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 SQL Database Compatibility

Future Trends and the Relevance of NonStop

Women in NonStop

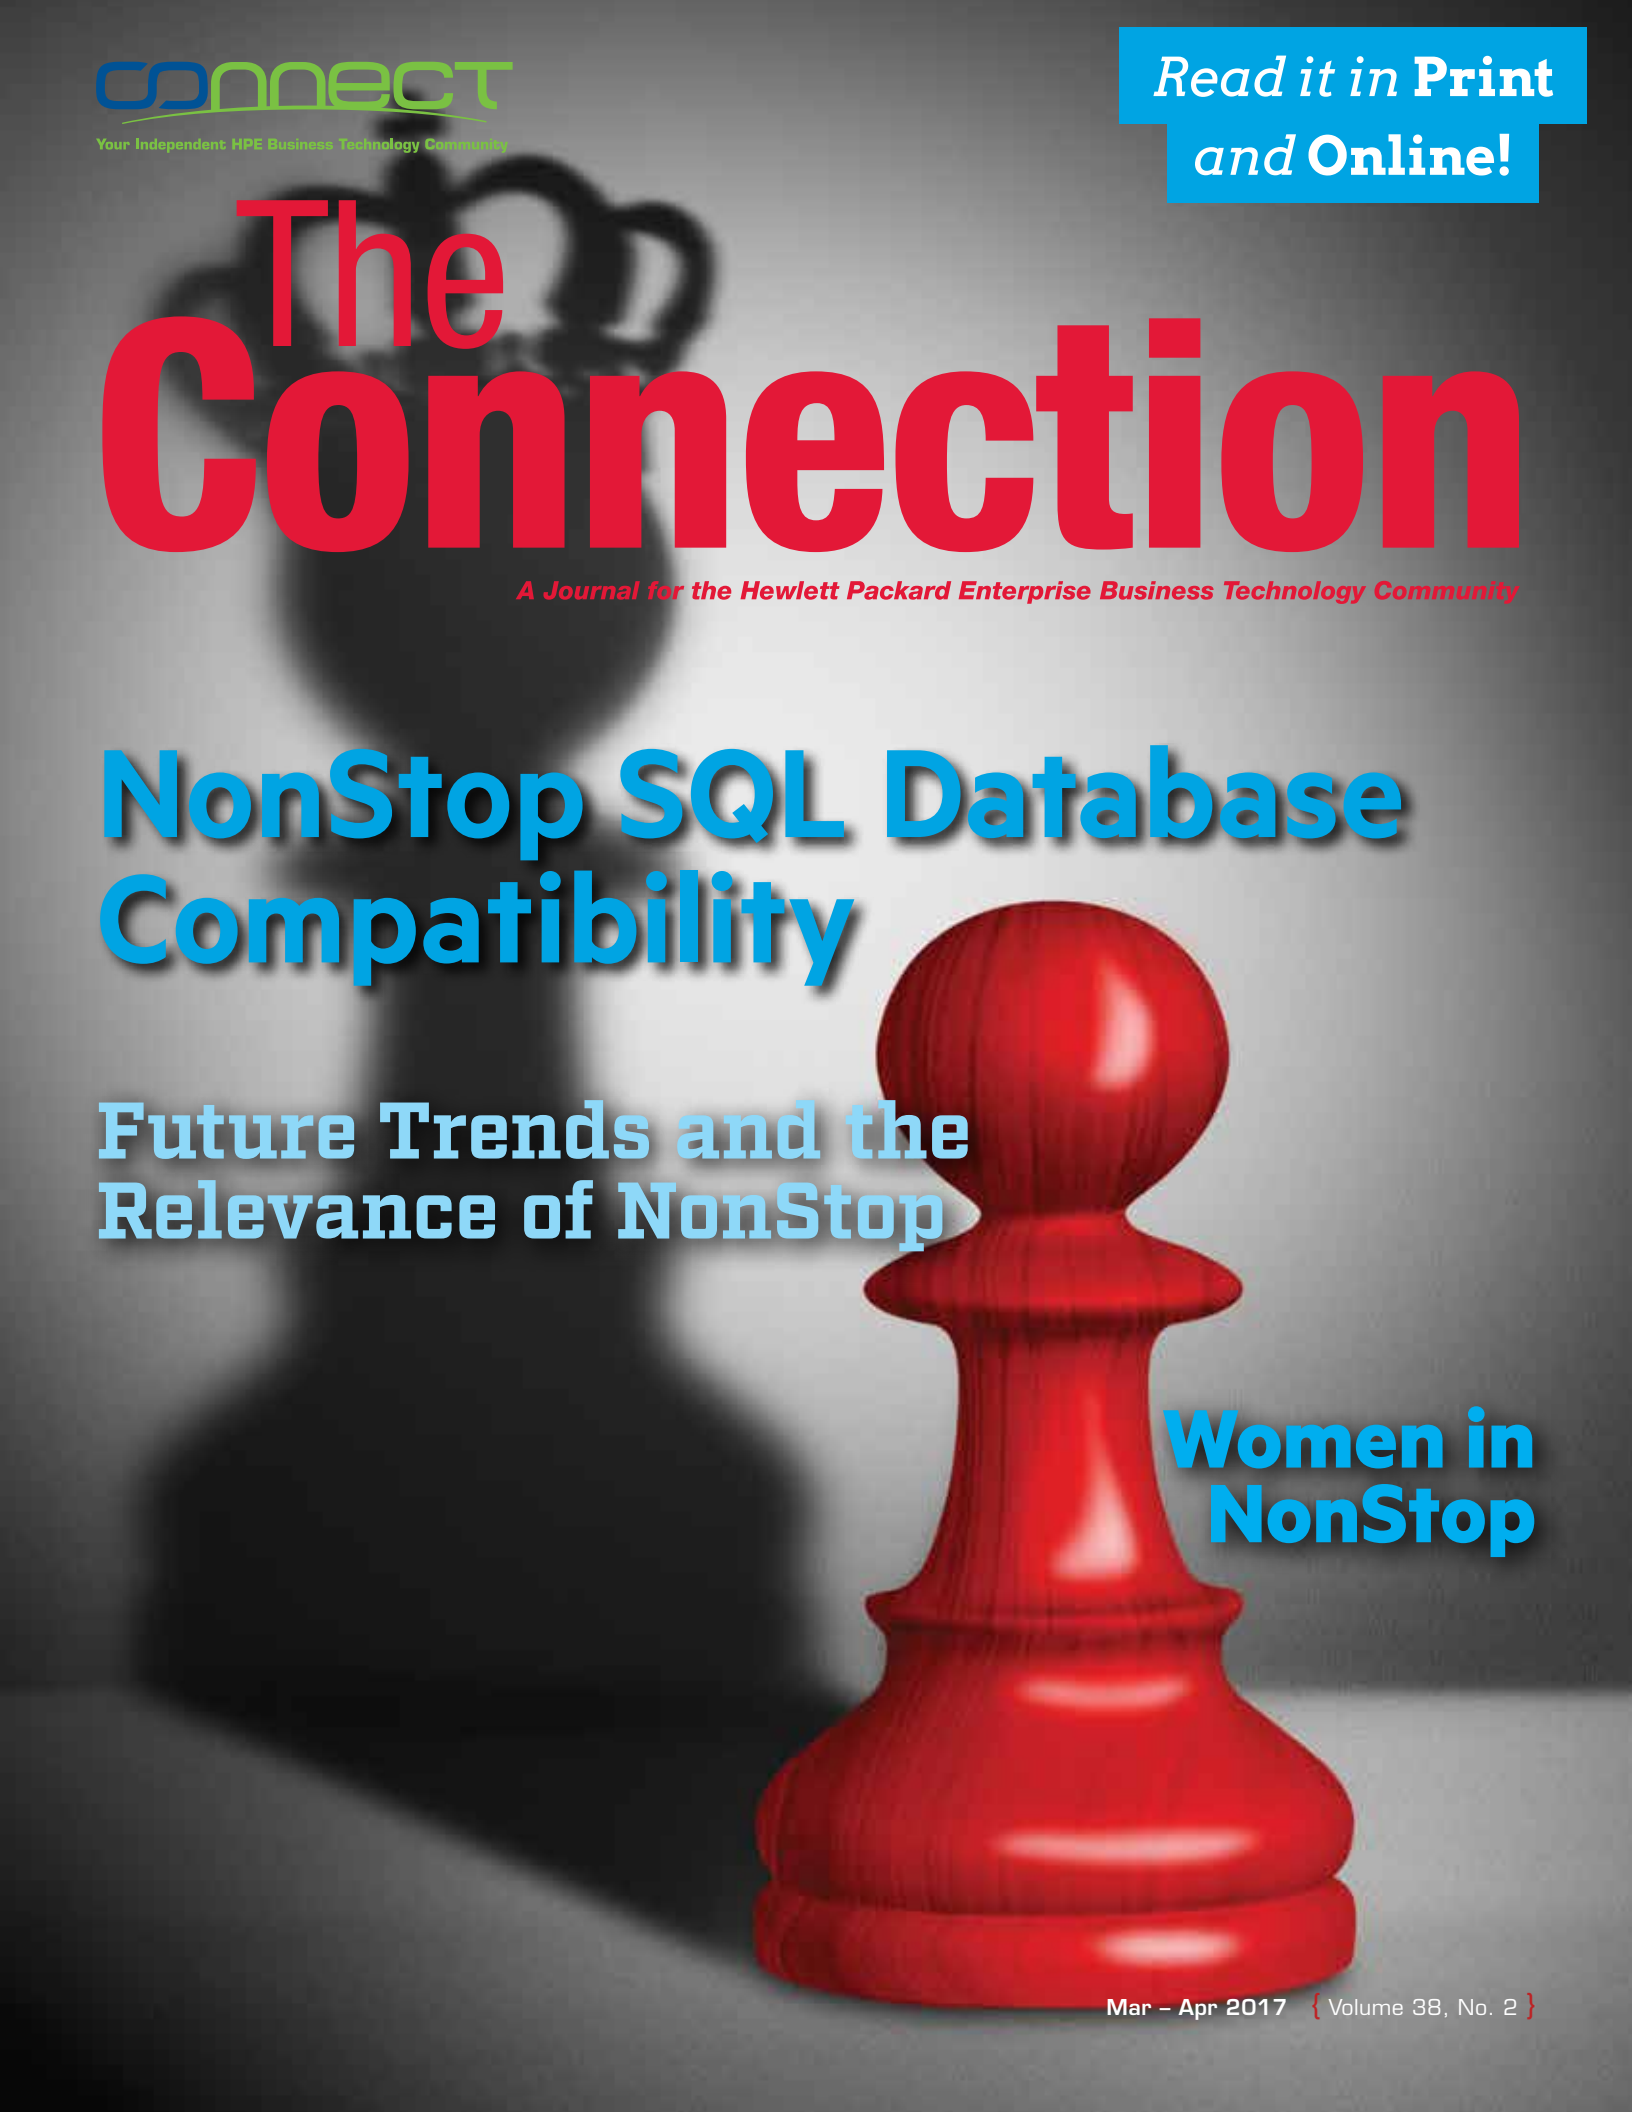

Mar – Apr 2017 { Volume 38, No. 2 }

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

Monitor 100% of Your Processes.

In 100% Real Time.

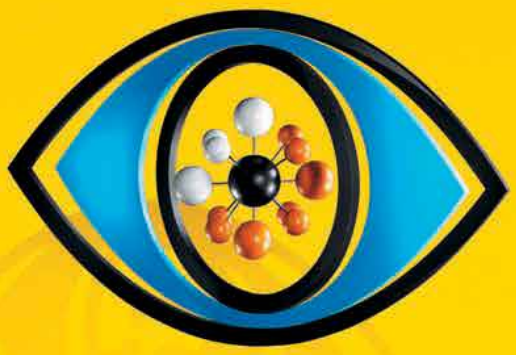

ALTERNATIVE THINKING ABOUT PERFORMANCE MONITORING:

Introducing HP's NonStop Real-Time Process Monitor (RPM), the new low-cost, efficient, real-time monitoring software that keeps you constantly aware of CPU and process resource consumption.

Engineered specifically for NonStop servers, NonStop RPM helps you monitor 1000s of CPUs and lets you know instantly when any CPU or process is using excessive resources. So you can fix bottlenecks before they become full-blown problems.

Here's your chance to keep your NonStop infrastructure more non-stop than ever before.

Technology for better business outcomes.

HP NONSTOP RPM

- Provides real-time process monitoring by CPU, Node or Expand super-cluster
- Color-coded alerts tell you instantly if any CPU or process exceeds limits
- Engineered specifically for NonStop servers providing linearly scalable monitoring from 1 to 1000s of CPUs

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

Europe's largest NonStop event: The eBITUG 2017 Summit

Why attend eBITUG 2017?

- Important insights from NonStop experts, solution providers and HPE
- Share your experiences with the NonStop community and learn from others
- Two-day event featuring multiple tracks and covering all areas in NonStop field
- Live Virtualised NonStop demos
- Network with your peers
- Valuable sessions by HPE, end-users and NonStop experts
- Vendor Fair featuring leading NonStop partners
- Unrivalled value for money
- Keynote by Jimmy Treybig

9-10 MAY 2017

DOUBLETREE BY HILTON, TOWER OF LONDON

EARLY BIRD PRICE: £150
REGISTER NOW AT WWW.BITUG.COM

CSP's eBITUG 2017 Evening at aqua shard

CSP presents this year's eBITUG 2017 Evening at aqua shard. This unforgettable event will be held at the world renowned aqua shard restaurant & bar on the evening of 9th May. At 310m tall, The Shard is the tallest building in Western Europe. aqua shard offers spectacular panoramic views of London from a 930sqm dining room, bar and entertainment area located on the 31st floor of this world famous building. CSP has reserved exclusive use of this facility for the enjoyment of eBITUG 2017 attendees. We hope to see you there!

CSP aquashard

EBITUG 2017 IS SUPPORTED BY:

CSP COMPUTER SECURITY PRODUCTS, INC.
network TECHNOLOGIES INTERNATIONAL
Hewlett Packard Enterprise
intel
FSS POWERING PAYMENTS
comforte better always on
GRAVIC Shadowbase
INSIDER TECHNOLOGIES
XYPRO
Ascert
connect
tcm
4tech software
OmniPayments
NuWave TECHNOLOGIES
BrightStrand INTERNATIONAL MISSION CRITICAL SERVICES
Canam Software
ETINET
esq
Paymark

Table *of* Contents

Columns…

The HPE Partner Ready for Technology Partner Program

Expand your customer base

Increase market share

Accelerate revenue growth

Join today
hpe.com/partners/technology

Hewlett Packard Enterprise

A Note from Connect Leadership

Hello "Connect"ed members!

Lots of things going on with the Connect Board of Directors these days, so I thought I would take this opportunity to share some updates. As you are probably aware, the NonStop customer base has always been one of the most active and vocal segments of the Connect community. I have never been able to figure out exactly why, but we seem to spend the most time with HPE as a group. It may be because we are the most concerned with how things go, or it may be that the HPNS platform just attracts the most vocal of advocates, but either way, we are what we are. And perhaps due to that, I am happy to announce that Connect is now being aligned under Randy Meyer and the Mission Critical Systems division and our HPE representative will be none other than Andrew Bergholz himself! And, just to keep things in perspective, I am happy to announce that Navid Khodayari of Idelji is joining the board as well.

Andrew will be the HPE representative to the Connect Board. He is taking over for our previous representative, Janice Zdankus who is taking on other responsibilities within HPE. Janice has been a great friend to Connect and we wish her good luck, she will be missed. Andrew has been an engineer with HPE for quite some time and reports to Randy and is very much in touch with what the Connect community.

Navid is joining us as a member of the Board. He is taking over for the position vacated by Jamil Rashdi, our representative from Dubai. Jamil has other duties that are taking him away from Connect as well. We also wish him good fortune in his endeavors. Navid is an active member of the Connect community. I am sure you have met him at any of a number of events over the years. He is also helping lead the charge to mentor younger HPNS professionals in the under 40 SIG.

Once again, there is so much going on in the NonStop world that it is good for us to have this level of representation on the Connect Board! Considering the place that NonStop holds in the enterprise, it is comforting to know that HPE views us as important enough to be in the Mission Critical Systems group.

This month we are focusing on empowering data-driven organizations. What better data driver is there than the NonStop?

Thanks to all of you. Do me a favor and the next time you see Andrew or Navid at an event, tell them how much you appreciate their help!

Rob Lesan

Thanks.
Rob Lesan
XYPRO Technology
Connect Worldwide President

2017 Connect Board of Directors

PRESIDENT
Rob Lesan
XYPRO

VICE PRESIDENT
Michael Scroggins
Washington St. Community College

PAST PRESIDENT
Henk Pomper
Plusine ICT

DIRECTOR
Trevor Jackson
SOCAN

CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide

HPE LIAISON
Andrew Bergholz
Senior Director of Development of HPE NonStop

DIRECTOR
Navid Khodayari
Idelji

The Connection

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

Kristi Elizondo.....CEO
Stacie Neall.....Managing Editor
Kelly Luna.....Event Marketing Mgr.
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.

Karen Copeland
Manager WW NonStop
Product Management
Mission Critical Solutions

News from HPE's NonStop Enterprise Division

In February, HPE announced the imminent availability of HPE Virtualized NonStop as a new trusted platform for mission-critical telecommunications applications like HPE Virtualized I-HSS. The announcement was made at Mobile World Congress in Barcelona, which is one of the largest telecommunications shows of the year.

Communications Service Providers (CSPs) are embracing Network Functions Virtualization (NFV) because of its promise of more efficient and effective service delivery. NFV allows them to be more nimble in responding to the demands of their increasingly diverse markets. Virtualization technologies and concepts have been proving their worth by delivering flexible network functions over the past few years, and CSPs are now looking to NFV to deploy key digital core network elements.

