connect
Your Independent HPE Business Technology Community

Read it in Print and Online!

The Connection

A Journal for the Hewlett Packard Enterprise Business Technology Community

NonStop X Opens Doors to Integration

Hardware vs. Software Data Replication for Business Continuity

Fast Track Your NonStop Integration with XYGATE and Active Directory

Sept – Oct 2016 { Volume 37, No. 5 }

XYGATE® SecurityOne™

Security Intelligence and Analytics for HPE Integrity NonStop™ Servers

Visibility

Faster Threat Detection

Improved Risk Management

Differentiate Noise from Actionable Incidents

Minimize the Impact of a Breach by Identifying it in its Early Stages

Reduce Mean Time To Detection

Learn more at xypro.com/SecurityOne

XYPRO®
Mission Critical Security

Get In-Depth Monitoring of NonStop Systems.

At Some of The Industry's Most Intelligent, No-Sweat Prices.

ALTERNATIVE THINKING ABOUT SYSTEMS MONITORING:

Alternative thinking isn't just about constantly monitoring your applications and systems.

It's about alerting you to issues and even fixing them – automatically – based on your specific business goals.

HP's NonStop Availability, Stats and Performances product (known worldwide as HP ASAP) provides automatic monitoring of your system and application objects, keeping you aware of the status of your processing environment at all times.

With HP ASAP, you can set goals that alert you if problems occur. You can also link goals to corrective actions such as initiating repairs or restarting failed components – automatically.

While many monitoring solutions track a few dozen system attributes, HP ASAP lets you monitor thousands simultaneously. What's more, ASAP consumes very few CPU cycles.

Now, you'll know instantly when CPUs are too busy, when disks switch paths, when applications are not performing well, or when there are a myriad of other problems that NonStop ASAP can isolate and resolve.

HP ASAP. For 10 years, the smart, affordable way to keep NonStop servers running non-stop.

Technology for better business outcomes.

HP NONSTOP ASAP

- Monitor applications, system & network resources
- Set user-defined goals
- Receive alerts via clients, events, email & mobile
- Automate recovery actions
- Monitor and resolve issues quickly and easily

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

Hewlett Packard
Enterprise

Discover 2016

London 29 November – 1 December

Explore the intersection of digital transformation, technology and ideas

hpe.com/discover

Table *of* Contents

Columns...

A Note from Connect Leadership

Integration...

I find it odd that it is now 2016 and we are finally getting around to talking about "integration" between platforms. Not that I am arguing, I think this is a huge step in the right direction. For the past 60 years of computing, systems have competed with each other to be different!

Being a NonStop supporter has really shielded me from what others consider integration. Trying to define what the term means in the technology sense has been a little troublesome for me. I think of integration as working well with others. Interfacing and interoperating. I have always thought of the NonStop as being the ultimate integrator. Consider the following:

NonStop has supported just about every communications protocol ever used in the enterprise and has been positioned as a gateway between systems since the platforms inception.

NonStop supports mainstream programming languages: PASCAL, FORTRAN, COBOL, C, even Python.

NonStop runs on commodity hardware: Intel CPU's, JBOD, RAM, standard chassy, etc.

NonStop even has a POSIX compliant interface, just like Unix/Linux.

Now things are getting even MORE integrated as new NonStop servers include features that allow us to integrate even more seamlessly than ever before:

TCP/IP v4 and v6
SSL/TLS and ssh
SIEM integration
ANSI compliant database services
JAVA
JDBC
LDAP authentication
etc.

What more could anyone ask for? I don't see any of the OTHER platforms trying to do what we do, or do I?

The more I think about integration, the more I realize that NonStop isn't just adding integration features, it has been integrated all along. NonStop is at the very core of every enterprise for the same reasons it always has been (scalability, availability, reliability) as well as something no one every mentions: it has always integrated with everything else.

Rob Lesan

Thanks.
Rob Lesan
XYPRO Technology
Connect Worldwide President

2016 Connect Board of Directors

PRESIDENT
Rob Lesan
XYPRO

VICE PRESIDENT
Michael Scroggins
Washington St. Community College

PAST PRESIDENT
Henk Pomper
Plusine ICT

DIRECTOR
Jamil Rashdi
wasl Asset Management Group

DIRECTOR
Trevor Jackson
SOCAN

CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide

HPE LIAISON
Janice Zdankus
Enterprise Group, HPE

TheConnection

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.
Joseph Garza ..... Art Director
Janice Reeder-Highleyman ..... Editor at Large
Dr. Bill Highleyman ..... Technical Review Board
Karen Copeland
Thomas Burg
Bill Honaker
Justin Simonds

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

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.

News from HPE's NonStop Enterprise Division

With autumn comes thoughts of the NonStop Technical Boot Camp

Summer is nearly gone! 2016 has been a busy year for the HPE NonStop team and here we are stepping into the fall of 2016. Fall brings visions of red, gold and bronze leaves on sidewalks, brisk walks through cooler air, the scowling faces of pumpkins, bales of hay and of course the annual NonStop reunion at the Connect NonStop Technical Boot Camp in the first half of November.

We are busily planning the details of this year's event located in San Jose at the Fairmont Hotel, November 13th – 16th. There will be three full days and 80+ breakout session slots, which will give us a chance to offer more repeats of popular sessions this year. There's an ever-growing list of breakout and vendor theater sessions posted on the Connect website at https://www.eiseverywhere.com/ehome/nonstoptbc2016/411952/.

In keeping with the changes coming in the landscape and the weather, Silicon Valley companies including HPE continue to announce changes. The Boot Camp event is your opportunity to hear directly from HPE leaders about the importance of the NonStop platform to HPE's new strategy. While you're there: talk directly with the engineers working on NonStop products, attend the IT Manager's Network and converse in a confidential environment with other customers about how they use the NonStop platform, catch up with NonStop partners over a cup of coffee and in general learn more about what's new and available now as well as what's coming next.

Looking beyond the Technical Boot Camp we are also planning our participation in the HPE DISCOVER London event scheduled for early December. As we demonstrated at DISCOVER Las Vegas earlier this year, we are planning to offer customers a glimpse of the virtual NonStop product in action. We hope you'll join us there as well.

This issue of Connection is themed "integration", which echoes HPE's hybrid infrastructure goals around helping customers make the best use of technology available today and build greater efficiencies into their data centers. The NonStop team continues to investigate emerging technologies that may have synergy with the platform. An example of this can be found in the article by Justin Simonds about Blockchain support and integration, which is being pondered [contemplated?] inside HPE. Richard Buckle's article on NonStop X's ability to integrate into the hybrid infrastructure within data centers also points to new capabilities available with the new platform.

NonStop customers continue to benefit from our rich Partner community and Partner companies continue to break new ground offering more open approaches to application development, including security aspects; an interesting article demonstrates that through XYPRO's Steve Tcherchian's look at integration with Microsoft's Active Directory. In this issue, HPE's Atalla and Data Security teams describe how they have integrated HSM with ESKM to make a new and stronger offering for customers. With PCI DSS 3.2 requirements looming, Andrew Price and Wendy Bartlett also remind you that we have already integrated the XYGATE User Authentication product into the systems we sell, which includes Multi-factor authentication capabilities.

We continue to look for opportunities to drive NonStop into new markets (such as the Internet of Things) and deliver our unique value proposition to more customers and applications worldwide. We look forward to seeing our customers and friends from around the world in San Jose in November and DISCOVER London in December!

Karen

Karen Copeland
Manager, WW NonStop Product Management
Mission Critical Solutions
Hewlett Packard Enterprise

Three Servers and Counting

OmniPayments Invests in NonStop X

Janice Reeder-Highleyman >> Principal >> Reeders & Writers

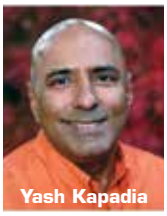
Yash Kapadia

Yash Kapadia loves, loves, loves the NonStop X. As the CEO of OmniPayments LLC, Yash and his company became the first NonStop partner to take possession of a NonStop X server. That was late 2014. Two more NonStop X servers later, Yash's projected business growth indicates a near-future need of additional NonStop X systems to enhance OmniPayments provisions for high-speed, multi-platform solutions of its standalone and cloud-based Financial Transaction Switch.

Yash is thinking six more NonStop Xs. What will he do with all that speed and performance?

Introducing OmniPayments LLC

OmniPayments LLC (www.omnipayments.com) is the product arm of Opsol Integrators Inc., a leading HPE NonStop system integrator. Opsol was founded by Yash in 1995, shortly after he left his role as a senior developer and project manager at Tandem Computers. Opsol is short for "Open Solutions," and its team of 100+ NonStop developers specialize in porting open-source solutions to NonStop servers running under OSS (Open System Services).

In 2008, ACI Worldwide announced the sunset of its BASE24 financial-transaction switch on NonStop servers. Ending as well was ACI's support for existing NonStop BASE24 applications. In the ensuing panic among BASE24 customers, Yash saw an opportunity to create a BASE24 replacement. Opsol already had a head start in that its existing business modules - OmniATM, OmniMessaging, OmniHub, and OmniCrypto - were installed in numerous locations worldwide. The four products, all flexible and customizable, formed the basis for what is now the OmniPayments Financial Transaction Switch. Additional modules were created to complete the SOA-compatible, layered design.

So large was the potential market that Yash established OmniPayments as a separate company focusing on payment transactions. Today, OmniPayments systems process 700 million transactions per month, generated by point-of-sales terminals and over 14,000 ATMs. A single OmniPayments system supports up to 10,000 transactions per second. Multiple OmniPayments systems can cooperate to provide any capacity required by an application. From seven worldwide locations, OmniPayments serves as a 24×7 managed services provider for remote production monitoring.

Migration Case Study

In 2009, the first BASE24 to OmniPayments migration was made by a large U.S. bank. It already was an Opsol customer through its use of OmniATM, which supported the bank's extensive ATM network. The bank also was a BASE24 customer. OmniATM is feature-rich and provided a level of tailored functionality that was not available from BASE24. It interfaced seamlessly not only with the BASE24 authorization system on NonStop for not-on-us transactions but also with the bank's IBM authorization system for on-us transactions.

When the BASE24 sunset was announced, the bank turned to Opsol's new OmniPayments switch as a possible alternative. OmniATM is a major component of OmniPayments, and its existing incorporation within the bank's debit/credit card authorization system eased what already would have been a smooth transition from BASE24 to OmniPayments. With no disruption for customers, OmniPayments went active within the bank's deadline and communicated effortlessly with ten interchange networks, including Visa, MasterCard, PULSE, and STAR.

The OmniPayments Financial Transaction Switch – A Modular, Flexible Approach to Payment Solutions

OmniPayments offers all the requisite functionality to manage credit/debit-card transactions. It manages multiple devices, hosts application interfaces, and interoperates with third-party products or other systems if required. OmniPayments easily expands to provide additional functionality when needed and supplies complete security functions for every financial transaction handled, including encryption-at-rest and encryption-in-flight. Available around the clock, OmniPayments will survive any single fault, requires no downtime for maintenance or upgrades, and supports a range of disaster-recovery products.

OmniPayments can be installed either as a complete, off-the-shelf solution or can be enhanced with customer-requested modifications and features to improve cost-effectiveness, efficiency, and risk mitigation.

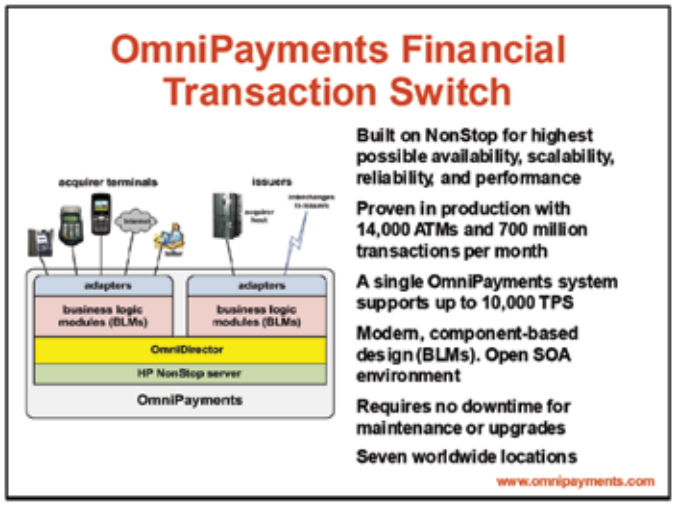

OmniPayments is a component-based software design that permits the creation of flexible business services, the kind that users can develop and deploy quickly to enhance the customer experience. The core of OmniPayments is a set of Business Logic Modules, or BLMs. As a whole, they comprise the OMNIs, whose open architecture (SOA) allows OmniPayments' building blocks and business modules to interact not only with each other but also with existing in-house systems and third-party solutions. While all components are included within OmniPayments, some can be purchased separately.

At the heart of OmniPayments is its Payment Engine, which supports ATMs, POS terminals, transaction authorization, stand-in authorization, settlement, and card management, among many other functions. The Payment Engine runs on NonStop servers, as does OmniDirector, a rules-based routing engine, and the OmniAuth Preauthorization Engine.

Other less critical OmniPayments modules run on Red Hat Enterprise Linux (x86 RHEL). The OmniOffender system monitor and the OmniDash business monitor offer 24/7 real-time views of all OmniPayments functions. Another module, the Customer Experience Hub, integrates with the Payment Engine and is implemented via a "party model" database. This furnishes the power to fuse many products and customer services, including loyalty modules, customized bill payment modules, etc.

The NonStop X Advantage

In the past, OmniPayments NonStop applications and the x86 RHEL applications ran on different blade systems. NonStop servers communicated over ServerNet, Linux applications communicated over Infiniband, and no intercommunication existed between the platforms. By placing the OmniPayments Financial Transaction Switch on NonStop X, OmniPayments can incorporate a NonStop X Payment Engine and the RHEL modules within the same cabinet. This significantly increases OmniPayments' transaction capacity and speeds up the Customer Experience Hub.

OmniPayments Optimizes Its Software Stack to Reduce TCO

No Pathway License	OmniPayments built its own NonStop checkpointed application monitor.
No Replication Software License	OmniPayments created its own replication software.
No External Monitoring Software	OmniPayments provides OmniOffender, its own monitoring software.
Discounted SQL/MX License	
Thin Client Model	

The OmniPayments 50%-Less Guarantee: How Do We Do That?

OmniPayments sells its popular financial transaction switch at a price that guarantees 50% off a company's current transaction costs. How does OmniPayments do that? Unlike other switch providers, who base costs on transaction volume, OmniPayments sells a one-time, perpetual software license. That pricing model is what makes the "50%-Less" difference. No volume fees. No transaction fees. No surprises. No penalties for company growth. All prices are presented up front.

The OmniPayments Financial Transaction Switch Lives in OmniCloudX

It's affordable, scalable, continuously available, and it's pay-as-you-go. OmniCloudX on NonStop X hosts numerous instances of the OmniPayments financial transaction switch at a price so attractive that mid-size retailers and financial institutions can enjoy the benefits of having their own high-capacity transaction switches, a luxury until now limited to their larger counterparts.

According to Yash Kapadia, OmniPayments' first private cloud was built in Northern California. It continues to serve as an active/active backup for several customers of the company's standalone switch and also is the host platform for ITUGLIB, Connect's library of user-contributed freeware and other software utilities. OmniPayments provides at no cost ITUGLIB's processing capacity, maintenance, power, and bandwidth.

In 2015, the company introduced OmniCloudX as a SaaS cloud platform implemented on NonStop X. It provides affordable payments services to a population of retail and financial businesses who cannot budget infrastructure investments that include a NonStop.

Each OmniCloudX customer pays only for the amount of CPU resources, storage, and networking that it uses. Backup systems are provided so that a system outage will be recoverable immediately via automatic failover to other geographically dispersed NonStop X servers. OmniPayments supplies the IT staff needed to manage the transaction switches running in its cloud. Also supplied are complete security functions for every transaction handled, including encryption-at-rest and encryption-in-flight.

Easily expandable to provide additional functionality when needed, OmniCloudX, like standalone OmniPayments, will survive any single fault, requires no downtime for maintenance or upgrades, and supports a range of disaster-recovery solutions.

OmniPayments Preauthorization Engine – We Call It the Fraud Blocker

The OmniPayments Preauthorization Engine is used by financial institutions in conjunction with the OmniPayments Financial Transaction Switch or as a seamless interface to other providers' switches via a custom support module (CSM). Customers call the Engine the *fraud blocker.* Modern and easy to manage, it preauthorizes millions of transactions far more effectively than its complex, compute-intensive competitors.

The Preauthorization Engine is one of the OMNIs – *OmniAuth*. It is flexible, SOA-based, sits on NonStop X, and interfaces easily with existing in-house systems and third-party solutions. Via complex rule sets, it identifies potentially deceptive transactions in real time and rejects them without having to send the transactions to issuing banks for authorization. The banks only see the validated transactions. In doing so, the "fraud blocker" mitigates the immense processing load demanded by the banks' own systems to perform preauthorization checks. It also minimizes the response times for transaction approval/rejection.

Migrate to OmniPayments

Pain-Free Transition, No Disruption to Customer Services

OmniPayments is an attractive BASE24 replacement. Its modular design permits gradual implementation for smooth migration to OmniPayments from existing payments infrastructures. A typical migration to OmniPayments averages four months because our team of migration specialists are based in time zones around the world. We work 24 hours a day. In addition to BASE24, our staff are experienced in migration from other transaction switches on NonStop. At a customer's request, we can enhance OmniPayments to address specific requirements.

- **Comprehensive payments solution for banks and retailers**
- **Built on NonStop for highest possible availability, scalability, reliability, and performance**
- **Proven in production with 14,000 ATMs and 700 million transactions per month**
- **A single OmniPayments system supports up to 10,000 transactions per second (TPS)**
- **Modern, component-based design (BLMs). Open SOA environment**
- **OmniPayments costs 50% less than our competitors because we sell you a one-time perpetual software license. No transaction or volume fees.**

Affordable OmniCloudX

- NonStop X hosts numerous instances of OmniPayments
- Allows mid-size OmniPayments customers to operate their own high-capacity transaction switches
- Offered on a pay-for-use basis, starts at $5,000/month
- Continuously available with automatic failover to other Non-Stop X systems
- Complete security functions for encryption-at-rest and encryption-in-flight. PCI-compliant

OmniPayments Fraud Blocker

- Modern and easy to manage
- Preauthorizes millions of transactions in real time and far more effectively than complex, compute-intensive competitors
- Sold as part of OmniPayments or as a seamless interface to other providers' solutions

OmniPayments is now an authorized reseller in Latin America of HPE NonStop servers and HPE Atalla security products. Contact us for more information.

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

Fraud Blocker Case Study

One of Latin America's largest suppliers of electronic transactions counts on the OmniPayments Preauthorization Engine, which seamlessly interfaces to the EPS (Electronic Payment Systems) provider's existing financial-transaction switch via an Opsol-created custom support module. The switch routes all financial transactions to OmniPayments for preauthorization prior to submitting the transactions to the issuing banks for final approval. For this EPS provider, that amounts to almost 200 million transactions per month.

