hadoop, as did Gawker with its Tomcat server. Cassandra, Hadoop, and Tomcat are all built with Java.

Older Motorola GPS receivers had a software bug that caused the receiver to revert to the previous day if a leap second had not occurred in 256 weeks. This bug hit on November 28, 2003. November 28th was repeated.

NTP includes a leap-second flag in its packets. The flag informs servers that a leap second is imminent. However, researchers have reported that not all of the NTP servers correctly set their flags, causing them to provide the wrong time for up to a day after the leap second is added. Security firms have suggested that hackers have been able to exploit this vulnerability.

Safeguard Your Finances, Credit and Good Name

LifeLock membership. Easy. Secure. Guaranteed.

SPECIAL OFFER:

10% OFF** for the lifetime of your membership

Enrollment is easy:

- Visit LifeLock.com or call 1-800-LIFELOCK (543–3562)
- Use promotion code: **CONNECT**
- Complete enrollment by providing your required information

connect
Hewlett Packard's Enterprise Technology Community

LifeLock®
Relentlessly Protecting Your Identity

** Offer is for new LifeLock members only.

The Status of Linux

The 2012 Bug

Perhaps the issue of most concern is whether Linux is now prepared for this year's leap second. Linus Torvalds, the creator of Linux, commiserates that:

"Almost every time we have a leap second, we find something. It's really annoying because it's a classic case of code that is basically never run and thus not tested by users under their normal conditions."

The 2012 problem was traced to a bug in the Linux subsystem, hrtimer. Hrtimer is used when the system is "sleeping," waiting for the operating system (OS) to complete some other task. In some cases, hrtimer sets an alarm for the sleeping applications. The alarm will go off if the OS is taking too much time with its other work. When the leap second hit, the hrtimers were suddenly a second ahead of the OS. They started generating alarms, waking up countless sleeping applications at once and overloading the machines' CPUs.

The State of Readiness for Red Hat Enterprise Linux

Red Hat has posted information on its Customer Portal concerning the readiness of its Red Hat Enterprise Linux (RHEL) operating system (https://access.redhat.com/articles/15145). It states in a January 10, 2015, blog posting that "Red Hat is aware of the upcoming June 30, 2015 leap second and is working to address it."

Red Hat offers the following suggestions. If a system is running NTP on any platform but an IA-64 architecture, RHEL should be upgraded to Version 7. Version 6 may work, but there is a chance that the leap second could lead to high CPU usage.

If a system is not running NTP or is running on an IA-64 architecture, the Linux operating system will not correct for the leap second. The system clock should be reset after the leap second occurs.

How Will NonStop Servers Do?

The following description of how NonStop servers will handle the leap second has been provided by HP Distinguished Technologist Wendy Bartlett:

> "Assuming the NonStop system is in sync with UTC via TIMESYNC or some other NTP client program, when the leap second is inserted into UTC, the system time of the NonStop systems will be 1 second fast when compared to the actual UTC time. When TIMESYNC next performs a sync operation with its NTP server, it will detect it is 1 second fast and will call the operating system to make a correction. Assuming the client is configured to apply a gradual correction, the NSK system will slowly bring the system time into sync with the NTP Server. The 1 second retardation will take about 42 minutes to be applied on the H, J and L series. It takes substantially longer on the S series.
>
> The operating system is unaware of the existence of a leap second. The system reckons date and time by counting microseconds since an epoch, with each day 86,400 seconds long. For example, if a program calls INTERPRETIMESTAMP(JULIANTIMESTAMP(0),dt) once during each second and displays the first six elements of dt, 2015/7/1 0:0:0 follows 2015/6/30 23:59:59.
>
> Conversion of a Gregorian date and time to a timestamp does not accommodate a leap second. For example, the COMPUTETIMESTAMP procedure rejects a value exceeding 59 for the number of seconds. If you pass an array containing { 2015, 6, 30, 23, 59, 60, 0, 0 } to COMPUTETIMESTAMP, the resulting timestamp value is –1, and if the errormask parameter is passed, its value is set to 0x0400: bit <5> (as numbered in TAL) is on, corresponding to element [5] of the date-n-time parameter."

Workarounds

It appears that Linux and Java have been patched to correct for the upcoming leap second. However, systems should be carefully monitored as midnight of June 30th approaches, just in case.

In 2011, Google devised a unique way to handle leap seconds, a method that could be adopted by others. On the day on which the leap second is to be added, Google periodically adds one millisecond to its clocks. This means that it adds one millisecond every 86 seconds or so, there being 86,400 seconds in the day. Thus, when leap-second time arrives, Google's clocks are already in synchronism. This strategy worked fine for Google for the 2012 leap second. Google dubs this method *leap smear*.