Andy Bergholz, the Director of HPE NonStop Engineering, attended Mobile World Congress to help telecommunications customers understand the enhanced capabilities of Virtualized NonStop. In the core network, specific components run on NonStop platforms in order to meet their continuous availability requirements. Now those applications, such as Home Location Region (HLR) and Home Subscriber Services (HSS) and associated subscriber databases, can be deployed on a rock-solid private cloud with Virtualized NonStop. The diagram below shows Virtualized NonStop's position in the VNF private cloud.

HPE Virtualized NonStop: NFVI Hardware and Software

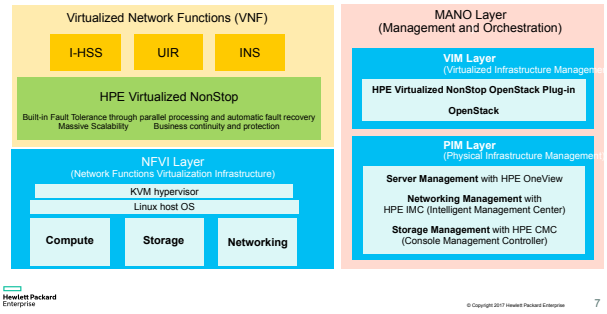

We are already working with a major telecommunications customer in Europe on an integration of Virtualized NonStop and HPE INS. We're looking forward to sharing this customer's experience later this year.

Virtualized NonStop isn't the only exciting development. This issue of Connection focuses on Database and the many investments we're making to offer greater compatibility with other market offerings. See Roland Lemoine's article on Database Compatibility in this issue to learn more about how we're enhancing SQL/MX to make it easier to port new solutions to the NonStop server.

Thinking about the vast databases kept by NonStop customers whose data is the lifeblood of the company, providing deep knowledge of orders, customers, histories, as well as supporting the analytics that businesses use more than ever today to rapidly predict trends, you realize how very important it is to protect your business by protecting the data. In this edition, you can learn more about protecting your data in two articles. Carole Murphy's "Protecting Sensitive Data In and Beyond the Data Lake" talks about how HPE SecureData fits into complex environments to secure Big Data and provides a nice diagram to help you understand. "A Large Financial Institution Migrates Datacenters with No Downtime Using HPE Shadowbase ZDM" by Keith Evans describes how to safely migrate your database while protecting the availability of your business with Shadowbase. We hope you find both of these articles enlightening.

2017 is shaping up to be another year with many NonStop and HPE events to choose from. There's probably a NonStop User's group meeting of some type coming to a location near you. Local meetings are a convenient and inexpensive way to check in with your NonStop colleagues, including partners and HPE speakers, to find out the latest about what's happening. We encourage you to go to the Connect WebSite to find out more about the following events coming up between now and June. www.connect-community.org/chapter-events-2017/. Do check back regularly for additional events.

Here's the current list:

DUSTUG (Desert Users of Tandem) – Mar 7th – in Scottsdale, Arizona.
SunTUG (Sunshine Summit) – Mar 17th – in Tampa, Florida. Andy Bergholz will be speaking. SunTUG will also hold the traditional and popular Golf Tournament the next day on Saturday, Mar 18th at the Heritage Harbor Golf Course.
METUG – Middle East Payments Processing Show - Mar 28 in Doha, Qatar and on Mar 30th in Dubai, United Emirates. Prashanth Kamath will be speaking
NonStop User Forum – Tokyo Japan – April 11th – Updates the Japanese customer base with latest news on NonStop. Jonathan Sechrist is speaking.
GTUG – German Tandem User's Group - Apr 25th in Hanover Germany.
MRTUG – Mid-Region Tandem User's Group – May 2nd, in Oak Brook, Illinois. Jeff Skinner and Wendy Bartlett will be speaking.
OTUG – Ohio Tandem User's Group – May 4th, at JPMC in Columbus Ohio. Jeff Skinner and Wendy Bartlett will be speaking
eBITUG – THE European NonStop Event for 2017 – May 9th – May 11th, London, UK at the Doubletree by Hilton, Tower of London. Word is that Jimmy Treybig will be speaking as well as Andy Bergholz who will talk about NonStop's future. Roland Lemoine will provide the latest updates on Database features underway.
N2TUG – Texas and Oklahoma NonStop User Group event – June 1st in the Dallas/Fort Worth area.

HPE DISCOVER 2017 – Jun 5th- Jun 8th at the beautiful Venetian Hotel in Las Vegas, Nevada – This is the biggest HPE Event of 2017. This year the event will include the HPE Global Partner Summit on the 5th and CIO Summit during the event. For more information check it out at: www.hpe.com/events/discover/

This is all just during the first half of the calendar year, more is still being scheduled. Events in New England, New York, Korea, Singapore, Canada and more are being planned.

We look forward to seeing many of you at these events during the year and we hope you enjoy this issue of The Connection!

Karen

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

Hewlett Packard
Enterprise

Discover 2017

Accelerating next

Las Vegas June 6 – 8

Prepare for a world where everything computes

The next wave of market leaders will possess the vision and technological agility to turn ideas into reality faster than the competition. Accelerate your digital transformation.

hpe.com/discover

Connect Introduces Tech Forums at HPE Discover

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

Connect Technical Forums deal in depth with specific topics of interest to the attendees of HPE Discover conferences. The forums are a new feature that have been added both to the annual Discover conference in Las Vegas as well as to the Discover conferences held in Europe or in the UK. They are presented by HPE employees or by customers and/or vendors who are experts in a specific topic.

The forums are yet another benefit for the members of Connect, which represents all HPE Technology users. Abstracts of the Technical Forums presented at last year's Discover conferences in Las Vegas and London can be found below. Be sure to participate in those that Connect has scheduled for conferences in 2017 and beyond. The forums are held in the Connect Community Lounge on the exhibit floors.

Tech Forums at the 2016 Las Vegas Discover Conference

Enterprise Networking

The "Enterprise Networking" Tech Forum was moderated by Steve Davidek, IT Manager of the City of Sparks, and by Miguel Olague of Summit Partners.

Abstract: Aruba, a Hewlett Packard Enterprise company, is a networking vendor selling enterprise wireless LAN and edge access networking equipment. Connect with the experts on what is going on in the Hewlett Packard Enterprise Networking world. Discussions include how ArubaOS, Aruba Wireless Access Points, Aruba AirWave, ClearPass and IMC can help you manage your HPE networking environment. We are also going to discuss Aruba Beacons and ideas on how to use them. Meet other HPE customers in an open roundtable discussion with HPE Networking experts.

Key Management: The "Key" to Successfully Adopting Enterprise Encryption

The "Key Management" Tech Forum was presented by Nathan Turajski and Farshad Ghazi of HPE Security – Data Security.

Abstract: For data encryption to be successfully integrated and deployed with enterprise IT, a best-practice approach to managing encryption keys needs to be an integral part of the strategy. But not all applications and use cases require the same approach. Join this chalk talk in which HPE Data Security product management experts will discuss the various approaches being adopted today based on the diverse data types, applications and systems requiring high-assurance security.

So You Want to be a CIO?

The Tech Forum "So You Want to be a CIO" was presented by Michael Scroggins of Washington State Community and Technical Colleges.

Abstract: Connect with the experts and HPE customers during this compelling session. The CIO is a high-risk position. There are many thoughts and much advice related to surviving as a CIO. You've got to get there first. This discussion will center on strategies and considerations that you can use to get there. Why would anyone want to be a CIO? It is the best job in the world... if you have what it takes.

Converged Systems

The "Converged Systems" Tech Forum was presented by Chris Purcell, HPE Manager, Influencer Marketing.

Abstract: Connect with the experts on HPE's Converged System team at this roundtable discussion about Converged and Hyper-Converged Systems. Learn where the data center is heading, how single integrated management is sweeping through the data center, and hear how customers are successfully deploying different types of key workloads across their infrastructure. This is a great opportunity to discuss your data center infrastructure questions directly with HPE technologists.

HP-UX

The "HP-UX" Tech Forum was presented by Jeff Kyle, Director, HPE MCS Product Management.

Abstract: Connect with global enterprise UNIX® peers for a Roundtable discussion about HP-UX, including the future of HP-UX, mission-critical computing, Integrity and beyond. Share best practices and provide feedback and enhancement requests directly to HPE technologists.

AIX to LINUX Migration Best Practices

The "AIX to LINUX Migration Best Practices" Tech Forum was moderated by Kyle Todd, Category Sales Manager, HPE Mission Critical Solutions, and Debbie Whitehurst, VP, HPE Data Center Consulting Service Line.

Abstract: Connect with the experts at this roundtable discussion on AIX to Linux migration and why HPE platforms are the ideal environment for mission-critical performance at a fraction of the cost. Learn best practices on successful migration implementations and what it means for your bottom line as you journey to

an open IT architecture. This is a great opportunity to speak directly with HPE technologists about migration successes or questions.

Protecting Underground Facilities with HPE Infrastructure

This Tech Forum was presented by Bill Kiger, President and CEO, and Jon DeMoss, Director of Technology, of Pennsylvania One Call System, Inc.

Abstract: Pennsylvania 811 (PA1Call) is part of a national network of call centers that monitor underground gas pipelines & electric lines. Their mission is to prevent damage and protect lives with an efficient communications network. In April 2015, a malfunction in the data center's fire suppression system caused damage to their mission-critical infrastructure. Learn how HPE Technology Services & All Lines Technology diagnosed the issues, developed recommendations, & implemented an upgraded & consolidated Datacenter with improved security and reduced risk.

Tech Forums at the 2016 London Discover Conference

Integrated Systems

The "Integrated Systems" Tech Forum was moderated by Chris Purcell, HPE Manager, Influencer Marketing.

Abstract: Connect with the experts on the HPE Software-Defined and Cloud team at this Roundtable discussion about converged management, hyperconverged appliances and Composable Infrastructure. Learn where the data center of the future is heading, how a simple integrated management experience is sweeping through the data center, and how customers are successfully delivering a cloud experience across their data center infrastructures. This is an excellent opportunity to discuss your data center infrastructure concerns directly with HPE technologists.

Enterprise Networking

The "Enterprise Networking" Tech Forum was moderated by personnel from Aruba Networks.

Abstract: Aruba, a Hewlett Packard Enterprise company, is a networking vendor selling enterprise wireless LAN and edge access networking equipment. Connect with the experts to learn what is happening in the world of Hewlett Packard Enterprise networking. Discussions will include how ArubaOS and Aruba's Wireless Access Points, AirWave, ClearPass and Intelligent Management Center can help you manage your HPE networking environment. You also will discuss Aruba Beacons and share ideas on how to use them. Meet other HPE customers in an open Roundtable discussion with HPE networking experts.

Enterprise Security

The "Enterprise Security" Tech Forum was moderated by Rob Lesan of XYPRO Technology,

Abstract: From the data center to the network, big data and social networking, security is now a crucial part of every IT conversation. Connect with the experts at this Roundtable discussion about enterprise security. Share best practices, offer feedback, and speak directly to HPE technologists in a small, intimate setting.

HPE Storage

The "HPE Storage" Tech Forum was moderated by Calvin Zito of HPE Storage.

Abstract: Connect with the experts at this Roundtable discussion about Hewlett Packard Enterprise storage products and solutions, including HPE 3PAR StoreServ, StoreOnce and MSA storage solutions, as well as software-defined storage. Learn how customers of all sizes implement and use HPE storage products. Connect with HPE technologists to provide feedback and enhancement requests directly to HPE.

Enterprise Cloud

The "Enterprise Cloud" Tech Forum was moderated by Steve Davidek, IT Manager of the City of Sparks.

Abstract: Join us for an intimate Roundtable discussion about cloud strategies and solutions. Learn from Hewlett Packard Enterprise customers and technologists, provide feedback directly to HPE, and share best practices with peers from around the world. Connect with the experts to learn how to support your workforce wherever it roams.

UNIX (HP-UX)

The "UNIX (HP-UX)" Tech Forum was moderated by Jeff Kyle, Director, HPE MCS Product Management, and Ken Surplice, HPE Category Manager.

Abstract: Connect with global enterprise UNIX® peers for a Roundtable discussion about HP-UX, including the future of HP-UX, mission-critical computing, HPE Integrity servers and beyond. Share best practices and provide feedback and enhancement requests directly to HPE technologists.

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.

Women in NonStop: *Karen Copeland*

Mandi Nulph >> **Marketing Coordinator** >> NuWave Technologies

In this new series, we will be chatting with some of NonStop's most influential female leaders to learn more about their professional journeys and their insights into the future of NonStop. In our first interview, we sat down with Karen Copeland, Manager of HPE NonStop Product Management.

Mandi Nulph: Please explain your current role in HPE NonStop.

Karen Copeland: I've done a lot of different things with NonStop in my career. Right now, I am managing the product management team for NonStop servers and product lines. We plan all new product deliveries, the introduction of new products, and the life cycle of all the NonStop products, including the hardware and the software, from cradle to grave. We work with the engineering teams to plan changes and enhancements, review what customers are looking for, and what we need to do to stay current within the market. Then we help plan the schedule and when the product comes out, we help to introduce it, and we support the marketing of it. We also assist with sales when the product is available, and as questions come up and problems occur, we drive the resolutions back to engineering. We also plan when products will be retired and taken off the market.

Mandi: What has been your professional journey so far?

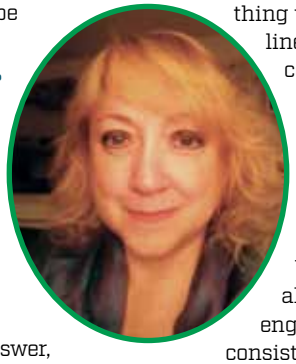

Karen: It's been a long journey! I've been working in the NonStop area now for 32 years, but I originally thought I would become a commercial artist or a writer. I came to California right out of high school, and my first job in the industry was working with IBM. I was 19 years old, and I worked in the support center taking calls from customers and people in the field and helping them resolve their issues. After a while I would hear the same questions over and over and I had the answer, so I started to say, "This is an answer we gave another customer. Does that make sense to you?" and they would say, "Oh yes, thank you!". Eventually, they started to automate that, so I began helping the engineering team to create an application that would do this, kind of like Genesis at HPE. I also did a little business travel to demonstrate this product to different groups around the country.

I left that job because the systems center in Palo Alto moved to Dallas, and I decided I didn't want to move. I took a data entry job, and then started working at Four Phase Systems, which used to be on the campus where Apple is today in Cupertino. I worked there for a few years doing data entry for manufacturing support, and then eventually moved into computer operations. Then they were purchased by Motorola, and that whole support team for manufacturing was laid off. The guy I had worked for there had a friend who worked for Tandem, so he contacted him and they said they had a lot of openings for people who could do computer operations, so I interviewed for that job, got hired, and that's how I started working for Tandem. That was probably in October 1983.

For my first couple of years I worked a swing shift where I helped users during the day, and then at night I did all the backups for four or five systems. Then I moved from that to a support role in software development. In those days, the commercial industry didn't really have a lot of SCM or software configuration management going on. It wasn't as evolved as the Department of Defense standards, which most companies had to meet if they were developing anything for the government. So, I came in and I was helping to take all the engineering fixes and compile the operating system and a lot of other products. I ended up organizing a process so that we could avoid having code collisions and avoid overwriting previous engineering changes when new ones were made. Soon, they started hiring new people to do that kind of job, they had me train them, and then I started managing them. And that grew until I was managing three departments across most Tandem product lines.

As things evolved, I was managing groups that did the actual release software. It would take us a year to put together a whole release version update (RVU), and it was partly because the system we were using to do it back then was manual. We had a team of eight people doing this at one time. Then I acquired some of the engineering teams that did the internal IT, and next thing you know I was a director with a staff of 70 with six line managers. We redesigned the whole tracking and control of the software release activities, and we built an application that could now control the different versions of the software revision package that we were putting together for customers.

I did that for a while, then there was a lot of reorganization when we were acquired by Compaq. I had an interesting style of management at the time. We were providing services to engineering, and I had all my people distributed and working closely with the engineering teams they needed to support. It provided a consistency of process across the teams, but it was kind of a matrix management to the extreme. I treated all those groups like they were my customer, and they found the service valuable and didn't want to lose it. So, somehow, we survived that transition.

After that, I decided it was time to make a career change. I had been line managing in Engineering support roles for 18 years. I looked around and I found they had lost several people in product management due to a very lucrative retirement program that was being offered to employees at the time. I spoke to Randy Myers, who was running Product Management, and I transferred to that team.. They took work that had been done by three different people and put them all in one job that I took on. At that time, I also started to work deals with partners, because I found that was a way to speedily bring something to market. There weren't too many people on the team who were comfortable doing that at the time. I wasn't afraid of working out deals because I had already done that with vendors in the past and worked with legal, so that didn't seem especially scary to me; it just seemed like another way of doing business. We already had a number of these types of deals on the books, but we began to do more of them. When we couldn't invest ourselves because of budget constraints and new products, we started to work with our partners who already had

products that had fulfilled a need in the market for 8-10 years. That put us on a road to reselling partner products to add to the portfolio of offerings.

I did that for two years on the manageability side, and then I moved into security for a couple of years. Then, the manager of the product management team decided he was going to retire. By then, I was at a point where I said, "Ok, I've worked in a lot of different depths and areas; maybe it's time to return to line management." So, I applied for that job, and I have been doing that now for about six years.

Mandi: That is quite the journey! I think it's always interesting to hear how people end up in the space.

Karen: It's interesting because one of the things I have appreciated about working with NonStop products and in the whole business space is the ability to do different things. Even though you have been there for a long time working around the product, your job itself evolves and changes. I never felt like I had to change companies to take on something new, and I found that if you worked hard and were smart about your decisions, that rewards would come and people would recognize that good things were happening.

Mandi: Would you say that's one of your favorite things about working in NonStop?

Karen: It was something that worked well for me, and I felt that was different from other places I had worked. If you had an idea for something, people might listen to it, and you could implement it. I came up with a lot of ideas, so out of 10 ideas I might throw eight away, but two would become something real that we could implement. It's neat now that I have been here for so long, when I interact with engineering and there are new people who tell me how things are supposed to work, I can tell them, "I used to do that job and this was the original intent, and here's why we did that or made this change." It's kind of interesting to see how things that you helped implement evolve as the business changes and as new people come in and put their influence on it.