New in 2016 - OmniPayments Goes Mobile

It was only a matter of time before OmniPayments customers in retail and finance asked Yash to develop mobile banking applications for use with their payment systems. Out of their requests came OmniPayments Mobile Banking, which brings real-time, anywhere banking and financial transaction convenience via a variety of digital devices.

Mobile Banking

The mobile banking application is intuitive, user-friendly, highly secure, fraud-resistant, and is architected for the Android operating system. Built on OmniPayments' HPE NonStop X servers, the app is backed by an automatic failover function to guarantee continuous availability. It readily can be customized to address any customer's unique requirements.

OmniPayments mobile banking seamlessly integrates with host systems and connects directly to Web Services via JSON or SOAP/XML. It provides instant updates, real-time reporting, includes a web-based UI Console/Dashboard for report analyses, and manages both debit/credit cards and loyalty/gift cards. Among numerous services, users can pay bills, transfer funds, recharge prepaid mobile devices, access digital wallets, deposit checks, make in-store purchases, and monitor balances.

Loyalty Card Management System

Loyalty programs are a marketing strategy designed to encourage repeat customer business via frequent flyer programs, rewards plans, club programs, points incentives, prescription plans, and gift cards. The OmniPayments Loyalty Card Management System provides all the necessary features to create, manage, and operate a variety of loyalty programs.

The Loyalty Card Management System is highly secure and fraud-resistant. It supports transactions made via POS terminals, cash registers, web portals, and mobile devices. Intuitive and user-friendly, the system connects directly to an OmniPayments NonStop X server and allows customers to check card statuses, confirm balances, view detailed transactions, and benefit from a host of other features.

Retailers can review, manage and maintain their own loyalty programs via the OmniPayments Online Web Console, which also provides access to detailed, real-time analytical reports.

Loyalty Card Case Study

Casa Ley is a longtime OmniPayments and NonStop customer. It is a large grocery store chain with hundreds of stores in over forty cities. Recently, the company decided to expand its customer incentive and retention efforts via new loyalty and reward programs. After a careful review of its options, the company selected the OmniPayments Loyalty Card Management System and Android Loyalty application, both directly connected to NonStop X via the store's OmniPayments Financial Transaction Switch.

OmniPayments is Now an Authorized NonStop and Atalla Reseller in Latin America

The OmniPayments presence in Latin America has been particularly successful, with a large installed base in several countries, including Colombia, Mexico, and the Dominican Republic. In areas where HPE does not offer 24x7 product support, Opsol Integrators provides managed services for remote production monitoring. Becoming an authorized reseller of HPE NonStop servers and HPE Atalla security products is a natural extension of the Latin America partnership that OmniPayments already shares with HPE. Although OmniPayments sells HPE solutions, Yash Kapadia has chosen to do so at $0 profit for his company. When Yash says "savings," he means it.

We're Here, We're There, We're Everywhere

Opsol Integrators and OmniPayments maintain a global presence. Locations include company headquarters in California; development facilities in India; offices in Houston, Mexico, and Colombia; and new representation in Europe. Our staff members are based in three time zones. As such, we work 24 hours a day and are known for rapid project turnarounds and meeting deadlines. It's also why migration from one switch to OmniPayments typically averages only four months.

The OmniPayments Crew

As for Yash, he keeps busy identifying new geographic regions for upcoming OmniCloudX purchases. Already, OmniPayments has seven data centers; but Yash sees a market for more. "The closer the data centers are to customers," Yash says, "the greater the mitigation of communication costs." Yet another savings that OmniPayments passes on to its customers.

We told you that Yash loved the NonStop X.

For further information about all OmniPayments solutions, contact us at +1 408-364-9915. www.omnipayments.com. We are Booth #48 in the Partner Pavilion at the upcoming NonStop Technical Boot Camp.

Janice Reeder-Highleyman is a contract writer and communications specialist. She has navigated the NonStop world for thirty years in a variety of volunteer capacities, including ITUG chairman. In her spare time, she is a flight instructor, which isn't nearly as demanding as is keeping abreast of the fast-paced changes in the technology industry. Contact Janice at jreederhi@gmail.com.

Connect 2016 Annual Survey Results Let Your Voice Be Heard

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

At Discover 2016 in Las Vegas, Connect launched an annual member survey. The survey was promoted in our digital and print publications (Connect Now, Connect Converge, and The Connection), via email to our member database, as well as via social media (Facebook, Twitter, LinkedIn). 1,508 people responded to the 2016 survey. We hope that you were one of those who responded to this survey, and that you will do so in the future. This is your one chance to let your voice be heard.

Survey Taker Demographics – 1,508 responses

Almost half of the survey takers were from North America (48%), and a large portion represented the EMEA region (37%). The largest group of survey takers (333) reported that their primary job function is "Project Manager/Programmer/Analyst/Engineer," which implies that the majority of survey takers are technologists. Directors/Managers (282) and System Administrators (251) were the second and third largest group of survey takers. Almost half of the survey takers reported themselves as "average" adopters of technology (47%).

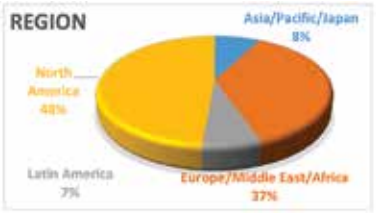

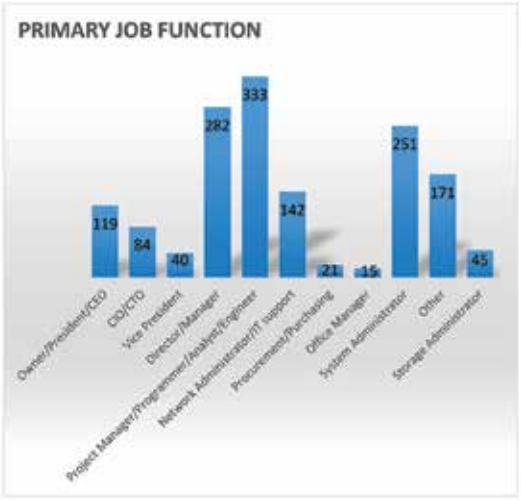

Organization / Employer Demographics

Many different industries are represented, with the majority of survey takers reporting "Other" as their primary industry (431). Financial Services (266) and Sales and Distribution (208) are also strongly represented in the survey. One-third of survey takers work at organizations with 1,000-4,999 employees, and 22% reported that they work at small to mid-size businesses with less than 99 employees worldwide.

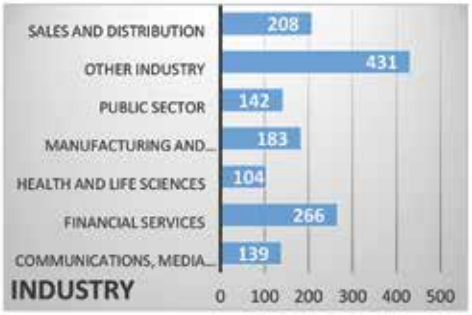

Relationship with HPE

Nearly half of all survey takers reported that their primary relationship with HPE was as an end user or customer (48%). When asked how much their organization planned to spend on HPE products or services in the next 12 months,

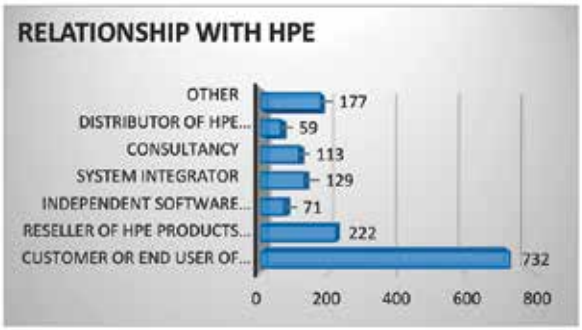

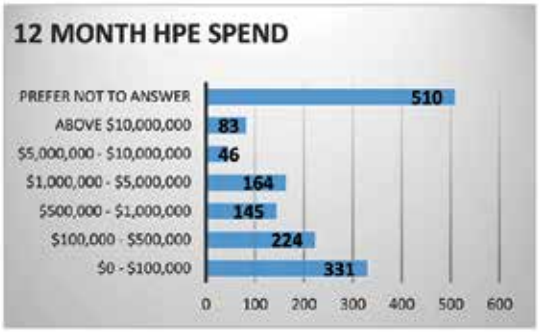

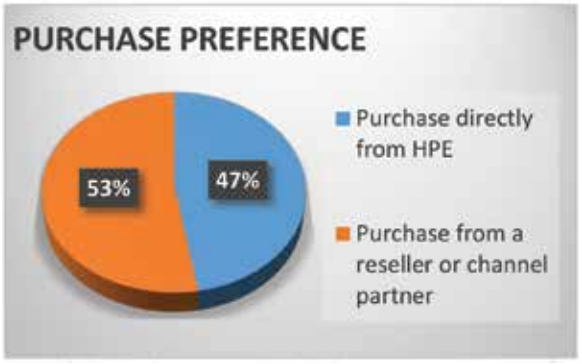

the majority preferred not to answer (34%), but 22% indicated that their organization would spend up to $100,000 this year on their HPE investments.

Of 732 people who responded to the question, more than half of the respondents (386) reported that they prefer to purchase their HPE products from a reseller or channel partner.

When asked what product/investment purchases their organization was planning to procure in the next 12 months, this open-ended question generated 850 responses from survey takers. 117 survey takers declined to answer the question, citing it was either unknown, confidential, or not under their purview. 32% (210 people) of survey takers said that they would be investing in servers within the next 12 months. Storage (13%, 82 respondents), Cloud (12%, 80 respondents), and Networking (12%, 79 respondents) were also commonly reported.

When asked specifically whether their organization was considering a Database or Enterprise Application upgrade within the next 12 months, 60% reported that their organization was not planning any substantial upgrades this year.

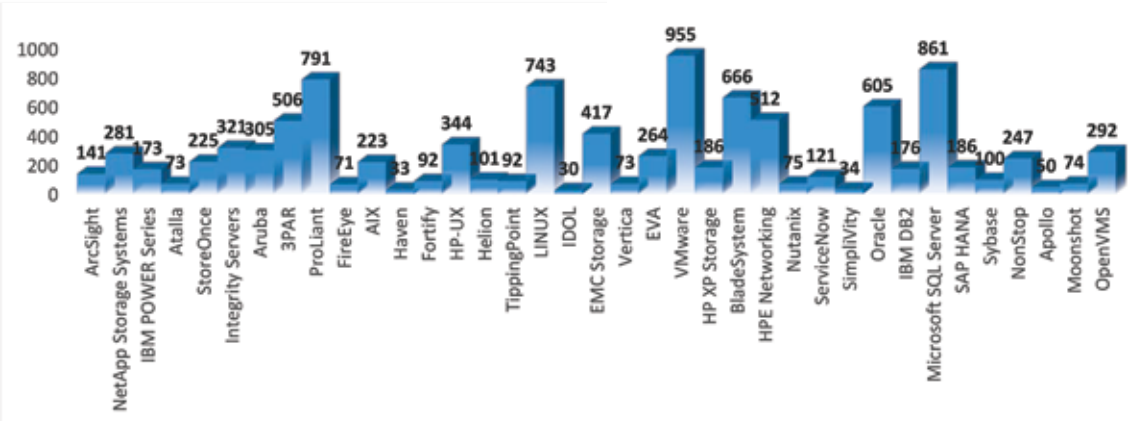

Technologies Used

When asked what HPE and non-HPE technologies survey takers used, 63.5% reported VMware as their main technology. Microsoft SQL server (57%), Proliant (52%), and LINUX (49%) were also commonly reported technologies.

When asked which technologies their company wishes to invest in, the 15% said their company would invest in servers. Storage (13%), Security (12%), and Cloud (12%) were also common responses.

Everyone has a moment when business continuity becomes real.

In the world of business continuity, there's no fire department to call before things get out of control. By then, it's too late. To protect your IT services from fire, or any one of a dozen other serious threats, you need protection in place *before* the worst happens.

Shadowbase business continuity solutions help ensure that no matter how toasty or damp your critical data becomes, there will always be an up-to-date copy available in another location to keep your business online. Don't wait for the fire department to arrive. Instead, contact Gravic today for more information on how Shadowbase software can protect your business now.

For more information, please see the Gravic white paper:

Choosing a Business Continuity Solution to Match Your Business Availability Requirements

GRAVIC® Shadowbase

ShadowbaseSoftware.com

Business Partner

Hewlett Packard Enterprise

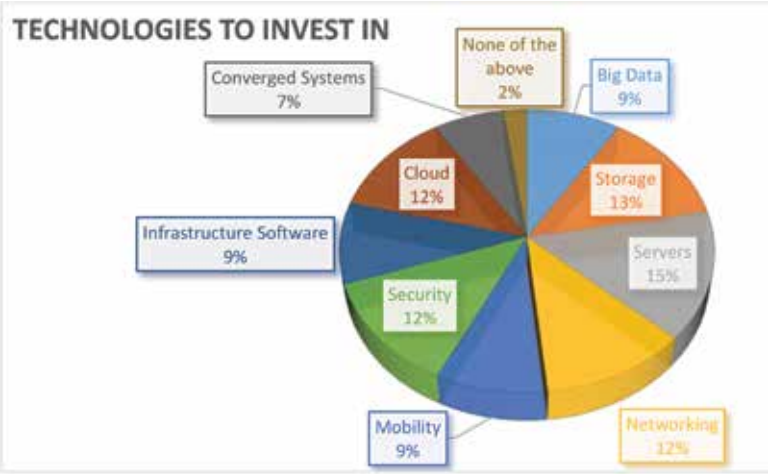

Education and Certifications

When asked about their preferred method for acquiring new knowledge for using, installing, configuring and maintaining new products, web training (17%), formal documentation (16%), colleagues (14%), and events (14%) were the most commonly used methods. 73% of survey takers (1,096 people) reported that they found certifications in general valuable, and 90% (985 people) said that HPE certifications are valuable. 81% of respondents (1,294 people) reported that they found online and classroom education valuable, while an overwhelming 94% (1,213 people) stated that HPE classroom and online training is valuable to them.

Discover Attendance

506 survey takers (34%) have attended an HPE Discover event within the past two years. Of those Discover attendees, the results indicate that technology sessions, general sessions and keynotes, product demos, and networking with HPE employees is of the most value to attendees.

When asked whether or not all survey takers had ever Livestreamed general sessions and keynotes from HPE Discover on the HPE website, 68% reported that they had never watched the Livestream. 478 survey takers reported that they had viewed a Livestream from HPE Discover.

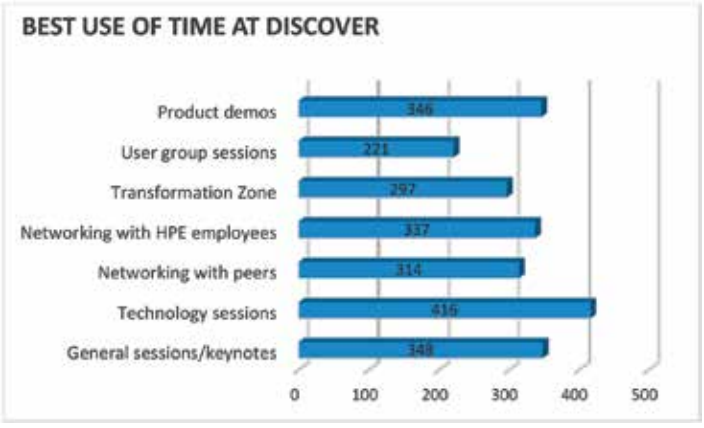

Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him 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.

Build vs. Buy:

The Challenge of Developing Homegrown Middleware for HPE NonStop Servers

Mark Roy >> **Senior Software Architect** >> NuWave Technologies **Mandi Nulph** >> **Marketing Coordinator** >> NuWave Technologies

HPE NonStop servers are arguably the best in the industry for processing mission-critical data. But to get the most value out of that data, you may want to exchange data with other systems. And to integrate with disparate systems, you'll need some sort of middleware. There are many commercial middleware products for NonStop that offer SOAP, REST, ODBC and proprietary APIs, among others.

When evaluating their options for middleware, some users explore the pros and cons of building their own solution. While this initially may seem to have benefits like savings for the IT capital budget, control over features, customization, etc., there are pitfalls to this approach.

Before you start developing your own middleware solution, consider some of the challenges that must be tackled, including standards compliance, security, multi-threading and scalability, data format conversion, diagnostics and statistics, administrative UI, and ongoing support.

Standards Compliance

If you want your middleware to work with third-party applications and development tools, standards adherence is critical. HTTP, XML, SOAP, and JSON are key standards you'll need to meet. And for security, SSL and TLS are essential. True, implementations for some of these are available through open source projects, but you need to understand them well to integrate them correctly and support them when something's not working. And, porting open source projects to NonStop is often difficult.

Security

All it takes is a quick glance at today's headlines to realize the importance of correctly implementing security on your systems, especially when you are dealing with sensitive customer information and mission critical data for your company. Not only must you implement standard security protocols to keep your data safe, but you must stay up-to-date with newly-discovered weaknesses and exploits. This means rebuilding and retesting your solution regularly to incorporate the latest security patches, which can be extremely time-consuming for developers who are expected to be working on multiple projects.

Multi-Threading & Scalability

It's easy to underestimate the complexity of building multi-threaded applications on NonStop. Without multi-threading, scalability will not be achievable. And a solution that doesn't scale is no solution at all. A related concern is memory management: if you're not careful about how you use memory, you'll also limit scalability.

Data Format Conversion

Another challenge you may face when developing middleware is dealing with data format conversions. NonStop has some unique data types. When handling data in- and outbound from NonStop, it has to be converted from the external format provided by the client/server/peer to the native format, so your NonStop application can use it. Differing byte-orderings, character sets, date formats, floating-point precisions, and non-standard data types (scaled integers, decimals) are all problems that must be dealt with. It takes a great deal of time, and skill, to identify each discrepancy and prepare to reformat all of the potential conversions, which means more time for your developers to create your solution.

Diagnostics & Statistics

A feature that's often overlooked when planning a software project is diagnostics. If you don't include diagnostics in your homegrown solution, you or your users will spend a lot more time, with a lot more frustration, chasing problems than you would if you designed-in a robust diagnostics framework from the start to tell you when your systems are operating properly (and to quickly help identify problems when they arise). And statistics? Think about all the transactions your NonStop is processing at any given time. You need to be able to analyze and do something with all of that data. That alone may require its own software solution, which takes up even more time and resources.

Administrative UI

Whether you develop your own solutions or you go with a pre-

packaged option, you need to be able to use the program. Having a user-friendly interface is important, especially in the event your developer or user is no longer with the company or is tied up with another project. In-house UI development is time-consuming. The homegrown solution developers usually skip it and configure by hand-editing text files and manual database manipulation, which is error-prone. Relying on third-party solutions can often save a lot of time and headache later on by providing out-of-the-box usability.

Build & Testing on All NonStop Platforms

Robust testing of the solution is a necessity. The business impact of software defects in production can be enormous, resulting in loss of revenue and user confidence after some time.