Summary

The leap second is such an erratic and infrequent occurrence that it is likely that many systems have not been built to account for it. Those that have may not have been thoroughly tested for the condition. This is the rational for everyone to monitor their systems carefully as the leap second approaches at midnight on June 30th.

With the Earth slowing down, do we have to worry about its rotation stopping? Probably not. At its current deceleration, it will take about 2.6 billion years to stop, if it stops at all.

Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him a popular seminar speaker and a sought-after technical writer. Dr. Highleyman is a past chairman of ITUG, the former HP NonStop Users' Group, the holder of numerous U.S. patents, the author of Performance Analysis of Transaction Processing Systems, and the co-author of the three-volume series, Breaking the Availability Barrier.

Back for More...

Richard Buckle >> CEO >> Pyalla Technologies, LLC.

I have to believe news of the HP Integrity NonStop X NS7 X1 will be liberally sprinkled through many of the articles and columns that make up this issue of The Connection. It's been a long time, given the short attention spans of today's IT leaders, since the news of the plans by HP NonStop to support the Intel x86 architecture were announced in late 2013. However, now that there's a new, complementary family of NonStop systems taking shape, with the first model reaching general availability stage, it's probably a good time to take stock and start thinking about the likely impact the HP Integrity NonStop X NS7 X1 will have on the community and indeed, the industry.

Not surprisingly, my own expectations are rather high – this being reflected in the many blog posts that I have written following the announcement. Starting with the very first post, that being to the NonStop community blog, Real Time View, Here comes NonStop X and here's to another decade or two, or four, of NonStop excitement! I opened with comments from HP executives on the start of this journey. Perhaps the most important quote of all coming from Martin Fink, Executive Vice President and Chief Technology Officer, HP. "It's been a long journey to where we see NonStop running entirely on industry-standard hardware," Martin Fink said. "Thanks to the significant investment HP has made in support of NonStop, customers will realize immediate value from the much improved price / performance that NonStop X now provides."

The word exciting can be overused at times, but the connection between NonStop and before that Tandem with exciting dates back to the unquenchable and enthusiastic support for NonStop that originated with the company's founder. And it isn't showing any signs of letting up. Look back at Randy Meyer, VP & GM, HP Mission Critical Systems, column in the Nov / Dec, 2014, issue of The Connection.

In that issue Randy said, "Our strategy is clear - HP is committed to providing choice, value and investment protection to our mission-critical customers. To that end, we are offering the 100% fault-tolerant solutions our customers rely upon on x86."

It's a long journey; customers will realize immediate value; redefining the availability and scalability of x86 application; and yes, offering value and investment protection to our mission-critical customers. Readers will string together their own summary from the above quotes by Martin Fink and Randy Meyer and yet, no matter how you cut and paste these statements, it really is all about choice. And for that, those in the NonStop community with significant investments in current HP Integrity NonStop systems – System i as the Itanium based systems have been called – will find support, and even potentially new models coming from HP. What the NonStop will see with the introduction of NonStop X is the "opening up a world of possibilities".

This came from the newest data sheet on NonStop X where you will also read that, "Mission-critical customers, like you, can continue to rely upon HP Integrity NonStop, deployed on either architecture, to deliver a resilient business foundation without compromise." The theme of this issue of The Connection is Mission Critical Computing and this segment of the market is growing rapidly, not the least because of the stepped-up awareness of the Internet of Things (IoT) and the volume (and velocity) of data spewing forth from any number of devices, be they out at sea, on our roads and railways or up in the air. And this segment is making a legitimate claim for the consideration of the Internet of Mission Critical Things (IoMCT) about which we will be hearing a lot more indeed. I would be surprised to find this topic not covered in this issue.

In her post of March 12, 2015, to the NonStop community blog, Real Time View, My financial advisor talks IoT? former ITUG and Connect board member, Margo Holen, wrote of how all she can think about is, "who is going to process, store and act upon the information all the things are providing. Most of the applications that my financial advisor mentioned, and most that I've encountered so far, are critical in nature, and it is unthinkable that the data could get lost, action delayed, or even worse, be compromised …Yes, it's all on the Internet!" Margo then observed that, "Almost every hospital went ahead and added a power generator so as not to be caught out by power outages in the middle of the medical procedure; have they all got a live backup system for the medical data they are collecting?"