Mandi: Did you have any role models or mentors who have helped you along your path?

Karen: I was lucky. I had some very good managers, but I also had some bad ones, and you can learn from both. I certainly worked for people who I thought to myself, "I would never do this to someone who worked for me. I would never speak to them this way. I would never handle a project or bad news this way." I took those lessons to heart when it comes to managing my own team.

The other thing that those of us who have been in this business for a while have the benefit of is that Tandem really invested in managers. Every year, there were one or two weeks where you would go to some sort of training. They would do a lot of what's known as "360-degree feedback". All the people you work for, all your peers, and all the people who work for you take a survey and give anonymous feedback about you. It gives you a picture of yourself that you may not have realized. It gives you an idea of how people perceive you that may or may not be true, and it gives you the unique opportunity to work on those things. I was lucky I could benefit from management training like that.

We did have mentoring programs, and I had several managers who were very positive influences and really helped me along in my career. I think it's very important for anyone taking a new position, to work for someone who is truly interested in them and their career. If you can find someone like that, they will help you see what you are good at, help you use and grow your skills. And if they are good, they will have a way of helping you work on the things you struggle with without breaking you down or making you feel demoralized about it, and turning it into a positive thing. Hopefully, when opportunities come along, that person will present them to you and help you decide if you should to take them or not. I think that's something that will move you along faster than anything else.

Mandi: What has been your experience working as a woman in the NonStop space?

Karen: The industry has changed a lot. From the early '80s when I came to Silicon Valley and began working, to today, the whole country has changed. When I first came into the business, there weren't rules about sexual harassment or people doing awkward or inappropriate things in the work place. When I was 19 or 20, there were men in their 30s or 40s who would make passes at me, or invite me out, or do things that made me very uncomfortable. I think every woman I know who has been in the industry for the past 25-30 years could tell you stories like that from their early days. But I think as the world has progressed and has become more aware of that, much of that has gone away. So as women, we don't face as much of that anymore.

For the last 20 years or so, I haven't felt too much like I was an outsider because of my gender. I haven't really found it as much of a struggle as I know some women have. Maybe I have just been lucky in the people I have had the pleasure to work with, or maybe I just think of it differently, so I haven't felt put aside because of my gender. It's how you approach things and how you present yourself that I think also helps determine what you get assigned to and what you show yourself as capable of doing. One of the wonderful things about Silicon Valley is that it becomes more about the idea than who suggested it. Fortunately, it hasn't been a huge thing for me in my career, and I haven't felt held back because I was a woman. There have only been a few rare cases when I have worked with people who seemed to have a real problem working with women.

Mandi: There has been a lot of talk recently about "the next generation of NonStop", as more people retire and younger people start working in the industry. Do you have any advice for young people coming into the NonStop space?

Karen: Come in with an open mind. There are still a lot of opportunities in the NonStop platform in fields beyond the traditional ones like telecommunications and finance. Think about new ways NonStop could be applied. Look for a type of position where you can have an impact and where you can have a manager who can mentor you and who cares about you and your career. If you have ideas or thoughts on where NonStop should be going or where something should happen, speak up! Maybe someone already had that idea and there's a reason why we aren't doing that, but maybe it's something new and exciting! Your first idea might not be the right one, but maybe your third or fourth is great and can be implemented. When you're first entering a career, you may have this idea of what you want to do or what you imagine you'll be doing in five years. Be open to new opportunities that present themselves. Maybe your real path is slightly to the right or to the left of where you thought you were going to end up.

Mandi: It's interesting to get the perspective from someone who has been in the industry for so long and learn about their journey. Thank you for sharing.

Mandi Nulph is NuWave's marketing coordinator. NuWave specializes in HPE NonStop middleware, including their newest product, LightWave Server™, which allows you to expose your existing Guardian or Pathway servers as industry-standard REST services. With a degree in Mass Communication and Journalism, she boasts 10 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.

PROTECTING SENSITIVE DATA IN AND BEYOND THE DATA LAKE

Carole Murphy >> **Global Product Marketing** >> HPE Security - Data Security

The need to secure sensitive data in Hadoop and IoT ecosystems

Hadoop is a unique architecture designed to enable organizations to gain new analytic insights and operational efficiencies through the use of multiple standard, low-cost, high-speed, parallel processing nodes operating on very large sets of data. The resulting flexibility, performance, and scalability are unprecedented. But data security was not the primary design goal.

When used in an enterprise environment, the importance of security becomes paramount. Organizations must protect sensitive customer, partner and internal information and adhere to an ever-increasing set of compliance requirements. But by its nature, Hadoop poses many unique challenges to properly securing this environment, not least of which include automatic and complex replication of data across multiple nodes once entered into the HDFS data store. There are a number of traditional IT security controls that should be put in place as the basis for securing Hadoop, such as standard perimeter protection of the computing environment, and monitoring user and network activity with log management. But infrastructure protection by itself cannot prevent an organization from cyber-attacks and data breaches in even the most tightly controlled computing environments.

Hadoop is a much more vulnerable target—too open to be able to fully protect. Further exacerbating the risk is that the aggregation of data in Hadoop makes for an even more alluring target for hackers and data thieves. Hadoop presents brand new challenges to data risk management: the potential concentration of vast amounts of sensitive corporate and personal data in a low-trust environment. New methods of data protection at zettabyte scale are thus essential to mitigate these potentially huge Big Data exposures.

Data protection methodologies

There are several traditional data de-identification approaches that can be deployed to improve security in the Hadoop environment, such as storage level encryption, traditional field-level encryption and data masking. However, each of these approaches has limitations.

For example, with storage-level encryption the entire volume that the data set is stored in is encrypted at the disk volume level while “at rest” on the data store, which protects against unauthorized personnel who may have physically obtained the disk, from being able to read anything from it. This is a useful control in a Hadoop cluster or any large data store due to frequent disk repairs and swap-outs, but does nothing to protect the data from any and all access when the disk is running—which is all the time.

Data masking is a useful technique for obfuscating sensitive data, most often used for creation of test and development data from live production information. However, masked data is intended to be irreversible, which limits its value for many analytic applications and post-processing requirements. Moreover, there is no guarantee that the specific masking transformation chosen for a specific sensitive data field fully obfuscates it from identification, particularly when correlated with other data in the Hadoop “data lake.”

While all of these technologies potentially have a place in helping to secure data in Hadoop, none of them truly solves the problem nor meets the requirements of an end-to-end, data-centric solution.

Data-centric security

The obvious answer for true Hadoop security is to augment infrastructure controls with protecting the data itself. This data-centric security approach calls for de-identifying the data as close to its source as possible, transforming the sensitive data

elements with usable, yet de-identified, equivalents that retain their format, behavior, and meaning. This protected form of the data can then be used in subsequent applications, analytic engines, data transfers and data stores, while being readily and securely re-identified for those specific applications and users that require it. For Hadoop, the best practice is to never allow sensitive information to reach the HDFS in its live and vulnerable form. De-identified data in Hadoop is protected data, and even in the event of a data breach, yields nothing of value, avoiding the penalties and costs such an event would otherwise have triggered.

The solution—HPE SecureData for Hadoop and IoT

HPE SecureData for Hadoop and IoT provides maximum data protection with industry-standard, next generation HPE Format-preserving Encryption (FPE), (see NIST SP-800-38G) and HPE Secure Stateless Tokenization (SST) technologies.

With HPE FPE and SST, protection is applied at the data field and sub-field level, preserves characteristics of the original data, including numbers, symbols, letters and numeric relationships such as date and salary ranges, and maintains referential integrity across distributed data sets so joined data tables continue to operate properly. HPE FPE and SST provide high-strength encryption and tokenization of data without altering the original data format.

HPE SecureData encryption/tokenization protection can be applied at the source before it gets into Hadoop, or can be evoked during an ETL transfer to a landing zone, or from the Hadoop process transferring the data into HDFS. Once the secure data is in Hadoop, it can be used in its de-identified state for additional processing and analysis without further interaction with the HPE SecureData. Or the analytic programs running in Hadoop can access the clear text by utilizing the HPE SecureData high-speed decryption/de-tokenization interfaces with the appropriate level of authentication and authorization.

If processed data needs to be exported to downstream analytics in the clear—such as into a data warehouse for traditional BI analysis—there are multiple options for re-identifying the data, either as it exits Hadoop using Hadoop tools or as it enters the downstream systems on those platforms.

To implement data-centric security requires installing the HPE SecureData infrastructure components and then interfacing with the appropriate applications and data flows. SDKs, APIs and command line tools enable encryption and tokenization to occur natively on the widest variety of platforms, including Linux®, mainframe and mid-range, and supports integration with a broad range of infrastructure components, including ETL, databases, and programs running in the Hadoop environment, and is available for any Hadoop distribution. HPE Security—Data Security has technology partnerships with Hortonworks, MapR, Cloudera and IBM, and certifications to run on each of these. HPE SecureData is integrated with the Teradata® Unified Data Architecture™ (UDA), and with the HPE Vertica Big Data Platform

Rapid evolution requires future-proof investments

Implementing data security can be a daunting process, especially in the rapidly evolving and constantly changing Hadoop space. It's essential for long-term success and future-proofing investments, to apply technology via a framework that can adapt to the rapid changes ongoing in Hadoop environments. Unfortunately, implementations based on agents frequently face issues when new releases or new technology are introduced into the stack, and require updating the Hadoop instance multiple times. In contrast, HPE SecureData for Hadoop and IoT provides a framework that enables rapid integration into the newest technologies needed by the business. This capability enables rapid expansion and broad utilization for secure analytics.

Securing the Internet of Things

Failure to protect sensitive data in the Hadoop environment holds major risk of data breach, leaking sensitive data to adversaries, and non-compliance with increasingly stringent data privacy regulations such as the General Data Protection Regulation (GDPR). Big Data use cases such as real-time analytics, centralized data acquisition and staging for other systems require that enterprises create a "data lake" — or a single location for the data assets.

While IoT and big data analytics are driving new ways for organizations to improve efficiencies, identify new revenue streams, and innovate, they are also creating new attack vectors which make easy targets for attackers. This is where perimeter security is critical, but also increasingly insufficient – it takes, on average, over 200 days before a data breach is detected and fixed.

As the number of IoT connected devices and sensors in the Enterprise multiplies, the amount of sensitive data and Personally Identifiable Information collected at the IoT Edge and moving into the back-end in the data center--is growing exponentially.

The data generated from IoT is a valued commodity for adversaries, as it can contain sensitive information such as Personally Identifiable Information (PII), payment card information (PCI) or protected health information (PHI). For example, a breach of a connected blood pressure monitor's readings alone may have no value to an attacker, but when paired with a patient's name, it could become identity theft and a violation of (HIPAA) regulations.

IoT is here to stay. A recent Forbes article predicted that we will see 50 billion interconnected devices within the next 5-10 years. Because a multitude of companies will be deploying and using IoT technologies to a great extent in the near future, security professionals will need to get ahead of the challenge of protecting massive amounts of IoT data. And, with this deluge of sensitive IoT data, Enterprises will need to act quickly to adopt new security methodologies and best practices in order to enable their Big Data projects and IoT initiatives.

New threats call for new solutions - NiFi Integration

A new approach is required, focused on protecting the IoT data as close to the source as possible. As with other data sources, sensitive streaming information from connected devices and sensors can be protected with HPE FPE to secure sensitive data from both insider risk and external attack, while the values in the data maintain usability for analysis

However, Apache NiFi, a recent technology innovation, is enabling IoT to deliver on its potential for a more connected world. Apache NiFi is an open source platform that enables security and risk architects, as well as business users, to graphically design and easily manage data flows in their IoT or back-end environments.

HPE SecureData for Hadoop and IoT is designed to easily secure sensitive information that is generated and transmitted across Internet of Things (IoT) environments, with HPE Format-preserving Encryption (FPE). The solution features the industry's first-to-market Apache™ NiFi™ integration with NIST standardized and FIPS compliant format-preserving encryption technology to protect IoT data at rest, in transit and in use.

You Know NonStop

You want to know that your information is accessible, yet secure. You also want to be able to get the most out of your hardware investment quickly and easily. We don't think that's too much to ask.

NuWave offers quick, easy, secure integration for your NonStop servers. Our products can connect your servers to virtually any platform, anywhere, using REST or SOAP services, allowing you to get the most out of your systems.

You choose NuWave for all of your NonStop middleware needs, because you know NonStop.

Why NuWave Middleware?

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

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

NuWave
THE MIDDLEWARE GUYS

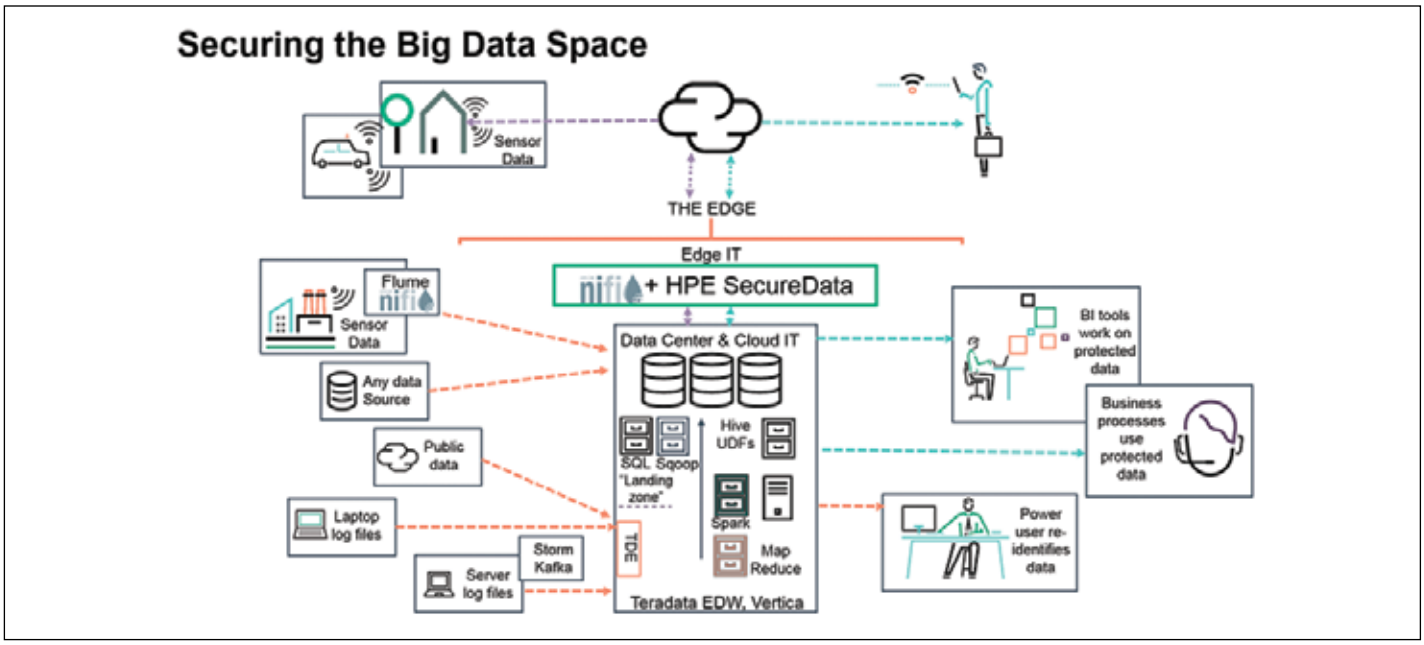

The HPE SecureData NiFi integration enables organizations to incorporate data security into their IoT strategies by allowing them to more easily manage sensitive data flows and insert encryption closer to the intelligent edge. This capability is included in the HPE SecureData for Hadoop and IoT product. In addition, it is certified for interoperability with Hortonworks DataFlow (HDF).

With this industry first, the HPE SecureData for Hadoop and IoT solution now extends data-centric protection, enabling organizations to encrypt data closer to the intelligent edge before it moves into the back-end Hadoop Big Data environment, while maintaining the original format for processing and enabling secure Big Data analytics.

Carole Murphy currently manages product marketing for HPE Security – Data Security, where she is responsible for developing market strategy for the HPE SecureData product line and Big Data/Hadoop, and IoT solutions, including go-to-market planning, product communication, strategic positioning and market awareness.

Improving Reliability via Redundant Processing

Dr. Bruce D. Holenstein >> **President & CEO** >> Gravic, Inc.
Dr. Bill Highleyman >> **Managing Editor** >> Availability Digest
Paul J. Holenstein >> **Executive Vice President** >> Gravic, Inc.

The three pillars of mission-critical systems are summarized by the acronym "RAS." It stands for Reliability, Availability, and Scalability. Reliability is the probability that the system will produce correct outputs. Availability is the probability that the system is operational. Scalability is the ability of the system to handle different size loads in a predictable manner.[1]

What is the difference between reliability and availability? If a system is available but is generating erroneous outputs, it is not reliable. Alternatively, if a system is generating reliable outputs while operational, but it is not always operational, then it is not always available.

Why would a system produce incorrect data? The problem could be a hardware fault, a firmware bug, or a software error. Bad data also can be generated if the system has been infected with malware. Depending upon its design, malware can corrupt any operation within the system. Malware is generally thought to be a software issue. However, malware also can be introduced into the hardware or the firmware of a system during its manufacture.

How can we determine if a system is malfunctioning before damage to the end user or environment takes place? One answer is to use a "Validation Configuration," a redundant system in which two or more sub-systems are running in parallel and are processing the same requests, presumably arriving at the same results. Preferably, the sub-systems are from different manufacturers to ensure they both do not contain the same faults, if any. Their results (or even intermediate processing states) are compared. As long as the results agree, it safely can be assumed that the sub-systems are operating properly. If one (or both) sub-systems have been infected with malware, the results will not agree; and the sub-systems should be taken out of service and checked to determine the problem.

Why Would There Be an Error?

There are several reasons why a sub-system that appears to be operating properly could be delivering erroneous results. Common reasons include a piece of hardware that is producing memory or disk read errors or a sub-system that has been infected with malware of some sort.