Commercial middleware vendors need to build and test their solutions on all HPE-supported NonStop platforms. Homegrown implementations may limit testing to just their current platform, only to get caught off-guard by unexpected problems when the platform is upgraded later on.

Even on a single platform, testing can require a lot of resources. A test plan must be created, from which individual test cases are designed and built. Automated testing is ideal, as it allows consistently repeatable execution, but adds significantly more time and planning for test development. Manual testing is less consistent and error prone. It is also time-consuming to execute, often leading to less overall testing being performed when done in-house.

Documentation

The development staff that creates the solution won't be with you forever, and even they won't remember the details a few years from now. This is why creating design documentation before beginning is important—just make sure it is updated to as-built specifications. In addition to this, you'll also need to create solid end-user documentation so you won't have to rely on oral history to train new resources.

Ongoing Support

Software requires constant updating to fix bugs, improve performance, integrate new technologies, and other routine maintenance. In addition to basic updates, troubleshooting when issues arise can often be a full-time job in itself, especially when you are talking about mission-critical applications that run on NonStop servers. There is also a risk when the original developer leaves the company for some reason.

Pre-developed middleware solutions often come with round-the-clock technical support to help should issues arise. They also stay on top of updates and release new versions of their software constantly so you don't have to.

Opportunity Cost

Opportunity cost is the loss of potential gains that you might have had by choosing the other option. While your resources are occupied with building, maintaining and supporting homegrown middleware, they could be used to implement new solutions that can't be bought and that have a more direct impact on your business.

The Bottom Line

When deciding whether to build or buy, you must consider all the costs. The decision to build software may be right when the unique features of the software directly impact the ability to differentiate your business from competitors, or to act with more agility. But for infrastructure software like middleware, it rarely makes sense unless you have time, expertise, budget and justification.

More than once we have had a prospective customer tell us, "We're thinking of 'rolling our own.'" However, they have either a) continually postponed the project, or b) eventually realized the savings just weren't there. And in one case, an enormous amount of time and money were wasted before they finally threw in the towel; their homegrown solution just couldn't perform to their own requirements. Don't be that guy.

Once you begin taking all of these factors into consideration, many people suddenly realize that they are saving neither time nor money by developing their own NonStop solutions. There are a variety of development companies who have already done the legwork for you and can provide on-call support and documentation for basically any configuration you need for your NonStop.

Diligent research and analysis of resources can help you to decide what the right approach is for your company's needs. However, when it comes to the long-term benefits of homegrown middleware solutions, you may find that it might not be worth it after all.

About NuWave Technologies, Inc.

NuWave Technologies was founded in 1999 and is well-known for providing exceptional middleware solutions for customers in the HPE NonStop space. Their middleware products include SOAPam® Server and SOAPam® Client, which use SOAP Web services to exchange data with NonStop, as well as LightWave Server™ and now their latest product LightWave Client™, which use REST APIs to integrate NonStop with modern devices and services. All NuWave middleware is high-quality, has a low total cost of ownership, and comes with excellent support from middleware experts.

Visit www.nuwavetech.com for more information.

Mandi Nulph is NuWave Technologies' marketing coordinator. With a degree in Mass Communication and Journalism, she boasts 9 years of professional experience writing and editing for a variety of publications, as well as an extensive career in marketing. Along with Gabrielle Guerrera, she volunteers to help interview companies making innovations in the NonStop space for a variety of trade publications.

Mark Roy has worked in and out of the NonStop arena for over 30 years, both in professional services and software development. His work has centered on integrating applications on NonStop with clients and servers on other platforms. He has been with NuWave for 15 years and another NonStop partner for 13 years prior to that. He is currently a senior developer for NuWave's LightWave products.

XYPRO's CISO on XS1 and Intelligent Data

Mandi Nulph >> NuWave Technologies

I recently had the opportunity to chat with Steve Tcherchian, Chief Information Security Officer (CISO) at XYPRO Technology about their newest product, XYGATE SecurityOne (XS1), their ongoing relationship with Merlon, and the role of intelligent data and security for the HPE NonStop server.

Mandi Nulph: Could you give us a little bit of XYPRO's history, as well as some of your professional background in the NonStop space?

Steve Tcherchian: XYPRO has been around for over 30 years in the mission-critical security space working on HPE NonStop servers. We provide security solutions for access control, data security, auditing, compliance, and our new product line that we will be talking about today is security intelligence and risk management; it is called XYGATE SecurityOne, or XS1 for short.

As far as my personal background in the company, I'm the CISO (Chief Information Security Officer) as well as the product manager and product owner of the XYGATE SecurityOne line. I have been with XYPRO close to 13 years now, and my background is mostly in security management, infrastructure and strategy in large global organizations.

Mandi: You recently released your latest product, XYGATE SecurityOne; tell me more about that.

Steve: Yes, we announced it at Boot Camp last year (November, 2015). XYGATE SecurityOne is our security intelligence and analytics solution for the HPE NonStop server and is XYPRO's largest development project to date. We are in beta right now, so we have a few customers to whom we have already deployed. We are working through that process at the moment. It's pretty exciting.

The product itself focuses on specific security indicators on the HPE NonStop server. We are able to bring multiple event sources together, correlate them, and most importantly contextualize them for a big-picture view of security and risk management in your NonStop infrastructure (see image below).

Mandi: How does XS1 differ from any of the other security intelligence products available for NonStop?

Steve: In most environments you're going to have SIEM (Security Information & Event Manager) solutions like ArcSight, QRadar, SPLUNK and others. The NonStop server can send its event data to SIEMs, but the challenge is that those solutions are not contextually aware of the NonStop. So, they receive events, parse them, put them in individual columns, and then the data just sits there in buckets, waiting to be analyzed. Unless someone writes a rule or some sort of specific action, that data doesn't go anywhere. All that really allows you to do is to check a box that says, "Yes, I'm compliant, I'm sending my NonStop data to the SIEM." XYGATE SecurityOne is different. It's a complementary solution to the SIEM. The NonStop server sends its data to SecurityOne and then SecurityOne differentiates itself with specific domain knowledge of the NonStop server environment. SecurityOne provides rich, contextualized data to the SIEM for the bigger enterprise security picture.

Other solutions position themselves as intelligent in analytics, but for the most part what they end up doing is taking event streams, or data sources, and correlating them. Correlation is more or less taking individual events and finding some sort of relationship between them and linking them together. What that ends up doing, however,

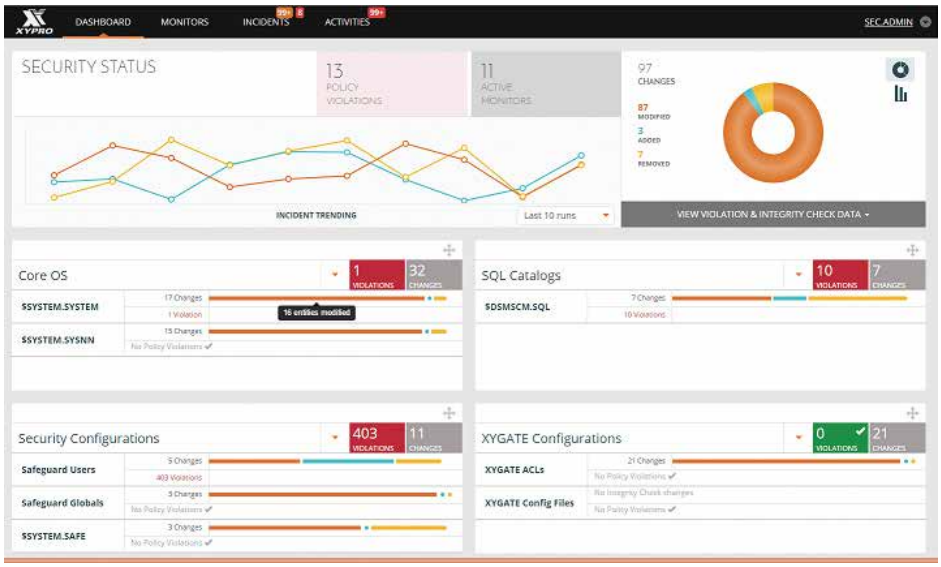

is creating more data and more noise. Now you've got to do something with all that data. What contextualization does is it looks at that new, enriched data and says, "Yes, based on what's happening in your environment and what we know about the NonStop, your specific system and your users, these are the incidents that are most important to your infrastructure."

Mandi: Are there any particular NonStop users that would benefit from a product like XYGATE SecurityOne more than others?

Steve: Everyone who needs to manage risk on their NonStop. So if you ask anybody, "Can you go home at night and sleep comfortably knowing your NonStop systems are 100% secure?" Anyone who answers "no" would be a perfect candidate for SecurityOne. It gives you the visibility you need to emphatically say you know what's happening on your systems as it's happening.

Mandi: How does SecurityOne work with other products in your portfolio? Do they operate as a suite to give NonStop users the full package?

Steve: The whole idea of SecurityOne came up because of the deep XYGATE suite of security products for the NonStop server. We are leveraging all of the data our products generate to feed SecurityOne. Without the history of XYGATE products, there wouldn't be a XYGATE SecurityOne. Products like XYGATE Merged Audit (XMA) and XYGATE User Authentication (XUA), which come shipped on every system, are integral to SecurityOne. XMA especially sits in a strategic place in the NonStop because of its audit capabilities. Because it comes as part of the NonStop security supplied with the operating system, it has unparalleled access to NonStop audit data. XMA gathers that data and is able to feed it to SecurityOne, feeding the product's intelligence and analytics. The same can be said about the rest of the XYGATE suite of security and access control software. We're able to do keystroke auditing and interpret those keystrokes for context. We know how the system is secured. We know how users are set up. We're able to generate login data and event data. We realized, this goldmine exists on the NonStop. Let's bring it all together and provide that big picture view that people need and clearly want. That's what SecurityOne offers.

Mandi: What is the potential for this solution?

Steve: Going forward, the more data that we feed SecurityOne, the more intelligent the decision-making process will become. That means if we can feed it data from other vendors' applications, or if we can give it data from different platforms and sources, the more data that can converge in SecurityOne, the better the results for the customer. The more data, the more intelligent that data becomes.

Mandi: That product sounds very exciting. I'd like to change subjects slightly and get caught up on one of your strategic partnerships. You have talked a lot about your partnership with Merlon at recent conferences. Can you tell us a little more about your relationship with Merlon and how that continues to evolve?

Steve: Our partnership extends just beyond Merlon as a company. The executive and development teams are made up of key people who have all worked with one another in the past and have very extensive NonStop security backgrounds. Our relationship with Merlon principals began way back in the early '90s as collaborative partners servicing the NonStop market with security solutions. So we have had quite a long, ongoing relationship with them. XYPRO has been the Merlon sales and marketing arm for years now, and we both have enjoyed success working together in that space.

With our strong existing relationship, the synergies between systems database management and security is something we continually bring together, because data integrity and data security on the NonStop is paramount. It's a very nice fit between what Merlon offers and what XYPRO brings to the table, and continuing to innovate in that area is something that we've absolutely got in our plans moving forward. We're enjoying long-term, cooperative success with Merlon by strengthening security and database management, and we will be servicing those needs in the NonStop space for a long time.

Mandi: You just announced SecurityOne not that long ago, but do you already have plans in the works for what's coming next at XYPRO?

Steve: Absolutely! On the software side, the functionality in SecurityOne right now is the tip of the iceberg. We see XS1 as being our flagship product going forward, and the platform itself will allow us to expand and integrate more and more solutions that will be beneficial to the NonStop and even beyond the NonStop space. We have quite a few things in the works.

Another benefit to the XS1 project has been the growth and injection of energy XYPRO is experiencing as a result of the team we're building for this project. Lisa Partridge, our CEO, and Andrew Price, our VP of technology, have given me the support and resources to assemble an amazing team of technologists. The XS1 team is comprised of traditional NonStop security experts who are joined by an energetic and invested group of engineers with amazing skills in a variety of areas that are not our traditional stomping ground. The addition of these resources to the XYPRO family has forever enriched our company, our development environment and the professional lives of everyone at XYPRO.

Mandi: That's all very exciting! We look forward to hearing about your new and improved products and partnerships in the future.

Steve: Thank you! So are we.

Check out the follow up interviews with XYPRO and Merlon on NuWave's NonStop Innovations Blog at www.nuwavetech.com/blog.

Mandi Nulph is NuWave Technologies' marketing coordinator. With a degree in Mass Communication and Journalism, she boasts 9 years of professional experience writing and editing for a variety of publications, as well as an extensive career in marketing. She volunteers to help interview companies making innovations in the NonStop space for a variety of trade publications.

Who are The Middleware Guys?

At NuWave, we know middleware.

We have specialized in HPE NonStop middleware for over 15 years, and our software architects have been developing NonStop middleware for decades. We eat, drink, and sleep middleware, so you can be sure you're getting the best.

Our products help you get the most out of your NonStop servers by integrating them with other platforms, and you won't find higher-quality, more intuitive NonStop software for the price.

Why NuWave Middleware?

- ✓ *HIGH QUALITY*
- ✓ *LOW TCO*
- ✓ *EXCELLENT SUPPORT*

Lightwave Server: Use REST services to access your NonStop applications from nearly any modern platform

LightWave Client: Access public or private REST services from your NonStop

SOAPam Server: Use SOAP Web services to access your NonStop applications from nearly any platform

SOAPam Client: Access public or private SOAP Web services from your NonStop

Learn more about NuWave middleware at
www.nuwavetech.com/middlewareguys

NuWave
THE MIDDLEWARE GUYS

Andrew Price >> VP Technology >> Xypro Technology **Wendy Bartlett** >> Distinguished Technologist >> HPE

XYGATE User Authentication (XUA) supports PCI DSS Multi Factor Authentication (MFA) requirement.

The Payments Card Industry Data Security Standard (PCI DSS) version 3.2 has recently been published. The previous version, 3.1, expires on Oct 31, 2016, at which time all new assessments must use PCI DSS 3.2. New requirements are considered best practices until Jan 31, 2018, at which point the new requirements become fully effective.

One of the main areas that has changed in scope from 3.1 to 3.2 is the requirement for Multi Factor Authentication (MFA). As of 3.2, MFA is required for all non-console administrative access to the Cardholder Data Environment (CDE). Simply put, and in NonStop terms, anyone who has access to the CDE (NonStop systems or applications) from anywhere other than the NonStop System Console is now required to be authenticated using MFA.

The good news is that many NonStop Security Administrators should be able to implement MFA using the tools they already have.

XYGATE User Authentication (XUA) has been included on all commercial NonStop Blade Servers shipped since September 2013. Others may order it separately using the PID QSN52 or QSN52U. XUA integrates NonStop authentication with a variety of off board authentication mechanisms, including Active Directory, LDAP, RADIUS and RSA SecurID. When configured correctly, these can provide MFA support, addressing the 3.2 requirement for all command line (TACL and OSS) based access "out of the box".

Application-level MFA can also be achieved in a Safeguard environment with XUA configured if the application already authenticates its users by calling USER_AUTHENTICATE_ or is changed to do so. The call to USER_AUTHENTICATE_ will invoke XUA automatically, which will perform MFA. Applications that have their own User store may need slightly more work, but can likely be modified to call USER_AUTHENTICATE_ to invoke XUA, rather than consulting their own User store for authentication.

To learn more about XUA, and to take advantage of a solution that you probably have access to already, click here, or contact your HPE or XYPRO account rep.

https://www.hpe.com/h20195/V2/GetDocument.aspx?docname=4AA3- 5035ENW&cc=us&lc=en
https://www.xypro.com/product/user-authentication-xua/

Andrew Price is VP of Technology at XYPRO. He joined XYPRO in 2011, and has over 25 years' experience in the mission-critical IT industry. Prior to joining XYPRO, Andrew was with ACI Worldwide for over 11 years, where he held roles in Product Management, Development and Architecture. At XYPRO, Andrew has engineering and product management responsibility for the XYGATE suite of products, ensuring that they continue to meet XYGATE users' stringent requirements for security and compliance on the HP NonStop. He can be reached at andrew.price@xypro.com.

Wendy Bartlett is a Distinguished Technologist in HP's NonStop Enterprise Division, and focuses on dependability – security and availability - for the NonStop server line. She joined Tandem in 1978. Her other main area of interest is system architecture evolution. She has an M.S. degree in computer science from Stanford University. Outside of work, Wendy is a dedicated choral singer and enjoys spending time hiking, in the gym, and just hanging out with her husband, Joel. She lives in the San Francisco area.

A great Chapter event ★ TEXAS style! ★

Bill Honaker >> XID Software Inc.

On May 5th of this year, N2TUG had another highly successful meeting. I could end the story right now, (and share these pictures), and most that have read of our previous meetings would probably accept it and move on.

However, this year we enhanced our meeting in 2 ways, once again raising the bar and making me proud to be part of this chapter!

NEW IDEAS

I'll start my story by going back to a February day this year when Diane Funkhouser, my planning committee partner and HPE Liaison, called me to float an idea and get my impressions. She talked to me about the Make a Wish® Foundation and suggested we leverage the generosity of our sponsors to do some good outside the group. That got my attention, and when she added another idea we had the basis of our plan. She said that there was a Reunion in the planning stages of people that used to work together at the Central Region offices of Tandem Computers, and said 'Why don't we plan the evening around getting old friends together?'

We picked a fairly simple venue for the evening (compared to years past), a nice local 'Southwestern-style' restaurant, Blue Mesa. We scoped out the big Hilton hotel next door for meeting space, and once we got the costs figured out, we started inviting our sponsors. I was very pleased to see they warmed up immediately to the charity idea and offered new highs of sponsorship beyond our hoped-for levels. Since most of them would bring 'door prizes' to our meetings to give away, we decided to leverage that as a resource which worked out extremely well.

So, come meeting day, we enjoyed a full day of presentations from the best our sponsors could offer, starting with a keynote presentation by HPE's new NonStop Database Product Manager, Roland Lemoine, who gave us a great roadmap highlights presentation. That was followed by Yash Kapadia, the CEO of OmniPayments, our Platinum sponsor. We were blessed with a beautiful spring day, and a great lunch around a pond right outside the venue to start our summer tans (and a few burns). As usual, throughout the day we had excellent technical content presented for our members. Lots of people took the opportunity to get some great networking done, too.

Then we all walked across the street for the combination of a Reunion and what Make a Wish® Foundation calls a 'Reveal'. They brought a child with a wish and his entire family to our event where he learned that his wish had been granted. They were going to Hawaii to swim with dolphins! All through the evening, there were a lot of items set out for a Silent Auction, including fantastic gifts like tablets, cameras, headphones, and gift certificates for dinners out, spas, and just general gift cards. Between that, a 'chance drawing', and other contributions, N2TUG members and sponsors donated $5,685 to help kids with life-threatening medical issues to have wishes come true!

Meanwhile, we had a lot of friends of N2TUG (and its predecessor, DRTUG) get together. All night, I saw people turn around and be excited to see a face they hadn't seen in a while. In all, there were around 100 people for the evening event, generously financed by OmniPayments. One attendee came all the way from Hawaii!

