CONNECT

Your Independent HP Business Technology Community

The Connection

A Journal for the HP Business Technology Community

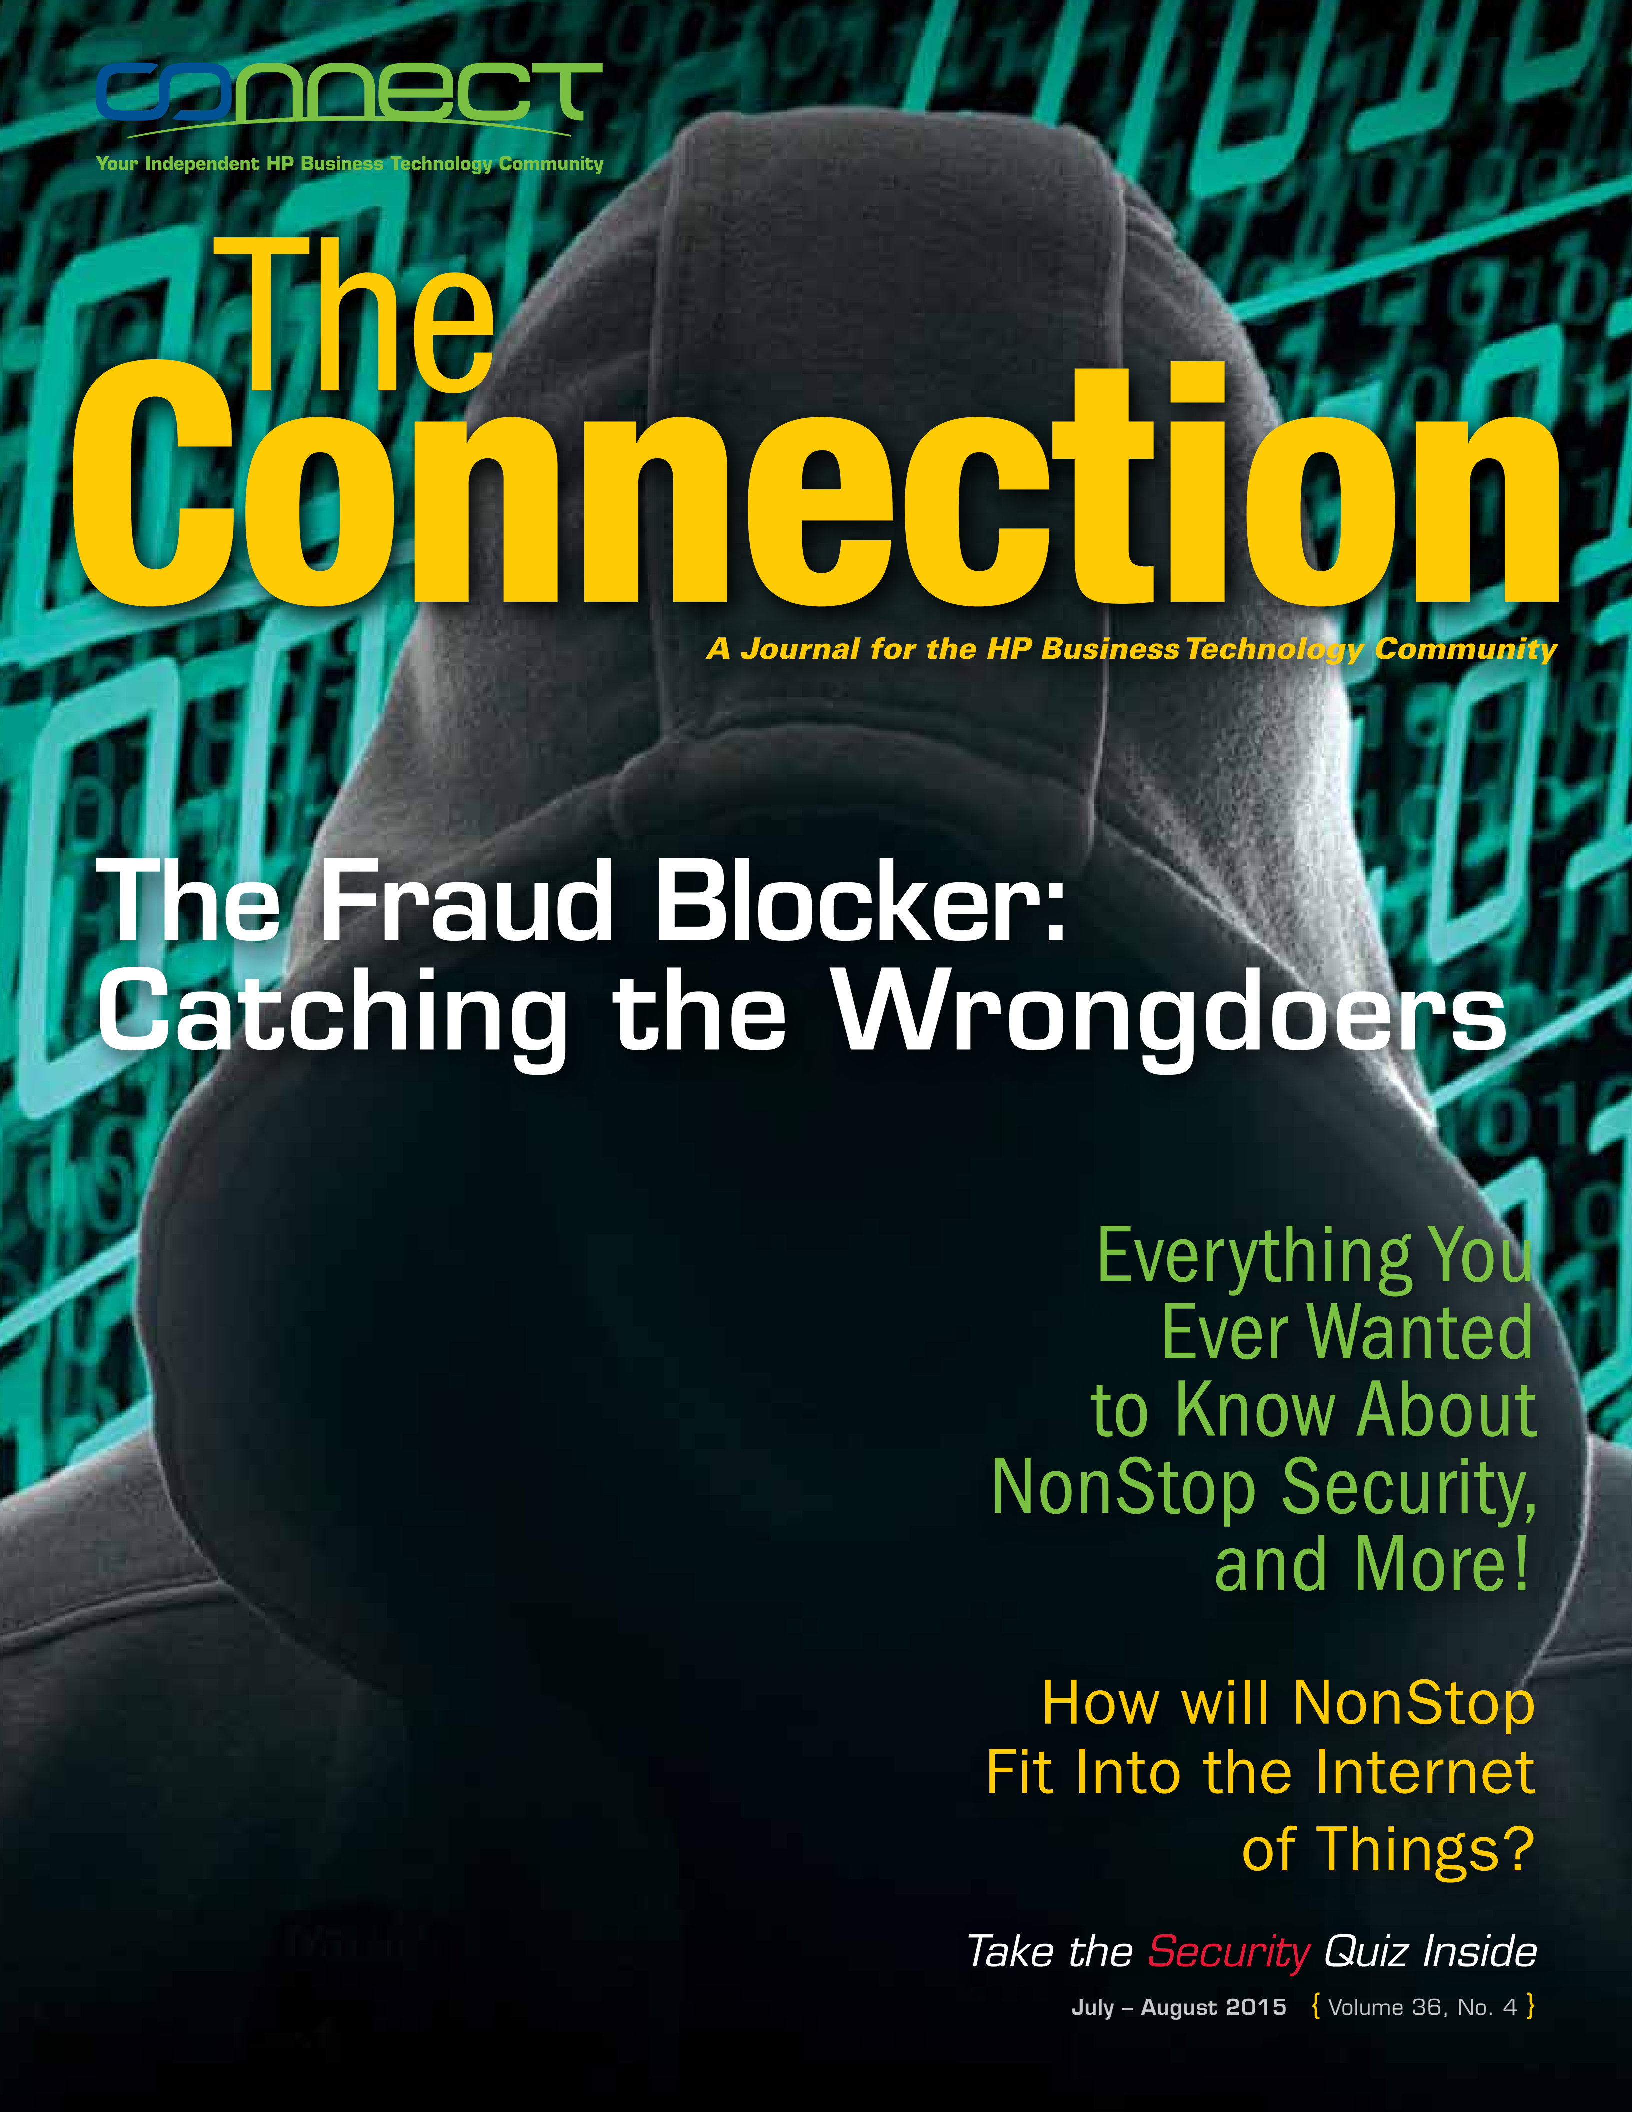

The Fraud Blocker: Catching the Wrongdoers

Everything You Ever Wanted to Know About NonStop Security, and More!