It also is possible (though unlikely) that the sub-system hardware or firmware could have been modified during manufacture to provide malicious results. In the most important mission-critical cases involving financial or health decisions or crucial process control sub-systems, care should even be taken that there is no common manufacturing point in the sub-systems, such as the printed-circuit masks being fabricated by the same company.

One challenge with malware is that it may be too late by the time the malware is detected. The system may appear to be operating properly, but the bad data that it is producing due to the malware infection may not be discovered for days, weeks, or even months, if at all.

An interesting example occurred several years ago in a banking application that calculated interest payments for banking customers. Typically, an interest calculation results in fractions of a penny. For instance, 3% of $167.58 is $5.0274. The customer is credited with an interest payment of $5.02. The remainder, $0.0074, is called overage. It is credited to the bank itself.

However, a hacker was able to install malware that took a portion of the overage and credited it to his account. The amount for each transaction was so small that it wasn't noticeable; but over hundreds of thousands of transactions, it amounted to a tidy sum. The malware was not discovered until the next annual audit of the system.

Running the sub-system in a Validation Configuration may expose these types of errors and corruption.

Comparing Results

Several methods exist with which the results of independent sub-systems can be compared. Two of those methods are Logical Synchronization Units (LSUs) and a new scheme that operates at the transaction level.

Logical Synchronization Unit (LSU)

An LSU, such as that used in the HPE Itanium S-series NonStop models, is a hardware device that compares two or more

[1] Originally, IBM called the "S" Serviceability, which is the speed with which a system can be repaired. However, Serviceability is characterized by the mean time to repair and is already incorporated into the term Availability.

values to ensure that they match. An LSU can be used to compare the outputs of two sub-systems (a DMR LSU – dual modular redundancy) or three or more sub-systems (a TMR LSU – triple modular redundancy). If all outputs agree, the result is published by the LSU. See Figure 1.

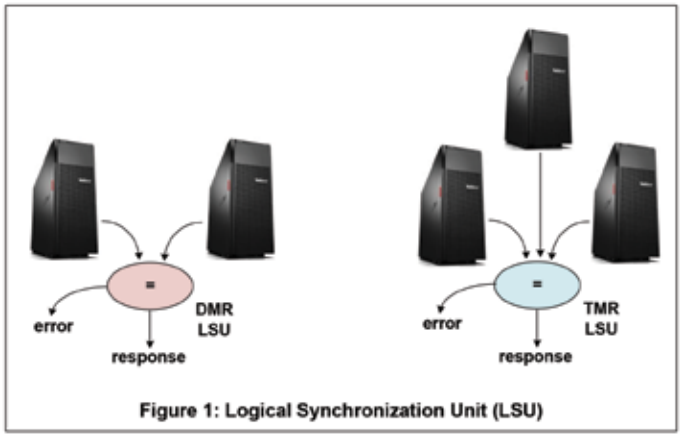

Figure 1: Logical Synchronization Unit (LSU)

In a DMR LSU, if the results do not agree, an error is returned by the LSU. In a TMR LSU, if one of the results is different, that result is ignored; and the majority result is returned. An error is generated identifying the sub-system that produced the erroneous output, and it can automatically be taken out of service.

An LSU was previously used by NASA for the space shuttle to compare the outputs of four independent computers (1, 2, 3, 4), as shown in Figure 2. The computers exchanged interprocess messages over four interprocessor synchronization buses (5, 6, 7, 8). The outputs of all four computers were sent to a voting LSU (9). The LSU outvoted the results of any failed computer and sent the correct output (the one generated by the majority of the computers) to the appropriate actuator in the space shuttle. The astronauts were instructed to turn off a computer that was generating false outputs.

An LSU can face several challenges:

- It represents a single point of failure.
- The LSU itself needs to be validated that it is free of infection.
- It must be simple in order to minimize the probability of failure.
- Therefore, it can vote only on simple inputs.
- If an error is caused by a malicious hardware, firmware, or software implementation, an LSU may not detect it because the same error will exist on all sub-systems.
- A major class of systems is transaction-processing systems, in which the different sub-systems and/or CPUs cannot operate in lock-step, and the outputs are thus not directly comparable.

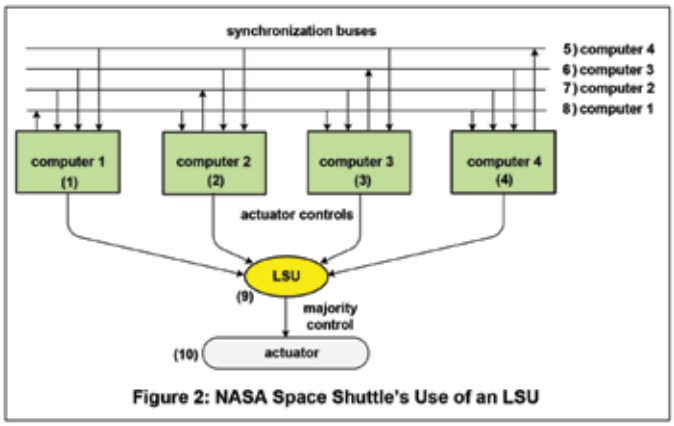

Figure 2: NASA Space Shuttle's Use of an LSU

Transaction Indicia Matching

The authors' ongoing research focuses on improving the data reliability of transaction-processing systems at the transaction level. The premise is that the same request fed into two independent sub-systems running the same application must produce the same transactional changes to the database, or else there is a reliability problem. The method, which we call *"transaction indicia matching* (TIM)," is described below and detects reliability problems, such as those caused by malware infections, operator malfeasance, or hardware or software errors.

The TIM method operates by matching "indicia" generated locally by the Validation Configuration. Indicia are measures of the sub-system state at any particular point in time and are primarily produced at the end of the DML (data manipulation language) operations made to a database. Indicia are calculated by a trusted piece of software or hardware called an *"Indicia Engine"* and may be arbitrarily complex. For instance, in a transaction-processing system, indicia may be the full set of DML changes that were made to the database by a transaction; or they may be a hash sum of the changes that were made to the database.

With two (or more) sub-systems running in parallel, the indicia calculated by each sub-system are compared in a Validation Configuration to ensure all sub-systems are operating correctly. It is optimal if the sub-systems are manufactured to the same specifications by different manufacturers and that no common point of manufacture exists (such as integrated circuit masks). This prevents a design error or a malware error maliciously introduced during manufacture from appearing in both sub-systems. The sub-systems are running the same application, preferably implemented by different software teams.

The indicia calculated by the Indicia Engines in each sub-system are compared. One means of transactional indicia matching is via direct communications between the Indicia Engines, as shown in Figure 3. Two sub-systems (1) and (2) are running versions of the same application (3). Each application version has an Indicia Engine attached or built into it (4).

At specific synchronization points, each application pauses; and each Indicia Engine calculates an indicium representing its application's current state. The Indicia Engines are connected via a communication channel (5, 6). They exchange their calculated indicia (7) and compare their own with that of the other sub-system. If the indicia agree, the Indicia Engines release the applications; and processing continues. If they don't agree, the applications are halted; and an error is posted.

In a DMR system, in the event of an error, both sub-systems can be shut down and tested. In a TMR system, the faulty sub-system is shut down and processing continues.

A Validation Configuration implementation that is designed to fit existing applications is shown in Figure 4. The Validation Configuration comprises two sub-system nodes (1, 2). An Indicia Engine is attached to each application that calculates appropriate indicia at one or more synchronization points within the application. The Indicia Engines communicate via a bidirectional synchronous data replication engine such as the Shadowbase® data replication engine from Gravic, Inc.

An application (3) sends an identical request (4) to both sub-

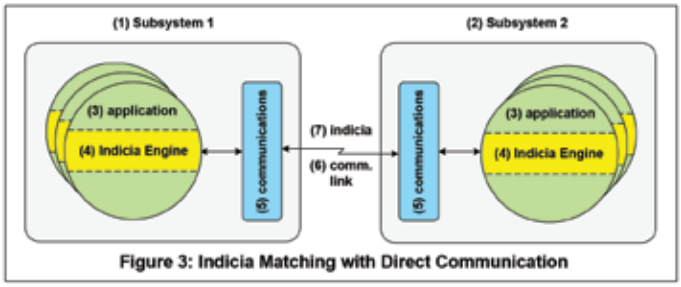

Figure 3: Indicia Matching with Direct Communication

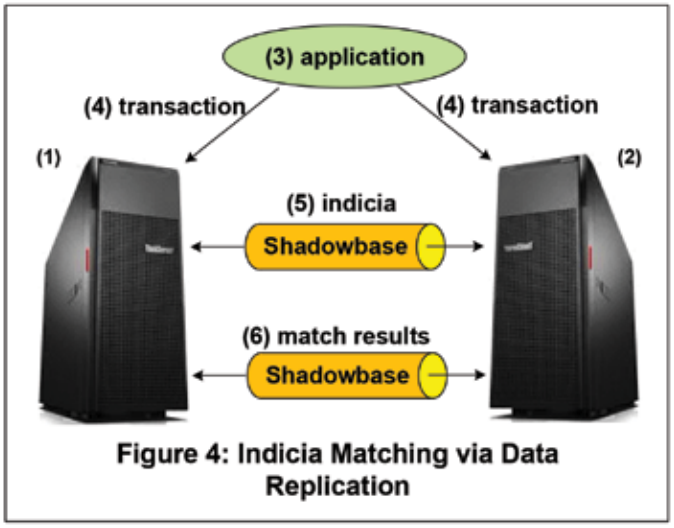

Figure 4: Indicia Matching via Data Replication

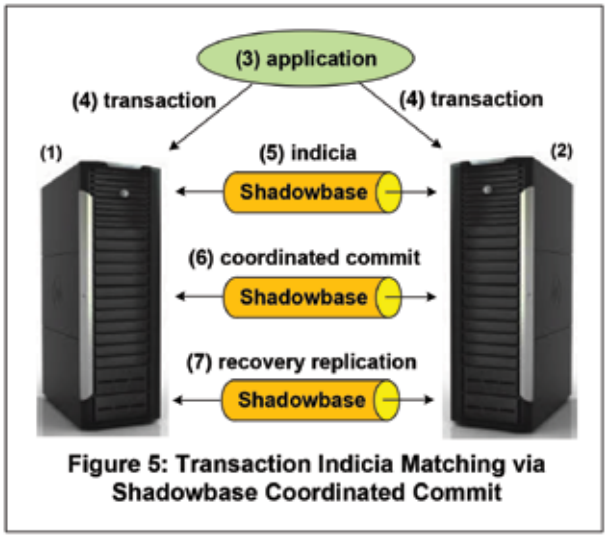

Figure 5: Transaction Indicia Matching via Shadowbase Coordinated Commit

systems. Each sub-system calculates indicia at the one or more synchronization points and sends its calculated indicia to the other sub-system via the bidirectional data replication engine (5). As an example, the indicia can be a hash sum of the changes to be made to the database by a transaction. The hash sum can be calculated by a User Exit[2] in the data replication engine. The indicia received from the remote sub-system are compared to the indicia calculated by the receiving sub-system.

At the end of request processing, each sub-system informs the other sub-system as to the results of its indicia matching (6). If all of the indicia have matched properly in both sub-systems, the result of the request processing is accepted by both sub-systems. In this case, the Indicia Engine at each node instructs the data replication engine to vote in favor of transaction commitment. If there is a mismatch in the calculated indicia by either sub-system, an error is posted and the transactions are aborted. In a DMR system, both sub-systems should be shut down and analyzed. In a TMR system, the sub-system that doesn't match with the other sub-systems is shut down; and processing continues with the operational sub-systems.

High Availability Transaction-Processing Systems

In transaction-processing systems, dual sub-system nodes (1, 2) often are deployed to achieve high availability (see Figure 5). Should one sub-system fail, all transactions can be sent to the surviving sub-system for processing. Advantage can be taken of this system redundancy to construct a Validation Configuration to detect data-reliability problems. The nodes exchange indicia at appropriate synchronization points to ensure that nothing has corrupted processing.

As shown in Figure 5, an application (3) sends a transaction (4) to both nodes in the Validation Configuration. The transaction-processing applications in each node make changes to the database at its node. User exits in the data replication engine

[2] A User Exit is a customized piece of logic that can be added to the Shadowbase data replication engine to perform selected processing on the various transaction components.

serve the purpose of the Indicia Engines to calculate the indicia. The calculated indicia is exchanged between the two nodes (5) via data replication, and each node compares its calculated indicia with that of the other node. If both nodes agree with the indicia calculated at the other node, the transaction is committed. Otherwise, the transaction is aborted.

If the Shadowbase data replication engine is employed, the transaction can be committed (or aborted) via the Shadowbase coordinated-commit[3] facility (6). Coordinated commits work in this case as follows. If a node's indicia have matched properly with the indicia sent by the remote node, each node will send a token to the other node to say that it is ready to commit its transaction. Each node will respond to the token with an indication that it is ready to commit the transaction. When a node receives a confirmation from the other sub-system, it commits the transaction. If either node cannot do so, it instead will send an abort request to the other sub-system. In this case, both nodes will abort their transaction.

Data replication used in this way to achieve both high-availability and data reliability has another important benefit if one of the nodes is taken off-line due a problem. In this case, the data replication tool will queue the changes on the running node and be able to forward them seamlessly to the recovered node once it has been brought back on-line (7).

Encryption

As an option, the indicia being exchanged between the systems can be encrypted. This can prevent a "man-in-the-middle" attack, in which an attacker can modify an indicium that does not match to one that matches in order to mask a malware infection. Alternatively, the attacker can change an indicium that matches to one that does not match to cause a sub-system denial of service outage.

Certifying a New Sub-system

The TIM approach can be used to certify the reliability of a new sub-system. The new sub-system is put into operation along with a known and trusted sub-system. As the sub-systems process requests, they both calculate indicia. The new sub-system sends its indicia to the trusted sub-system, which compares them to its own indicia.

If the indicia should not match, the new sub-system can be taken out of service for further diagnostics. The error could be in the hardware design, the firmware, or the software. Alternatively, the error could be caused by malicious malware that has infected the new sub-system.

Summary

The proper operation of a sub-system can be verified via a Validation Configuration that compares the operation of one sub-system to another sub-system running the same applications. Verification is accomplished by comparing indicia generated by the two sub-systems at specific synchronization points. If the indicia agree, the sub-systems are operating properly. If the indicia do not agree, one of the sub-systems is misbehaving.

Comparing sub-system outputs via TIM is a significant improvement over the use of an LSU for validation purposes in a transaction processing system because of the set of challenges faced by an LSU, not the least of which is that it represents a single point of failure in the system. There is no single point of failure when using transaction indicia matching.

Our research shows that TIM can be implemented with a synchronous data replication engine such as Shadowbase from Gravic, Inc. We would appreciate feedback if you are interested in exploring this concept with us for your real-world application.

Dr. Bruce D. Holenstein, President and CEO. Dr. Holenstein leads all aspects of Gravic, Inc. as President and CEO. He started company operations with his brother, Paul, in 1980, and is presently leading the company through the changes needed to accommodate significant future growth. His technical fields of expertise include algorithms, mathematical modeling, availability architectures, data replication, pattern recognition systems, process control and turnkey software development. Dr. Holenstein is a well-known author of articles and books on high availability systems. He received his BSEE from Bucknell University and his Ph.D. in Astronomy and Astrophysics from the University of Pennsylvania.

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.

Paul J. Holenstein is Executive Vice President, Gravic, Inc. He has direct responsibility for the Gravic, Inc. Shadowbase Products Group and is a Senior Fellow at Gravic Labs, the company's intellectual property group. He has previously held various positions in technology consulting companies, from software engineer through technical management to business development, beginning his career as a Tandem (HPE NonStop) developer in 1980. His technical areas of expertise include high availability designs and architectures, data replication technologies, heterogeneous application and data integration, and communications and performance analysis. Mr. Holenstein holds many patents in the field of data replication and synchronization, writes extensively on high and continuous availability topics, and co-authored Breaking the Availability Barrier, a three-volume book series. He received his BSCE from Bucknell University, a MSCS from Villanova University, and is an HPE Master Accredited Systems Engineer (MASE). To contact the author, please email: SBProductManagement@gravic.com . Hewlett Packard Enterprise directly sells and supports HPE Shadowbase Solutions (www.ShadowbaseSoftware.com); please contact your local HPE account team.

[3] Coordinated Commits are a special form of processing that allows the results of two or more transactions to be coordinated to the same conclusion before they complete. Refer to the Breaking the Availability Barrier book series for more information.

NonStop SQL Database Compatibility A Key Element of the NonStop Database Services Strategy

Roland Lemoine >> NonStop Product Management >> Database, Languages and ID

With the recent proliferation of new Open Source databases, many companies are re-evaluating the TCO of their current solution and looking at alternatives. Customers who are unhappy with high license costs or lack of scalability often mention Oracle RDBMS as a product they'd like to replace. In addition, while Open Source solutions may include free-license software and therefore appear as an obvious choice for lower TCO, a closer look or experience shows that often the free versions lack many of the Enterprise class features. The gap is even more apparent when the target is a mission critical application. Creating scalability and availability, for example with Open source based solutions, in the end requires so much re-work and application redesign to cover new unmet needs that it defeats the purpose of the migration. This gap amplifies the perception that switching database is a high-risk task. Furthermore, Open Source databases are rarely multi-purpose databases which rules them out as a viable choice for handling data as part of the digital core of the enterprise. Instead, their use promotes duplication of databases throughout the enterprise with the same data sometimes needing to be propagated to 10 different databases. The TCO in that case is exploding as the customer needs a different set of skills for each database and also must invest in data federation and data virtualization software and complex ETL solutions.

To make this picture even more unpleasant, SQL as a standard is a goal that has been abandoned since the end of the 1990s. This resulted in removing incentives for vendors to stay close to the ANSI standard, to the detriment of application portability. While portability has made great progress in programming languages such as Java or Python allowing an application to just "run everywhere" unchanged, the same cannot be said for the database access part of the application. Such lack of compatibility immediately impacts the porting effort and is yet another obstacle to switch databases.