Diane and I have to thank all of our sponsors for making it possible. They are: BlackWood Systems; Bowden Systems; comForte; Hewlett Packard Enterprise; Idelji; IR; Merlon; NTI; NuWave; Odyssey; OmniPayments; Tributary Systems; XID Software Inc.; and XYPRO.

I hope you enjoy browsing a few of the pictures from the day. We all had a great time and are looking forward to next year's event!

Signage


Registration


Impromptu Networking


Keynote


Food and drink


Yash Kapadia and Bill Honaker


Session Attendees

Silent Auction


Group picture at Blue Mesa


Chance drawing


Evening venue (Blue Mesa)


Reveal

Enterprise Key Management for Hadoop Encryption

Joe Achett >> **Principal Software Engineer** >> Hewlett Packard Enterprise
Jim Christensen >> **Senior Technical Writer** >> Hewlett Packard Enterprise

Overview

Hadoop provides a rich set of technologies for storing and analyzing Big Data, but securing that often-sensitive data has not been addressed in a robust and enterprise-level way before. Encryption for HDFS has been added only recently, and its key management is lacking in many respects, including scalability, configuration, and auditing.

The Hewlett Packard Enterprise SecureData Key Server provides a stateless, robust, and highly scalable replacement for the Hadoop Key Manager, bringing true enterprise-level key management to Hadoop.

Encryption in Hadoop

As Hadoop adoption grows, companies are faced with the challenge of securing the data in their Hadoop clusters. Often large amounts of potentially sensitive and unstructured company data is moved into Hadoop for analytics, but without a standard/robust way to secure that data.

To address this shortcoming, the Hadoop ecosystem has recently introduced the concept of **Hadoop Transparent Data Encryption (TDE)**.[1] This relatively new technology allows files in the Hadoop Distributed File System (HDFS) to be transparently encrypted/decrypted, using keys provided by a Key Management Server (KMS).

Unfortunately, the default KMS that comes with Hadoop is not robust, easily configurable, or scalable; it generates keys that must be stored on the server, and authorization is configured manually in a text file. Basically, stateful key generation has management/distribution/configuration issues for enterprise deployments. As a result, companies that want to implement Hadoop TDE for their Big Data clusters are struggling with managing their KMS, which makes the overall TDE solution problematic for securing their sensitive data efficiently.

To understand this, let's start by looking at how files in HDFS are encrypted and decrypted in the Hadoop TDE architecture.

Hadoop TDE — Client-side

From the Hadoop user's perspective:

1. A regular user on the Hadoop system asks the KMS to create a key. This is done through the "hadoop key create" command, providing the name of the key and the location of the KMS.
2. After the key is created, the special Hadoop administrative user creates an **Encryption Zone (EZ)**, by assigning this key to a directory. This is done through the privileged "`hdfs crypto -createZone`" command, providing the name of the key and the path in HDFS to turn into an encryption zone. An encryption zone is basically a directory within HDFS that is protected through TDE, such that any files added to it are automatically encrypted and any files read from it are automatically decrypted, provided that you have access to the associated key.
3. A Hadoop user runs a regular hdfs command to copy a file into a directory in HDFS. For example:
   ```
   hdfs dfs -put myfile.txt /ezone_1
   ```
   If the file resides in a directory that is an encryption zone, TDE automatically encrypts the file contents, using a key managed by the KMS.
4. Later, when the file is read from HDFS, TDE automatically decrypts the file contents, using the same key associated with that file.

As described above, Hadoop users must run special commands to set up the zone, but once that's done they can deal with files in HDFS exactly the same way as they always have: encryption and decryption are handled transparently, hence the "T" in "TDE".

But what's happening on the KMS side of things during all this? That's the more interesting part, so let's take a look at that next.

Hadoop TDE — Server-side

From the server's perspective:

1. When a user asks the KMS to create a key, the server authenticates the request (through Kerberos, which is

1 Transparent Encryption in HDFS: https://hadoop.apache.org/docs/r2.7.0/hadoop-project-dist/hadoop-hdfs/TransparentEncryption.html

beyond the scope of this discussion) and generates the key. The key is then stored by name on the KMS, for later use.

2. After the encryption zone is configured, the first file copied to that directory triggers a request to the KMS for **data encryption keys** (DEKs), which are the keys used to encrypt and decrypt the file contents. The KMS randomly generates these DEKs, and encrypts them with the base key created in step 1. Because the Hadoop system must store these keys securely, the KMS returns them as encrypted DEKs (EDEKs). And because Hadoop uses a different key for each file in the zone/directory (for enhanced security), it requests such EDEKs in batches (typically 150 at a time) to store locally, associated with the zone.
3. To encrypt a file new to the EZ, the Hadoop system associates the next available EDEK with the file, and calls the KMS again, but this time to decrypt this specific EDEK into a DEK. The KMS uses the base key created in step 1 to decrypt the key, which is sent back to the Hadoop system. The Hadoop system then uses this DEK to encrypt the file contents.
4. For each new file copied to the encryption zone, the next unused EDEK from the initial batch is sent to the KMS for decryption. And whenever an encrypted file is read from the encryption zone, the specific EDEK associated with that file is sent to the KMS for decryption.

"Encryption is easy; key management is hard"

This scheme presents several problems.

Everyone who's dealt with securing data through encryption has at some point come to the realization that dealing with keys is the tricky part. There are plenty of libraries available to perform encryption, with no heavy lifting needed on your part; all you have to do is provide the key. But therein lies the problem: how do you create and manage the keys securely and efficiently?

Traditional key management systems generate and store the keys, thus involving some state (often a lot of it!). This can be problematic for distributed and high-availability systems, where keys are managed by a cluster of servers. If one server goes down or is under heavy load, another server in the cluster must handle the requests.

But what happens if a key request is made to one key management server in the cluster, and then a few minutes (or even just a few seconds) later, a request is made for that same key to another server in that same cluster? The same key must be returned from both servers for the client to work properly. This is not a trivial thing to implement efficiently. Some approaches involve real-time replication of the keys to all the other servers in the cluster, or storing the keys in a central database. Data replication requires network communication across the servers in the cluster, and using a central database introduces a single point of failure.

Basically, distributed state is a known problem in server clusters, requiring often complex and difficult-to-manage solutions.

Authorization

And what about authorization for the keys, in addition to authentication? One of the primary cases for TDE is to create encryption zones that only certain users can access. For example, an encryption zone (directory in HDFS) might contain sensitive financial information that only the company's finance team should see; another encryption zone might contain sensitive employee information that only the HR team should see. Since each zone has its own key, an authorization scheme must make sure that only appropriate users can access keys: only finance people can access the "finance" key, and only HR people can access the "hr" key. Kerberos handles the authentication part, but that only proves that users are who they claim to be. How then do you also prove that users are allowed access to the requested key?

Problem — default KMS

The default KMS that comes with Hadoop TDE is not enterprise-grade. It is usually co-located on the same Hadoop cluster as the encryption zones themselves, creates and stores keys locally within that KMS instance, and manages authorization for specific named keys using a hand-edited text file. Additionally, events are written to local log files, without a strong auditing infrastructure.

This is why some Hadoop distributions either do not provide full TDE support at all yet, or do so with the disclaimer that it is not meant for production deployments, due to the limitations of the default KMS provided by Hadoop.

Solution — HPE SecureData

Since the problem with Hadoop TDE is not on the client side (the Hadoop cluster), but on the server side (the KMS), the solution is to replace the KMS.

The Hewlett Packard Enterprise SecureData Appliance can serve as a stateless Key Management Server for Hadoop TDE. Providing a stateless KMS, managed by the robust and easily configurable HPE SecureData Appliance, solves all of the management/distribution/configuration issues inherent in the default Hadoop KMS, and adds enterprise-level scalability to this critical key provider component of Hadoop TDE.

The underlying technology behind this approach is the HPE SecureData Appliance, which uses **Identity-Based Encryption (IBE)**[2] to statelessly derive keys on-the-fly as needed, using a provided identity and a set of master secrets (base keys) stored securely on the server.

The HPE SecureData Appliance supports the **Hadoop KMS REST API specification**[3] that Hadoop TDE uses to retrieve keys and metadata from the KMS, allowing it to seamlessly replace the Hadoop KMS as the key provider for Hadoop TDE.

How it works

The HPE SecureData Key Server supports Hadoop TDE through several REST APIs, with underlying KMS functionality to perform the following TDE-requested operations:

- Derive an encryption zone key (EZK), which is the base key used to encrypt other per-file keys.
- Generate and encrypt random data encryption keys (DEKs), in batches.
- Decrypt individual DEKs.

[2] HPE Identity-Based Encryption: https://www.voltage.com/technology/data-encryption/identity-based-encryption/
[3] Hadoop Key Management Server (KMS) - Documentation Sets: http://hadoop.apache.org/docs/stable/hadoop-kms/index.html

Figure 1 below shows how the HPE SecureData KMS integrates into the Hadoop TDE flow:

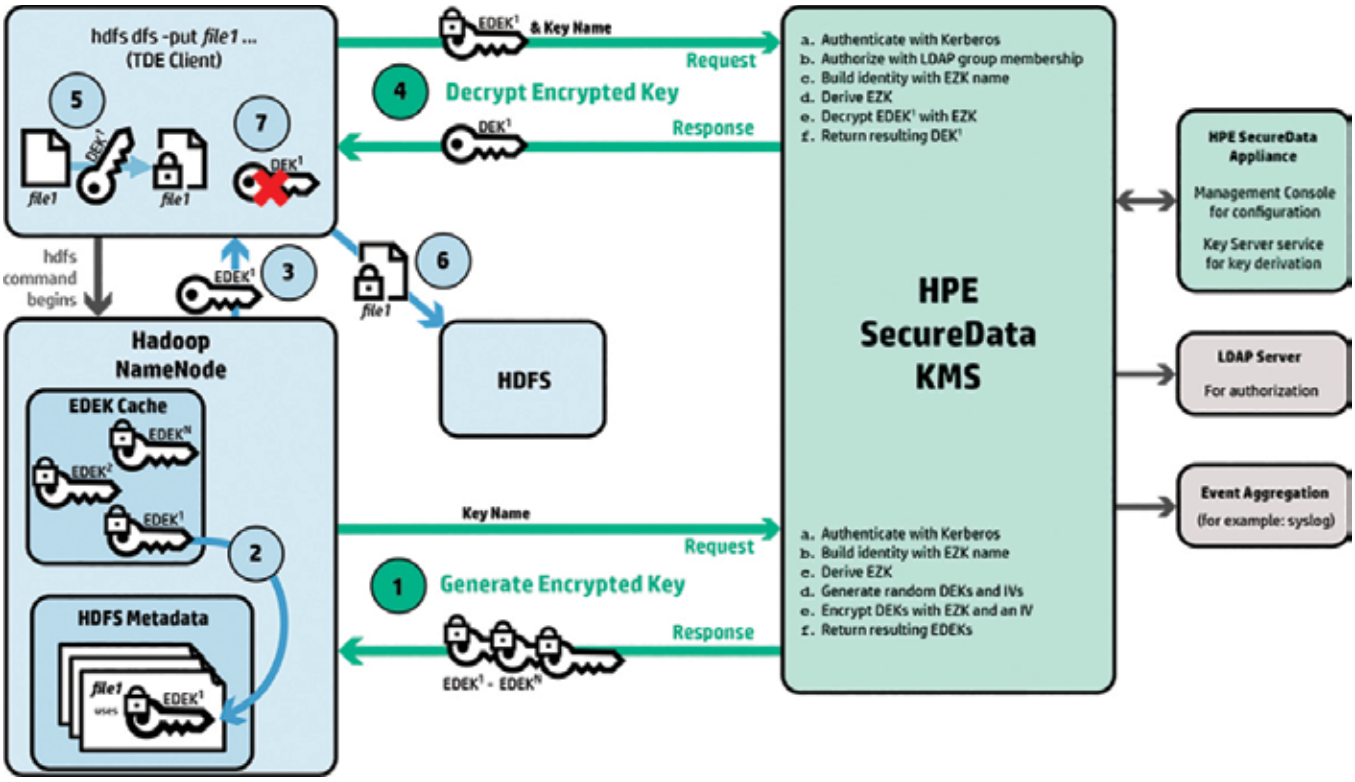

1. When the first new file is added to an encryption zone in HDFS, TDE requests a batch of encrypted DEKs (EDEKs) from the KMS for the associated EZK name. The **HPE SecureData KMS** statelessly derives an AES key for that key name, and uses it to encrypt the randomly generated DEKs in the response.
2. The Hadoop NameNode caches the returned batch of EDEKs, and associates the first one with the new file.
3. The NameNode provides the EDEK to the Hadoop hdfs command.
4. The hdfs command sends the EDEK to the KMS for decryption. **The HPE SecureData KMS** statelessly derives the EZK and uses it to decrypt the EDEK, returning the DEK to Hadoop. Authorization is via LDAP.
5. The Hadoop hdfs command uses the DEK to encrypt the file.
6. The encrypted file is written to HDFS.
7. Hadoop securely disposes of the DEK.

The main advantages in this solution are as follows:

- **Stateless Key Management and rollover** – Generating and storing EZ keys, as in the base Hadoop KMS implementation, requires state to be backed up any time a new key is created, and that those generated keys be distributed/replicated to the other servers in the KMS cluster; this distributed state can be problematic for consistency within a cluster of servers. But in the HPE SecureData KMS, these EZ keys are derived on-the-fly using IBE technology, and need not be stored or distributed. Any number of KMS machines can be configured behind a load balancer, for high availability and redundancy, with no state distributed between those machines.

  For enhanced key rollover support, the HPE SecureData KMS encrypts the DEKs using a specific approach within IBE: **Identity-Based Symmetric Encryption (IBSE)**. This innovative technology uses embedded identity information to make the ciphertext self-describing in terms of key versioning, allowing simplified key rollover.

  If, for regulatory compliance, you need to roll over the base key associated with an encryption zone, you can do so by creating a new set of master secrets through the central Management Console and deploying them to all the KMS nodes in the cluster. Multiple key rollovers can be performed over the years, without the Hadoop system on the client-side ever having to deal with managing which data encryption keys were encrypted with which "versions" of the base encryption zone keys; the HPE SecureData KMS handles that automatically through IBSE.
- **LDAP Authorization** – In the base Hadoop KMS, users are authorized for EZ keys by manually editing a text file, which lists which users are allowed access to which keys. This is error-prone, difficult to manage efficiently, and requires that this configuration file be copied to each KMS in the cluster. But in the HPE SecureData KMS implementation, EZ keys are authorized for users using LDAP group membership, leveraging an enterprise's existing directory service. In most cases, the company would already have an LDAP server (such as Microsoft Active Directory) that manages employees and business units/groups. So, by naming the encryption zone keys based on the business unit/group names in that LDAP infrastructure, access to encryption zones can be limited to authorized users only, without having to explicitly configure or manage an additional set of access

controls. When group membership in the LDAP server changes, as new team members join or leave, access to the encryption zone keys in HDFS automatically reflect that, without having to synch up these changes in other configuration files.

- **Central Configuration** – The base Hadoop KMS is configured by editing configuration files, and started via command line scripts. In contrast, the HPE SecureData KMS is configured through the HPE SecureData Appliance Management Console, a robust web-based UI with strong access control and auditing infrastructures. Leveraging this management UI allows administrators to easily configure the HPE SecureData KMS for Hadoop TDE service, and to deploy this central configuration to any number of load-balanced servers in a cluster configuration, with administrators authenticated via LDAP and actions written to audit logs. Additionally, the configuration can be locked down using dual-control protection, requiring approval from a second administrator before any configuration changes can be made.
- **Reporting/auditing** – The HPE SecureData Appliance comes with a robust logging/reporting system, making it very easy for administrators to see all the incoming service requests, for auditing purposes. With the integration of the Hadoop KMS API to the HPE SecureData Appliance, that same reporting system is leveraged for auditing Hadoop TDE requests to the KMS as well. Administrators can run an included report to view all the KMS requests from the Hadoop cluster, including details such as client IP addresses, users making requests, and encryption zone key names. This provides a powerful way to easily monitor the Hadoop TDE key requests in a distributed cluster deployment, not available in the base Hadoop KMS implementation.

Conclusion — HPE SecureData and Hadoop Distributions

The HPE SecureData KMS for Hadoop TDE has been successfully tested with the client-side TDE implementations on Hortonworks Data Platform (HDP) 2.2, 2.3, and 2.4 initially, and then later on other Hadoop distributions as they started providing initial TDE support, including Cloudera CDH and IBM BigInsights.

The HPE SecureData Appliance is an industry leader in stateless key management, and those capabilities are leveraged in the KMS implementation built on top of it. Hadoop TDE is a relatively new technology in the Hadoop ecosystem, so there are few competitive approaches for the KMS implementation. Other major Hadoop distributions, such as MapR, are taking a wait-and-see attitude before integrating this feature in their production releases. Introducing the stateless HPE SecureDataKMS for Hadoop TDE which has the potential to really energize this technology for real use in production enterprise deployments, and to strengthen HPE's position as a forward-thinking leader in the Big Data/Hadoop security space.

Joe Achett is a Principal Software Engineer - Master Technologist for Hewlett Packard Enterprise - Data Security, responsible for enterprise security integrations with Big Data technologies such as Hadoop. Prior to that, he was one of the main technical leads who developed the HPE SecureData Appliance, including its stateless Key Server for FPE (format-preserving encryption). Joe's career spans over 20 years of programming experience, primarily in Java and web-based technologies, in industries including online banking, VoIP, healthcare, computer security, human resources, and credit reporting. He can be reached at joseph.achett@hpe.com.

Jim Christensen is a Senior Technical Writer for Hewlett Packard Enterprise – Data Security. His responsibilities include producing the documentation for protecting PCI and PII using a variety of Data Security APIs, and he has innovated solutions that bring the code and documents together into a cohesive whole. Jim spent the first half of his career as a programmer, working for a variety of Silicon Valley start-ups, ending up at Microsoft in Seattle at the end of the 90s. Not long after that, he switched to writing documentation for developers and has been doing that ever since.

NonStop X is opening the doors to even greater NonStop integration with modern architectures!

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

It takes very little encouragement for me to travel to user events anywhere in the Americas. This year I have participated in numerous NonStop user group meetings, from America's west coast on through to the mid-west. I have also participated in the BITUG event held this past spring even though this required me to cross the Atlantic. Whenever promotions for upcoming events of any kind appear in my inbox I do take a moment to check what they are all about and oftentimes, I respond in the affirmative. For events outside of the NonStop community, particularly when they involve payments, I have been fortunate to obtain a press pass and that goes a long way to ensuring that I will show up.

Just a few days ago I saw a promotion by Gartner for an upcoming summit that will be held at the end of the year. I am not a Gartner customer, but even so, when I read about this event, **Gartner Application Strategies & Solutions Summit 2016**, billed as the premier event for accelerating engagement, driving customer experience and delivering digital business innovation I have to admit, I was sorely tempted to register. What came close to tipping the scales was the description Gartner provided for one of the tracks, integration.