How will NonStop Fit Into the Internet of Things?

Take the Security Quiz Inside

July – August 2015 { Volume 36, No. 4 }

XYGATE® Data Protection

All the benefits of SecureData™
No application changes

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

XYPRO®
Mission Critical Security

HP Security Voltage

Learn more at
xypro.com/XDP

©2014 XYPRO Technology Corporation. All rights reserved.
Brands mentioned are trademarks of their respective companies.

Get In-Depth Monitoring of NonStop Systems.

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

ALTERNATIVE THINKING ABOUT SYSTEMS MONITORING:

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

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

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

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

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

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

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

Technology for better business outcomes.

HP NONSTOP ASAP

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

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

©2009 Hewlett-Packard Development Company, L.P.

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

With 0% Headaches, Hassles or Human Error.

ALTERNATIVE THINKING ABOUT DATABASE REPLICATION:

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

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

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

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

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

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

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

Technology for better business outcomes.

HP NONSTOP SDR

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

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

©2009 Hewlett-Packard Development Company, L.P.

Table of Contents

Columns...

Special Feature Brought to You by

OmniPayments

OmniPayments' OmniCrypto – Bulletproof Security for HP NonStop

Look no further than OmniPayments Inc. for bulletproof security. Our OmniCrypto suite of security products is valued by retail and financial networks that manage and deploy unique keys for automatic teller machines (ATMs) and point-of-sale (POS) devices. Flexible and customizable, OmniCrypto delivers role-based and permissions-based authentication and authorization, hardware and software encryption for data-at-rest and data-in-flight, management and migration tools to simplify and speed implementation, and digital certificates and signatures for e-commerce applications.

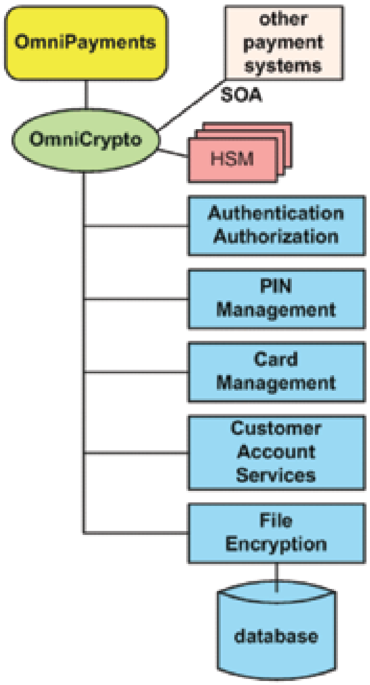

In addition to its role as a critical component of the OmniPayments Financial Transaction Switch, a popular BASE24 replacement, OmniCrypto is sold separately and seamlessly integrates with other payment systems. HP NonStop, Windows, UNIX and Linux are supported platforms.

Based on a modern, component-based design and an open SOA environment, OmniCrypto includes modules for ATM/POS transaction security • ATM key management • Public Key Infrastructure (PKI) • encryption and tamper proofing of database/files, passwords and credentials-based user authentication • and role-based Access Control Lists (ACL) for user authentication. It offers complete card and PIN management services that are fully compliant with ANSI X9.8 (3DES) and ANSI X9.24 (Unique Keys), works with third-party products to provide two-factor authentication, and supports the CVV, CVV2, iCVV, ARQC, CSC, and CVC verification codes.

OmniCrypto is PCI-DSS compliant and conforms to all industry standards. It can be accessed from other computers via standard SOA application servers so that companies can use OmniCrypto as their centralized authentication security service. Via a common interface, OmniCrypto offloads encryption to Hardware Security Modules (HSM) such as HP Atalla and Thales. For high volumes, OmniCrypto will route transactions to a pool of HSM devices to achieve fault tolerance and load balancing.

OmniCrypto's SecureChannel allows client applications to communicate securely with server applications via SSL (secure socket layer). SecureFile Transport encrypts any type of file for transfer from one place to another.

OmniPayments is a comprehensive architecture by which financial institutions acquire, authenticate, route, switch and authorize transactions across multiple input channels such as ATMs, POS terminals, kiosks, IVRs and the Internet. It supplies a full set of functionalities to support payment transactions. Based on a modern Service Oriented Architecture (SOA), OmniPayments consists of several service modules. The critical payment components are built on NonStop. OmniCloudX, running on NonStop X, hosts numerous instances of OmniPayments at a price so attractive that mid-size retailers now can enjoy the benefits of their own high-capacity transaction switches. Starts at only $5,000 per month.

OmniPayments Inc.

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

About Opsol Integrators and OmniPayments

OmniPayments is a switching solution for the financial industry. It is deployed on NonStop for the highest availability and offers all the requisite functionality to manage credit/debit-card transactions. It manages multiple devices, hosts application interfaces, and interoperates with third-party products or other systems if required. OmniPayments easily expands to provide additional functionality when needed and supplies complete security functions for every financial transaction handled. OmniPayments will survive any single fault, requires no downtime for maintenance or upgrades, and supports a range of disaster recovery solutions. Now available on NonStop X and OmniCloudX. Call us today!

*Paid Advertisement

2015 Connect Board of Directors

PRESIDENT
Henk Pomper
Plusine ICT

VICE PRESIDENT
Rob Lesan
XYPRO

PAST PRESIDENT
Steve Davidek
City of Sparks

SECRETARY»TREASURER
Glen Kuykendall
Xerox

DIRECTOR
Michael Scroggins
Washington St. Community College

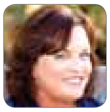

CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide

HP LIAISON
Janice Zdankus
Enterprise Group, HP

The Connection

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

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

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

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

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

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

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

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

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

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

© 2012 by Connect
All company and product names are trademarks of their respective companies.

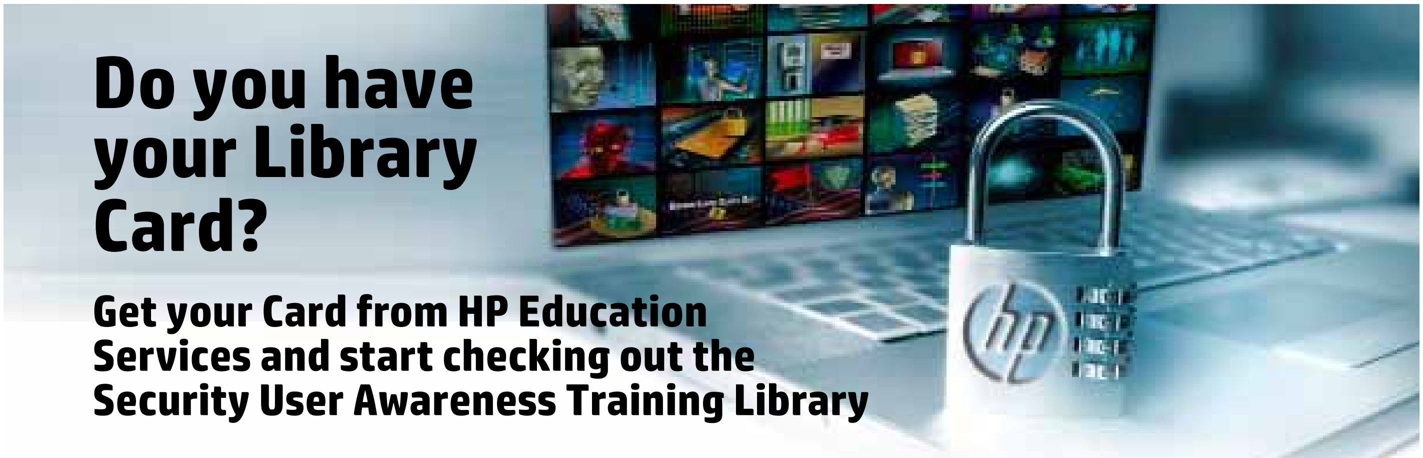

Do you have your Library Card?

Get your Card from HP Education Services and start checking out the Security User Awareness Training Library

HP offers computer-based training (CBT) that has advantages like:
- The ability to scale the training across your organization.
- Users can take training as per their schedule.
- It ensures that your program communicates a standardized message.
- It is easier to track who took the training, which is often required for compliance purposes.

Free 21 Day Trial Available Now
Library Card Available January 2015

Why choose HP Security User Awareness training?
- More than 40 engaging modules.
- Available in 28 languages.1
- Sharable Content Object Reference Model (SCORM) compliant.
- U.S. Federal 508 compliant for compliance with the Americans with Disabilities Act.
- Regularly reviewed and updated.
- Global content for global enterprises.

Get your Card...Get Secure.
Learn more at
hp.com/learn/securityawareness

Delivering protection with every transaction

HP Integrity NonStop X ensures the confidentiality, integrity, and availability that you and your customers require—with sophisticated protection of resources and data.

System-wide security is built in, giving you the ability to control and monitor user access to data and system resources and rapidly detect anomalies.

HP Integrity NonStop X.
When security matters.

hp

Make it matter.

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

© Copyright Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

News from HP's NonStop Enterprise Division

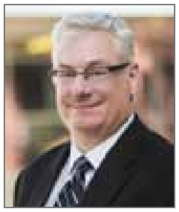

System-wide protection with every transaction

2015 is a breakthrough year for HP NonStop. Earlier this year, we launched our new family of HP Integrity NonStop X systems and we'll continue to round out the portfolio with additional products and enhancements. Stay tuned for exciting news on this front.

HP Integrity NonStop systems have delivered mission-critical computing across global industries for four decades. And as I look to the future, I see nothing but opportunities for new workloads in industries that recognize the value of a fully-integrated, x86 fault-tolerant system. Running on industry-standard hardware building blocks, including both Intel® Xeon® processors and InfiniBand system interconnect, NonStop X is that system.

Independent research tells us there is significant demand for such systems in industries with new workloads where the risk of downtime and data loss is rapidly becoming unacceptable—to both the business and to its end customers. These same industries require the highest levels of protection, and security is a key in those environments. HP NonStop provides it in a number of ways.

Enterprise-class security

HP and our partners provide a secure computing ecosystem to help you defend against ever-increasing security threats and comply with industry-specific laws and regulations such as the Payment Card Industry Data Security Standard (PCI DSS) and the Health Insurance Portability and Accountability Act (HIPAA).

We take meeting this demand for sophisticated protection of resources and data seriously—with continuous investments in new and enhanced products that provide the security capabilities your customers require.