There is good news though and it is coming from HPE NonStop! For one thing, NonStop is a well-known leader in Mission Critical applications with its unrivaled scalability and availability. Second, HPE NonStop SQL is a multi-purpose database with best-in-class data integrity features making it a perfect candidate to play a central role to maintain the database of record. Third, NonStop can also claim a very low TCO based on multiple factors such as the fact that NonStop is a Single system Image and therefore only requires one Operating System, one database, to setup and maintain regardless the number of nodes (up to 4080).

Introducing Database Compatibility

Database compatibility is an important factor in minimizing the impact of database conversions. HPE NonStop SQL is introducing Database Compatibility to remove the last obstacle that prevents taking an application and moving it unchanged (or with very few changes) onto the target platform. Database Compatibility enables low risk, economical port of applications to NonStop SQL. Investment in the application and associated skills is preserved. Porting cost is kept at the minimum and HPE NonStop SQL's existing record for mission critical applications is insurance that the benefits will be there without any bad surprises at the completion of the project.

But wait, that's not all, so far we've discussed that moving away from Oracle could be motivated by reducing cost reasons. It may also worth noting that NonStop SQL performs better than Oracle in other areas than TCO.

Scalability

Oracle RAC is the Oracle solution to scale out instead of scaling up. As it is now known in the industry scaling up has performance limitations and an inherent lack of fault tolerance. However while Oracle RAC adopts the multi node approach it does not mean that it does scales very well. In fact after 4 nodes scalability becomes questionable as shown in this real application benchmark:

NonStop SQL/MX can claim a linear scalability meaning more than

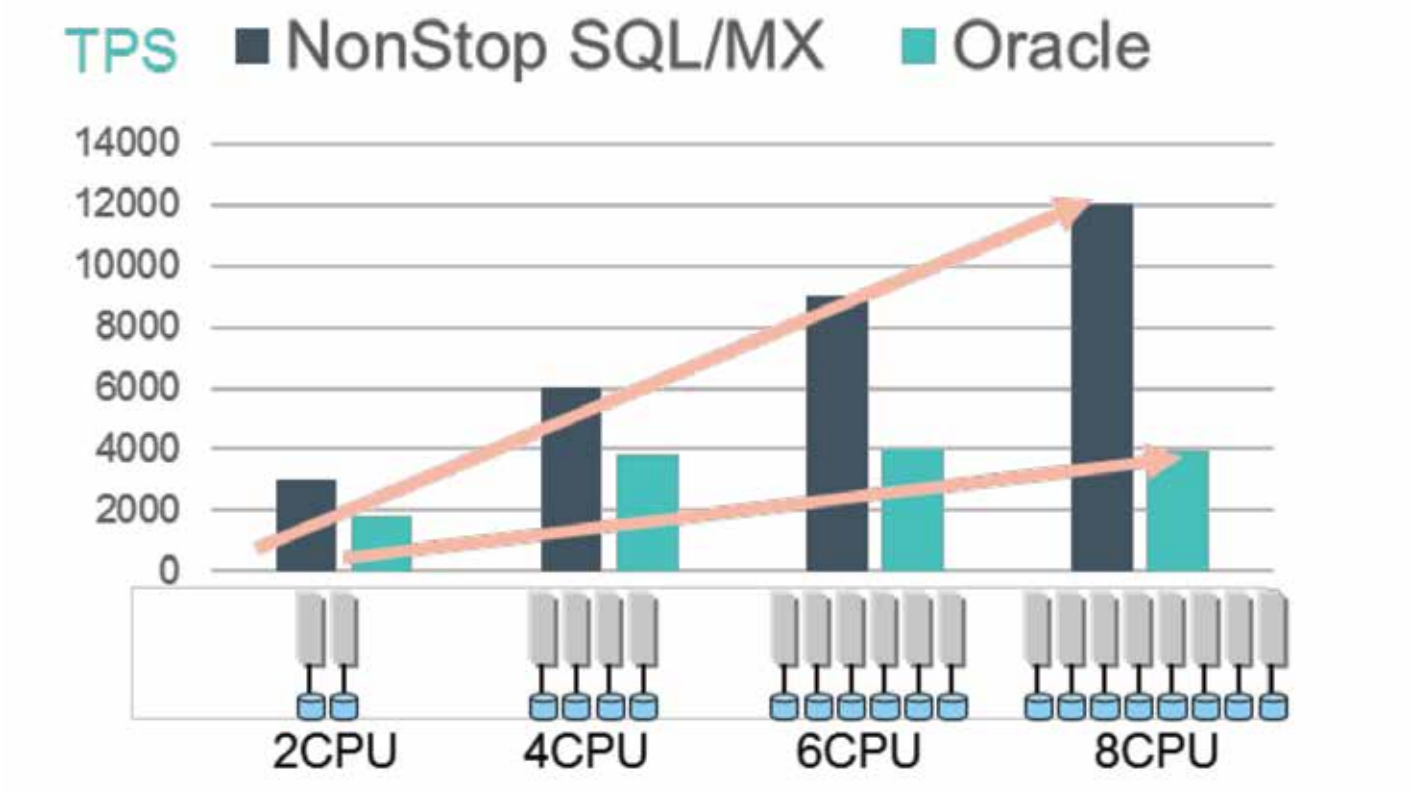

This benchmark graph shows how Oracle RAC scalability is limited to a few nodes.

98% of added processing capability translates to increase in throughput. Lack of scalability for Oracle RAC lies in its design. Originally an SMP architecture, Oracle RAC is still exposed to the bottleneck of having to sync the disk cache in memory for all nodes. NonStop being a shared nothing architecture is not exposed to this bottleneck.

Simplified architecture

The Single system Image mentioned earlier not only reduces the TCO because it is much simpler to administer but also brings a much simpler architecture as opposed to Oracle RAC where the whole cluster is made of a manual setup of multiple pieces that constantly need to be re-certified together as Linux versions are changing and requires IP management for load balancing purposes. This complexity for Oracle RAC increases the chances for failures or misconfiguration. The NonStop simplified architecture is the required recipe to achieve the best reliability. Oracle has recognized this issue and released their Exadata offering, which is a set of only four pre-assembled Oracle RAC system configurations. While Exadata may solve the initial setup issue, it still is equally complex to maintain when problems happen because it is still based on the same complex Oracle RAC solution.

Availability

Despite all its best efforts Oracle RAC does not achieve the same levels of availability as NonStop SQL. An Oracle RAC setup, assumes very skilled DBAs and architects participating in the setup (unlikely in general), may at best achieve what IDC qualifies as "AL3" (element failure is not transparent to the application). This requires additions of multiple Oracle options, such as partitioning, compression, which more than doubles the license costs. The highest level of availability, "AL4", can only be obtained in a one unique scenario for Oracle which is by using replication with 2 Exadata systems. NonStop does not require replication to achieve AL4. And replication, while it is perfect for disaster recovery purposes, is not as ideal for availability purposes (replication incurs added latency and can require manual resolution of conflicting database updates or collision).

Database compatibility evaluation

From a high level point of view, NonStop SQL and Oracle RDBMS have a lot in common:

Feature	NonStop SQL	Oracle
Database type	Both are multi-purpose, relational DBMS supporting static and dynamic SQL, ACID transactions, referential integrity, concurrency, distributed consistency, access control, partitioning, replication and online backups	
Application domain	OLTP, OLAP, Mixed workload, Big Data	
Other Features Connectivity	Triggers, Publish/subscribe, Sequence generators, Multi-tenancy ODBC, JDBC	
Programming Languages User environment Development environment Built-in functions	C, C++, Cobol, Java, Python, Node.js* Shell, bash Eclipse support 60 same functions (aggregate, string or date manipulation)	
Schemas	Both database can name tables as schema.table	

What is the same in NonStop SQL and Oracle.
*Node.js off platform.

Database Compatibility contents

However to achieve compatibility, it is imperative to implement the same features across the whole software stack as are

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

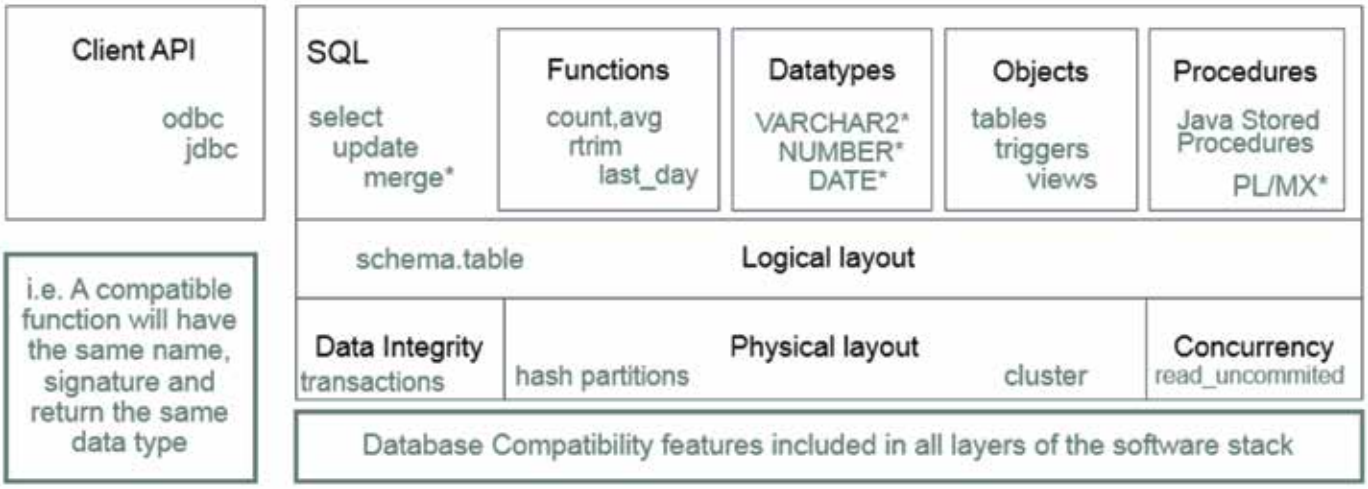

Database compatibility features. * Most recent additions.

> **Database compatibility enables a low risk, economical port that overcomes the common burden of switching databases. The customer is not locked in with one vendor anymore and can take advantage of the unique, unmatched, NonStop value proposition of scalability and availability.**

implemented by Oracle. The diagram above shows example of Oracle-compatible features implemented in NonStop SQL.

PL/MX is being introduced as part of SQL/MX to provide compatibility with PL/SQL language. Oracle PL/SQL is a vast language and has many features, so additional compatibility will be implemented across multiple releases of SQLMX.

Database Compatibility is available on both NonStop X and HPE Virtualized NonStop offerings.

Conclusion

Moving from Oracle to NonStop SQL does not just bring a better TCO but true linear scalability which makes it possible to start with a small application and overnight grow it to be a mission critical application without having to re-architect it. Database compatibility enables a low risk, economical port that overcomes the common burden of switching databases. The customer is not locked in with one vendor anymore and can take advantage of the unique, unmatched, NonStop value proposition of scalability and availability. In its strategy to play a role as the digital core of the enterprise, delivering Database Services in a cloud computing environment, NonStop SQL database compatibility allows customers to easily migrate their applications to NonStop SQL and truly reduce their database sprawl. NonStop SQL ultimately delivers better service levels for end users while allowing IT departments to control their costs with a best-in-class multi-purpose database.

Roland Lemoine has been working on NonStop for 22 years and is currently product manager for database, languages and development products. Previous experience includes customer support for middleware products, Open Source advocacy and a strong UNIX background.

Data-driven organizations now realizing the value of NonStop – data has to be fresh and accurate!

Richard Buckle >> Pyalla Technologies, LLC

There has rarely been a time when I have considered myself anything other than a "computer guy." Of course, over the years, there have been other distractions. On one occasion, following regional trials for the Australian schoolboy Rugby League team, I came home from practice with a contract to join a professional League team as a junior – the Balmain team, as I recall. I also loved to sail competitively, so much so that day-sailing just doesn't do it for me anymore. I was part of the crew when we won the Middle Harbor (Sydney) Division 1 in-harbor championship a number of times. As for the League contract, there were two sticking points – I was under 18 years of age and needed my parents to sign and oh yes, we had to move to Balmain to qualify on basis of residency. Many years later, in 1979, I almost bought a yacht but mostly, I am a "car guy." That is, when I'm not a "computer guy!"

In a recent motor magazine editorial column, an emotional journalist wrote, "The fevered rise of automated car research is inextricably tied to the rise in distracted driving, both of which are enabled by high-speed data. And data is the operative word here." The editor then concluded, "Instead of solving the distracted-driving problem by simply freezing driver cellphone activity within a moving car, the big data companies encouraged self-driving technology, sensing an opportunity to collect / deliver data from / to the occupants … I think the automated car serves the data companies vastly better than it does the cause of safety." As a "car guy" who is also a "computer guy," I see that this editor may have a point and yet, when I look back at my career as well as at the cars I have driven, I can compartmentalize and there should be no confusion about which device takes priority when I am behind a keyboard or behind a steering wheel. Consuming data while driving, other than listening to the radio, isn't a good idea but just the recognition that drivers have a need to consume data while driving probably correlates highly with why so much goes wrong out there on our highways.

There is a small group on LinkedIn called Data-driven Apps that I have supported for many years. Membership the last time I checked was less than twenty but I have my fingers crossed that in time it will attract more and more members of the NonStop community. However, the data-driven organization is the subject of considerable focus of late and the complexity of data as we save text, video and voice in tiered storage is driving discussions about the future of computing and specifically, the future of the data center. And yet, among the NonStop community, it remains very much business-as-normal for the simple fact that the pursuit of a data-driven organization must still deal with transactional systems. With as much talk as is coming from HPE about transforming to hybrid infrastructure perhaps the encroachment of cloud computing onto the data center is only part of the story. Perhaps upon further consideration, when the worlds of transaction processing meet the data-driven organization, there is a need for hybrid world of a different type that may or may not include clouds. But what is a data-driven organization?

One of the better explanations of the data-driven organization I have come across was a May 24, 2016, article in InfoWorld, What is a data-driven company? "The surge of a data-driven culture has also had a significant impact on how companies are structured. The complexity of data forces companies to merge different department to harness their individual strengths to make the most of data. Being data-driven means making use of massive quantities of unstructured data – text, video, voice." The introduction of a data-driven culture helps explain why there is as much coverage of data-driven organizations as there has been of late. For many of its supporters, it is the only way to look at data and these supporters have proved evangelical in their zeal and their disquiet over traditional data center approaches! "To make sure data is not only in the hands of IT and other data enthusiasts, organizations need to embrace a switch in culture. Most experts agree that business intelligence needs to be in the hands of every decision maker in the company to make sure the entire staff is aligned and fighting the same battles."

When it comes to the value proposition for customers, Forbes Insights (financed by SAS) wrote that "Data-driven customer experience is critical to the future growth and development of organizations, particularly in today's hyper-competitive economy. A Forbes Insights survey of 357 executives of large organizations finds that the benefits of evolving to data-driven customer experiences are wide-ranging, including enhancing revenue generation and enabling cost reduction, as well as accelerating process efficiencies and quality improvements." Echoing what InfoWorld referenced, Forbes notes that in order to achieve this, "there needs to be greater alignment of people, processes and technology across enterprises—involving not only sales and marketing teams, but also other key players behind customer experience, including information technology, purchasing and production." Such alignment has also been recognized by other industry pundits as "evidence-based decision making."

While some of this may not yet strike a chord with every member of the NonStop user community there are numerous NonStop vendors already working to better support data-driven organizations. The variety of approaches and the number of products that support the movement of data, whether it's part of data recovery strategies, simple back-up procedures of ensuring data can be consumed by end users, customers, or processes as mandated by governments or industry associations, or to more precisely target sales opportunities, they all empower the data-driven organization where NonStop has a presence. When I approached DataExpress Inc. CEO,

Billy Whittington, about the underlying business driver for DataExpress his response was as relevant today as it was when DataExpress was founded, “Moving NonStop data for consumption? This is exactly what we have been doing since the early 90’s, successfully and without fanfare!”

Another company that might not come to mind is Striim, Inc. Just look at how quickly Striim has applied its real time data analytics to addressing some very simple problems that arise with data. At a time when many NonStop users are pursuing data replication and disaster planning strategies that call for replicated databases, how can the data-driven organization know for sure that the data is really in synch. While there are options to license replication products for replication to be configured in support of active-active, data-driven organizations still need to operate within agreed-upon SLAs – and Striim has now successfully implemented a capability that highlights when SLAs established to ensure databases don’t fall behind. After all, there is no lessening of the need to move files even as there is no letup in the requirement to protect information – yes, the data-driven organization still operates under the premise that data will be stored globally and yet be fresh and accurate when it comes time to meet local needs.

To become a data-driven company the belief in the importance of the integrity and quality of information needs to permeate the culture of the company at all levels. It is not enough to start a formal data governance program, becoming data-driven requires a disciplined shift in the mindset of all employees towards maintaining the integrity and quality of their data.

"To become a data-driven company the belief in the importance of the integrity and quality of information needs to permeate the culture of the company at all levels. It is not enough to start a formal data governance program, becoming data-driven requires a disciplined shift in the mindset of all employees towards maintaining the integrity and quality of their data," said Chris Jennings, vice president of technology services at Collaborative Consulting, in the Infoworld article. While Striim is providing tools to ensure such shifts in mindsets are supported there is still some very basic needs to be addressed when it comes to availability. Of course, Striim assumes that data is being replicated and indeed backed up and in every situation where Striim plays a role you will find traditional D/R and Data BackUp solutions already deployed.

When we think of data-driven organizations the conversation oftentimes turns to clouds and in particular, public clouds. After all, the elasticity they provide when it comes to storing data allows many companies turn to them when the requirements for storage aren’t well understood or adequately defined. This is the battleground for competing solutions from Amazon Web Services (AWS), Microsoft’s Azure as well as Google with its Cloud Platform services. However, just this week, news quickly spread of the massive outage that took place at AWS. According to the February 28, 2017, article in USA Today, Massive Amazon cloud service outage disrupts sites, “While not all AWS clients were affected by the outage at one of AWS's main storage systems, some experienced slowdowns, after a big portion of its S3 system went offline Tuesday afternoon.” Furthermore, “‘this is a pretty big outage,’ said Dave Bartoletti, a cloud analyst with Forrester. ‘AWS had not had a lot of outages and when they happen, they're famous.’”