"Connecting the dots across and beyond the enterprise," says Gartner

Gartner proffered the explanation about how, "Integration is no longer a complex cost center that companies have to endure. The explosion of mobile, fragmentation of application portfolios across the cloud and rise of the Internet of Things mean that integration is a pervasive issue that demands upfront thinking. We examine the changing technologies and practices required to meet today's digital business demands." As we all have come to observe through the years, Gartner, as Gartner so often does, is connecting the dots for us between Mobility, Cloud and IoT.

Gartner realizes that the message coming from many of the vendors including HPE, center on Cloud, Mobility, Security and Big Data. Or, as HPE reminds us every chance it gets, Transform to a hybrid infrastructure, **Enable** workplace productivity, **Protect** your digital enterprise, **Empower** the data-driven organization. By now few of us have missed the emphasis HPE places on these key themes – the fact that Gartner is now taking these themes to enterprise customers with, of course, the usual Gartner embellishments and "unique interpretation" doesn't come as a surprise.

However, the story of integration is vitally important for all of these four themes to come together in a cohesive manner and provide enterprises not just renewed value from their IT investment but the opportunity to be better aligned with the expectations of their users, even as they remain competitive in the marketplace. Integration of the systems as we transform to a hybrid infrastructure, integration of the networks as we deploy clouds and edge products and communicate among the many processors involved and yes, integration of the data. The underlying message coming from all these recent vendor proclamations is that the older model of "best of breed" is dead and the myriad number and variety of application silos is long past its prime.

"At the beginning of the 21st century, the earth needed to find a way to keep up with the data from over 30 billion connected devices, which changed the basic architecture of computing."

There are two significant and not entirely unrelated drivers for greater integration and the dismantling of application silos – the impact of consumer-driven engagement and the rise of IoT as one by one, each and every device we have ever interacted with, be it a toaster, a lawn sprinkler or a car, comes equipped today with multiple sensors. It didn't shock anyone who saw the promotional video clip in support of the new movie, ***Star Trek Beyond***, which was unveiled at HPE Discover 2016 to see HPE collaborating in this manner with the film producers.

It was part of the first HPE marketing launch in support of The Machine but the message wasn't lost on anyone who saw it, including actors in the video clip. What, just 30 billion connected devices? Any thoughts among the NonStop community that somehow they would be shielded from the explosive growth in mobile payments are beginning to subside. It's not just payments but almost all applications as well as the majority of third-party middleware supports mobile devices. Perhaps even more alarming is the potential for sensor information streaming into NonStop and if that becomes a reality, even more attention will need to paid to just how soon 30 billion devices will be connected.

Already we have seen some NonStop vendors step up to support protocols like HTML5 for better integration with mobile devices, including smartphones and tablets. While essentially a further development out of Web services in that it's a consumer-to-provider relationship its benefits center on the HTML5 protocol itself accommodating different form factors. Whether its simple alert messages out of Prognosis, or part of a larger application supported by comForte Client Server Link (CSL), NonStop solutions haven't been immune from the rush to

supporting mobile devices. comForte, ACI, NuWave and a couple of other vendors are all succeeding in providing middleware to facilitate integration with mobile devices, even as they all supply slightly different solutions.

However, I have serious doubts about the long term prospects for NonStop middleware based on SOA and Web services given the sheer number of mobile devices, in the hands of consumers, should they be directly connected to NonStop solutions. Likewise, I haven't heard any software architects warm to the idea of flooding NonStop systems with sensor information as part of any IoT initiative. In both cases, where such integration is anticipated, talk quickly turns to cloud computing and here, the transformation to a hybrid infrastructure begins to make a whole lot more sense to the NonStop community. With the expressed goal of integrating traditional IT with the new world of clouds, hybrid infrastructure offers the prospect of having a massive scale-out private cloud capability front-ending mission-critical solutions.

And nowhere in the slideware HPE has shown of late do I see any indication of NonStop systems not being a part of the hybrid infrastructure.

Indeed, quite the contrary is in evidence as numerous times this year, starting at GTUG but continuing on at HPE Discover 2016, HPE IT has been giving presentations on just how they are going about integrating vNonStop with the rest of the HPE IT infrastructure whereby vNonStop, with its newfound ability to participate rapidly in any provisioning undertakings, supports NS SQL as a Data Base-as-a- Service – yes, SQL/MX is open for access from any applications that are part of the transformed hybrid infrastructure underpinning HPE IT on the basis of a DBaaS. With the arrival of a virtual NonStop running as a guest of the open stack Linux / KVM environment, the opportunity for NonStop to be integrated with the rest of HPE IT is inescapable and will be showcased to everyone who is escorted around the new HPE IT datacenters. Not all of HPE's 100,000 plus employees will be directly connected to vNonStop at any one time but whenever critical data is needed, they will arrive at NonStop from essentially a private cloud front-ending vNonStop DBaaS.

Clouds and SOA aren't mutually exclusive. Indeed, for those NonStop users who already have started down the path towards private cloud integration, they are very likely leveraging SOA albeit in a slightly different manner. There are positives to using SOA but there are also liabilities that are exposed when private clouds enter the picture. According to a June 16, 2016, article in ***The Edge Technical Trends***, How Cloud Computing Has Reshaped SOA, there are "disadvantage associated with SOA in that it isn't the best option for software that utilizes GUI functionality, because those systems become unnecessarily complicated when designed using SOA".

"In traditional SOA, coupling of systems could be loose or tight," it states before adding, "There can be thousands of cloud services within an architecture. Tight coupling would mean that an application would be completely stopped if a particular service goes down or if the connection were to be lost. Loose coupling is essential to allow for flexibility and independence, but the concept of coupling systems is still very much alive in the cloud model." For more than a decade the NonStop community has had a choice in SOA and Web services implementations and the good news here is that when it comes to front-ending NonStop solutions with clouds, that integration between NonStop and the cloud can leverage SOA as it always has. It's just that the integration with mobile devices and indeed sensors needs to be addressed with a light-weight solution.

Before settling on how this may be implemented, two new middleware offerings come to market

Two highly significant middleware offerings, NonStop Application Direct Interface, or NSADI (previously known by its internal project name, Yuma), and the runtime platform, Node.js (supporting Server Side JavaScript) give applications and middleware developers new tools to better integrate with the hybrid infrastructures they encounter. These two offerings are only support NonStop X but both of them bring some serious heavy-lifting to applications and middleware developers targeting NonStop X.

I have to admit that I was excited to hear of NSADI and of the reception it was given by the NonStop vendor community. Having removed the last proprietary components from the NonStop hardware, NonStop systems shared the same packaging as HPE provides for Linux and Enterprise Windows users. SuperDome X running Linux was indistinguishable from NonStop X running right alongside it – something that was in evidence last year at HPE Discover 2015 when a HPE salesman, explaining the properties of SuperDome X, stood in front of a NonStop X system. Of course, I couldn't let that continue without pointing out an obvious case of mistaken identity

While HPE hasn't gone so far as to package NonStop with Linux in the same chassis, with the advent of NSADI I can't help but wonder whether solutions vendors will recognize that this is an option and with NSADI connections in place between NonStop and Linux (switched via a Mellanox InfiniBand switch), it is bound to happen and a recent conversation with OmniPayments, Inc. CEO, Yash Kapadia, leads me to believe at least one vendor shares this observation and is working towards such a goal. When asked, Yash wasn't shy about telling me how, "vNonStop and NSADI are very much a part of our roadmap for OmniPayments."

Node.js is every bit as important as NSADI and has been deployed within the NonStop Advanced Technical Center and been accessed numerous times for demonstration purposes by numerous solutions architects. Whereas in the more traditional world of client / server computing, the code we wrote wasn't quite the same on a client device as it was on the server and it took the two development teams to get together to discuss real world integration whereas with JavaScript, what runs on the client, now runs on the server greatly simplifying the integration of the client with the server and a first for IT.

More importantly, the issue concerning dependence upon kernel-level threading has become an issue of the past with JavaScript, as it's almost as if the architects that provided the initial implementation of JavaScript had the loosely-coupled,

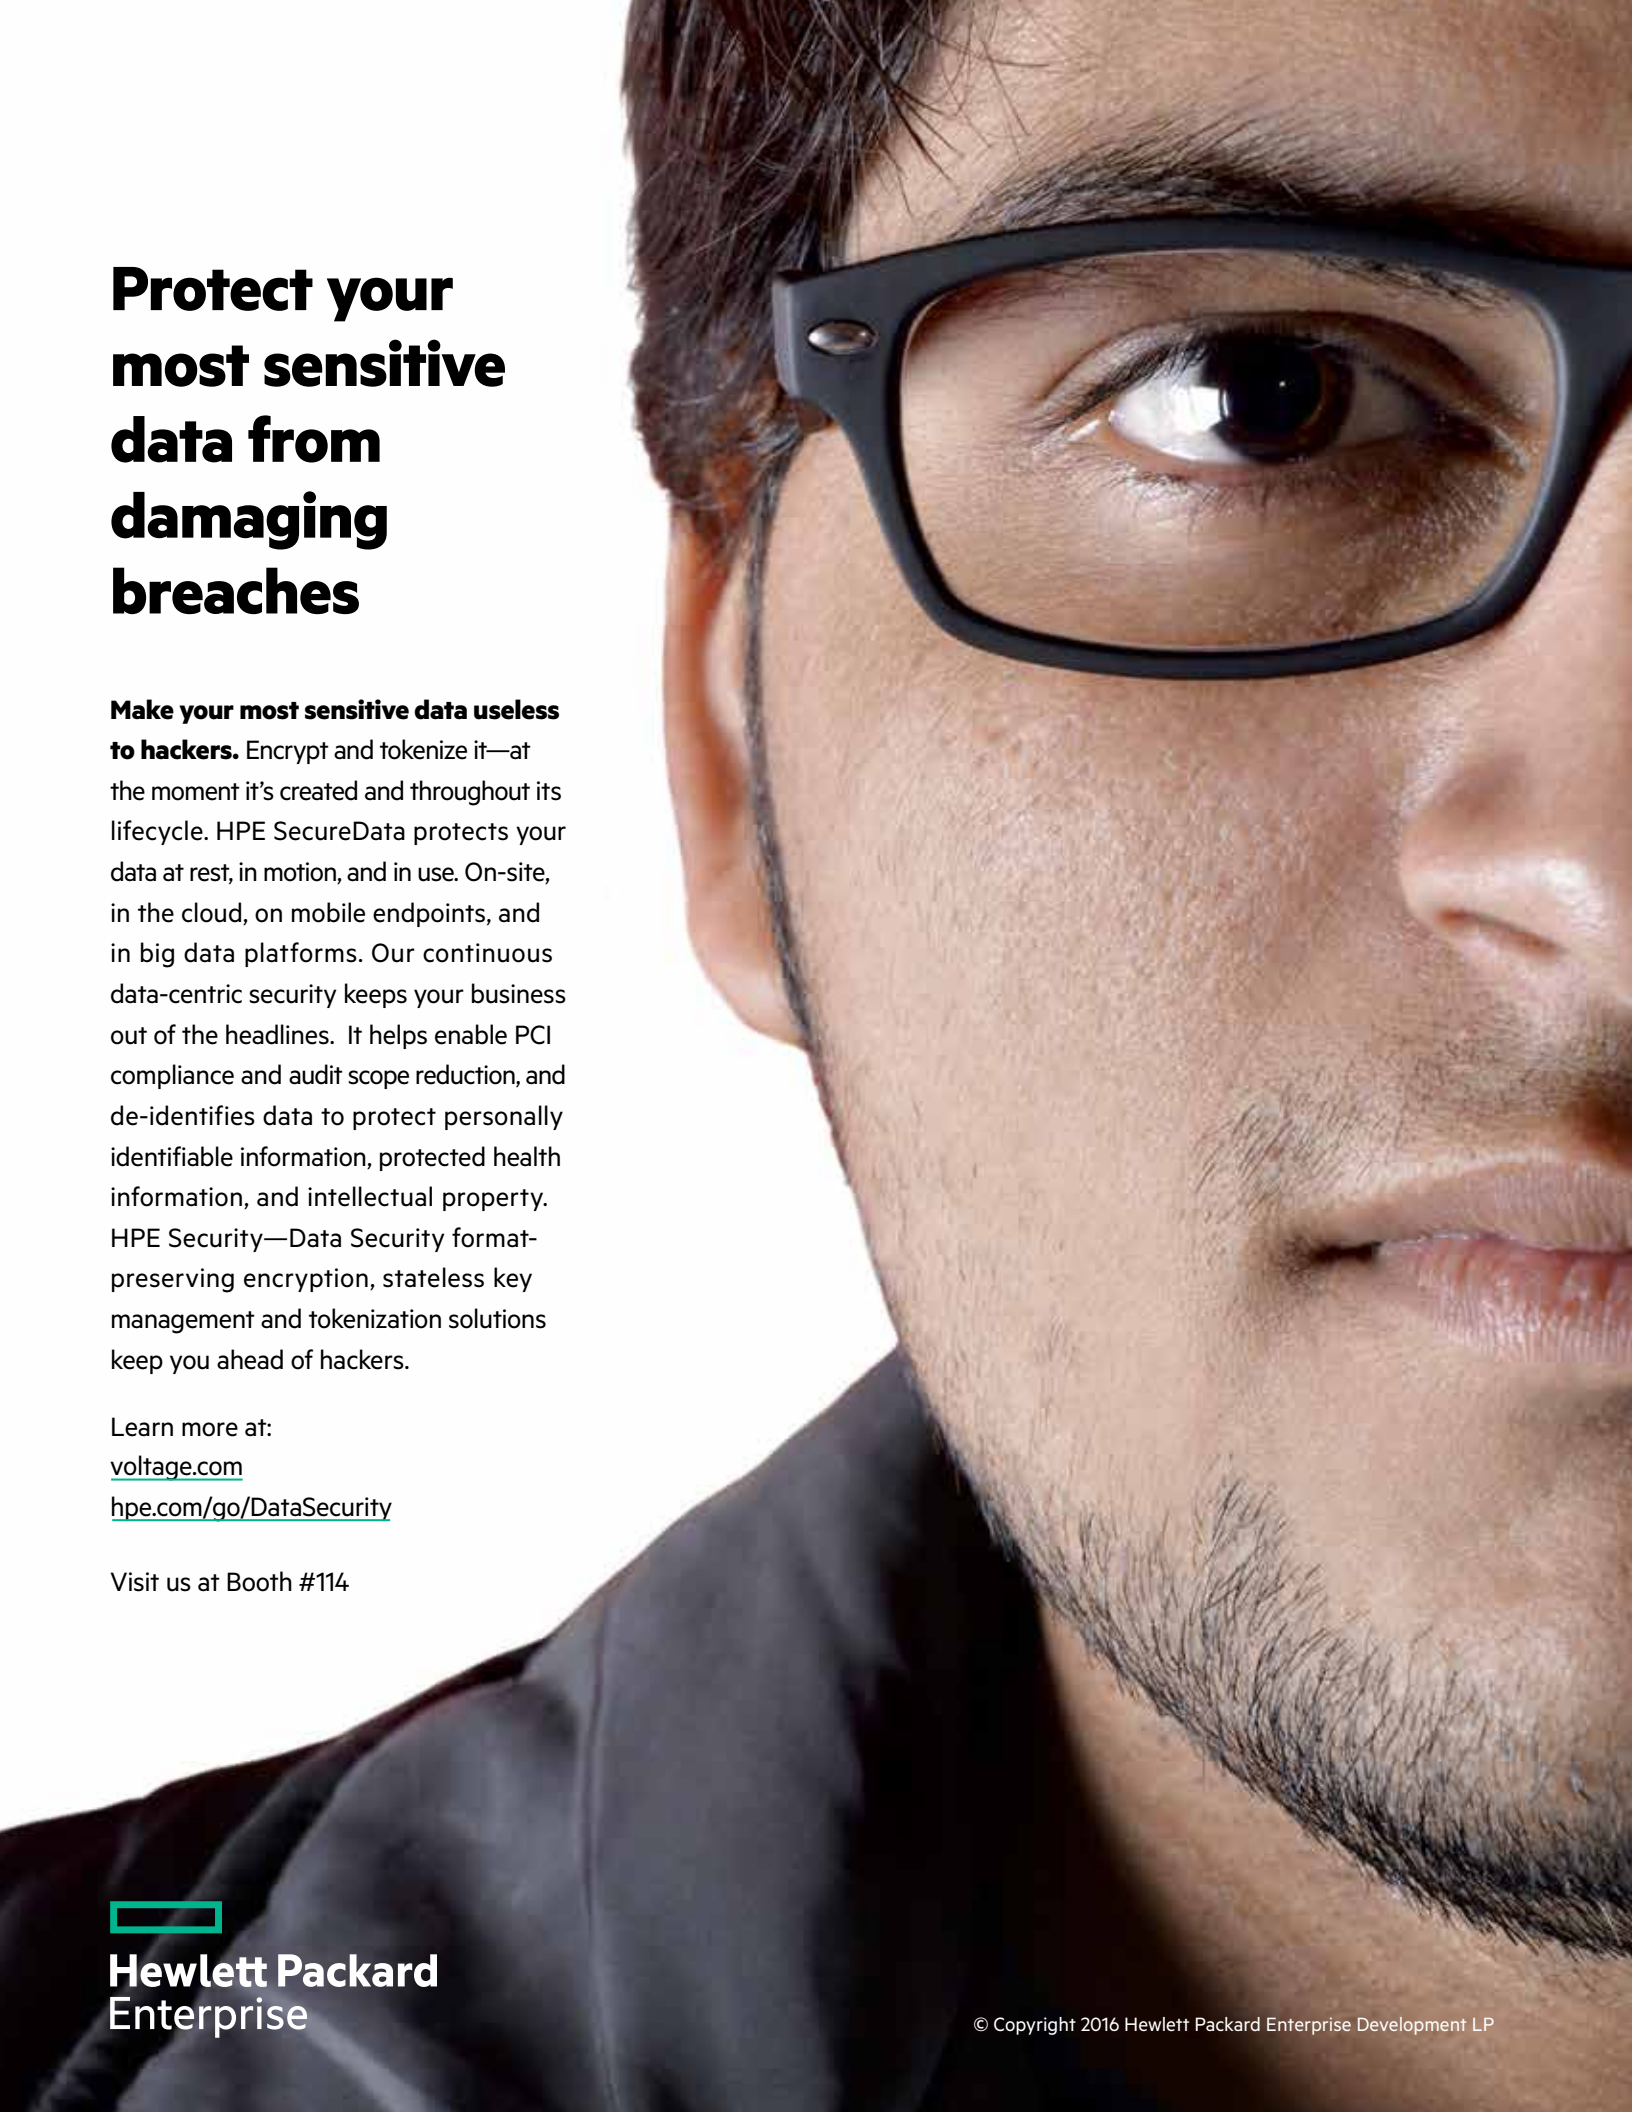

Protect your most sensitive data from damaging breaches

Make your most sensitive data useless to hackers. Encrypt and tokenize it—at the moment it's created and throughout its lifecycle. HPE SecureData protects your data at rest, in motion, and in use. On-site, in the cloud, on mobile endpoints, and in big data platforms. Our continuous data-centric security keeps your business out of the headlines. It helps enable PCI compliance and audit scope reduction, and de-identifies data to protect personally identifiable information, protected health information, and intellectual property. HPE Security—Data Security format-preserving encryption, stateless key management and tokenization solutions keep you ahead of hackers.