- **On-platform security**
  HP Integrity NonStop platforms come with a security subsystem built into the environment that protects access to the system and data kept on the system. Add-on products are offered to extend built-in system security—from additional options for managing user access to configurable keystroke logging.
- **Data-in-motion security**
  HP Integrity NonStop systems include out-of-the-box SSH and TLS/SSL capabilities to secure data in transit between the NonStop server and other servers or workstations.
- **Data-at-rest security**
  HP Integrity NonStop offers the optional Volume Level Encryption product to encrypt data being stored on CLIM-based disks or tape. With the recent acquisition of Voltage, HP now offers SecureData on NonStop systems, to give customers the option of an industry leading solution to tokenize sensitive data subject to PCI DSS. Our partners offer security solution frameworks and field-level or column-level encryption.
- **Security compliance**
  HP offers products specific to the NonStop platform to help customers verify that security policies are being enforced and meet compliance requirements — avoiding breaches, fines, customer loss, or public embarrassment.

We live in volatile times where cyber criminals are growing ever more sophisticated in their attempts to break into systems and steal financial and private information. HP and our partners are taking steps to offer more security capabilities on our systems every year. To learn more about NonStop security, see Wendy Bartlett's article in this issue on security reference material, and visit the website at www.hp.com/go/nonstop/security.

Finally, I'll close with a quick wrap-up of HP Discover Las Vegas. Once again it was a fantastic event, and provided an opportunity to see the full depth and breadth of what HP has to offer. Our Mission Critical portfolio was featured in the HP Servers keynote – for a replay of that session and many others visit www.hp.com/go/discover. I look forward to seeing you at the London event this December, or at the Connect NonStop Technical Boot Camp in San Jose this November.

Randy Meyer

Randy Meyer
VP & GM, Mission Critical Solutions
HP Servers

Go Easy on Your CISO – They're Under a Lot of Pressure

Steve Tcherchian, CISSP >> CISO >> XYPRO

Threats are everywhere. And it's no secret cybercriminals are getting more organized and working more patiently to accomplish their objective: stealing your data, access to which is provided via an infinite number of channels. There are currently an estimated five billion devices connected to the internet today and Gartner estimates that number to grow to 25 billion + by 2020. That's almost 4 devices for every man, woman and child on the planet. Look around your desk. How many connected devices do you count? 5? 6? More? Further, just how many of these devices are being brought into and connected to the enterprise at work, exposing not just you, but your company to further risk? A recent article I read described the exponential increase of connected devices - including devices such as smart ice cubes that pulse to the beat of your music and monitor how much you're drinking and smart diapers that can tell you when the baby needs to be changed! Every one of those devices pose a risk at home and work. This risk increases the strain on security resources that now have to be responsible for plugging up every hole, even ones they don't know about.

The odds are against your success

Security professionals in today's landscape have to be right 100% of the time to stop criminals, whereas the criminals only need to be right once. Those are pretty scary odds considering all the different devices they need to take into account. You would get better odds if you were to wager that I'd end up being the President of the Moon. I kind of like the sound of that.

I've been in the security space for nearly my entire career. Back in the old days, security was nothing more than installing a small firewall in a locked room that also housed the cleaning supplies because someone said a firewall was the right thing to do. Back then, the only time sensitive part of the job was ensuring the antivirus software on everyone's system was up to date. Synching your phone was simply putting your Palm Pilot or Windows CE on its cradle to sync your contacts and calendar and working from home meant printing out your spreadsheets to take with you. As a CISO now, the number and magnitude of security risks that have to be factored into the day-to-day monitoring and overall corporate strategy are mind blowing. BYOD, IoT, VPN, ISO, PCI are all acronyms CISO's lose sleep over. As a profession, we not only need to worry about the outside attackers intentionally threatening everything in our organizations, but we also have to make sure our legitimate inside users are informed enough so they're not going to accidently open our systems to a vulnerability or circumvent security controls because "It gets in the way of doing real work". We have all received those tempting emails offering an all expenses paid trip to Tahiti or texts saying that we are entitled to a ten million dollar Nigerian inheritance; all we have to do is send them our SSN and corporate password to collect our riches. How do you ensure your users are armed with enough information so they know not to click or respond to obvious phishing scams as well as the more sophisticated ones designed to look perfectly legit?

Couple all that with the 100 different security tools we need to deploy from 50 different security vendors all with their own proprietary implementation and it's a mystery to me why anyone would want to willingly work in the cybersecurity space. Don't even get me started on physical security and regulatory compliance. Good luck getting any sleep at night! (Have I mentioned I'm thinking about switching careers to become a fisherman?)

You're only as strong as your weakest link

Dear Customer,
This is your bank. We forgot your social security number and password. Why don't you send them to us so we can protect your money.
Sincerely,
I. B. Banker

You can spend millions of dollars on the fanciest security hardware with cool flashing lights and engage every vendor for their "Next Gen Security Whatever" solution, but all that aside, one gigantic vulnerability still exists in every organization. PEOPLE. Users are the single largest vulnerability when it comes to cybersecurity. In fact, studies show that 95% of successful security attacks are the result of human error - that is a scary number. Users can be manipulated into giving up sensitive information. Users can forget proper protocols and passwords, they can even forget that they're not supposed to click that link!

A proper security awareness program with frequent reinforcement messages that advocate vigilance will help arm your users with the knowledge needed to protect your organization. Most regulatory compliance and security frameworks incorporate and in fact require a security awareness program for users. There is no use locking all the doors and windows when the users are going to continuously open them again. Informing your users why and how they should keep windows and doors locked empowers them with the information they need to turn them from our biggest vulnerability to our greatest asset.

Security is no longer an IT problem, it is a business problem and security professionals at all levels need to work together to minimize the risk of the "people factor" and maximize the success of their security posture. As such, consider the following points to assist in the success of your security program.

1. Executive Support – A security awareness initiative without executive support will not get much traction with the rest of the organization. If the executives don't see the value, other key departments you need to work with won't either. Getting

this level of support can be difficult, even though there is a correlation between compliance and awareness efforts and reducing corporate risk.

2. Peer and Interdepartmental Support – Everyone is busy, but just as important as getting executive support is getting departmental support. If people don't see the direct value of your program, it's not likely to succeed. Tailor your message to the specific department. Partner with them. Make key departments such as legal, HR and finance understand they have a vested interest. As security practitioners, we know we have to work up, down and sideways to get the support needed for our initiatives to be successfully implemented.
3. Walk before you run - When it comes to security awareness, there isn't a one size fits all solution. Depending on your industry and company culture, you will need to evaluate your audience and their level of expertise and cater your program to allow your audience to extract the most value out of your message.
4. Have a Plan – I cannot stress this enough. After getting everyone to support your program, you need a way to execute and measure the success of the initiative. Compliance programs, such as PCI have their suggested methods, but these should be used as a baseline or framework to build on and customize for your organization. Create a 90 day plan and identify key performance indicators to keep your program successful and progressing.
5. Reinforce, Reinforce, Reinforce – It doesn't have to be weekly stern emails telling people not to click on links in emails sent from Romania or not write their corporate passwords on a sticky-note and place it under their keyboard. Be creative and be consistent. Make people want to join and engage in the program. Create a cybersecurity week and celebrate a theme. Have others share experiences. Give out prizes for participation. Create posters. Get everyone engaged. Employees feel engaged in the program if they can relate personal experiences to the message.

We all wear multiple hats when it comes to cyber protection and security awareness. We must protect ourselves against both internal and external threats, inform our legitimate users about what not to do, sniff out those looking to harm the rest of us. It's an ongoing effort that requires teamwork. Make sure your team is educated, motivated and armed with the knowledge and tools necessary to do their part and they will.

Steve Tcherchian, CISSP, is the CISO for XYPRO Technology. With almost 20 years in the cybersecurity field, Steve is responsible for overseeing XYPRO's risk, compliance, infrastructure and product security to ensure the best security experience to customers in the Mission-Critical computing marketplace.

Test your Security Knowledge!

Try a short security quiz to test how well you and your users' cyber security knowledge stacks up.

1. Multi-Factor authentication consists of something you know and
A. A secret question challenge
B. Your email address and password
C. Something you are or something you have
D. A properly configured authentication service, such as LDAP

2. When a system permits access to a file or program, what is it doing?
A. Authorizing C. Auditing
B. Authenticating D. All of the above.

3. Password rotation is the practice of
A. Rotating passwords among co-workers
B. Changing passwords after a certain period of time
C. Using variations of the same password

4. Once an encryption key is rotated out of service, all copies of it should be destroyed immediately
A. True B. False

5. When you add idle timeouts to an interactive environment (e.g. TACL), what have you accomplished?
A. Caused batch to be more complicated.
B. You've inconvenienced the users.
C. Prevented users from leaving their terminals while they are logged on.
D. Provided limited protection against the fraudulent use of that user's session
E. All of the above.

6. What is the concept of "single signon" intended to provide?
A. A single use userid for temporary access to a system, good for only one session logon.
B. A single userid for all computer systems that provides access for a number of users in a job function (e.g. Operations Users).
C. A userid for use on all computer systems in a network for a single user.
D. A userid and authentication method that provides access to and correct authorization (without re-authentication) to all computers systems for all functions for which the user is authorized.

7. What is the concept of "least privilege" intended to provide?
A. No access beyond what is required to do your job.
B. Authorized access to resources of the system.
C. All the access that is required to do your job.
D. B and C.

8. Data Tokenization is the concept of
A. Replacing sensitive data with a non-sensitive surrogate value
B. Using a mathematical computation to derive a new value of the data
C. Encrypting data
D. None of the above

9. POODLE is a vulnerability in which protocol?
A. SMTP
B. SSL
C. HTTPS
D. RADIUS

10. Spear-phishing is what form of an attack?
A. SQL Injection
B. Dictionary Attack
C. Social Engineering
D. None of the above

11. Database Auditing can be used for
A. Investigating suspicious activity
B. Deterring users from performing inappropriate actions
C. Monitoring access
D. All of the above

12. Connecting to coffee shop or other public WiFi network is completely secure as long as the WiFi network uses SSL or another form of data encryption
A. True B. False

13. Firewalls, encryption and antivirus provide protection against social engineering attacks
A. True B. False

14. How can you determine if an email is fake or spam and you likely shouldn't click on it?
A. Asking for account numbers or passwords
B. Poor spelling and Grammar
C. Offering you a trip or reward.
D. All of the Above

15. Spyware and malware are usually distributed through
A. Your friend's USB stick
B. Free games
C. Free screensavers and toolbars
D. All of the above

Answers on page 43

NonStop Innovations Deep Dive

A Breakdown of EMV and Payments Security with **Terence Spies**

Gabrielle Guerrera >> NuWave Technologies

Terence Spies

The United States is barreling towards the Europay, Mastercard, and Visa (EMV) standard to help bolster card payment security. These new (for North America) requirements demand sweeping technology changes across the payments, retail, and financial industries, making it a must for many NonStop companies. How will card payment technology be altered? What does each industry have to do to prepare for EMV, and should they already be going beyond these requirements to bolster security?