What stands out in this article for anyone familiar with NonStop are the closing observations by USA Today. “Companies have been steadily moving storage to the cloud because it is cheaper, easily accessible and more resilient. But the downside is that when there are problems, there's a cascade effect,” it said. “It's possible to contract with multiple companies to avoid potential problems, but that strategy is pricey, so many companies make peace with the knowledge that on rare occasions they're going to have a very bad day. ‘Only the most paranoid and very large companies, distribute their files across not just AWS but also Microsoft and Google, and replicate them geographically across regions - but that's very, very expensive, Gartner's [cloud analyst, Lydia] Leong said.” When I had a conversation later in the day with StrongBox Data Solutions, Inc. CEO, David Cerf, his immediate response was “A great reason for tape! You know it’s secure!”

Data-driven organizations need layers of redundancy even as they need to build out their strategy for empowering their employees. Data has to be delivered and data has to be accessible and every data-driven organization must be confident that the data being relied upon is fresh and accurate. And for the NonStop user this often turns to conversations about the database and this increasingly involves NonStop SQL/MX. “These organizations realize that their businesses depend on providing tools allowing 24 X 7 data access to everyone in the company,” said Gravic Inc. EVP, Paul Holenstein. “Whenever customers run solutions on HPE NonStop that access databases, such as NonStop SQL/MX, availability becomes an important consideration. There is little value in the contents of the database if the database is no longer accessible. Moving data to a second system, and indeed to a second site, helps ensure availability from a business continuity perspective – if the primary site fails, the backup site provides continuous access.”

The data-driven organization needs to go beyond purely ensuring the data is available after a failure. “They realize the value in unlocking and sharing their key data assets to provide enhanced value-add to their information portfolio. Sometimes this value-add takes the form of feeding data changes from an operational system into an online and real-time data warehouse for both tactical decision making as well as strategic planning support. At other times, it is integrating their information from one environment to another to support real-time business intelligence, such as performing real-time fraud detection for a financial authorization switch,” noted Holenstein. “At Gravic, we have witnessed an increasing priority on ensuring that nothing takes the data offline for any reason, and we have continued to provide the features and tools in our Shadowbase line of data replication products to ensure the databases and data sharing on NonStop systems are no exception.”

Part of ensuring data is available involves backing-up data critical to the functioning of the business – yes, there are multi-level tiering options readily available but at some point, there is still the need to back-up to virtual tape and this isn’t lost on anyone in the NonStop community, paranoid or not. ‘’Our customers that use Shadowbase also use BackBox replication in order to protect their database integrity against data corruption or against an accidental record lost. With the backup also replicated, the customer can always go back to a specific backup and restore the file with integrity,’’ said Sylvain Tétreault, COO, ETI-NET.

“Shadowbase customers are pushing more and more for synchronized replication of the database along with the more advanced active/active architectures, in order to take over transparently,” noted Holenstein. “The BackBox Catsync synchronizes the NonStop backup catalog DSM in real time in

GRAVIC®
Shadowbase

Would you bungee jump without knowing it was safe?

Then why take chances with your Business Continuity solution?

Many business continuity solutions are difficult or even impossible to effectively test. You'll never know if they work until you really need them. And by then, it's too late. Eliminate these risks by using an **HPE Shadowbase Active/Active** or **Sizzling-Hot-Takeover** business continuity solution. Then, you'll know for sure that your safety net will work.

Don't cling to the cliff, contact us!

For more information, please see the Gravic white paper: *Choosing a Business Continuity Solution to Match Your Business Availability Requirements.*

ShadowbaseSoftware.com

Hewlett Packard Enterprise
Silver Partner
Technology Partner

order to complement Shadowbase and allow the customer to switch transparently the backup/restore on their DR system," added Sylvain Tétreault. Acknowledging this aspect of making sure data is always available empowering the organization, Holenstein explained how they "spend a lot of time advising people to go beyond virtual tape for DR because it is complementary and provides better Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the database."

With Gravic and ETI-NET providing value with respect to ensuring data is always available, Holenstein then said, "We also strongly advise customers to use virtual tape to take periodic backups of their critical data because, as you say, if you 'accidentally' (or employee malfeasance) delete something on the source, the data replication engine will dutifully delete it on the target. Of course, you 'may' have that data in an audit trail and can use a recover/roll forward from TMF to mine the deleted data from audit and reinsert it." Highlighting one specific instance when it all can go horribly wrong, even at the best managed NonStop site, Holenstein added, "It is actually much more urgent than that as some operations, such as a PURGE FILE or DROP TABLE do not generate audit; they merely whack the file or table, and if you are configured for replication to replicate these events, the data replication engine will dutifully comply and whack the file or table on the target node. Now, you are really in trouble if you don't have a copy of it 'somewhere' (e.g., on tape)." And that's "what a virtual backup solution like BackBox is addressing," added Sylvain Tétreault.

Who knows what may have transpired if my parents had looked more favorably upon my opportunity to play professional rugby. Then again, I am a "computer guy" and yes, I am a "car guy." And the two seem inseparable of late to the point where more recently, it's been hard to tell whether a convention is about cars or electronics. According to former GM Vice Chairman, Bob Lutz, "Detroit tries to counterpunch Las Vegas's CES, which is supposedly a consumer electronics show but is attracting more and more major car introductions. It might be another deplorable sign that we are at a point of convergence: Are these cars with onboard computing power, or are they computers with four wheels and an engine thrown in?"

Collecting and delivering data to the occupants of a car may be the best example yet where moving data directly to the consumer. However, for the NonStop community, where applications have been moving data to consumers standing in line at a shop or in front of an ATM or even scrolling through their smartphone, without knowing all the processes working behind the scenes to make sure the data is fresh and accurate, this is not surprising. Perhaps there is a small element of paranoia within the NonStop community but then again, given the state of technology today and the most recent incidents being reported in the media, it's good to know that there as many NonStop vendors as there are today all working to better support our data-driven organizations.

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.

A Large Financial Institution Migrates Datacenters with No Downtime Using HPE Shadowbase ZDM

Keith B. Evans >> **Shadowbase Product Management** >> Gravic, Inc.
Paul J. Holenstein >> **Executive Vice President** >> Gravic, Inc.

HPE Shadowbase Zero Downtime Migration (ZDM) Overview

How is HPE Shadowbase ZDM able to migrate systems with no application downtime? An overview of the ZDM process is shown in Figure 1 through Figure 4. (Note that these figures do not reflect the subject company's actual system configuration, but demonstrate the general principles involved.)

Step 1: Set up the new system (Figure 1).

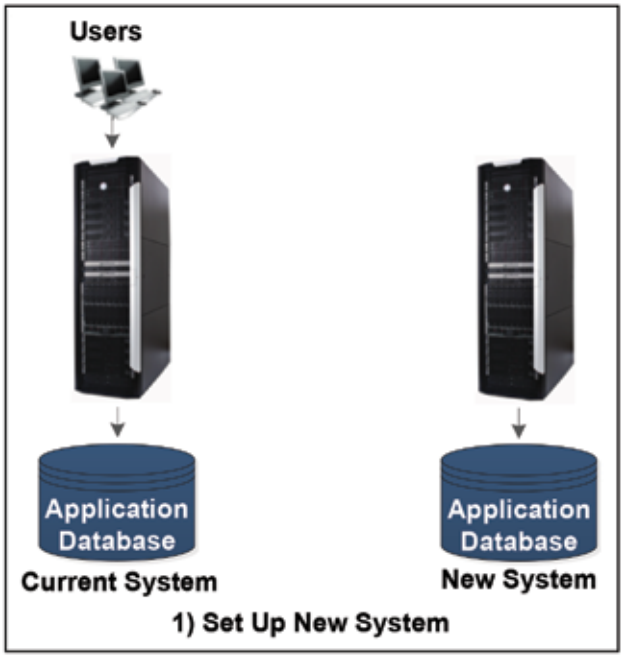

Figure 1 – HPE Shadowbase ZDM Step 1

Step 2: Load the existing source application database onto the new (target) system (Figure 2).

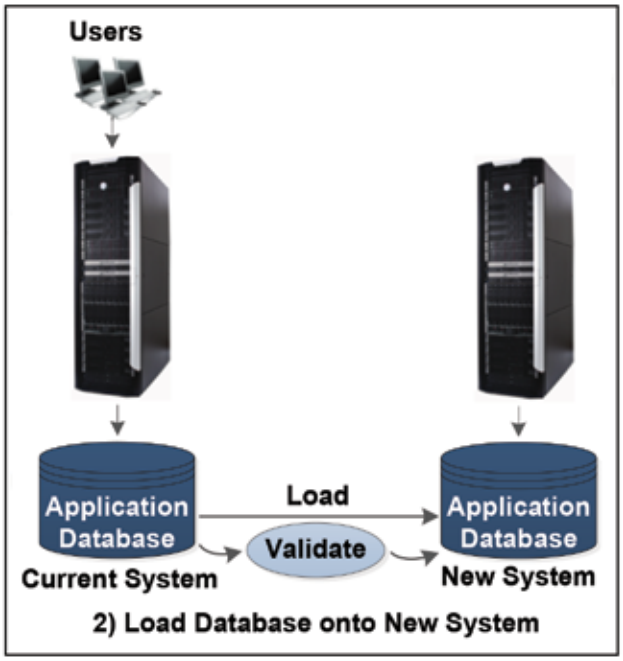

Figure 2 – HPE Shadowbase ZDM Step 2

Step 3: Once the load is complete, it is good practice to validate that the target database is an exact replica of the source database before proceeding. (The HPE Shadowbase Compare product may be used to perform this verification.) Uni-directional data replication is enabled so that the database of the new system will remain synchronized with the current system while the current

system continues to process transactions. Step 3 allows the testers to work on a current copy of the production database to verify proper functioning with 'real' data.

Step 4: Test the new system (Figure 3).

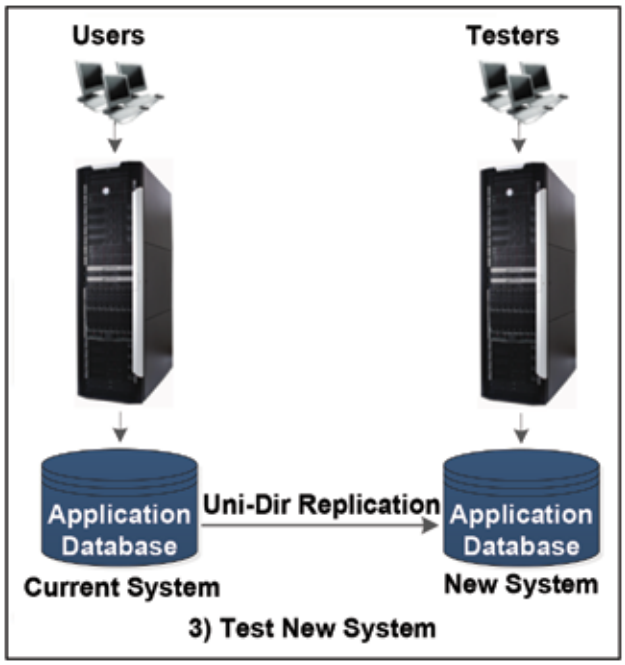

Figure 3 – HPE Shadowbase ZDM Step 3

Finally, when the testing of the new system is complete, bi-directional replication is enabled in Step 4 and all users are moved to the new system (Figure 4). Using bi-directional replication allows the original system's database to remain current when first using the new system, which greatly speeds a failback to the original system without losing any data.

Step 4: Enable bi-directional replication and move users to the new system.

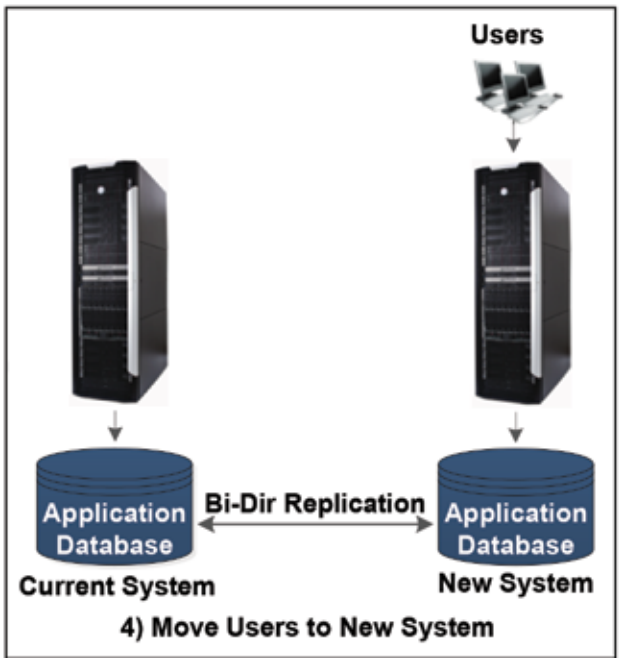

Figure 4 – HPE Shadowbase ZDM Step 4

After an appropriate period of time passes (to ensure the new system is properly functioning), the original system can be taken down (perhaps for its own upgrade).

Migration of users to the new system is therefore accomplished with no application downtime. Furthermore, the normal risks of migration are eliminated. Users are moved to a known properly functioning system, and the existing, unchanged system is always available, providing service. If there is a problem, users can quickly be returned to the original system until the problem is corrected. And, with bi-directional replication being deployed, any data modified on the new system is reverse-replicated to the original system so that no data will be lost if a failback occurs.[1]

In summary, even the most challenging of migrations (hardware, software, and location) can be undertaken efficiently and with low risk using HPE Shadowbase ZDM technology and methodology.

The Company's Migration Objectives

The company's original active/active system included a twelve-processor, four-core NonStop NB54000 system, DC1 (in Datacenter 1), and a similar system, DC2 (in Datacenter 2). The two systems were interconnected via the company's existing bi-directional data replication engine, called "Old Bi-Dir Replication," as shown in Figure 5.

The company wanted to accomplish several objectives with its migration:

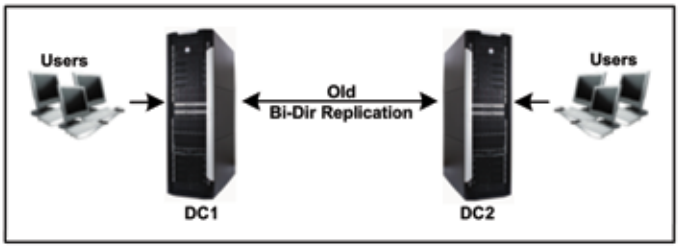

Figure 5 – The Company's Original Active/Active System

1. Retire DC1, a twelve-processor four-core NonStop NB54000 system, located in Datacenter 1, replacing it with a new twelve-processor four-core NonStop NB56000 system, called DC3, installed at a new datacenter location (Datacenter 3).
2. Move the DC2 system from Datacenter 2 to a new datacenter location (Datacenter 4), and rename the system DC4.
3. Replace its current data replication product ("Old Bi-Dir Replication") with a HPE Shadowbase active/active solution. (Its original data replication product was expensive, functionally stable and not being enhanced.)
4. Train its entire systems and operations center staff on HPE Shadowbase using a hands-on approach while the project occurred.
5. Achieve all of these goals with no application outage and no decrease in business continuity protection during the entire process. At all times, there must be at least two copies of the application and database available on two separate nodes.
6. Complete the project in a few short weeks as leases were expiring.

Figure 6 shows the company's final active/active system configuration:

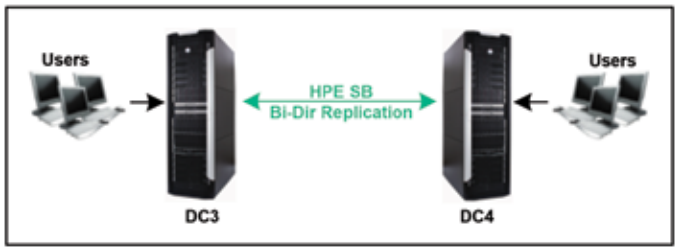

Figure 6 – The Company's Final Active/Active System

[1] For a more detailed description of HPE Shadowbase ZDM, please see the Gravic white paper, *Using HPE Shadowbase to Eliminate Planned Downtime via Zero Downtime Migration*.

All of these objectives were met with the use of HPE Shadowbase ZDM. The migration proceeded as follows:

Step 1: Install and Load the DC3 System
Step 2: Migrate Users from DC2 to DC3
Step 3: Shutdown and Move the DC2 System to Datacenter 4 (Rename it to DC4)
Step 4: Migrate Users from DC1 to DC4
Step 5: Retire the DC1 System

The following sections drill down into the details of each step.

Step 1: Install and Load the DC3 System in Datacenter 3

The company's first step was to obtain its new NonStop NB56000 (DC3) and install it in Datacenter 3 (Figure 7). The company timestamped the Audit Trail in the DC1 system and loaded the DC3 database from the DC1 database with data up to the timestamp. The company next enabled HPE Shadowbase bi-directional replication between DC1 and DC3 and began replication to DC3 from the DC1 timestamp, keeping the DC3 database synchronized with the DC1 database. The DC3 system was then exhaustively tested and verified before being put into service.

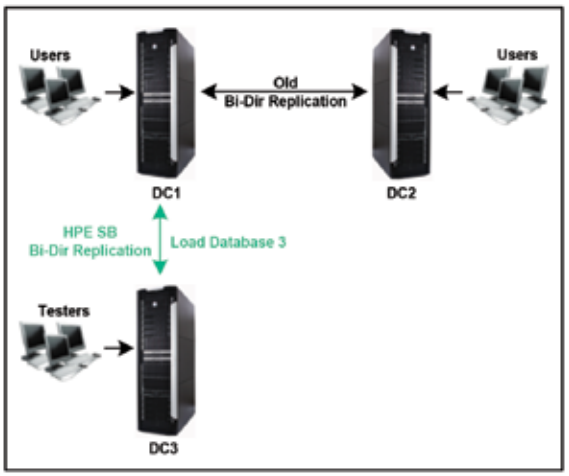

Figure 7 – Add the DC3 Node

Step 2: Migrate Users from DC2 to DC3, Isolating DC2 for the Shutdown/Move

When the test of the DC3 system was successfully completed, the migration of users from the DC2 system to the DC3 system began (Figure 8).