Learn more at:
voltage.com
hpe.com/go/DataSecurity

Visit us at Booth #114

Hewlett Packard Enterprise

shared-nothing, architecture of NonStop in mind. As the team behind the deep port to NonStop of Node.js (a team that included NonStop solutions architects as well as developers from InfraSoft Pty. Limited) explained it all to me more than a year ago, "The Node.js model of event-driven, non-blocking I/O is particularly suited to I/O bound applications". They stated, "It may have almost been dictated by the fundamental concepts of writing a high-performance OLTP application running on the NSK operating system."

From the perspective of languages and development frameworks, as far as I am concerned, it just doesn't get any better than that.

The NonStop community now has access to a language that is optimized for performance on NonStop and yes, run your application on NonStop X, run it on Linux connected to NonStop via NSADI, run it on a cloud also connected to NonStop via NSADI (and shortly, according to Andy Bergholz, Senior Director of Development, HPE NonStop, run it via RDMA over Commercial Ethernet – RoCE), run it on mobile devices including smartphones and tablets and yes, run it out on the latest "edge products" now being introduced to better handle IoT.

The very fact that there is as much talk these days about the best ways to integrate NonStop into hybrid infrastructure masks the importance that HPE is giving NonStop. Who could have imagined just a short time ago that NonStop would gain such prominence? If there's any takeway from discussion on integration it's that NonStop still provides the value enterprise users salivate over – the gradual acceptance of three nines (8 hours of downtime / year) is beginning to subside as recent failures among airlines and stock exchanges have highlighted, something that has been highlighted numerous times in the LinkedIn group hosted by Dr. Bill Highleyman, the Continuous Availability Forum.

IBM is once again flaunting just how many nines they support, but eventually, talk returns to NonStop. If HPE IT is placing as many chips as it does on NonStop and on NonStop integration, isn't it time we all followed suit and talked to our own development teams? I may not make it to the ***Gartner Application Strategies & Solutions Summit 2016*** this time, but I think the message is very clear. HPE has focused on the right themes – Transform, Enable, Protect and Empower with the transformation to a hybrid infrastructure very important for the long term prosperity of NonStop; integration, however, is the key aspect of any transformation that NonStop users pursue and yes, NonStop can play a key role in even the most complex of transformations as HPE IT so clearly demonstrates today.

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, Striim, Inc., InfraSoft, and OmniPayments, Inc.

HPE SecureData Integrates with HPE Atalla HSM

Trish Schaefer Reilly >> Hewlett Packard Enterprise

The Challenge

The volume of data, the sophistication of abundant computing, and the borderless flow of data are outpacing the ability to understand how personal data is being used in enterprises businesses. In this data-driven economy, the number of cyber-attacks continues to grow in frequency and severity. Heightened regulatory requirements and increasing risk of cardholder data breaches create a feeling of uncertainty for companies, merchants, payment processors, and acquirers who need to protect payment data anywhere it moves, anywhere it resides, and anywhere it's used. Companies of all types are being targeted—anyone with valuable data for attackers to steal.

Integration allows end-to-end data protection

HPE Security – Data Security provides a comprehensive data-centric approach to data protection that addresses the security and privacy needs for data at rest, in motion, and in use through the combined, integrated solutions of HPE SecureData and HPE Atalla Hardware Security Module (HSM). By joining data-centric data protection with a tamper-reactive hardware security module, companies are able to neutralize data breaches by protecting data and rendering it useless to attackers.

Unique Capabilities

HPE SecureData provides an end-to-end data-centric approach that enables you to protect data over its entire lifecycle—from the point at which it's captured, throughout its movement across your extended enterprise, all without exposing live information to high-risk, high-threat environments. HPE SecureData enables companies to neutralize breaches and render data useless using breakthrough, proven encryption, tokenization, and stateless key management solutions—protecting sensitive

data such as personally identifiable information (PII), protected health information (PHI) and credit card and social security numbers. HPE SecureData leverages HPE Format-Preserving Encryption (FPE) and Secure Stateless Tokenization (SST) to enable companies to securely protect data while maintaining its usability and referential integrity for data processes, applications, and services preventing costly database schema changes. HPE SecureData has a unique key derivation and management infrastructure called HPE Stateless Key Management that dynamically derives keys on demand after authorization as needed for data protection. HPE Stateless Key Management eliminates the need to store or manage keys and seamlessly integrates with existing identity management and authorization systems to provide policy-based access to data.

HPE Atalla HSM integrates seamlessly with HPE SecureData to enable enhanced protection of the underpinning cryptographic secrets and key materials needed for key derivation, data encryption, de-identification, and masking. HPE Atalla HSM manages the system-level keys used for key derivation within a hardened FIPS 140-2 Level 3 device, eliminating the risk of exposure or compromise. Cryptographic encryption or decryption and key derivation are performed within the secure boundary of the HPE Atalla HSM inside the tamper-reactive security module environment.

HPE Atalla HSM provides differentiated capabilities for the security market such as a flexible approach to HSM configuration and key management. This is especially important where flexible and compliant solutions are required to manage HSM configuration in lights-out facilities while meeting Payment Card Industry (PCI) Dual Control Requirements. Additionally the Atalla HSM offers robust backup and restore capabilities where a policy can be set to specify that M of N cards must be required for restore. This approach provides increased robustness and policy control around recovery of sensitive encryption keys and configuration data.

The integrated solutions of HPE SecureData and HPE Atalla HSM enable organizations to quickly pass audits and additionally implement full end-to-end data protection. The integrated solution helps to reduce risk impact of data breaches, all without the IT organization having to completely redefine the entire infrastructure and IT processes or policies. It protects information in compliance with PCI Data Security Standard (PCI DSS), Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA), state and national data privacy regulations, as well as the European Commission's General Data Protection Regulation (GDPR). The integrated solution enables companies to comply with the more stringent PCI DSS v3.2 requirements on transport encryption.

The Benefits of a Complete HPE Security Solution

The integrated HPE SecureData and HPE Atalla HSM solution provides a comprehensive data security offering to address the challenges of end-to-end data protection through data-centric security, and ensures the security from an HSM by safeguarding and managing system-level encryption keys. By selecting the integrated HPE Security solution, companies are able to reduce deployment and configuration time by streamlining setup through a centralized management console, while reducing risk and demonstrate PCI DSS compliance and PII or PHI data protection with a single vendor solution. Additionally, Hewlett Packard Enterprise has more than 35 years' expertise in data protection, security, and cryptographic performance.

Customers demand data-centric security that is powerful yet simple to deploy and administer. This HPE Security solution offers the power of HPE SecureData plus HPE Atalla HSM in a simple, integrated, easy-to-buy, easy-to-install-and-configure, and easy-to-administer solution.

Unique Differentiators

- ✔ Single point of purchase – HPE Security.
- ✔ Centralized configuration for management of HPE FPE keys.
- ✔ Industry leading data-centric security combined with root of trust to store your most sensitive secrets.

Trish Schaefer Reilly has over 18 years of product marketing and management experience. She has a broad range of expertise in marketing, defining and managing varied technology platforms including: security, data storage, encryption, key management, big data, analytics, virtualization and cloud services for the enterprise and channel for multiple industry verticals. Trish has played a prominent role in building demand and resources within database, security and licensing organizations focusing on the protection of data. Trish brings a unique, broad perspective both to the challenges facing the industry today and the difficulties that experience making critical technology decisions.

Hardware vs Software Data Replication for Business Continuity

Keith Evans >> Shadowbase Product Management >> Gravic Inc.

Introduction

Today, businesses with access to real-time online transactional data have a competitive advantage. To gain the greatest benefit from this data it must be current and available at any given time. The counter to this advantage is that the inability to access or update current data denies service to users and carries significant business costs, possibly measured in many thousands of dollars per second. These requirements necessitate an IT infrastructure that is continuously available.

Business continuity encompasses activities that an enterprise performs to maintain timeliness, consistency, and availability of its data, operations, and services. Application availability depends upon the ability of IT services to survive any fault, whether it is a server failure, a network fault, or a data center disaster. Data availability depends on the existence of up-to-date backup data copies. *Data replication* is an enabling technology for achieving high or continuous availability for application services and the timely backup of important data. There are two primary data replication technologies, hardware replication and software replication. Each of these technologies, and the differences between them, are discussed in this article.

Data Replication – The Fundamental Force Behind Business Continuity

Improving availability via data replication depends upon having at least two nodes (or disks), each capable of hosting data. As shown in Figure 1, the purpose of data replication is to keep target data synchronized in real-time with source data that is being updated by a source application.

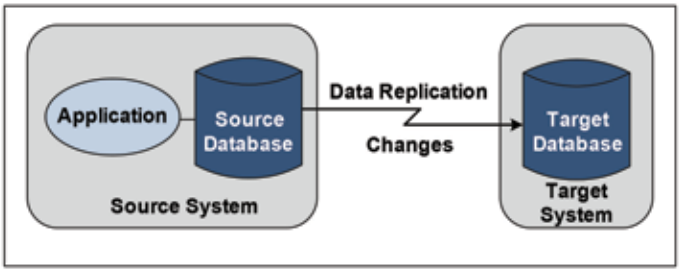

Figure 1 – Data Replication

The *source database* is hosted by the source node and the *target database* is hosted by the *target node*. The two nodes comprise the distributed data processing system. As an application makes changes (inserts, updates, and deletes) to its local source database, these changes are immediately sent by some means to the target system, where they are applied to the target database, which typically resides on another independent node. Because the target database is kept synchronized with the source database, if the source system becomes unavailable, processing can continue using the target system, maintaining service availability. The means by which the data is replicated between the disks on the source system and the disks on the target system falls into one of two categories, hardware replication or software replication.

Hardware Data Replication

Hardware replication is usually implemented in the storage system controller, which replicates disk blocks to a target disk as they are written to the source disk (Figure 2). If a failure occurs in the

source system to which the source disk is attached, a backup system to which the target disk is attached can take over processing.

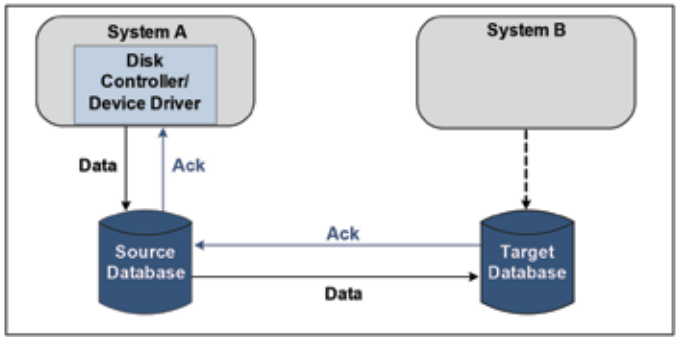

Figure 2 – Hardware Replication – On Disk Write

However, disk blocks are typically only written to disk when they are flushed from the disk's cache. There is no logical order to the disk-write sequence since other factors control cache flushing. The recent least used disk blocks are flushed to disk when cache space is needed for new blocks that must be read from disk. As a consequence: the target disk is not guaranteed to be consistent with the source disk; target disk blocks may be partially split; indices may exist without the rows or records to which they refer; and children may exist without parents. The data is consistent in cache, but the target disk image is generally useless. As a result, applications cannot use the target database for any application processing. Because of this inconsistency, if the source node fails, a lengthy recovery process is required to bring the target database into a useful, consistent state, which extends the period of service unavailability. Additionally, if synchronous replication is used, large amounts of data may be lost due to a source-system failure, as any data still in cache will not have been flushed nor replicated at the time of failure.

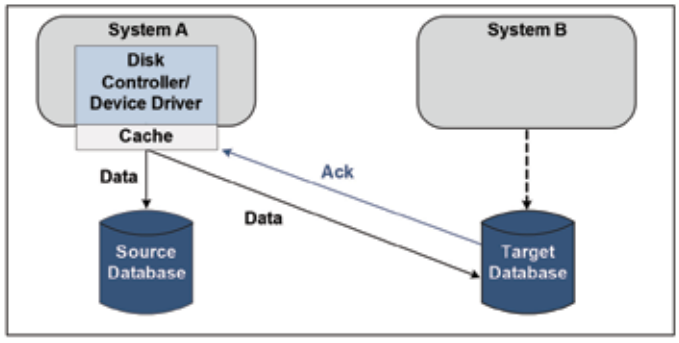

Figure 3 – Hardware Replication – On Cache Write

Some storage controllers replicate changes as they are made to a disk's cache regardless of whether or not they have been physically written to the source disk (Figure 3). The replication of cache updates ensures the logical consistency of the target database, since changes are replicated to the target system as soon as they are made at the source system. If synchronous replication is used, no data will be lost following a source system failure. However, due to other limitations (as described next), the target database may still not be useable by applications while replication is occurring.

Whether based on disk flushing or cache updates, hardware replication typically sends blocks of changes to the target. In some cases, the controller compresses data to only the changed bytes. In other cases, entire data blocks are sent, which requires high communication bandwidth and co-location of source and target disks. Both hardware replication techniques typically do not replicate current data locking protocols nor transaction end-state information (commits and aborts). Hence, the target database typically contains many "dirty records" and cannot be used by applications, even for read-only activities.

Hardware replication generally requires identical storage technology, including the version level, to be used at both the source and the target. This requirement means it cannot be used to integrate diverse systems and applications, to eliminate islands of information and implement new business services. Further, if any component is required to be upgraded (for example, to fix a fault), all components must be upgraded at the same time (or else the fault would bring down both components).

Hardware replication also typically does not allow the target database to be opened by applications at the same time that replication is taking place, thus preventing the use of hardware replication for active/active systems[1], or for enabling read-only activities to be performed on a backup system. Consequently, hardware replication is not an option in order to achieve recovery times measured in seconds or minutes, or for maximizing system utilization.

Another issue with hardware replication is that the maximum distance between the source and target disks is limited. Having the target disk insufficiently distant from the source disk increases the chance that a local area incident (e.g., a flood) will affect both the source and target disks and prevent a timely recovery.

Software Data Replication

Software replication may take place by: event, transaction (several events all treated as a single unit of work), request, or *log-shipping* (which is discussed further in the following section):

- *Event replication* replicates data-manipulation language (DML) events as they occur. DML events include insert, update, and delete operations. In some cases, event replication may also replicate data-definition language (DDL) operations that affect the database's data structure and schema.
- *Transaction replication* replicates entire transactions, either one operation at a time as they occur, or as a group of operations once the transaction has committed on the source. When replayed at the target, the transaction is either committed or, if the entire transaction is not received, is aborted.
- *Request replication* replicates the entire application request, which is reprocessed in its entirety by the application running on the target system.

In this article we are concerned about achieving the highest levels of replication performance, application service availability, data consistency, and minimizing data loss when a failure occurs. Of the various modes of software replication, *transaction replication* best meets these requirements. In transaction-level software replication, a data replication engine running on the source and target systems performs the replication task. The *data replication engine* is driven by a queue of database change events which are read on the source system and sent to the target system and applied

[1] Active/active systems spread the processing load over multiple environments, allowing for instantaneous takeover when one of the systems or environments fails, and lead to the highest Recover Time Objective (RTO) attainable. Active/active architectures are referred to as continuously available and disaster tolerant, and achieve the best RTO possible when implemented with transactional software data replication. See the Gravic white paper, *Choosing a Business Continuity Solution to Meet Your Availability Requirements* (http://shadowbasesoftware.com/white-papers/) for more information.

to its database. This form of replication is able to replicate at the transaction level because transaction control information is also replicated, and updates are applied to the target as transactions. That is, either all updates in a transaction are applied, or none are, thereby preserving source transaction consistency at the target database. In addition, the updates are applied to the target system in the same order as they were generated on the source system. As a result, the target database can always satisfy all of the requirements of referential integrity and database consistency, and consequently can be used by other applications for both read and write operations (the latter in an active/active architecture).

Unlike hardware replication, software replication works between heterogeneous systems and databases. This capability enables the integration of diverse applications and data, elimination of islands of information, and implementation of new business services, such as real-time business intelligence.

As with hardware replication, software replication may be asynchronous or synchronous. An *asynchronous data replication* engine is completely transparent to the applications running in the source node. As shown in Figure 4, it extracts changes made to the source database from a database change queue[2] and sends them after-the-fact to the target database.

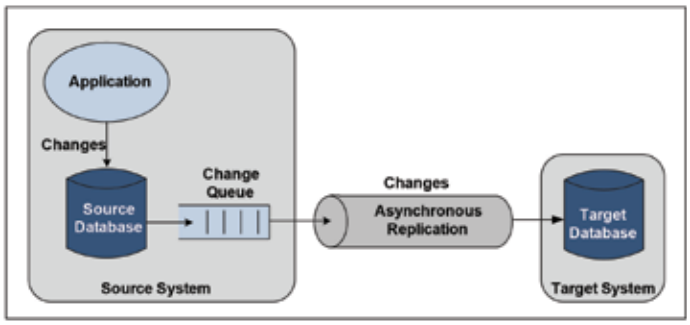

Figure 4 – Asynchronous Software Replication Engine

Unlike some forms of hardware replication (disk flushing), *synchronous software replication* guarantees that no data will be lost after a failure. Using a technique known as "coordinated commits"[3] (Figure 5), synchronous software replication makes no permanent changes to any database copy unless these changes can be applied to both source and target database copies. With coordinated commits, the replication engine participates in the source application's transaction, and at commit time, it votes "yes" or "no" dependent upon whether all the updates in the transaction have been replicated to the target system. If "no," then the source transaction aborts. It is guaranteed that all participating databases, source and target, received and/or applied the same updates, or none did; therefore no data will be lost in the event of a source system outage. Another major benefit of coordinated commit technology is that while it guarantees no committed data will be lost, impacts to application throughput are also minimized as synchronization only occurs at transaction commit time, and not on every database change event.

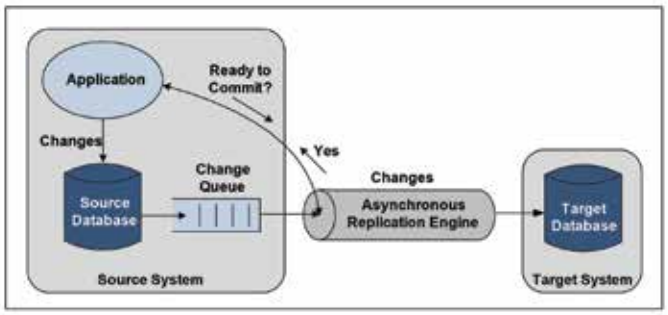

Figure 5 – Synchronous Software Replication using Coordinated Commits

Software Data Replication – Log-Shipping

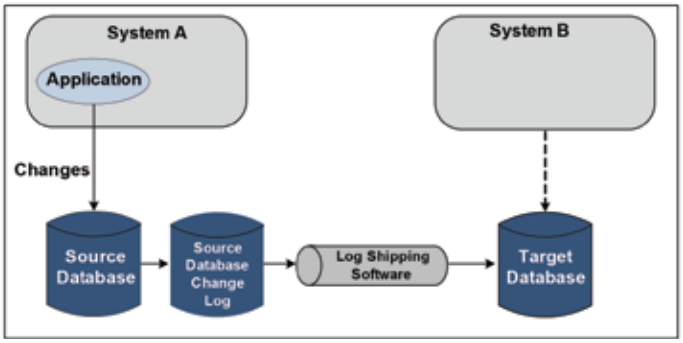

Figure 6 – Software Replication – Log-Shipping

Log-shipping (Figure 6) is a form of software data replication that operates more like hardware replication. Log-shipping sends the entire source database change log periodically to a target system. On the target system this change log is read and the data and index blocks are applied against the physical database structure, which is analogous to how hardware-based data replication works; it also has all of the same issues.