I spoke with Terence Spies, the chief technologist at HP Security Voltage (formerly Voltage Security) and an expert on cryptography, about the importance of the upcoming EMV mandate. We also discussed the importance of mobile payments, and the changes they will provoke for any company that handles payments.

Terence Spies has over 19 years of security and systems software development experience. Prior to joining Voltage, Terence worked at Asta Networks as the director of development and VP of engineering. Before Asta, Terence was with Microsoft for almost nine years, where he started the public key cryptography group and led the development of Crypto API. While there, he also designed the SSL server and client side implementations for Internet Explorer, participated in the PCT/TLS protocol design, led the development team for the Certificate Server and led the integration of the certificate server and active directory.

Terence is active within the standards community and currently serves as chair of X9F1, the cryptographic tools group of X9 whose charter is to draft cryptographic algorithm standards for use in the financial industry. Prior to the acquisition by HP, Terence served as CTO for Voltage Security, overseeing the expansion of Voltage technology into new application areas such as mobility, payments and other areas where application data security is required. Terence will maintain these responsibilities at HP Security Voltage.

Gabrielle: Since HP Security Voltage is now part of HP Atalla, what can you tell us about the history of Atalla and payment security?

Terence: Payment security has been evolving for a long time. The credit cards we use today are based on technology that was invented in the 1960's when somebody had the great idea of taking a piece of magnetic tape, bonding it onto the back of a cardboard card, and using it to transport the information that would be used to start payments. This technology was originally used in situations where there would be dial-up connections to transport that data--back in those days people assumed things like dial-up connections were secure because they were all managed by your friendly telephone company. These things started transitioning onto more open networks with more terminals and automation, which meant suddenly the security assumptions people had made around how these cards operate started getting invalidated because the market was growing.

One of the innovations that Martin "John" Atalla, the original founder of the Atalla Corporation, came up with was the idea of a personal identification number, or PIN, which was basically an extra level of security for the card. There are principles of authentication based on who you are, what you have, or what you know. You can think of a card as an authenticating mechanism that is something you have, and you can think of having a PIN as an additional factor that is something you know. John Atalla came up with this idea of a PIN and it became common practice as part of ATM and debit card transactions.

A big thing is you want to be able to transmit that PIN in a secure fashion--it has to be encrypted so someone looking at the transmissions across the network cannot simply steal PINs by recording them. This entails something that is called a PIN block, which is an actual piece of cryptography used to encrypt information when it is entered either in an ATM or at a terminal where you are doing a debit transaction. This information is encrypted and sent to a trusted endpoint, where ultimately the encryption will be translated for the issuing bank of the card. The bank verifies the PIN, but the PIN should not be accessible anywhere else in the process.

Managing cryptographic keys and the encryption process for PINs has been the centerpiece of HP Atalla's flagship product, Atalla NSP. Atalla NSP is a piece of cryptographic hardware that enables people to manage the secure keys used to predict PIN data as it traverses the network. One of the focuses for Atalla and Voltage is to go to the next level and encrypt even more transaction information, rather than just the PIN. We are now looking into encrypting the primary account number (PAN) in flight and tokenizing that PAN data since it ends up getting stored in backend systems.

The advent of large-scale database systems that are used to hold transaction data has created a new form of risk, which is seen in the volume of credit card breaches when people go into the backend of these systems and pull out the PAN data. Voltage's technology enables us to compliment what Atalla had done with PIN security; we can start using these methodologies to secure the PAN data both as it is being transmitted during the transaction and also in the system it is being stored in after the transaction.

Gabrielle: In terms of the EMV mandate, are these security standards enough, and if not, what do companies need to do to fill in the gaps?

Terence: EMV, which is also sometimes known as "chip and pin" or "chip and signature", is a technology that uses more than just a magnetic stripe on the card: there is an actual chip on the card. If you have gotten a credit card issued in the past year or so you probably noticed it came with a metallic chip on the front of it, and underneath that is actually a small cryptographic microprocessor that contains keys that will never leave that device. When you insert the EMV card into an EMV reader, it is going to send the credit card number, plus a dynamically-created value that uses the cryptographic keys from the microprocessor. What this means is there is a unique validator for every single transaction, so if somebody steals that credit card number off the wire they are not going to be able to get the cryptographic keys needed to create that dynamic authenticator. This prevents them from using that stolen card data to build counterfeit cards, which is what happens now.

EMV is a mechanism that has been around for a long time. Almost every card that has been issued in Europe and other regions outside of the United States utilizes EMV and has been used to prevent what some call "card present fraud" or fraudulent activities where the actual plastic card is being used. This is not sufficient for protecting the entire payment infrastructure: there are still "card not present" transactions where the card number is used without the authenticator. What we have seen in places where EMV has been deployed is fraud shifts to these "card not present" transactions.

I would say there are two areas where there are security gaps: one is "card not present" transactions, and the other is card number repositories on the backend. Merchants or processors that are handling card data as part of transactions will want to retain that transaction data, typically including the card numbers. They use this information for analytical purposes after the transaction has been completed for fraud detection, customer loyalty, and other uses. The card numbers will often be retained in backend databases, so EMV is going to protect the card from outright counterfeiting in "card present" situations, but with regard to "card not present" transactions and the PAN data being held in these backend databases, there are other aspects of card security large merchant processors have to worry about.

To fill in those gaps, the HP security portfolio includes technologies that enable people to do what is called "tokenizing": taking the card number in those backend databases and replacing it with something that looks like a card number and shares some characteristics with a card number, but is not the actual card number itself. This actually negates the threat of fraud and removes the incentives to create data breaches. Backend databases are no longer going to be full of genuine credit card numbers; instead they will be full of tokenized data, which is useless to an attacker when trying to do "card not present" fraud or printing counterfeit plastic.

The other thing I would say in terms of what companies need to do to fill in security gaps is there is never going to be one technology that is an "end all" solution. The simple fact is security is not a particular state that you get to, but a practice that you have to keep on refining. Attackers are going to keep getting better, and you, as a person involved in the payments industry, are going to have to keep up with those attacks. Continuing to implement practices and technologies that are going to prevent people from stealing card data is essential.

Gabrielle: How are Apple Pay and similar payment methods changing the payments industry?

Terence: What Apple Pay uses is a combination of creating an EMV-style dynamic identifier with card tokenization. When you enroll a card with Apple Pay, you are not going to store that card number on the phone itself; instead, a token is generated. The token gets placed on the phone and will be paired with a secure element on the phone. When you do a transaction with Apple Pay, this will generate a per-transaction tag, which will be sent along to a token service provider, who will then take the token and transform it back into the card number deeper in the network. What this means is that if the phone is stolen, the thief will not get the card number; instead, they only get the token which can be deactivated independently of the card. If you lose your phone, you do not have to reissue all of the credit cards you have enrolled on your phone.

The significance of this in the payments industry is it drives this idea of tokenization forward very quickly, because essentially every company that issues cards or processes cards wants to be able to deal with these tokenized card numbers. There is a lot of interest in having tokenization as a security technology when these card numbers are being built. It utilizes several technologies that have been created to secure transactions, and it deploys them in one place in a particularly secure way.

Gabrielle: How do you think this move to higher security and mobile payments will affect HP NonStop users?

Terence: The chances are if companies using NonStop as their payment host are doing debit transactions, they already have an Atalla NSP as a way to process PIN transactions. What they are going to see is that the transaction stream coming in is going to be more diversified in terms of people looking for support for P2PE (point-to-point encryption). A piece of software will actually be put on their NonStop to enable the channel from the POS device to the NonStop host to keep card data encrypted.

There will be some additional complexity in terms of the protocols for transmitting authorization messages, and other parts of the payment scheme will get more complex and will have encryption added to them. You may see NonStop situations where there are calls out to a separate tokenization server or where the tokenization is actually done on the NonStop system. There are new technologies that will start running on the NonStop platform in order to make these security technologies work for a host using NonStop as their transaction processing platform. We certainly have lots of customers who are using NonStop as their payment host and fully support point-to-point encryption and tokenization on the NonStop platform.

Do you want to learn more about EMV and the upcoming payment security changes? The most recent article on the NonStop Innovations blog dives in deeper with Terence Spies to discuss the technology behind EMV; along with tokenization, encryption, and the differences between the two.

To read more, access the NonStop Innovations blog now at www.nuwavetech.com/hp-nonstop-innovations.

Gabrielle Guerrera is the author of the NonStop Innovations blog, which is in its second year of publication. The blog discusses current issues in the HP NonStop space and features the industry's movers and shakers, as well as the latest products and services that are available to NonStop users. Thought leaders like Justin Simonds and Joe Androlowicz have participated, as well as companies like WebAction, XYPRO, 3Qube, and Oracle.

Got 30 minutes? Then start INTEGRATING.

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

NuWave Technologies specializes in HP NonStop middleware. It's what we do.

- Enable any application to communicate with Nonstop server applications through standard, secure SOAP or RESTful APIs
- Enable NonStop applications to communicate with SOAP-based Web services within the enterprise or in the cloud

Why NuWave middleware?

- ✓ *INTUITIVE*
- ✓ *QUICK IMPLEMENTATION*
- ✓ *MINIMAL OR NO CHANGES TO APPLICATIONS*
- ✓ *BUILT-IN SECURITY*
- ✓ *RUNS IN GUARDIAN*
- ✓ *SUPERIOR TECHNICAL SUPPORT*
- ✓ *LOW TCO*
- ✓ *NO ADDITIONAL SOFTWARE REQUIRED*

Learn more at:
www.nuwavetech.com/30min

Scan to learn more.

NuWave
TECHNOLOGIES

Everything You Ever Wanted to Know About NonStop Security, and More

Wendy Bartlett >> Distinguished Technologist >> HP NonStop Division

- Where do I go to learn about NonStop security?
- What are my options for securing data in motion?
- We've got a new auditor who isn't familiar with NonStop servers – how can I best educate him/her?
- We have well-established security practices, how can I benchmark them against NonStop best practices?

There are many resources that can help you answer these and related questions, including two relatively new white papers from HP. Here's a quick overview of what's available and where to find it.

HP Resources

For overviews, product data sheets, and similar materials, visit www.hp.com/go/nonstop/security. Be sure to check out the Technical white papers section. It includes a NonStop security overview, which is an excellent starting point for anyone new to the NonStop system who needs a technical overview and also serves as a refresher for the experienced crowd.

You can find product-specific security details in the manuals collections at www.hp.com/go/nonstop-docs, along with a more general NonStop Security Management Guide.