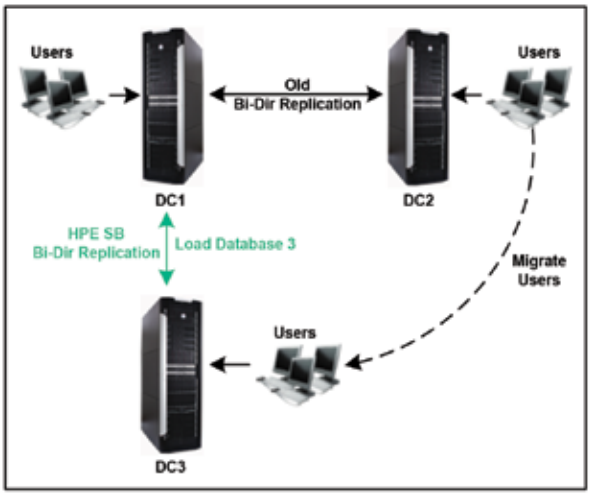

Figure 8 – Migrate Users from DC2 to DC3

Once the migration was completed, the DC2 system no longer had any users attached to it. At this point, the old data replication product that had kept the DC2 database synchronized with the DC1 database was decommissioned, and bi-directional replication between DC3 and DC2 was established using HPE Shadowbase (Figure 9). During this step, the Shadowbase configuration was verified before moving the DC2 system.

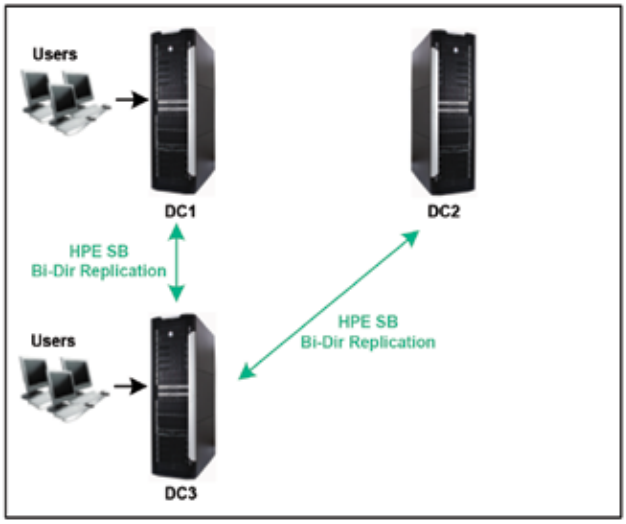

Figure 9 – Migrate Users from DC2 to DC3

Step 3: Shutdown and Move the DC2 System to Datacenter 4 (Rename it to DC4)

Replication between the DC3 and DC2 systems continued until the DC2 system move was ready, and then was paused with a timestamp in order to create a replication restart point. The DC2 system was then shut down and moved to Datacenter 4 (Figure 10).

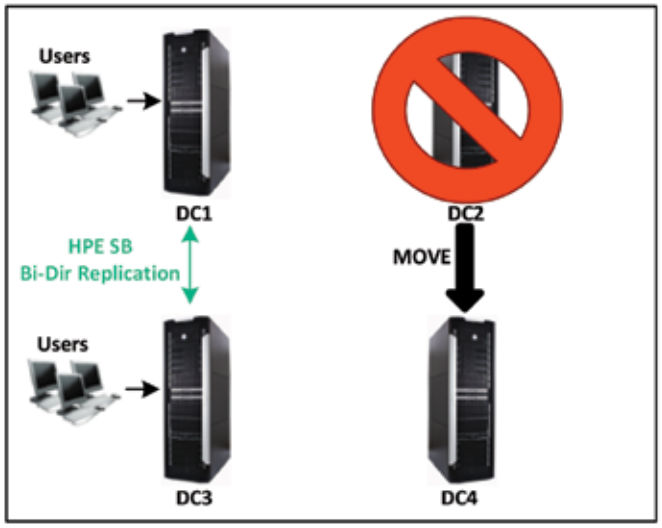

Figure 10 – Shutdown, Move, and Rename DC2 to DC4

The previously validated HPE Shadowbase data replication configuration was then deployed between the DC3 system and the DC4 (previously DC2) system (Figure 11)

The DC3 queue of database changes that had built during the move was then flushed to the DC4 system by the HPE Shadowbase replication engine from its timestamp replication restart point, thereby making the DC4 database current. The two systems were continuously kept synchronized via the HPE Shadowbase data replication engine.

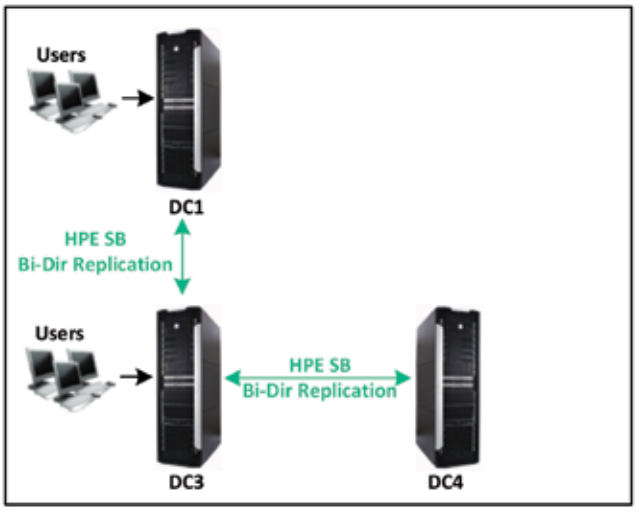

Figure 11 – Synchronize DC4 with DC3

Step 4: Migrate Users from DC1 to DC4 to Isolate the DC1 System

Next, the users served by the DC1 system were incrementally moved to the DC4 system (Figure 12).

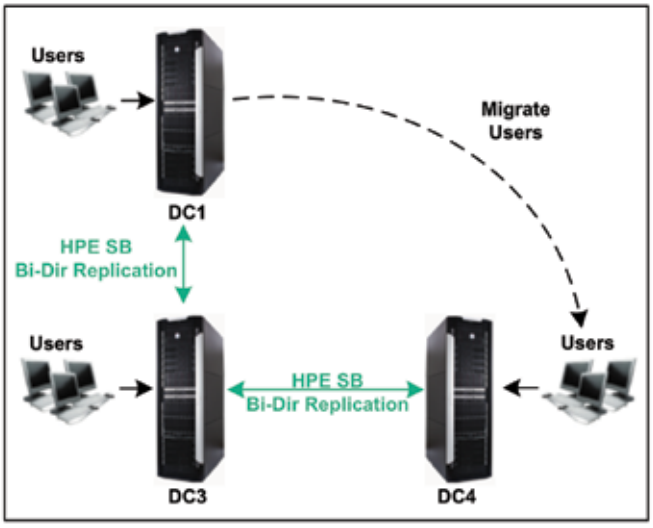

Figure 12 – Migrate Users from DC1 to DC4

Step 5: Retire the DC1 System to Create the Final DC3/DC4 Active/Active Solution

After all the users were migrated to the DC4 system, no users were left on the DC1 system. So, the HPE Shadowbase data replication engine connecting the DC1 and DC3 systems was decommissioned, and the DC1 system was retired, resulting in the final configuration (Figure 13).

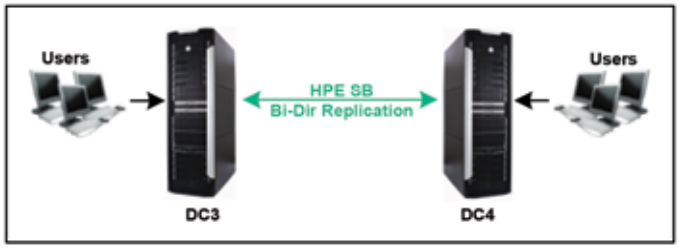

Figure 13 – Retire the DC1 System, Final Configuration

Summary

This migration presented a number of challenging parallel requirements:

- Complete within a few weeks
- No application downtime
- Full business continuity protection maintenance at all times (meaning at least two systems continuously available in the event that one system failed)
- Physical replacement (upgrade) of one hardware system
- Replacement of the existing data replication software product with a new product (HPE Shadowbase)
- Relocation of all systems and networks to new datacenters in other cities
- Minimal risk, only proceed to the next step once the current step is tested and proven
- Easily accomplish failback of each step if necessary

These requirements were undeniably challenging, but the fact that they all were met is testament to the power of the HPE Shadowbase ZDM approach, which allowed for the migration of systems with no application downtime. The company used HPE Shadowbase ZDM to change its data replication software product, to move its active/active system to two new cities, and to perform a system upgrade in the process. During all of these activities, the company suffered no application downtime.

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.

Paul J. Holenstein is Executive Vice President, Gravic, Inc. He has direct responsibility for the Gravic, Inc. Shadowbase Products Group and is a Senior Fellow at Gravic Labs, the company's intellectual property group. He has previously held various positions in technology consulting companies, from software engineer through technical management to business development, beginning his career as a Tandem (HPE NonStop) developer in 1980. His technical areas of expertise include high availability designs and architectures, data replication technologies, heterogeneous application and data integration, and communications and performance analysis. Mr. Holenstein holds many patents in the field of data replication and synchronization, writes extensively on high and continuous availability topics, and co-authored Breaking the Availability Barrier, a three-volume book series. He received his BSCE from Bucknell University, a MSCS from Villanova University, and is an HPE Master Accredited Systems Engineer (MASE). To contact the author, please email: SBProductManagement@gravic.com . Hewlett Packard Enterprise directly sells and supports HPE Shadowbase Solutions (www.ShadowbaseSoftware.com); please contact your local HPE account team.

"NonStop strives to find areas in which to contribute as we march into the ever-changing future."

Future Trends and the Relevance of NonStop

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

2017 and beyond promise a lot of change. This continues from the accelerated pace of change we've been experiencing for quite some time. NonStop strives to find areas in which to contribute as we march into the ever-changing future.

A brief retrospective of NonStop and change. Those that follow NonStop know that early capabilities were demonstrated in cloud computing back in 2011 and 2012 (see Figure 1). GuardianAngel later released as MaRunga demonstrated NonStop's ability to both support cloud burst – now known as hybrid cloud computing and an ability to provide some NonStop attributes off platform. By off platform we mean a NonStop capability to watch over and respond to events occurring on Linux systems or even, as demonstrated, applications running in the Amazon public cloud.

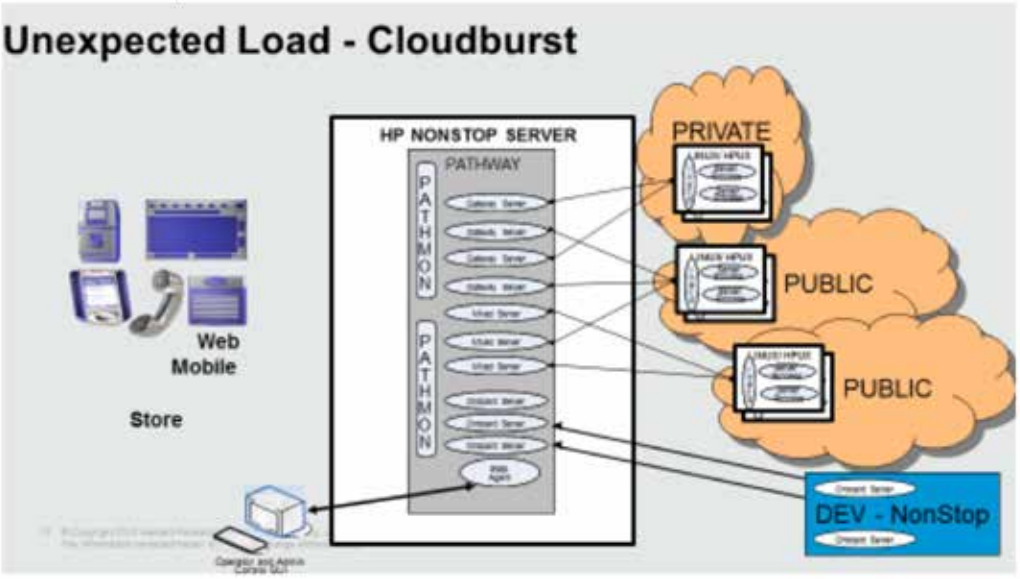

NonStop Cloudburst - Figure 1

In Big Data, NonStop was presented in the exciting area of real-time data-in-motion processing as we teamed up with STRIIM and Network Kinetix to demonstrate real-time analytics on data in flight. This directed NonStop into the Internet of Things and what will become a massive data-in-motion avalanche of information (see Figure 2).

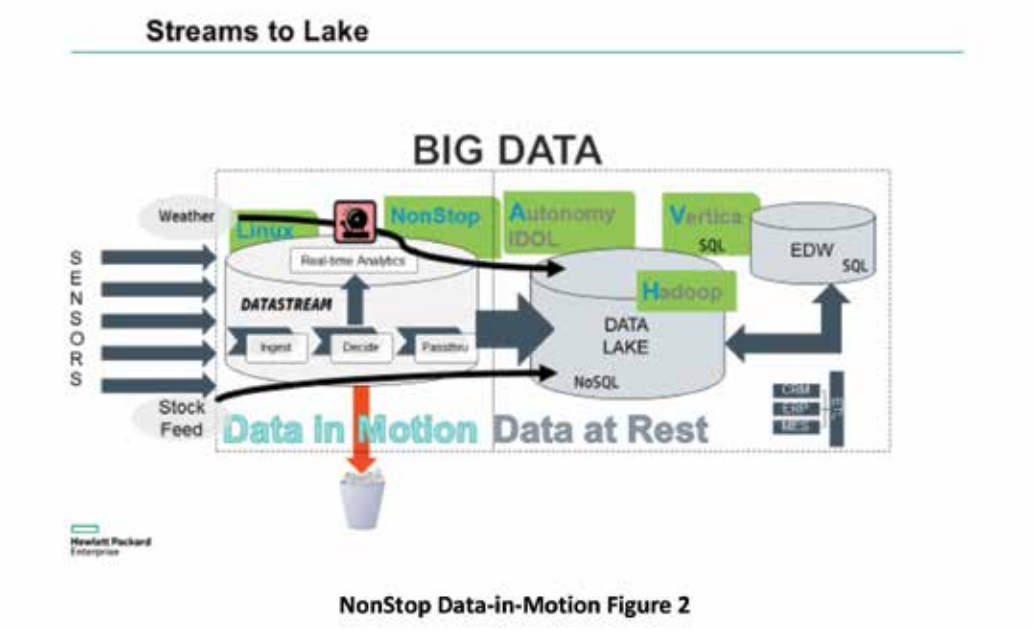

NonStop Data-in-Motion Figure 2

Of course the best place to observe, analyze and react to this information is at the edge, where the data is created. In 2015 NonStop demonstrated direct connectivity to sensors and an ability to collect this information using open-source middleware (Mosca and Redis) and an ability to communicate back to the sensors thanks to our node.js (Bombora) capabilities (see Figure 3).

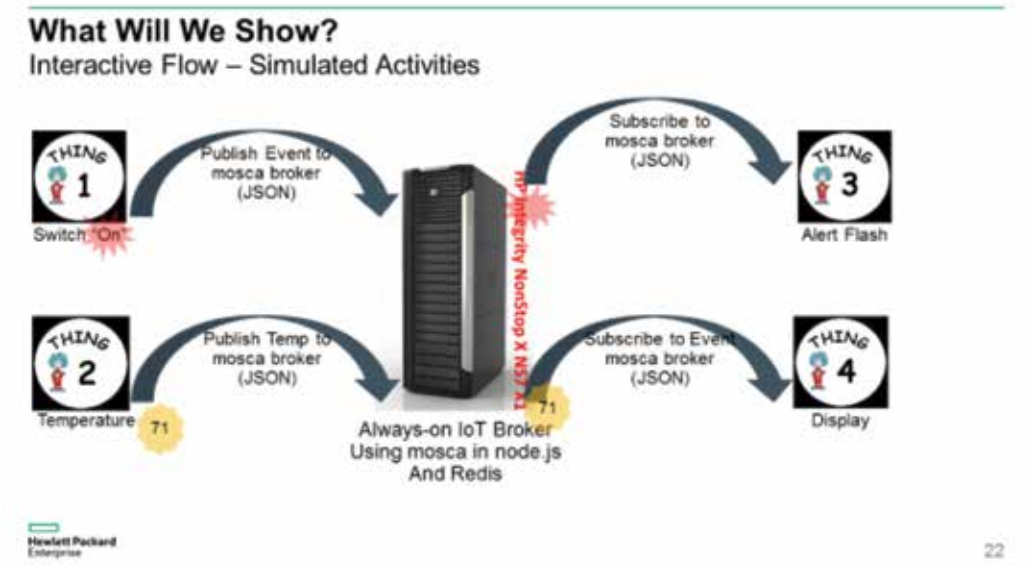

NonStop Internet of Things Sensor Demo Figure 3

This demonstrated NonStop as an edge-capable Operating System which could support Hybrid Transactional and Analytic Processing (HTAP – a 2014 Bootcamp talk).

The demonstration clearly revealed a capability for NonStop in the Industrial Internet of Things and the capability to integrate IT with Operational Technologies (OT – which was the subject of a 2013 Bootcamp talk). At this year's Bootcamp there were talks on Blockchain/DistributedLedger Technology and Field Programmable Gate Arrays (FPGA) and NonStop (see Figure 4 on page 34).

Of course the big news was the virtualization of NonStop and the capability of running on almost any x86 device. NonStop has demonstrated an ability to quickly understand and add value to industry trends and directions. So what are we seeing for 2017 and beyond?

Mobile payments, after much hype in previous years saw some uptake in 2016. For example Apple announced that 1 million Apple Pay users were being added per week. Most of the top retailers rolled out payment capabilities to their mobile apps (CVS, Kohl's, Walmart, etc.). As mobile

NonStop has demonstrated an ability to quickly understand and add value to industry trends and directions.