Prevailing wisdom is that the recipients of the data coming from the IoMCT will wind up on systems based on the Intel x86 architecture and so, bestowing on x86 systems the power, and indeed excitement, of a fault-tolerant system like NonStop will ensure advocates across the industry will be quickly found. The challenge will continue to be gaining traction with modern solutions vendors but this is one aspect of the launch of the NonStop X family of system I believe will prove more likely than we have witnessed previously with the launch of any other NonStop family of systems – coverage of the arrival of Node.js and with it, support for Server Side JavaScript (SSJS) for me will be a key contributor to generating such an interest among the vendors delivering modern solutions.

Exciting can be used to excess and even today there are those within the NonStop community who cringe each time they hear that word. However, few words provoke the kind of pent-up expectation as much as the word excitement does – this time, is NonStop about to cut lose and carve out new market segments? Choice; mission-critical; IoMCT; and not to be ignored, fault tolerant – all reverberate well within the IT community at large. After all, who wouldn't want a system as robust and resilient as a NonStop system if only it ran on x86 and if only it supported InfiniBand (IB) and yes, if only it's TCO made it eligible for more serious consideration! Well, it's happened, it's arrived and yes, it's exciting!

Data Backup, Replication and DR Solution for the New HP NonStop X and All Other Server Platforms

STORAGE DIRECTOR®

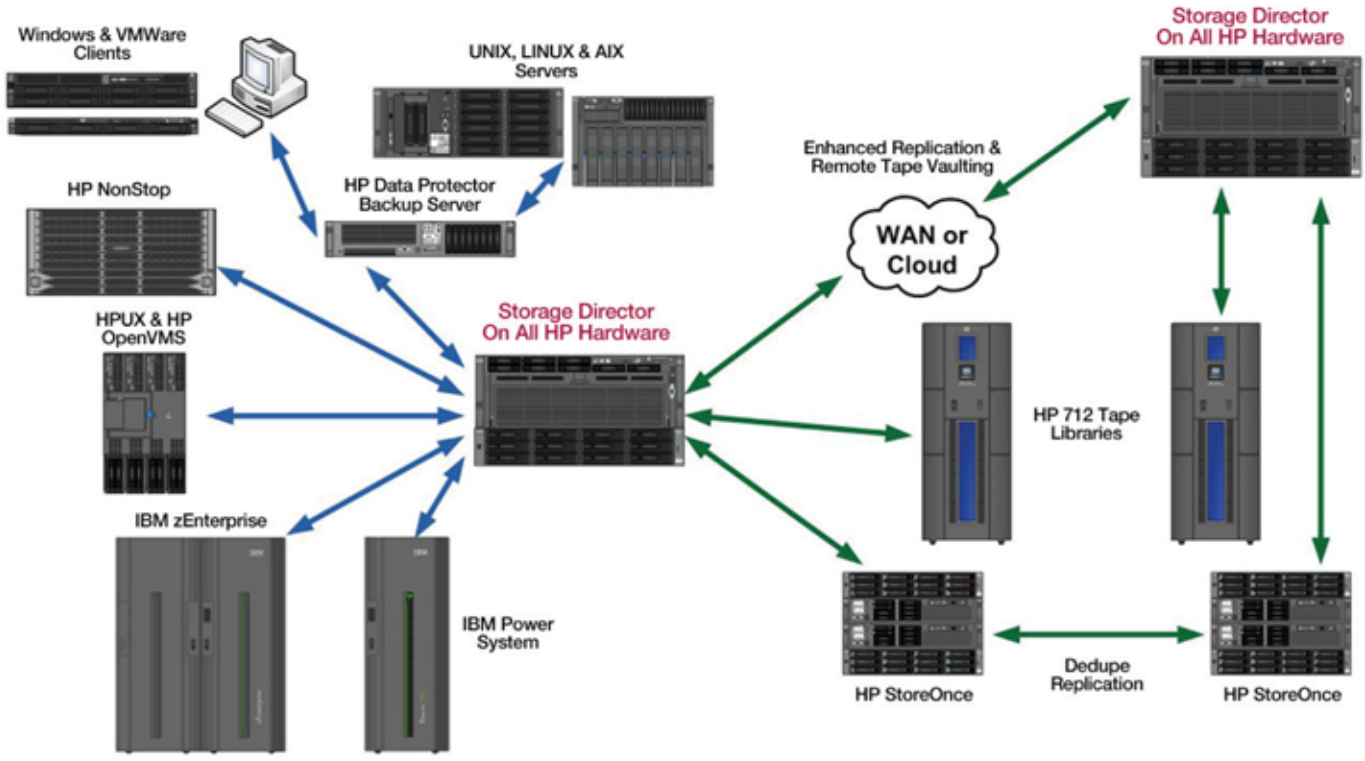

Storage Director Implementation in an Enterprise Environment

Storage Director is a policy based solution with Any-to-Any connectivity (any host platform with any backup application to any storage technology or media).