The newest member of the family is the NonStop Security Hardening Guide. Like the security overview, it is useful for multiple audiences. To make it more accessible to auditors and other readers who are not familiar with NonStop servers, the guide is structured as a mix of commentary/tutorial ("background") and how-to information ("best practices"). Its intent is to highlight what areas need attention both from a general security perspective and for individual products, so in many places it includes pointers to specific manual sections for further details. Here's a peek at the current table of contents:

- HP references
  - Documentation
  - HP security products
- Security hardening implementation approaches
- Initial system hardening overview
- User management: best practices
- Guardian file security
- Open System Services (OSS) file security
- Additional security considerations for individual subsystems
- Remote system access
- Data at rest protection
- Data sanitization
- NonStop Cluster I/O (CLIM) security
- NonStop System Console (NSC) security
- Securing the NonStop Virtual Tape Server (VTS)
- Compliance monitoring and reporting
- Hardening against known vulnerabilities
- Viruses and Virus protection
- Additional references

I'm the primary author of the hardening guide, but its development has been a collaboration with NonStop security experts in development, the field, and our partners and customers. It's a living document, so please send me corrections and suggestions for added material and new topics to include in the next version. You can reach me at wendy.bartlett@hp.com. The current version of the guide, 1.2, is now available in the manuals collections at www.hp.com/go/nonstop-docs.

Books written by partners

XYPRO has published two books about NonStop security:

- HP NonStop Server Security
- Securing HP NonStop Servers in an Open Systems World

These are very thorough references, with coverage of both Safeguard and the security aspects of other products. You do need to be aware that both books were published some time ago and do not reflect newer security features in Safeguard and other products or current thinking about certain best practices such as minimum/maximum password length.

comForte has written an e-book that includes a brief security section:

- NonStop for Dummies, comForte

On the web

There are many NonStop-specific white papers on security available from our partners, and an ever-growing flood of articles on various security topics of more general interest. See the references section of the hardening guide for some suggested reading and links to partner websites.

A word about the Payment Card Industry Data Security Standards (PCI DSS)

Actually, more than one word…

If you are anticipating a PCI DSS audit, as noted above you can use both the technical overview and the hardening guide to help educate your auditors. Before starting discussions with your auditors, you should take a close look at both the current standard and the supporting material that the PCI Council has produced, which you can find at www.pcisecuritystandards.org. In addition, several partners and consultants have written NonStop-specific white papers on PCI compliance that you should review. You'll find links to these in the references section of the hardening guide.

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

Hold On Tight – It's a Complex World and It's a Bumpy Ride!

Monitoring gives us the eyes we need to mount a better defense

Richard Buckle >> CEO >> Pyalla Technologies, LLC

When it comes to security there's a need to know about security breaches but just as importantly, there's a need to alert IT operations and business managers. However, what is becoming every bit as important as alerts is the early detection of security attacks, as they develop, from first testing of your defenses to much deeper probing. For the NonStop community this is particularly important as so much of the world's money touches a NonStop system somewhere along the transaction path and whereas other systems involved in the transaction flow may be vulnerable, there's no excuse to leave the NonStop solution's "keys to the kingdom" unprotected.

In this edition of The Connection is a number of articles on building defenses to better protect from intrusions and this includes an article from comForte. Very much focused on ensuring data isn't simply left lying around, unencrypted and in the open, comForte exhibits considerable competence both in its products and more lately, its services, when it comes to ensuring data cannot be easily stolen or otherwise compromised. "For many, tokenization is becoming a best practice for adding an extra layer of data protection to any type of data at rest," advises comForte CTO, Thomas Burg, in his article. "When access control measures fail and intruders manage to get unauthorized access to files and databases, tokenization is your last line of defense making sure that the data will be of no use to the criminal. As a result, the compromise of confidentiality, integrity or availability is averted."

This article by Burg published elsewhere in this magazine should be a good read and I encourage all to check it out. However, there's a lot more to security than pursuing the type of depth of defense that comForte energetically promotes. While I totally agree with the need for as deep a defense as we can muster, even as I am fully onboard with making life difficult for the bad guys, when it comes to the topic of security I am just as interested in the monitoring aspect of security. In other words, the consoles and dashboards where flashing red lights emanate and where spikes and other unusual activities are depicted. In a world where attacks perpetrated on an unsuspecting financial institution (FI) can range from nuisance and opportunistic probes to straight out fraudulent use of stolen cards to well-funded criminal and even government-sponsored attacks, an automated email delivered over lunch may be our first indication that an attack is under way. Yes attack; where once such adversity was described as hacking, as recently as June the Canadian government openly referred to latest breaches as an attack!

The history of NonStop is liberally supported by many who venture into management and monitoring over the years and today the NonStop community is well served by vendors solely focused on providing NonStop users with insight into all that is happening on their NonStop systems. Furthermore, there are new contributions coming from within the NonStop vendor community, that include support of Clouds and Big Data. This I acknowledged in a blog post I wrote last November, published on the WebAction web site as, Time for an Ever-watchful Guardian… In the post I proposed that with NonStop interacting with big data analytics, much like the military depends on having a God's Eye view of all that is transpiring in a theater of operation (via hi-tech radar representations et al), is just as important for today's data center managers to have a God's Eye view of their systems given that they too are under constant attack. Whether it's for systems, or networking components, or the entire data processing ecosystem (business partners included),to be alerted of developing patterns more quickly and to recognize possible fraudulent attacks, even as they appear, the integration of big data with monitoring is only going to become more important as tangible results are realized and promoted to all within the NonStop community.

"What WebAction is particularly good at providing is 'actions' generated whenever selected criteria have been met," said WebAction Cofounder and EVP, Sami Akbay. "We know that security is a priority for many enterprises and it's a complex situation to effectively monitor – there are so many moving parts in a data center these days, with systems spanning many generations - where even the simplest device possesses incredible intelligence." What strikes me about the development of WebAction is that it includes support for NonStop systems even as it provides an end-to-end platform delivering actionable insights from big data (e.g. the sum of all event logs being generated 24 X 7). Other products are likely to appear that support NonStop systems but, as of now, WebAction is the first and in so doing, is well-positioned to catch the slightest variance that in and of itself may be the first clue that an enterprise is under attack. Data center managers with a NonStop presence will be well-served by taking a good look at the integration of big data and monitoring.

Today the NonStop community is well served with monitoring solutions. My own path to Tandem Computers came about as I championed NET/MASTER and I found a home within the Distributed Systems Management (DSM) group. Looking for better ways to monitor NonStop, and indeed, integrate the monitoring of NonStop into larger enterprise offerings, I thought a deep port of NET/MASTER might just do the trick, but it was the efforts of several independent developers where traction really developed across the installed base. Products like IR-Prognosis, ESQ PNA (and follow-on features) and even the occasional free-spirited Tandem developer (e.g. WebViewpoint) that now has a a home at Idelji even as it is sold by HP as part of the HP NonStop product offering. Gresham's TOP, created by former ITUG Chairman, Tony Bond (with the help of Rod Falk and friends), is now a part of comForte, and Insider Technology, too, have monitoring products and turning to just two of these vendors, IR and Idelji – at different ends of the security monitoring spectrum, but a good representation all the same – provided additional insight not only into the challenges but also the benefits that can come with the greater monitoring of security.

If as yet you haven't caught up with what Idelji has been doing of late you may want to check out their web site and look what they are doing with Remote Analyst – you will find it under the Cloud tab. This is what caught my attention as I wasn't anticipating quite as proactive promotion of cloud support as I was – big data, yes, but clouds? "What we have been doing is to build capabilities into our product offerings that support the storage of any type of event log off-platform, into a Cloud. In so doing we can consume data from practically any source – any time there's a change to data that in turn creates an event that is logged we capture it," said Idelji Corporation President, Khody Khodayari. "These extensions to Remote Analyst have traditionally focused on performance monitoring and subsequent analysis but increasingly, customers are looking at every tool they have to use in the fight against criminal activities and even the sudden increase in disk or process activity unrelated to processor loads can hold a clue to a possible imminent attack."

When I asked IR about their take on the role Prognosis plays in overseeing security, newly appointed Chief Solutions Officer, John Dunne, was quick to respond. "Clearly, applications monitoring is very important and IR Prognosis has been providing a way for users to set thresholds and then automate the process all the way through to resolution. In so doing we are seeing Prognosis play a key role in ensuring compliance according to industry standards and mandates." Dunne then added that, "For those customers with larger enterprise monitoring products deployed then yes, Prognosis contributes much-needed alert information that helps in providing greater insight as to what's happening across the enterprise. Over the last couple of decades, IT infrastructure may have reduced in size but its complexity has expanded exponentially."

No vendor that I talked to was reticent about discussing just how difficult the task of monitoring security had become. And for Dunne, this meant, "Taming this complexity is a big challenge. Not only does Prognosis integrate well with other platform monitoring solutions, but with the work we have done on integration and support under a single pane of glass, any application or device in your entire network, for example web services, SNMP, WMI and ODBC can be monitored and you can then use this information to drive alert notifications of threshold breaches. Just as importantly, the integration of HP NonStop systems with enterprise management frameworks, such as Tivoli, Remedy, CA-Unicenter and HP Operations Manager can give those data centers where Prognosis has been set up for such a span of monitoring, additional capabilities and when it comes to security, monitoring of the whole is pretty much the only way to get results."

No NonStop solution is more at risk than those solutions offered to FIs. It would be simple to just admit that they are more at risk because that's where the money is, but nevertheless, payments solutions providers are either partnering with established monitoring companies or developing integrated solutions themselves. "When it comes to NonStop, and payments solutions," said Dunne, "IR has a version of Prognosis tailored to meet the needs of payments solutions providers and as such, we are now a partner of ACI Worldwide and provide unique capabilities when it comes to both BASE24 and BASE24eps. As we present Prognosis to ACI customers and prospects, then to have any chance of staying one step ahead of the sheer number of fraudulent activities, you need to be able to monitor all of your transactions, payments, ATMs and EFTs in real time. With IR Prognosis Payment Transactions Solution it becomes easy to spot patterns in data and to stop problems in their tracks, respond to issues faster, and prevent outages enabling any BASE24 / BASE24eps user to get back to doing what you do best."

One payments solution vendor that has developed an integrated solution is OmniPayments, Inc. For OmniPayments it's a matter of convincing prospects that their vulnerability (to attack) is lessened with an established product already accepted by FIs worldwide. "When it comes to security there's the aspect of security concerned with detecting potential fraudulent activity and then there's the more general threat we associate with hacking and that is, detecting potential infiltrators looking to retrieve personal information for large-scale financial gain, said OmniPayments, Inc. CEO Yash Kapadia. "In other words, securing individual transactions as well as securing sensitive personal information files. In both cases, we have implemented tools as part of OmniPayments to detect and thwart such attempts together with monitoring tools to quickly alert operations staff."

As for these monitoring tools that are a part of OmniPayments, "The architecture is based around a central hub, but there's two distinct points of monitoring – the Console in support of the application itself and the Dashboards that are customizable to meet individual user requirements," explained Yash. "Anything to do with trying to penetrate sensitive files and data bases would be displayed on the Console whereas attempted fraud would be detected and reported via a security-focused dashboard." Again, it's just so important these days that those flashing red lights not only are visible to those responsible for the application but also, can feed automated escalation routines so indeed, those in charge of security can be notified immediately by whatever media they prefer. In the exchanges I had with all of the vendors referenced here what struck me most was the willingness to talk openly of the effort being made to stem the rise in fraudulent activity committed on NonStop systems. It was if the well-published attributes of availability, scalability, data integrity, networking, database etc. were all part of the "nice-to-have" category, all things considered, and that the IT world had rotated significantly on its axis to where all that mattered was security. Or to express it in another way, unmonitored and poorly secured systems have no place in today's data center even as prices drop and options multiply – in today's complex IT world, we just have to get our arms around our systems and hold on tight!