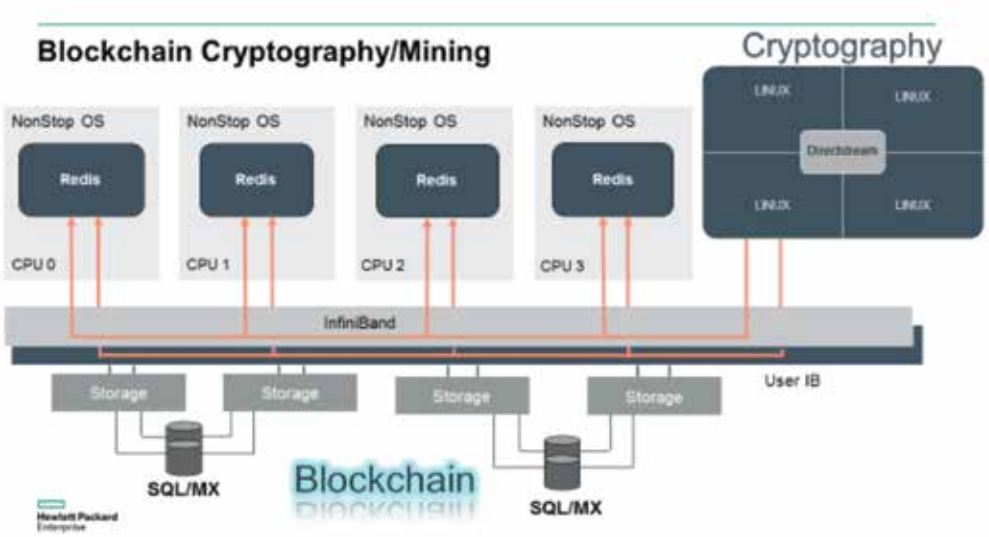

NonStop and DLT/Blockchain with FPGA cryptography – Figure 4

payment options multiply application developers will be interested in creating a simple and unified experience for mobile payments such as Google's Hands Free, an app that allows owners of newer Android and iPhone devices to make in-store payments without reaching for their phones or wallets. What can't be sacrificed is speed, security and accuracy in the transaction. A fast growing frontend mobile payment market will require a solid, reliable backend infrastructure. So NonStop will continue to play a key role in the financial infrastructure while supported new payment methods. As always, behind the scenes. As always, indispensable (see Figure 5).

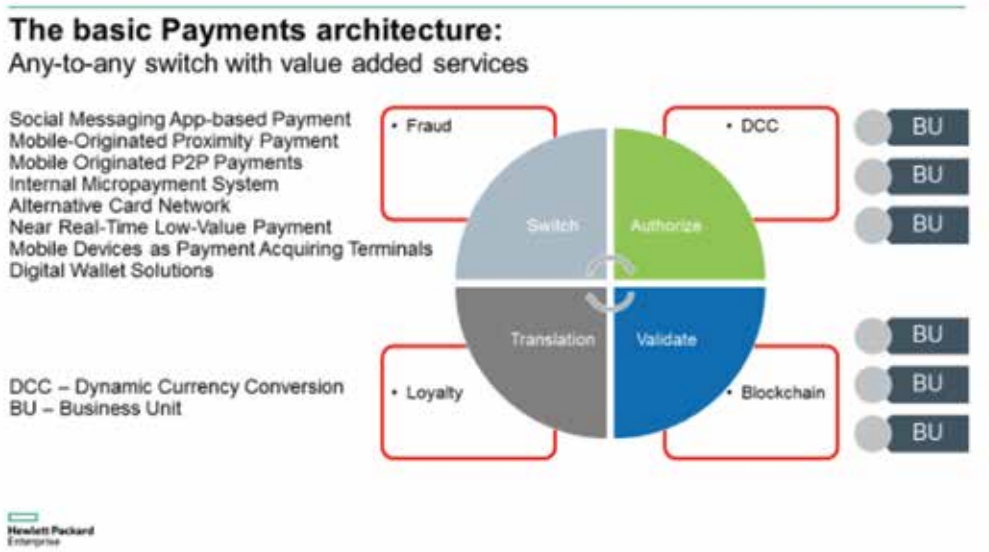

NonStop Payment Hub Figure 5

The Industrial Internet of Things is being used by manufacturers to increase efficiency through the use of automation, "many decisions are now starting to be automated based on data and analytics, often in real time," according to a December 2016 study by research firm IoT Analytics GmbH:

"Cheaper sensors and integrated information are now making shop-floor entities smart agents which can process the information to [make] autonomous decisions. In this context, we may see smart processes and smart products that communicate within this environment and learn from their decisions, thereby improving performance over time."

The basis for the decisions and the automation will be predicated on Machine Learning. This has become the great hope of the future and much effort is being spent on machine learning. At a recent lecture I attended by a professor from the University of California, Irvine who was specializing in this area, caution was recommended. In an experiment using machine learning algorithms a team developed a 'machine' that could tell the difference between a picture of a wolf and a picture of a dog with 95% accuracy. After much back slapping and praise, they dug into what the machine really

'learned' and was actually evaluating. It turned out that the learning pictures had led the machine to conclude that any picture that had snow in it equaled 'wolf'. Any picture without snow therefore equaled 'dog'. So as it turns out the machine wasn't even considering the animal itself. Once they changed the learning pictures the accuracy dropped dramatically along with their spirits. The point is we have a long way to go before machine learning is really productive. That being said a gateway on the edge applying these rules and analyzing incoming data streams for automation and decision making will need to be very fault tolerant. With shop floor control systems making their own decisions, a bad one could be extremely costly, if not disasterous. As those of us with experience know, many times computers don't fail gracefully. NonStop was designed from the start to fail fast (without damage) and keep the application up and accurate (see Figure 6).

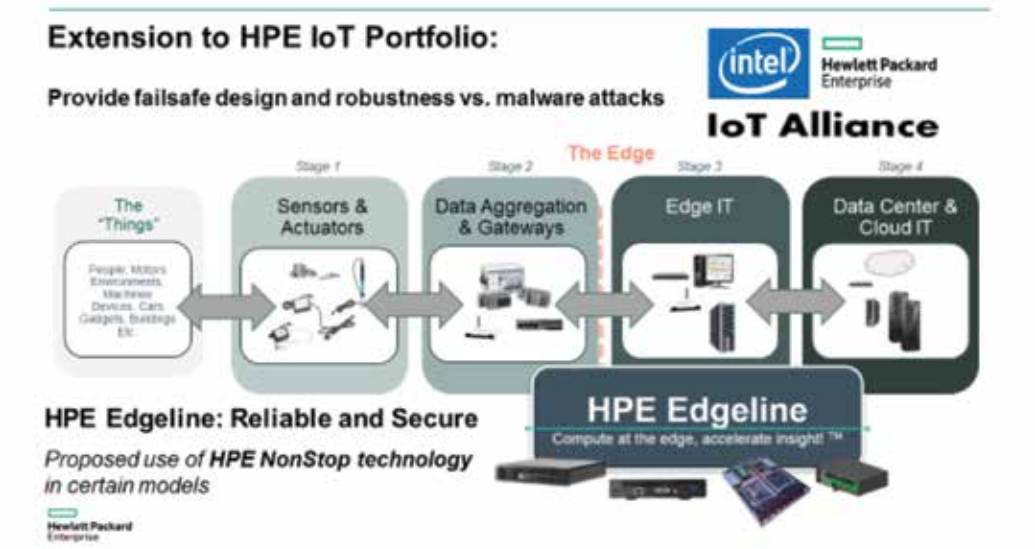

NonStop on the edge – Figure 6

The September 2016 distributed denial-of-service attack used the IoT (1.5 million closed-circuit TVs, cameras, and other devices) as a launching pad. Unfortunately, the hackers then released the code (Mirai) that was used for the botnet attack which produced other IoT-based attacks. Security continues to be a massive problem since analysis indicates hackers are growing at much quicker and more alarming rate than security trained individuals. The good guys are being overwhelmed. Here again machine learning is put forth as a potential solution and likely will be at some point in the future. Until that time Linux, Windows and now IoT sensors are open to attack.

Hopefully not but the 2016 Botnet attack might just be the warmup act for 2017. An October 2016 report on IoT risk identified common IoT devices that can be hacked in as little as three minutes (ForeScout Technologies). Consider an edge Operating System immune to standard attacks that could oversee sensor traffic and if a DDOS attack is launched stop transmission and reboot the sensors. By allocating some NonStop cores on edge gateway systems, such as the HPE Edgeline series, a bulletproof hybrid transactional and analytic processing operating system can maintain the edge and avert major attacks. NonStop will be required on the edge.

HPE is a leader in converged and now hyper-converged infrastructure. This is a method of consolidation and efficiency. Servers, networking and storage – once all separate areas of expertise – have been consolidated for greater efficiencies. As most of us learned when you consolidate systems you also consolidate risk. As more and more is put into fewer things, those things become increasingly critical. The other area of risk is the leading cause of outages, humans. Systems can be converged but we have not succeeded in hyper converging the people that run those systems. The recent 'major' AWS outage is being blamed on an operator who gave the wrong command while trying to down a few servers for maintenance. On NonStop many tasks have been simplified and automated but the problem, as it has always been, is when things go wrong. Specialists in the various areas are still needed for these times. NonStop systems have been designed for minimal involvement and can withstand a number of mistakes without compromising the application. As we roll into the disruptive future NonStop and it's 40 year mission of making computers that don't fail becomes increasingly relevent.

"As those of us with experience know, many times computers don't fail gracefully. NonStop was designed from the start to fail fast (without damage) and keep the application up and accurate.

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.

As someone who has maintained an interest in motor vehicles for nearly all of his life, it's not surprising to know that my passion for cars has dimmed very little through the years. There are many others within the NonStop community who share this passion and for a while those in the know knew that many of the engineering leadership of NonStop development held a pilot's license and raced cars. From the very first time I turned the steering wheel in earnest on a race track, there was at least one NonStop engineer standing on the sidelines providing encouragement.

As far as hobbies go it is a true measure of your commitment to your pastime if you arrive with a supporting cast. In the beginning you arrive at a race track with just your car and as the day progresses, you do all the work yourself from the simple tasks of making sure you have gas in the car (surprise; gas disappers quickly and forgetting to top up the tank has happened to many racers at the least opportune moment) to changing wheels and even doing track-side wheel alignments. However, as your commitment grows and checkbook permitting, you stop driving your car to track, preferring to trailer your "race car", even as your entourage of helpers grows. Looking at any NASCR, Indy or F1 event, the ecosystem in support of a team reaches numbers unimaginable to the hobbyist and yet, that's what it takes to win. When it comes to computer systems, the ecosystem that develops in support of a computer system ultimately leads to a winning computer vendor and there shouldn't be any surprises when reading this.

From the earliest days of Tandem Computers, there have been programs in support of independent vendors who chose the NonStop platform to build and then offer their products. The original Alliance program saw a compendium printed that was about an inch thick as it detailed every vendor and its products that ran on Tandem. Today, after many years of uncertainty, indeed waning enthusiasm, the ecosystem in support of today's NonStop systems is gaining strength. Perhaps the first signs came with the successful Boot Camp event of just a few years ago as today, having moved out of the earlier hotel into the conference center of much bigger hotels, the annual NonStop Technical Boot Camp is indelibly marked on the calendars of the majority of NonStop users and vendors. Ask any of the NonStop community stakeholders about this and the response is pretty much the same every time – HPE is making investments in NonStop and the results are proving to be just what the NonStop community needed.

If the NonStop Technical Boot Camp is the premier event for all stakeholders within the NonStop community, then the NonStop Partner Symposium is quickly becoming the premier event for all NonStop vendors. By the time this issue of The Connection becomes available, the event will be over but looking ahead at what to expect at the Partner Symposium, it doesn't take a whole lot of imagination to assume that there will be major updates on the key programs – NonStop X and Virtualized NonStop together with further insights into the progress being made to ensure NonStop SQL/MX helps drive new business to NonStop

If you missed the column by Andy Bergholz in last month's issue of The Connection, News from HPE's NonStop Enterprise Division, you may want to go back and read it. It spells out much of what are the goals for NonStop SQL/MX in the near term and is something everyone in the NonSotp community should be familiar with – "take a close look at what it would take to start migrating those applications to SQL/MX."

This was Andy's admonition even as he wrote of how "Application portability is key to continuing to grow the NonStop business." And growing the NonStop business is the sure-fired way to further build out the NonStop ecosystem. The Partner Symposium was first held last year and so this will only be the second time the NonStop development and product management teams have hosted such an event but they really do need to be congratulated for investing in this event and yes, it is in direct response to what they heard coming from the NonStop vendor community and a clear indication that the HPE NonStop team really does listen!

I briefly referenced Virtualized NonStop, the new name for what many of us had written about as vNonStop. The progress being made with Virtualized NonStop is simply amazing and as impressive as the work has been done to transform NonStop to where it is really the best software platform on the planet and this is beginning to spark conversations everywhere I go. For this program to get funded, HPE identified a key market where, in time, virtualization would become the only game in town. And the reason for this comes down to industries electing to transform their own usage of technology in ways that now look to be bellwether of yet more shifts towards virtualization to come.

If you have seen any of the recent NonStop product roadmap presentations of late, you won't have missed the reference to the Telco marketplace. "It's not hyperbolic to say that telecom trends in 2016 will reveal the direction of operator services for years to come. Three revolutionary models for service-building - the cloud, software-defined networking and network functions virtualization (NFV) - have emerged and been tested enough in labs to validate them at the functional level," as one writer to the tech publication, TechTarget, noted in the post of late 2015. However, to meet the needs of this key NonStop market, HPE needs its partners and a thriving ecosystems surrounding NonStop to ensure that "the latest technology partners in key markets" know NonStop.

We may not all like to drive cars on race tracks but whether you like to sail, play hockey or indeed participate in any organized activity, it's all about teamwork. And today, NonStop needs the team and the last couple of years have provided ample proof that there's a big team out there and they are focused on NonStop. Ecosystems are important and the fact that HPE is supporting NonStop vendors is proof enough that the future looks good for NonStop and whether your penchant is for physical or virtual deployments of NonStop, the variety of offerings should lead to even more markets opening to NonStop in 2017 and beyond.

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

SecurDPS Enterprise Data Protection

Data-centric approach for enterprise data protection across heterogeneous platforms

- **Protect sensitive payments and customer data**
- **Includes tokenization and encryption**
- **Neutralize impact of data breaches**
- **Reduce compliance audit scope and costs**
- **Enable PCI compliance and tackle GDPR legislation**
- **Integrates transparently with no software changes required**

Already protecting millions of PANs for customers today

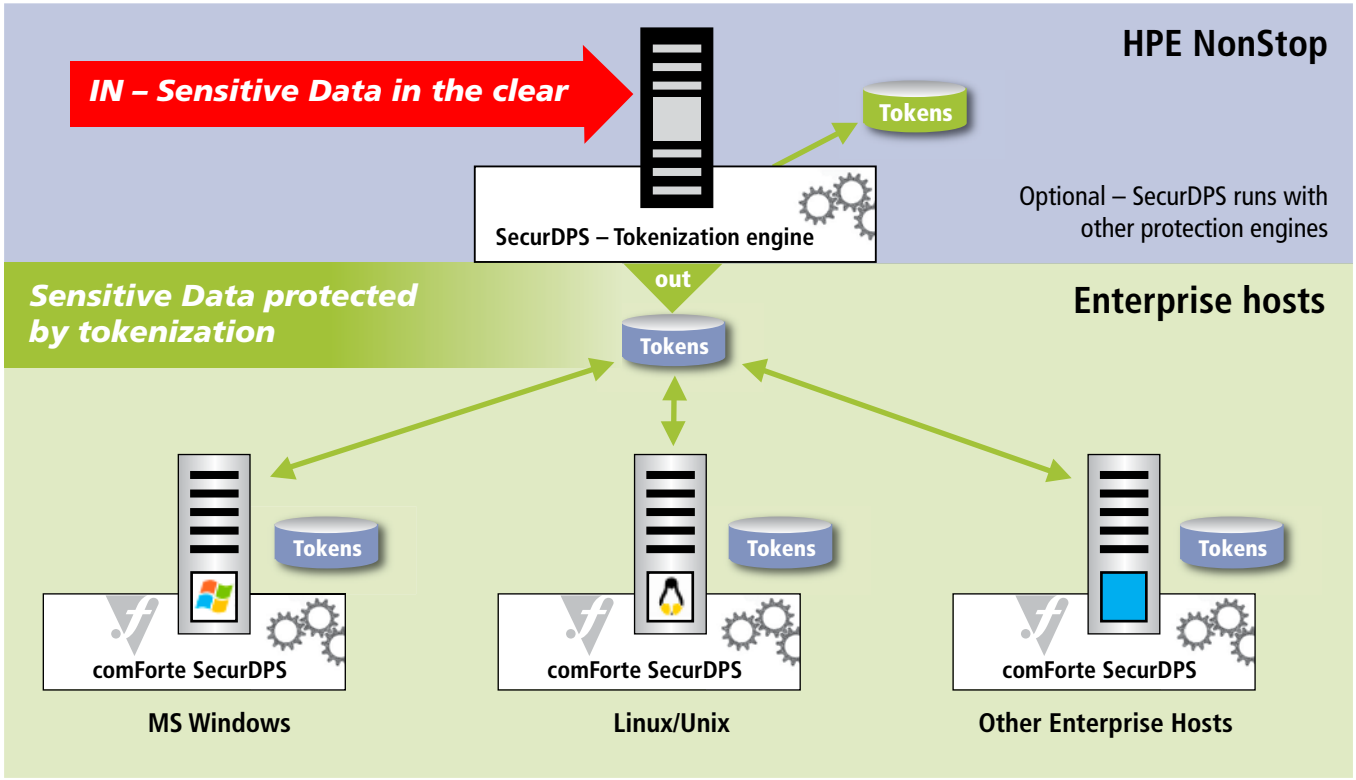

Learn more at comforte.com/SecurDPS

comforte® better always on