Storage Director is unique in supporting de-duplication AND tape in the same environment.

Tributary SYSTEMS, INC
www.tributary.com

Contact us today for a no-obligation assessment.
info@tributary.com
+1 (817) 354-8009

The guiding light for your mission critical business
Improve your NonStop'ness. Better always on!

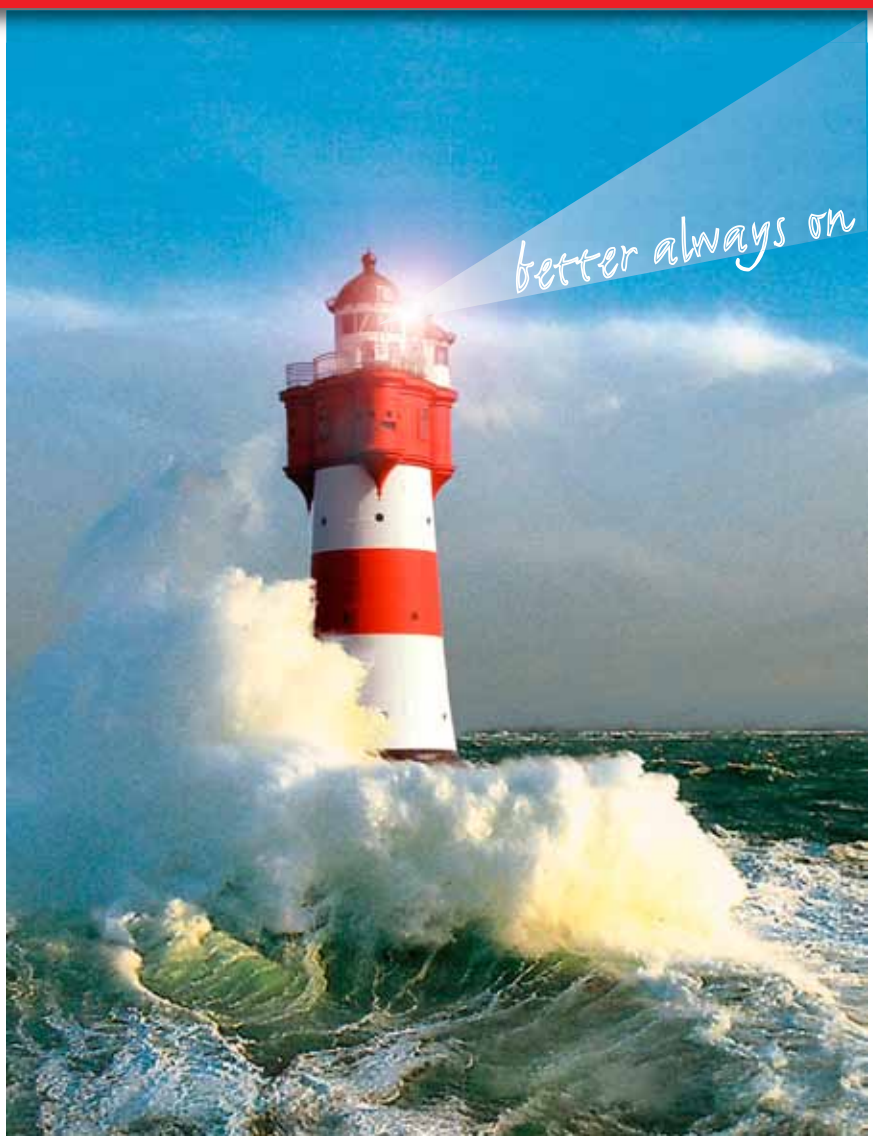

www.comforte.com/better_always_on

Today's demands of mission critical businesses and customers are ever increasing. Unreliable and unavailable systems and applications are not an option. Minimizing downtime whilst maximizing security and operational efficiency is therefore paramount for the IT department. If your light is going out, your business and your customers can get in trouble. Systems and applications can't stop; they must be on, always!

comForte „better always on" solutions help you gain …

Better Infrastructure
Make the most of best in class communications and connectivity solutions by providing end users and system administrators with high performance, secure and reliable access to NonStop systems.

Better Security
Protect your mission critical data-in transit and at-rest. Improve your overall security posture on NonStop and achieve compliance with industry standards and regulations.

Better Applications
Modernize your legacy applications from the database layer, through better integration in the enterprise all the way to refreshing the application's Graphical User Interface.

Better always on with comForte's unparalleled solutions for HP NonStop.

comforte®

better always on