Having excellent visibility to all that was taking place today of transactions passing through NonStop had become an opportunity many vendors had stepped in to enhance the NonStop users capabilities and it was the final thoughts of Yash that impressed me the most. "To date our monitoring of security has proved effective and has reduced the financial loss from fraud for many of our customers. Whereas security had been just another check item in a list of features, today it's too important to simply be left buried in a list," concluded Yash. "For many of the customer situations we face today, we often start out with security as so often we hear that the three most important components of payments solutions is security! security! security!"

Richard Buckle is the founder and CEO of Pyalla Technologies, LLC. He has enjoyed a long association with the IT industry as a user, vendor, and more recently, as an industry commentator. Richard has over 25 years of research experience with HP's NonStop platform, including eight years working at Tandem Computers, followed by just as many years at InSession Inc. and ACI Worldwide. Well known to the user communities of HP and IBM, Richard served as a Director of ITUG (2000- 2006), as its Chairman (2004-2005), and as the Director of Marketing of the IBM user group, SHARE, (2007-2008). Richard provides industry commentary and opinions through his community blog and you can follow him at www.itug-connection.blogspot.com, as well as through his industry association and vendor blogs, web publications and eNewsletters.

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

"Editor's note: Quotations in this article are largely those of clients of the author. In addition to the NonStop security monitoring products mentioned in the article are XYGATE Merged Audit (XMA) from XYPRO, which collects entries from various logs and stores them in a SQL database suitable for sending to a SIEM system, and Audit-Plus from CSP, which provides real-time intrusion detection."

Committing to Git: Integrating a Software Change Backbone Into Your Organization

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

Preamble

It was March 2015 and the NULL DOS vulnerability was discovered in OpenSSL. A number of vendors and customers were clamouring for a fix, but how was that going to happen? Who was going to fix it? When? Fortunately, a group of LINUX and NonStop heroes were on it from the moment the news broke.

It took the development team some time to come up with and certify the correctness of a fix. Then, the wheels started moving. The change was pushed out to the LINUX repositories, then to GNU and GitHub. Within minutes of the commits and tags showing up at GitHub, the ITUGLIB team pulled the change down to their repository. Five minutes after that, the platform changes were merged in, built and the tests started. Forty-five minutes after that the code was packaged for deployment to the website.

What you have just read was entirely based on true events, but is the exception rather than the rule. The integration between external entities, development repositories, and production is informal at best, and denied at worst. Tracking of changes between different contributors still mimics paper-based accounting, not state-of-the-art policies. Our IT infrastructures do not generally have a solid backbone to support themselves. This article will help you understand how you can be efficient, successful, and all those good things, using distributed version control systems as communication and transport method for your software assets. With a solid supported backbone, your organization can withstand radical technology changes, including to your backbone itself.

Operational History

In 1972, an early version control system called SCCS was built at Bell Labs for an IBM 370 system to track changes on mainframes. This system was quickly adopted into the UNIX project and became a standard. In those days, there were no client-server systems; no workstations; no Internet, no distribution methods. Even disks were new. Most production programs existed as punch cards, so version control involved physically managing boxes and boxes of cardboard. Sharing of code was like sharing books in a library. You had to borrow bits and pieces. The idea of a communication backbone existed only as a concept in the minds of researchers.

In the years that followed, there were few advances in version control technology. There was little need. Better mouse-traps were built, but all had the same basic notions: track changes on a computer's disk instead of on punch cards; record who made the change and why; store differences so you can recover old versions just in case. Along came Tandem in 1978, which was fundamentally a client-server machine. EXPAND was there too and we had production-hardened systems that communicated with each other. UNIX still had a way to go to come up to speed with that concept and ARPANET was still being built. For change managers, this presented challenges, because code could be moved between systems easily, but processes to do that were not really well understood.

By the late 1980s, people started to see the importance of code movement. Workstations were becoming pervasive. Program editing started happening on desktops, and the need for central repositories became important. Fortunately for IT, there was an intuitive understanding of the centralization that came with the origins of SCCS back on mainframe systems. Companies adopted SCCS-like solutions that supported centralization, and we still tend to operate using policies that restrict solutions along those lines.

Once such product, RMS, from the beginning was designed to move code between systems. This was intended for multiple purposes: first, to allow development code to move to production through releases; second, to provide a means where vendors could send releases of code to customers for integration into their environments. This capability was revolutionary at the time, being one of the first distributed version control systems (DVCS). The (overly ambitious) intent of that capability was to build a form of code migration middleware. If NonStop had been more pervasive in organizations, or RMS available on other platforms, perhaps that would have happened.

Time passed, the Internet happened, and thousands of developers started collaborating on joint projects like Linux, GNU, Tomcat. Products like Subversion and CVS were eventually replaced by DVCS systems including Mercurial and git because of the ability to migrate code from system to system and to identify the path code took over time to arrive at fixes. This ability has enabled, for example, the ITUGLIB team to be extremely responsive to find out about a bug in OpenSSL, receive the multi-file fix, automatically merge changes into ported code, test, and deploy the fix to the Connect Community website within about a day.

The Software Change Backbone

What is really interesting about how DVCS systems evolved is a seemingly incidental requirement: to be able to interact with other types of DVCS systems. There are bridges between Mercurial and git, Subversion and git, even Team Foundation Server and git – git being one of those common bits of enabling systems. A lot of effort has been put into these connectors. The Subversion connector even has a CVS variant that allows conversion from a CVS repository to git. A normal reaction would be to say "Oh good, so you're saying I can migrate to git. That's nice." That point of view is fine but is based on the ideological need for a central repository, which is actually no longer necessary and actually problematic for many IT departments. What you should really ask now is: "Randall, where are you going with this?" Good question.

The ability to link DVCS systems together creates a mindboggling capability. With it, we can build a Software Change

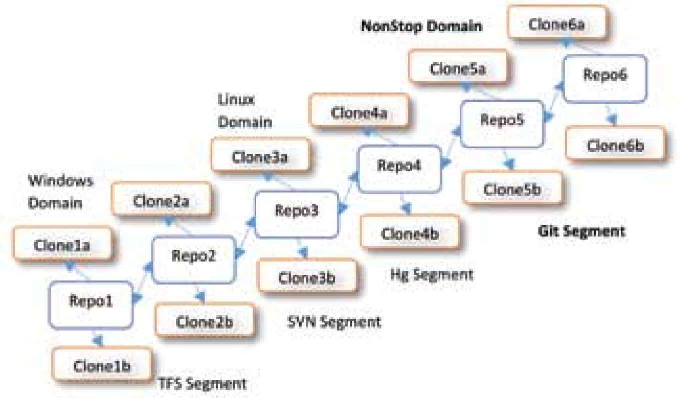

Backbone – a structure where repositories are linked together to share changes.

In a Change Backbone, developers interact with their official Repository of Record (RoR). This repository contains all of the change history that the department wants to keep for posterity. It may be before or after the Quality Control part of the organization. It may also not be the final destination for the changes. The Repository of Record can be used as a source for mirror sites that pass the changes along to other repositories on other platforms or other products. The integration between DVCS products should allow that, although you may have to do some automation work to make it happen in your shop.

A practical Change Backbone for broad use of Open Source for NonStop should include git, Subversion, and Mercurial.

The Need

In today's development environments, code is shared between platforms. Whether it is a JSON library that runs partly on a mobile device and partly inside an NSJSP 7 TS/MP server, the code needs to be managed effectively. How do we do this particularly with different development groups participating in the effort? With a Change Backbone, keeping cross-platform development in sync is simple. Let's take a look at an overly complicated example to illustrate the point:

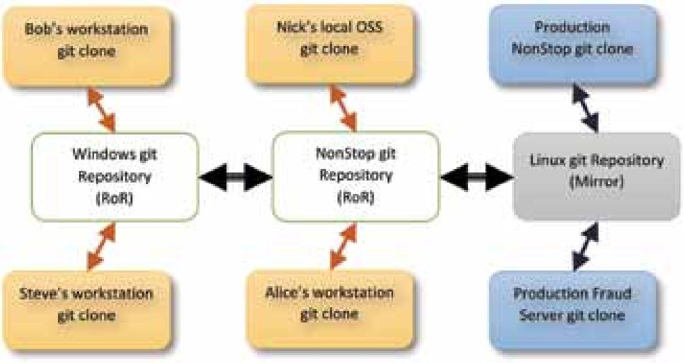
Bob's workstation git clone
Nick's local OSS git clone
Production NonStop git clone
Windows git Repository (RoR)
NonStop git Repository (RoR)
Linux git Repository (Mirror)
Steve's workstation git clone
Alice's workstation git clone
Production Fraud Server git clone

In our project, Bob and Steve are collaborating on a bit of mobile development. They are building and using JSON libraries that Nick and Alice need for their server development. The Windows and NonStop Repositories of Record automatically keep the main development branches in sync. Nick is working on server configuration definitions so is doing his work in OSS. Alice is building code using NSDEE on his workstation. Jan is coordinating merges on both the NonStop and Windows boxes, and is pulling production releases from the Linux mirror into various production environments that need the packaged products.

Movement between the Windows, NonStop, and Linux servers is automatic and continuous. From a git point of view, this can be done by hooks that are invoked when a push function occurs. As a result, Steve and Bob can publish their changes to the backbone through a simple push. This can be selective so that their works in progress, or topic branches, are not published. Nick and Alice can pick up their changes through a rebase off their own repositories, and really do not need to interact much with Steve and Bob at all; although, if they have to fix anything in the JSON libraries, they can commit and push to the NonStop repository. This will cause updates to the Windows repository that Steve or Bob can integrate into their projects. For anyone but Jan, and her support group, having the backbone in place is really not visible.

The decision to pull production releases from the Linux mirror is really a verification step for knowing the repository backups are in place. Having a mirror is a really important part of repository management and provides an active backup. There is no need for replication software to do this function in the DVCS world – it is a built-in bonus.

Implications

This brings up another really important effect: once we have a working Change Backbone, we are no longer restricted by needing to have the same product on every platform or even in every department. As long as a department's DVCS has a solid connector to the main repositories they can participate. This means that one department may use git, while another uses Mercurial, and a third uses Subversion. The limiting factor on what is available is based almost on staffing availability and budget to support the products. Even migrating from one product to another or one platform or another does not really involve major technological efforts. On a Change Backbone, migration involves setting up a new participating mirror as a destination for changes. Running multiple repository products in parallel as a transitional step becomes almost mechanical.