Compared to software transaction-based data replication, log-shipping has a higher amount of data loss on failure (a higher Recovery Point Objective, or RPO), because the target database is only current to the point of the most recent log-ship, which may only be when the source log file has been closed. Any subsequent changes on the source system are lost. Log-shipping data is usually applied to the target as index and data blocks of changes, which results in an inconsistent target database, rendering it practically useless while replication is taking place, as is the case with hardware replication. Even if the change events are extracted from the log and applied to the target, source transaction consistency is typically not maintained; events are applied without regard to source transaction bracketing. The result is an inconsistent and unusable target database.

Comparison of Hardware vs Software Data Replication

Although hardware replication (including software-based log-shipping) appears to offer a simple and cost-effective approach to maintaining data and service availability, it is often not a good business continuity solution for the reasons previously discussed. To summarize:

- Since there is no concept of transaction boundaries, database consistency or referential integrity, hardware replication cannot be used for active/active systems (i.e., it cannot provide continuous availability).

[2] A database change queue is a DBMS-maintained list of transactional insert, update, and delete operations that have been performed by applications against the database.
[3] For more information on synchronous replication and coordinated commits, see Dr. Bill Highleyman, Paul J. Holenstein, Dr. Bruce Holenstein, Chapter 4 – Synchronous Replication, Breaking the Availability Barrier: Survivable Systems for Enterprise Computing, AuthorHouse; 2004.

- The backup disk is highly inconsistent due to missing data not yet replicated, and is consequently not usable for query/reporting or other functions.
- The primary and backup systems must generally use identical database hardware and software. All upgrades must be simultaneously applied to all hardware components, thereby increasing the risk that a fault in one component will affect all components.
- The source and target systems must be homogeneous, and integration of diverse systems, applications, and data is not possible.
- Recovery from a failure is a complex and lengthy task, requiring data "fix-up" on the backup disk, which leads to long recovery times and service unavailability.
- A significant amount of communication bandwidth is required since whole disk blocks rather than individual rows are typically replicated.
- Data corruption of the source database is replicated to the target, perhaps preventing it from being opened for recovery.
- Distance between disks is physically limited (typically about 100 km), which increases the likelihood that a local area incident could affect both disks and prevent recovery.
- If synchronous replication is used, data can still be lost when a failure occurs due to the cache flushing issue.

By contrast, none of these issues affect software replication:

- Primary and backup databases are consistent and maintain transactional and referential integrity. Software replication therefore enables active/active systems which deliver continuous service availability.
- Because software replication maintains database consistency, backup systems can be used for query/reporting and other online activities, while replication is taking place.
- Primary and backup systems can be completely different (heterogeneous). The platforms, operating systems, database software, and database structure can all be different. The data replication engine takes care of managing the necessary data transformations.
- Software replication can be used for the integration of different applications and data, enabling the implementation of new business services.
- Because the backup database is transactionally consistent and ready for use at any time, there is no need for a complex takeover process (database "fix-up" to bring it into a consistent and usable state), and recovery times as low as sub-seconds are possible.
- In an active/active architecture many users will not see an outage, and recovery is simply a matter of re-routing users from a failed node to an active system.
- Because only row change data is sent between systems, much less communications bandwidth is required for software replication.
- Software replication replicates changes described by a source transaction log. Such changes are executed completely independently on target systems, thereby avoiding the mirroring of corruption from the source database to the target database.
- Software replication has no physical distance limits between nodes, which can be positioned sufficiently far apart to ensure continued service, including an outage incident that affects a wide area.
- Synchronous software replication guarantees that all data associated with committed transactions is replicated, and hence no data will be lost in the event of failure.

Summary

On the surface, hardware data replication appears to offer a simple and cost-effective solution to the problem of maintaining data and service availability in the event of a system outage. Scratch the surface however, and it becomes clear that hardware data replication suffers from many significant issues which make it unsuitable for this task for mission-critical applications. Simply put, the likelihood of a timely recovery from an outage with minimal data loss is very low, and perhaps not possible at all. Furthermore, backup system capacity is wasted since the replicated data is inconsistent and unusable. It will only take one incident for it to become apparent that hardware replication is not cost-effective and does not enable the highest levels of service availability with minimal data loss. While suitable for some tasks, hardware data replication is inadequate for supporting business continuity of mission-critical applications.

Conversely, software-based transactional replication suffers from none of the issues which afflict hardware replication, including:

- distances are not limited;
- source and target databases are consistent;
- recovery is simple, fast, and repeatable;
- continuously available active/active architectures are supported;
- zero data loss is provided when running in synchronous mode;
- backup databases can be used for productive work;
- source and target systems can be completely heterogeneous.

For mission-critical applications, the highest levels of service availability and protection against data loss is required. A business continuity architecture built on software-based transactional data replication is the only viable solution to meet this requirement.

Shadowbase Software Data Replication

The Shadowbase product suite from Gravic, Inc. provides the full range of software data replication features to satisfy the most demanding IT business continuity and other replication requirements, including:

- active/active continuously available business continuity architectures to eliminate *unplanned downtime*;
- synchronous replication technology for zero data loss (ZDL) when disasters occur;
- zero downtime migration (ZDM) capability to eliminate *planned downtime* for upgrades and migrations;
- data and application integration between heterogeneous systems.

The Gravic white paper, *Choosing a Business Continuity Solution to Meet Your Availability Requirements* contains more information on the subject of hardware versus software data replication, and the requirements to consider in choosing the best solution to meet your business continuity needs.

Keith B. Evans works on Shadowbase business development and product management for Shadowbase synchronous replication products, a significant and unique differentiating technology. Asynchronous data replication suffers from certain limitations such as data loss when outages occur, and data collisions in an active/active architecture. Synchronous replication removes these limitations, resulting in zero data loss when outages occur, and no possibility of data collisions in an active/active environment. Shadowbase synchronous replication can therefore be used for the most demanding of mission-critical applications, where the costs associated with any amount of downtime or lost data cannot be tolerated. For more information and the availability of Shadowbase synchronous replication, please email sbproductmanagement@gravic.com.

HILTON WORLDWIDE | CONNECT+

HILTON ELEVATED DISCUSSIONS by Hilton Worldwide

We're very pleased to announce the long-anticipated launch of **Hilton Elevated Discussions**, a growing collection of short videos of customers and Hilton Worldwide subject experts sharing insights on specific topics that are relevant to the Connect+ audience. Currently we have 25 clips that include the ROI of meetings and events, procuring grants and sponsorships, negotiating Wi-Fi, effective contracting, planning hybrid live/virtual meetings, and much more.
To view the content, please visit **Hilton Elevated Discussions**

To learn more about Hilton Worldwide's unique suite of customer solutions, please contact:

Rocco LaForgia
Director of Sales Technology Associations
p. 212-820-1715 e. rocco.laforgia@hilton.com

Hilton HOTELS & RESORTS · WALDORF ASTORIA HOTELS & RESORTS · CONRAD HOTELS & RESORTS · CURIO A COLLECTION BY HILTON · DOUBLETREE BY HILTON · EMBASSY SUITES · Hilton Garden Inn · Hampton · HOMEWOOD SUITES BY HILTON · HOME2 SUITES BY HILTON · Hilton Grand Vacations · HILTON HHONORS

Fast Track Your NonStop Integration with XYGATE and Active Directory

Steve Tcherchian, CISSP >> **CISO** >> XYPRO Technology

Over the last 15 years, I've worked in a variety of environments - from dining room startups with a handful of systems, to multinational corporations with tens of thousands of systems. One thing I could almost always be sure of is Microsoft Active Directory (AD) is somewhere in the mix.

For most organizations, Microsoft Active Directory is the main corporate directory service for managing access to information systems across the enterprise. Active Directory serves as the central authority for network security, distributed resources, systems, users and services. When a user logs into a Windows workstation, their credentials are submitted to and evaluated by AD to determine if they are authorized for access, to which groups they belong, to which resources they have access, and which policies are enforced, along with other controls. This process enables a user to have a single set of credentials controlling their access to corporate systems their resources.

What makes Active Directory even more powerful and a staple in nearly 90% of Fortune 1000 organizations is its extensibility, scalability and integration capabilities. AD can communicate over Lightweight Directory Access Protocol (LDAP) and Kerberos, two standard application protocols. These standards extend AD's capabilities by integrating non-Windows platforms and enterprise applications. This centralizes user information in a single repository, making it available to multiple platforms and applications.

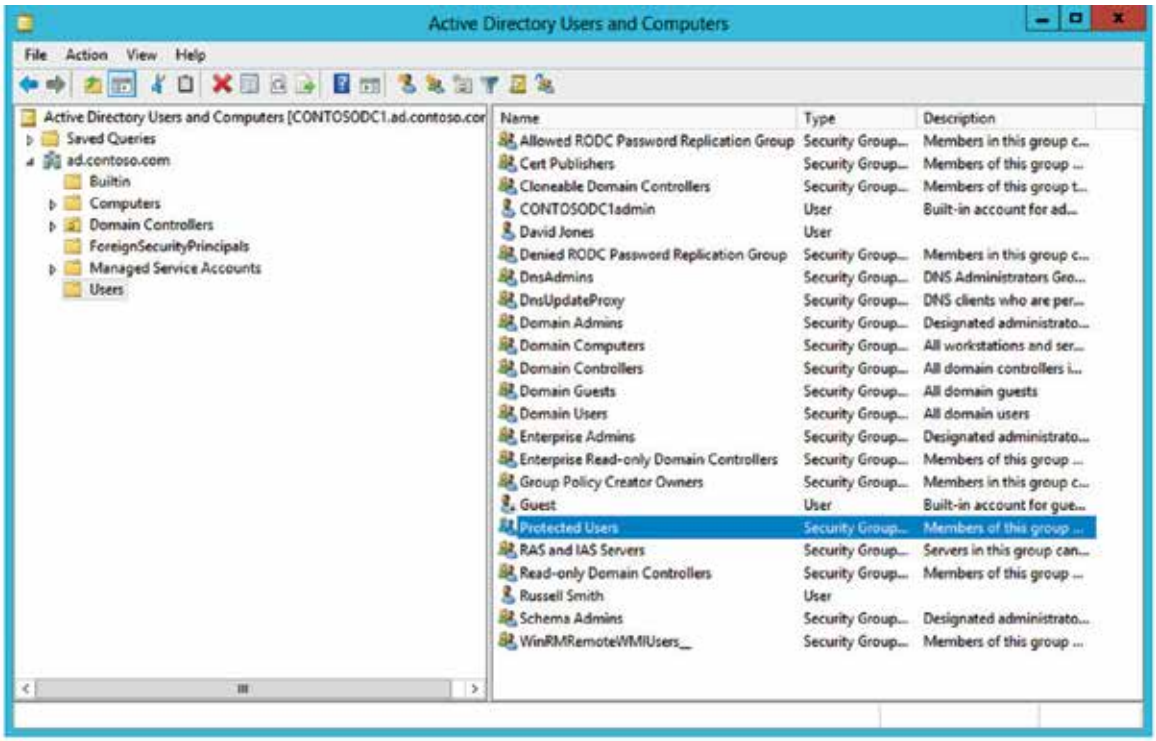

This central management not only comes with a remarkable reduction in administration costs, but facilitates compliance with security frameworks and enhances overall security right out of the gate. Security is at the forefront of everyone's mind, therefore centralized user management is one of the necessary steps towards proper risk management.

The Challenge

For years, the HPE Integrity NonStop server has been the box in the corner humming along, reliably and effectively handling its mission critical workload. Identity and access was managed by an administrator who would manually add users to each NonStop system, assign individual user ids and passwords, and manage user updates on a case by case basis. This not only created operational overhead but potentially exposed organizational risk to users, administrators and the enterprise. Users had to remember (or store in an unsecured Excel file on their desktop) yet another set of credentials. Administrators were burdened with managing user accounts separate from the rest of their enterprise, and corporate policies had to be enforced on multiple platforms with different parameters to configure.

The challenge doesn't stop there. When an employee leaves a company, organizational workflows disable a user's Active Directory credentials. This in turn disables access to any system or application integrated with AD. In most cases, the NonStop is often still handled via a separate, most often, manual process. Unless the NonStop administrator was notified of the departure and manually acted on the request in a timely fashion, the risk of a valid and active user ID remaining on the system long after the user has been removed from the main corporate directory exists. I don't think I need to explain why this is a major security concern.

NonStop in the Modern Enterprise

HPE recognized the need for integration of the NonStop server with the rest of the enterprise. To address this evolving requirement, XYPRO and HPE partnered to offer XYGATE User Authentication (XUA) on all HPE Integrity NonStop servers shipped since 2013. XUA enables enterprises to efficiently integrate their NonStop user environment with their Active Directory infrastructure. Other platforms such as Linux, Unix and Windows were already capable of integration with the corporate directory. XUA integration with AD strengthens user authentication on the NonStop and supports enforcement of corporate password policies while reducing the costs of user provisioning and management.

XUA features include:

- Enterprise SSO participation through LDAP and Active Directory client interfaces
- Log-on controls based on ancestor program, requester program, port or IP address, time of day or day of week, or current logged-on user
- User impersonation support to reduce the need for sharing sensitive user passwords, for example—ability to log on as SUPER.SUPER but provide the individual user's password
- Authentication controls customized at the user or group level
- Enhanced log-on event audit collection and reporting capabilities
- Integration with SIEM solutions through XYGATE Merged Audit

Policy Management

In most environments, the corporate password policy is typically defined through the Active Directory Group Policy Object (GPO) and pushed out to users, systems and applications. Any application integrated with Active Directory will have their password policy governed through this central authority.

With XUA, NonStop userIDs take advantage of the same password policies set in a single location. Active Directory is the system of record for all password policy management.

Active Directory – Getting Started

Preparing Active Directory for NonStop integration requires very little setup on the Active Directory side. There is no client or agent to install. You will only need the following 4 items.

1. Domain Controller and LDAPS Connectivity
2. LDAPS Certificate
3. NonStop Users Container
4. LDAP Bind User and Password

Domain Controller and LDAPS Connectivity - Don't be Lazy!

As tempting and easy as it may seem to use traditional unencrypted LDAP for connectivity – DON'T! Integration as critical as this should always be done using LDAP over SSL (LDAPS). This ensures no credentials are sent from the NonStop to Active Directory in clear text. The default port for unencrypted LDAP is 389, while encrypted LDAPS uses 636. In most cases, LDAPS will already be enabled, but in case it is not, a server with the Certification Authority role is required within your AD environment. LDAPS connectivity can be verified by using an LDAP browser like Softerra or JXplorer.

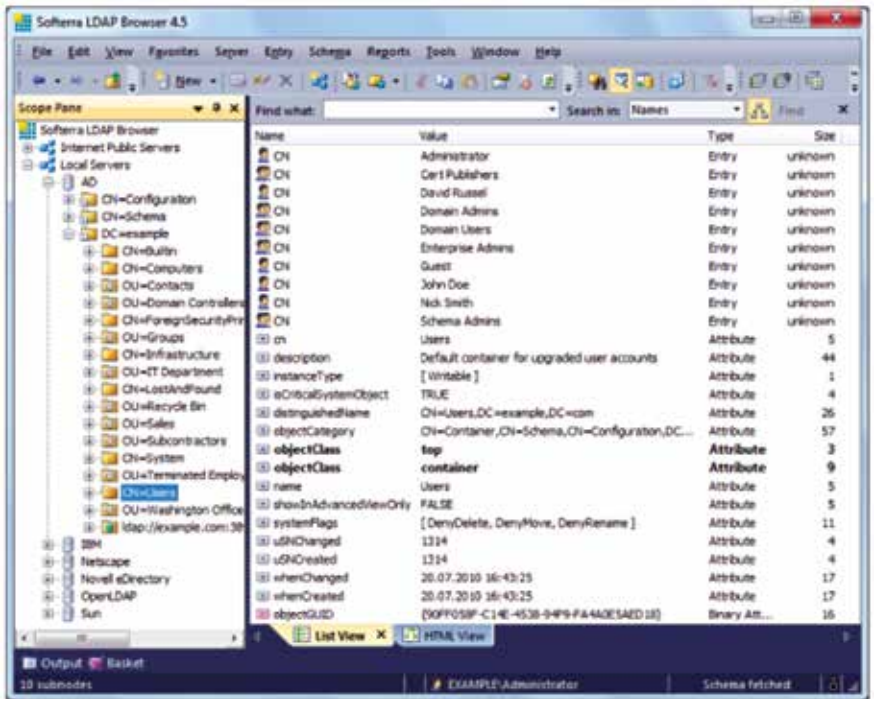

LDAPS Certificate

Once connectivity is verified, you will need an export of the entire certificate chain used for LDAPS. This is typically provided by the Active Directory administrator. This important step enables LDAPS and generates the certificate chain that will be required by the NonStop.

NonStop Users Container

User IDs within Active Directory are organized within Organizational Units (OU) and Containers (CN). This allows for separate GPOs to be assigned to different OUs. The Active Directory administrator will need to identify which CN or OU the NonStop User IDs are located within the directory. For example, to locate the OU for user John Pierce (jpierce), the AD administrator would run the following command in a Windows command prompt

```
dsquery user -name jpierce
```

which returns

```
"CN= jpierce,OU=NonStop Users,DC=MYDOMAIN,DC=COM"
```

This shows us that our NonStop users are located in the "NonStop Users" OU in the MYDOMAIN.COM Active Directory domain.

LDAP Bind User and Password

The last requirement is to have access to an Active Directory user that will be used for directory lookups. This is typically known as an LDAP BIND user. It is recommended that the bind user be given just enough privileges to perform the required lookups within the directory and nothing more. This should NOT be an administrator account.

HPE NonStop Integration with Active Directory

The installation of XYGATE User Authentication (XUA) on the NonStop is not a difficult task, but one that should be undertaken with the knowledge and understanding of what is being accomplished. XUA is installed as the Safeguard Event Exit Process (SEEP) for authentication events. If an existing authentication SEEP is installed, it must be

NonStop Technical Boot Camp

SAVE THE DATE

November 13-16, 2016

Fairmont San Jose Hotel | San Jose, CA

- 50+ Technical Breakout Sessions
- Executive Keynotes, Roadmaps and MORE!
- Partner Pavilion and Partner Networking Events
- Downtown San Jose Venue

REGISTER ONLINE TODAY!

www.connect-community.org/nonstop-technical-boot-camp

removed prior to installing XUA.