From the backbone's point of view, even changing the pointers to the repository of record is a very simple task that can be automated. If for some reason, you need to move or replace a repository server, you can either change its name via DNS or modify the upstream identifier (in git). Even better, because the change identifiers and access keys can be made global across segments of the backbone, developers will likely not be impacted by migrations.

Built-in Availability

We all know that NonStop is seriously available. Some other notable platforms are not; and yet, we may need those to participate in the backbone and have their code managed. This is really important when you are trying to keep track of your company's DNA. With NonStop, the backbone is always available.

Release management

Possibly the most powerful capability of a Change Backbone is to provide a method of moving code from system to system and platform to platform in a consistent manner. Imagine being able to identify a production file back to the developer who made the change across four or five distributions. With a Change

Backbone, a commit to one repository is preserved no matter where it is distributed in the organization. This allows releases to be built and identified tying source and object together in one immutable package. Production machines, regardless of platform, can connect to the backbone to pull releases, and the need to copy (and potentially miss) files between systems vanishes. Operations should start planning for a day when installation and fallback will be as clean as pulling the appropriate branch into the working production area from a Repository of Record clone.

Vendor Management Intake

Using an industry standard DVCS, vendors can participate in your Change Backbone. I can hear your thoughts: "Wait. What? How could that possibly work? I'm not giving vendors access to my network." Many DVCS systems are symmetrical, meaning that you can choose to replicate content either using a push model, or a pull model, or both. Let's suppose that your vendor has migrated to git. They make their code available on an SSH server behind their firewall, and have given you access. Your backbone can periodically pull branches from their server into a server, which can then publish content internally. The commit identifiers would be consistent from the vendor's environment right through to your production machine. Even if you had to apply customizations, those are still merge operations off of the vendor's commits. Your changes and the vendor's become part of the history that is contained in your repositories. I think the word you are now looking for is "Nifty".

But wait, there's more. Suppose you find a problem in production in one of the scripts supplied by a vendor and have to do that scary 3 a.m. fix thing. Committing and pushing that change to production's repository can initiate a sequence of events involving pushing the changes back to the developer's copy of your production branch – so that developers can see the change – and can even push all the way back to the vendor – assuming you choose to do that. The vendor can then take the change and merge it into their next release, without anyone having to email the code around. Emergency changes become part of the company's DNA through the same mechanism as any other change and now include vendor code bases.

NonStop Production

There are a few variations for NonStop, depending on how involved you get with a Change Backbone. The following is a very simple picture that assumes an RoR on a NonStop server that is not in a backbone, but has git and NSDEE involved. Because the repository is visible over EXPAND from /E, you could use standard security for access.

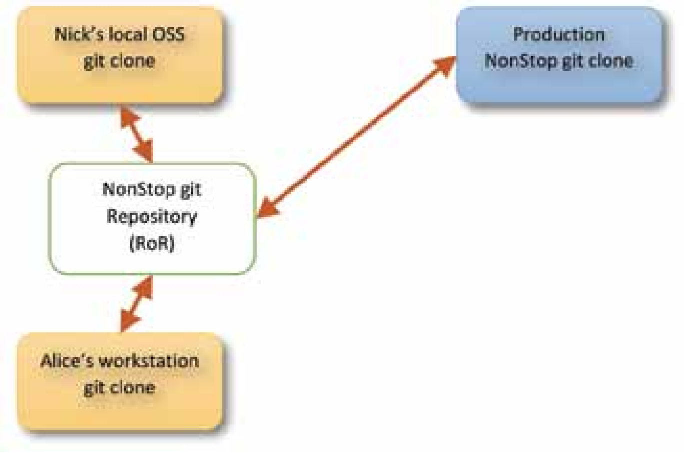
Nick's local OSS git clone
Production NonStop git clone
NonStop git Repository (RoR)
Alice's workstation git clone

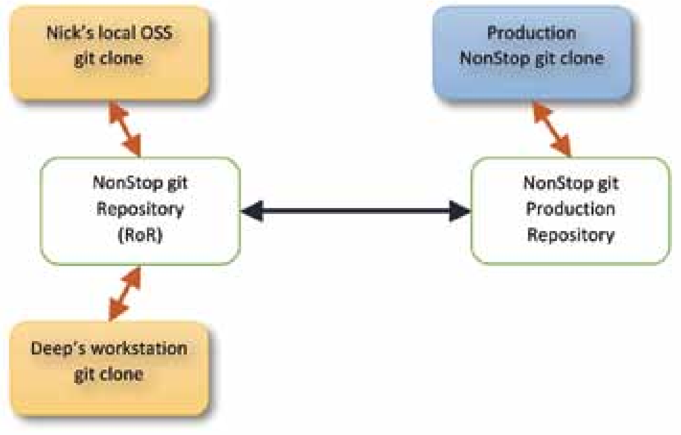
Nick's local OSS git clone
Production NonStop git clone
NonStop git Repository (RoR)
NonStop git Production Repository
Deep's workstation git clone

Of course, no one with a separate production system should do this. Another picture keeps the RoR on development, although it could easily be in production. The interesting point about this picture is that it is the first step in building a backbone. The next picture shows the introduction of a production repository. This repository would pull from the development repository to avoid any security issues. Only people with access to the production server would need access to its repository. If a production fix is made, it could be committed and pushed back to development.

The separation of repositories by security rules is actually a really important concept to the Change Backbone. Not only does each segment in the backbone have a specific purpose, platform, and product, it also has potentially distinct security rules. Where a DVCS is critically different than traditional VCS systems is that the production repository does not have to be local to the host where it is being used. In the above picture, the production repository could sit on any platform in the backbone. If we add a new repository for Quality Assurance, the picture starts to take shape:

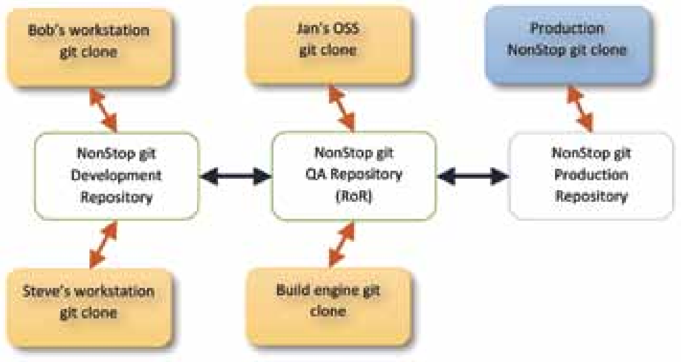
Bob's workstation git clone
Jan's OSS git clone
Production NonStop git clone
NonStop git Development Repository
NonStop git QA Repository (RoR)
NonStop git Production Repository
Steve's workstation git clone
Build engine git clone

Once QA is involved, it generally makes sense to move the RoR to that environment. It can still be on the same development box, but under a separate secure environment. Another interesting capability of a DVCS is maintaining multiple copies on the same box. Git has an advantage over some other systems in that you do not need an underlying database engine to maintain multiple instances. Repositories are very easily moved and replicated to help with virtualization of this type.

Integrating with ITUGLIB

The ITUGLIB team is creating its own backbone that can be integrated into your own backbone. Let's take a look at how this might work for the OpenSSL project:

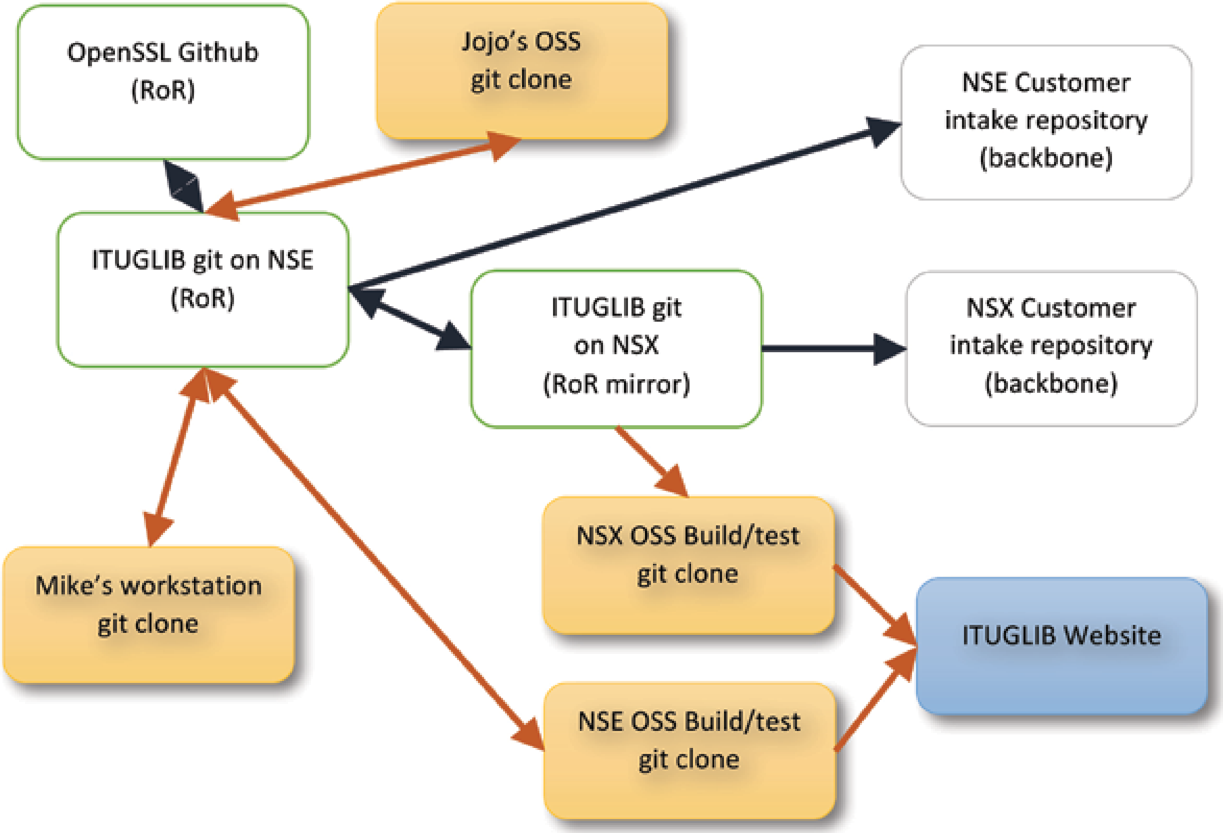

ITUGLIB has already integrated with the OpenSSL repository. Changes are pulled on demand from Github. When the team decides to put together a new release, they create a new branch for the release anchored from the appropriate commit; for example, the 1.0.2c version. The specific changes needed for NonStop are then merged into that branch, which involves Mike's workstation repository and Jojo's local OSS clone. The release is tested, and when ready, the ituglib_release branch is updated with those changes. The plan is to trigger an automatic update of the NSX repository and automatic build/test in that environment, which then goes to the ITUGLIB website to allow downloads. The NSX mirror will then have an updated **ituglib_release** branch that customers can pull into their environments with all of the changes, and the original history from Github. Customers can automatically pull the **ituglib_release** branch to keep up to date and use their own branch identifier to make that version live within their own backbone, after their own review and approval process.