The following information is required to complete the installation of XUA/Active Directory integration on the NonStop for secure communications:

1. The IP address or name of the Domain Controller running the LDAPS service
2. The LDAPS service version (must be 3 or later)
3. The LDAPS service type (Windows for Active Directory)
4. The name of the TCP/IP process on the NonStop to be used for Active Directory communications
5. BIND information (outlined above) for fully qualifying user names
6. The SSL certificate chain
7. The port number of the LDAPS service (if different from the default)
8. Access to SUPER.SUPER (or an alias) to complete the installation

Once you have this information, run the XUA installation process to capture this information and setup the NonStop configuration.

```
LDAP_HOST                dc01.mydomain.com
LDAP_TYPE                WINDOWS
LDAP_PORT                636
LDAP_PROXY_IP_PROC       $ZTC0
LDAP_PROXY_LOG           /G/log/location
LDAP_PROXY_CACERT        "BOTHCERT"
```

Additional entries needed if LDAP lookup is being used:

```
LDAP_LOOKUP              ON
LDAP_USERS_CONTAINER     "dc=mydomain,dc=com"
LDAP_LOOKUP_ATTRIBUTE    "samaccountname"
LDAP_DN_ATTRIBUTE        "dn"
LDAP_SEARCH_USER         "cn=ldap,cn=NonStop
                           Users,dc=mydomain
                           dc=com"
```

Once the install process is complete, users can be added to the configuration for optional or required Active Directory authentication.

In Summary

Confirming and controlling the identity of users accessing your system is crucial to protect your systems and data. The native HPE Integrity NonStop operating system and its Safeguard security infrastructure provide unique identification for users through Guardian user IDs and aliases, both with 64-character—strong password and passphrase support. However, until XUA became available through HPE they did not offer direct integration into SSO environments.

XYGATE User Authentication (XUA) is normally ordered with all new HPE Integrity NonStop i systems and is included in all HPE Integrity NonStop X systems. XUA allows for integration into an enterprise's Active Directory and SSO environment, simplifying provisioning and management of NonStop users. Users can access all of their authorized systems, including NonStop servers, using a single user ID and password. The user benefits from the simplicity of authentication while the administrator benefits from a reduced user ID and password maintenance burden. Overall, security is improved and risk and costs are reduced for the enterprise with XUA and Active Directory.

Steve Tcherchian, CISSP, PCI-ISA, PCI-P is the CISO and SecurityOne Product Manager for XYPRO Technology. Steve is on the ISSA CISO Advisory Board and a member of the ANSI X9 Security Standards Committee. With almost 20 years in the cybersecurity field, Steve is responsible for XYPRO's new security intelligence product line as well as overseeing XYPRO's risk, compliance, infrastructure and product security to ensure the best security experience to customers in the Mission-Critical computing marketplace.

Blockchain and NonStop

>> Justin Simonds
>> Master Technologist
>> Americans Enterprise Solutions

A little over 18 months ago a customer was asking about Blockchain and if NonStop was investigating the technology. At that time almost everyone equated Blockchain with Bitcoin. Of course Bitcoin remains the preeminent example for Blockchain but equating the two is somewhat like equating TCP/IP with the applications that run over it. Blockchain is the foundational technology that makes Bitcoin work but has applicability in various industries and verticals. The big deal about Bitcoin was that it was issued and backed not by a central authority but by automated consensus among the networked users. Furthermore it did not require those users to trust each other. The blockchain portion was the enabler. Since that customer enquiry members of the ES&A (Enterprise Solutions and Architecture) team have been having spirited debates over NonStop blockchain use cases.

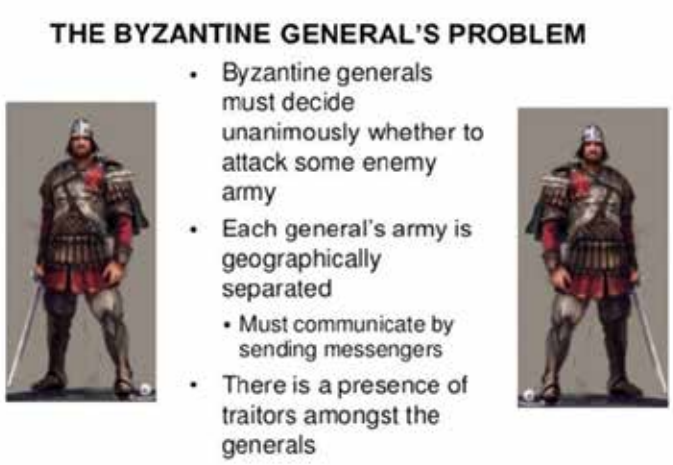

So what is Blockchain? The simplest answer to that is "a distributed ledger" but that probably won't be sufficient. The original blockchain definition was published by someone using the name Satoshi Nakamoto in 2008 and was implemented in 2009 as the original source code for Bitcoin. For one thing blockchain purportedly solves a very sticky problem known as the Byzantine Generals Problem (see figure 1). The problem is how do a group of generals, each with an army, agree to a consistent plan, such as, they all attack or all retreat? It doesn't sound like too hard a problem but remember this was in the days before instant communication. The generals receive orders to attack or retreat from messengers. The orders may be legitimate or they may have been altered by the messenger (spy) or a disloyal general. So how do they determine if an order is legitimate? For a good overview of the problem and issues go to http://research.microsoft.com/en-us/um/people/lamport/pubs/byz.pdf . You may recognize that some of these same issues applied in the design of the Tandem systems, which had to solve similar concerns of consensus and trust.

Blockchain also solves another very interesting issue, which is decentralization. One of the fascinating things about Bitcoin is that there is no financial institution associated with the currency. It is completely decentralized and its very openness (anyone can join) creates the trust. Bitcoin operates using miners who create (mine) blocks to be added to the chain. This intentionally requires a lot of computational work and if the miner is successful, by being the first one to create a valid Block, he is paid a percentage of a Bitcoin. The work of the miner must be verified by additional miners (consensus) before it is written to the chain and therein lies the trust. There would have to be collusion on a pretty massive scale (51%) to successfully corrupt the blockchain. Also each Block is chained to the previous Block by an elaborate cryptographic hash so any tampering with any previous blocks would invalidate the chain. Many (and any) user may copy the entire blockchain and since there are so many in circulation a malicious hacker changing one instance would be found out quickly and that blockchain invalidated. So again, it is a method to allow trust between two parties that do not know each other and do not use an intermediary or central authority. For a very complete description of this I suggest the book "Mastering Bitcoin" by Andreas Antonopoulos.

So far these things all apply to public chains and are not a good fit for NonStop since availability is based on many copies of the chain. Integrity is based on the mining efforts used on the block creation and the work is verified/validated by several other miners. These were the necessary items needed for Bitcoin to launch its Internet cryptocurrency.

What has become intriguing to many is the idea of a trusted ledger which can facilitate transactions quickly for little to no cost perhaps even across borders. Could blockchain technology become the new Internet transaction settlement architecture in today's markets? Let me quote an article from Bloomberg "When investors buy and sell syndicated loans or derivatives or move money around the world, they must cope with opaque and clunky back-office processes that rely on negotiated contracts between buyers and sellers, lots of phone calls, lots of lawyers, and even the occasional fax. It still takes almost 20 days, on average, to settle syndicated loan trades." See www.bloomberg.com/news/features/2015-09-01/blythe-masters-tells-banks-the-blockchain-changes-everything for the complete article. From the same article "A June report backed by Santander InnoVentures, the Spanish bank's fintech investment fund, estimated the blockchain could save lenders up to $20 billion annually in settlement, regulatory, and cross-border payment costs." There are a lot of monies that cross borders as breadwinners work in remote countries. The transfers can take days and often incur very high fees. Would a private blockchain be a better, cheaper solution?

There have been various and numerous startups looking beyond finance to smart contracts (contracts embedded in a blockchain that are self-activating when specific events occur), proof of existence and provenance for all sorts of assets (cars, houses, real estate, diamonds, art, etc.), personal identity (social security, health records, work history) and smart property (proof of ownership). The potential uses for blockchain are crossing all verticals: Finance, Insurance, Healthcare, manufacturing and energy to name a few. It may also be involved in major forces such as the Internet of Things.

With so much interest by major institutions and so many start-ups being funded by venture capitalists what's wrong with, or what can be improved upon the Bitcoin implementation?

Bitcoin's capacity and speed are governed to a large degree. Bitcoin transactions are confirmed every time miners create a new block for the networks chain. Each block takes about ten minutes to mine, and can hold 1MB of information. The 1MB size will generally hold a number of transactions which will vary since transaction size varies but generally Bitcoin adds about 300,000 transactions per day. We can easily calculate blockchain growth assuming a block is added every 10 minutes. Since there are 1,440 minutes in a day then 144 1MB blocks could be added each day to Bitcoin's blockchain. If we compare these figures with a high-speed credit card processor that may be handling 2,000 to 4,000 TPS or about 150 million transactions per day, the Bitcoin implementation just won't work. It should go without saying but those rates are currently being run on NonStop systems. Bitcoin's proof-of-work mining is well-suited for a general purpose decentralized network, it just wouldn't work for more conventional users and currencies. As mentioned the delay is somewhat unpredictable but at least 10 minutes for a confirmation. Although low, there is also a risk that miners could refuse to confirm transactions for ideological or economic reasons. Again a low risk but the potential exists for a 51% attack, where a group of miners controlling over half of the network's computational power conspire to rewrite the blockchain's history and invalidate and/or steal Bitcoins.

As mentioned, the network is made up of anyone who would like to join and numerous copies of the blockchain are held by many systems. There are a number of miners competing to add blocks (for which they are paid). So how much does a transaction cost? To the users of the system not much at all, generally a small percentage of the transaction itself and not dissimilar to credit card transaction fees and generally lower. What is not calculated are all the participating computers which are quite a massive network. In a manner not dissimilar to Uber and Airbnb, Bitcoin has allocated its IT to the participants of the system. Bitcoin's actual compute resources may be fairly sparse, the real work is being done by the miners. So the true costs are a bit elusive. Currently it is based on the miner's willingness to accept a rather low fee per transaction but that could change and since the miners and not Bitcoin itself control the costs it is unpredictable.

By agreement, all bitcoin transactions are visible to all network nodes and therefore to the entire world since anyone can participate and get a copy of the complete blockchain. This aspect of Bitcoin is somewhat alleviated by the fact that bitcoin addresses cannot easily be connected to real-world owners. But the existence and rate of transactions are not protected, they are public record, and participants run the risk of their identities being revealed at some point in the future. Which means that their entire transaction history could be made public. This could cause some unease in more conservative, security minded thinkers or philandering individuals.

What if the benefits could be achieved and most of the risks reduced or eliminated? This is where we take a look at private chains and where NonStop begins to make a lot of sense.

Private Blockchain technology lends itself to fast, documented, low-cost transactions creating a permanent ledger. This is good for customers, institutions and regulators. The mining and governance done by Bitcoin is there so that peers do not have to trust each other. But if a private chain is established between trusted parties then the elaborate mining exercises and compute intense cryptography would not need to be performed. Simpler, format preserving encryption could be employed to ensure privacy and support PCI and other regulations at a much reduced compute cost.

The unknown costs of replicated computer systems supporting copies of blockchains in Bitcoin could be replaced with a continuously available system hosting the 'golden' or master chain. Additional copies could be made for continuity and analytic purposes but with limited and predictable costs.

What is required for NonStop to participate in this exciting new market? On figure 2 the components are broken out with many of those components already existing on the system. Our node.js port allows access to Restful services and to anything in the open-source arena – which may include some blockchain (BC) APIs, smart contracts and even some rules engines. For the rules engines to execute in minimal time perhaps a programmable gate array (FPGA) would make sense. Items in blue

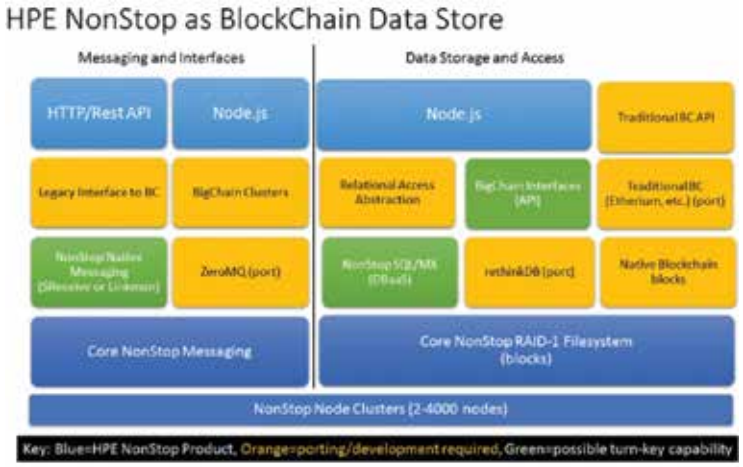

Indicate what we already have and once again the NonStop architecture having a single system image, massively scalable, parallel processing system creates a compelling fit. And why wouldn't you want to host your private master chain on the system with the highest out-of-the-box availability in the industry?

NonStop could address the speed and capacity concerns, the security concerns and the scalability concerns. It could be perfect for enterprises considering the implementation of a private blockchain. If your company is interested in exploring blockchain capabilities, the ES&A team would be very interested in participating in the discussions to explore a possible NonStop fit. Contact your NonStop sales specialist for more information.

Justin Simonds is a Master Technologist for the Americans Enterprise Solutions and Architecture group (ESA), a member of the HPE IT Transformation SWAT team.. His focus is on real-time, event-driven architectures, business intelligence for major accounts and business development. He is involved with HPE Labs on several pilot projects. He has worked on Internet of Things (IoT) initiatives and integration architectures for improving the reliability of IoT offerings. He has written articles and whitepapers for publication on Helion cloud, TCO/ROI, availability, business intelligence, Internet of Things and the Converged Infrastructure. He is a featured speaker at HPE's Technology Forum and at HPE's Executive Briefing Center, at industry conferences such as the XLDB Conference at Stanford, IIBA and the Metropolitan Solutions Conference.

Back for More

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

The theme of this issue is integration and for many in the NonStop community perhaps this will just produce a yawn. Integration? I cannot recall ever hearing of a solution running on a NonStop system being described as just another application silo in the data center. What I have heard openly discussed is the wealth of connectivity options available for NonStop – being a fault tolerant system, where users routinely report having never seen a single outage (for any reason), todays NonStop systems are ideal runtime platforms when it comes to hosting mission-critical applications. Even as the business world is waking up to needing to support an always-connected, always-on, "I need it now" consumer-driven world, 24 X 7, where practically every application has become mission-critical, no one system runs every mission-critical application.

Whether for historical reasons or simply prioritizing some mission-critical applications as even more critical than others, NonStop remains a system that is integrated with the world around it. Within the data center, as you look up and down rows of servers, it may be getting harder and harder to recognize a NonStop system but it's there and it's plugged into a variety of supporting cast members. And where we are headed on this integration journey is to the world of clouds.

However, clouds have their own integration issues. While they support massive scale-out requirements and can provide High Availability (HA) through redundancy and a level of parallelism such that minor outages, limited to a failed server, can be easily masked, shifting up another gear and providing Continuous Availability (CA) in the manner NonStop provides today, well, that will require integration with NonStop, where NonStop becomes an integral part of the cloud providing availability as part of any Infrastructure-as-a-Service (IaaS) offering. Think again about the potential here, for those businesses planning for CA in IaaS, think vNonStop.

There's some interesting reading out in the blogosphere on HA / CA for IaaS. From one Microsoft blog, going by the name of Carpe Datum – yes, seize the data (why didn't I think of that name) - "IaaS involves Virtual Machines (VNs) – so in effect, an HA strategy here takes on many of the same characteristics as it would on-premises. The primary difference is that the vendor controls the hardware, so you need to verify what they do for things like local redundancy and so on from the hardware perspective."

And then comes the surprising conclusion that will generate some amusement amongst the NonStop faithful. "As far as what you can control and plan for, the primary factors fall into three areas: multiple instances, geographical dispersion and task-switching. In almost every cloud vendor I've studied, to ensure your application will be protected by any level of HA, you need to have at least two of the Instances (VM's) running." You need to have at least two VMs – go figure. In all my discussions with NonStop development the gating factor concerning the release of vNonStop is ensuring guidance, possibly templates, to be provided to help avoid configuring vNonStop within just one VM.

The more I read the more I see the market coming to NonStop. Integration? I will give you integration – physically and logically. On real systems and virtual systems! And the value of having the option to run virtually is to ensure no matter the underlying hardware and infrastructure, CA can be supported and where I am seeing the initial interest in vNonStop focuses on running NonStop SQL within the cloud as a DBaaS, this just makes sense. Integration? Businesses that capitalize on building better clouds through HA IaaS featuring vNonStop will be able to service their customers a whole lot better than they are today.

But don't rush to clouds as simply the be-all and end-all solution to infrastructure problems already plaguing your IT. Don't be lulled into a false sense of security over replacing the physical wreck you may have with a logical wreck – virtualization will expose the weakest link in a heartbeat and for most businesses, they won't even see it coming. Farm it out to AWS, Azure and Google and no, you will not see magic – you still have to do the work to ensure CA that matches what so many of us take for granted today with NonStop. Yes, you can roll-your-own NonStop given enough servers, switches and know-how but you cannot maintain it past just one release of the software stack. Top to bottom. It just won't happen, so get over it – install vNonStop and work with it!

Perhaps the industry with the worst infrastructure and yet depending so much on integration with its customers, its employees and many layers of retailers and financial institutions, is the airlines industry and look what has happened to it of late. And why don't airlines run NonStop you ask? A happy mix of Arrogance and Ignorance, I contend – just the latest assignment given to the two letter acronym, AI. A little history too I suspect but ultimately, as USA Today reported on August 12, 2016, in the article Travel trouble? Here's why your airline flight is delayed it's all about infrastructure and integration – layers of integration done poorly.

"While the basic foundation of many airline systems has been in use for decades, complexity, not age, is the real problem, says Lance Sherry, director of the Center for Air Transportation Systems Research at George Mason University. 'So many systems are layered on top of each other that we don't always know who's talking to whom," Sherry said … airlines run multiple, intersecting systems which must flawlessly interact with each other. In the industry, there are at least six — ticketing reservation, aircraft assignment, flight crew scheduling, airport gate assignment, air traffic flow management and irregular operations systems. Often they come from different vendors and use different software languages. And yet they must be synchronized, and timing is split-second and critical."

Integration – we all need it; we all want it; we just don't know much about the how and ignoring the NonStop will not be wise. And everyone I know wonders why on earth I drive rather than catch a plane!

XYGATE® Object Security

A single solution for complete control of Safeguard and OSS security

Sophistication

RBAC for HPE NonStop

- Minimize Security Administration Overhead
- Maximize Security, Control and Audit

"Simplicity is the ultimate sophistication" Leonardo da Vinci

Manage Access to Information

Learn more at xypro.com/XOS

XYPRO Mission Critical Security

The guiding light for your mission critical business
Improve your NonStop'ness. 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 HPE NonStop.

www.comforte.com

comforte®
better always on