An effective integrated backbone depends on humans to review incoming changes from suppliers

Security

The elephant in the room for using DVCS continues to be perceptions about security. There are a few questions where this is important, and this will probably be the subject of the next article in this series:

1. Who has read and/or modify access to the code?
2. Are historical records secure and how visible should they be?
3. How is the Repository of Record managed?
4. Which branches need to be kept secure and protected?

These are interesting questions, and strangely not really any different from traditional central VCS systems. The core difference is whether history visibility represents a vulnerability. If you have code that needs to be protected using different access rules, put that in its own repository and lock it down. Consider however, that security differences should define a potentially distinct segment in your backbone, with different audit and production requirements.

Summary

Building a Software Change Backbone enables a new level of resiliency in development and production environments. Being able to close the loop between departments, divisions, and vendors on change history allows unprecedented tracking of change history that pushes the boundaries of current requirements for software auditing in our community, and normalizes how code is managed. Ultimately, the move of integrating NonStop code management into the entire enterprise ecosystem can only be a good thing.

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

Compliance Does Not Equal Security: What the EMV Mandates Mean for You

Cynthia Leonard >> Marketing Manager >> HP Security Voltage

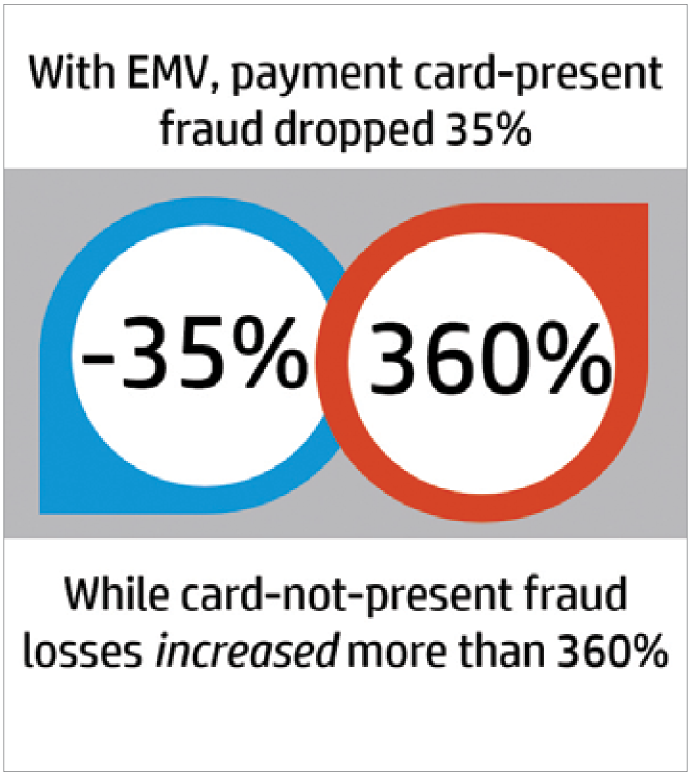
With EMV, payment card-present fraud dropped 35%

While card-not-present fraud losses *increased* more than 360%

We've all been reading and hearing a lot about EMV and the fast approaching EMV mandates. This article will lay out some basic EMV background, and then present the case for why compliance to the EMV mandates, while important, will not equal security.

A quick background if you are not as familiar with it: EMV, which stands for Europay, MasterCard, and Visa, is a global standard and is also referred to as "chip and pin" cards. The card has an embedded computer chip and usually a personal identification number (PIN) is required as well.

Currently, the United States uses the "swipe and signature" point-of-sale (POS) devices, utilizing the card's magnetic stripe. The magnetic stripe contains consumer information including the card number. Once stolen, this information is easy to duplicate on a magnetic stripe of a fraudulent card. Chip and pin cards are harder to duplicate for data thieves. The fact that the card also requires a PIN to work, introduces a second layer of authentication, making it harder for data thieves to duplicate.

For EMV cards to work as designed, the POS device must be enabled with dual-interface (contact/contactless) terminals that are capable of processing complete chip transactions. Dual-interface terminals are able to process chip transactions and mobile devices and wallets, as well as magnetic stripe cards. Instead of swiping your card and the POS device reading the magnetic stripe, you put the card into a terminal, then enter a PIN or sign your name. Further complicating the decision-making for shop-owners switching to these new terminals, is the fact that there are three basic merchant payment system environments: Standalone, Semi-integrated, and Fully Integrated. The choice will depend on your current system for transaction and could get expensive. There are an estimated 15.4 million POS terminals currently in use worldwide, and to replace them all would cost around $8.65 billion.

EMV or chip and pin technology has been around for a long time, Europe has been using it as a standard for years. The way chip and pin technology works is rather than have a magnetic stripe on the back of a card, the chip contains a key only for that card. When accessed by the POS device, or reader, the reader sends a crypto key with an identifier for the card and for the transaction. Therefore, every transaction is unique. Because the chip embedded in the card generates a different, single-use code for every transaction, the data is useless to attackers.

The U.S. is finally migrating to EMV technology–driven in large part by the increased severity and never-ending news of cyber-attacks that result in data breaches. October, 2015, is a significant date and is often referred to as the "the liability shift" or "EMV Mandate". After that date, when credit card fraud takes place, liability for the costs will pass to the entity using the lesser technology. If a merchant is using the old "swipe and signature" POS and the customer has the chip and pin card, the liability for the costs will fall on the merchant. If the merchant does have the updated POS device and the bank has not upgraded the customer's credit card to the new EMV technology, the bank is liable for any fraud charges. If the bank has issued the chip and pin card and the merchant has in place the EMV reader and fraud happens, the credit card company will be liable, as it is presently.

So, we have established that EMV helps protect from what is called "card present" fraud. But what about "card not present fraud?" In France, for example, payment card-present fraud dropped by 35% between 2004 and 2009 after the implementation of EMV, but domestic card-not-present fraud losses increased more than 360% in that same time span. Card-not-present fraud happens with purchases via the Internet, and EMV does nothing to protect these transactions or the sensitive data as it transits web tiers and enterprise systems.

I cannot stress this enough: Compliance does not equal security. The fact remains that while PCI DSS delivers excellent guidance to address a broad range of risks, any compliance assessment, no matter how frequent, only measures risk at a single point in time and provides no guarantee of protection. It is important to think of both compliance and security as ongoing processes and not as a state achieved or a matter of simply checking boxes. EMV will protect the data at the point of origin of the credit card transaction, however, data thieves will still go after non-encrypted information stored on systems or being transmitted. Moreover, while EMV protects against counterfeit card fraud, it does nothing to protect against malware and hacking–the exact methods used in the most notorious data breaches over the past 2 years.

To protect against malware attacks and online fraud, as well as credit card fraud, retail merchants, consumer-facing enterprises and processors need to rethink security and take a data-centric approach that uses field-level, format-preserving encryption and tokenization technologies to protect data in transit, in use and at rest–from the card-swipe or browser to the trusted host.

This data-centric security approach calls for protecting data at the source, by transforming the sensitive data elements with de-identified, yet usable equivalents that retain their format, behavior and meaning. Credit card numbers, track data and other types of structured data are protected without the need to change the data format. Merchants can preserve existing business processes while protecting sensitive digits from the browser or card reader to the payment processor.

Three innovative data-centric technologies can be used to effectively neutralize the threat of data breach, enable PCI compliance, and close critical protection gaps that remain even when EMV has been implemented. These are point-to-point encryption (P2PE), page-integrated encryption (PIE), and secure stateless tokenization (SST).

Point-to-point encryption (P2PE) protects credit card transactions data in payment acceptance systems from the card-swipe through the authorization and settlement processes. Page-Integrated Encryption (PIE) encrypts payment and personal data in ecommerce browser-based transactions from the moment data is entered into a web browser and all the way through the web tier, the application tier, cloud infrastructure, and upstream IT systems and networks to the trusted host destination.

Tokenization works by replacing sensitive data like credit card numbers with a token, or random equivalent, enabling processing on the token while removing the data value to attackers. This protected form of the data can 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. Applications handling the tokenized data–including back-end applications such as fraud analysis and loyalty programs—may be removed entirely from PCI audit scope, delivering significant reductions in management and compliance costs.

A major benefit of these format-preserving technologies is that a majority of analytics and business operations, usually handled by back-end processes, can be performed on de-identified data protected with data-centric techniques. Business users and data scientists do not need access to live payment card data or personally identifiable information in order to achieve the needed processes and business insights.

By using proven, data-centric encryption and secure stateless tokenization technologies, in combination with the new EMV chip card technology, businesses will have a comprehensive data security solution to protect customer credit card and online transactions, end-to-end across the multi-platform enterprise.

Note: Voltage Security was acquired by HP in February 2015, and is now HP Security Voltage, part of the HP Enterprise Security Products Group. HP Security Voltage delivers data-centric encryption solutions and HP SST as a native, stateless tokenization solution running on HP NonStop servers, the platform of choice for payments processing for more than 40 years.
www.voltage.com/breach

Cynthia Leonard joined HP Security Voltage in April 2015 as marketing manager for Communications and Brand. Cynthia provides writing, editing and content creation expertise across all of HP Security Voltage's marketing and internal communications channels, including PR, social media and support for product and market collateral.

Save the Date

NENUG 2015
Monday, October 19th
HP Offices
Andover, MA

Welcome to all NonStop users in the northeast!

Did You Know?

Augmented Reality is Real In HP ExpertOne Courseware

New in HP courseware- images come to life using Augmented Reality (AR) technology giving the user the ability to trigger an experience from an image .These other realities, can help learners, learn quickly by watching a short video during a standard break or after other classroom work. To access the learning, download the HP ExpertOne App from the Google Play Store and the Apple App Store (coming soon), open the app and point the viewfinder at the page wherever you see the HP ExpertOne AR icon to launch your digital content. Check out the HP FlexNetwork Fundamentals, Rev. 15.21 course to see this new style of learning for yourself.

How Will NonStop Fit Into the Internet of Things? Part II

Justin Simonds >> Master Technologist >> Americas

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

New NonStop Architecture Fundamentals

In Part I we discussed the IoT (Internet of Things) market with some general examples in the automotive and energy markets. We discussed how this expanding market is similar to OLTP and how data stream processing requirements were a good fit for NonStop. In this section we'll discuss some ideas for enhancing NonStop based on the x86/InfiniBand announcement that was made during the 2013 Boot Camp.

What InfiniBand Brings to the Party

OLTP transactions typically have an all-or-none nature that is specifically tied to a database. IoT data will not involve a rigid definition of transactions, but will require fault-tolerance – the ability to survive any single point of failure. With storage systems implemented using flash memory (see www.hp.com/3PAR and NonStop's use of SSD) the simple truth is that accessing secondary storage across a fiber network implemented with InfiniBand (IB) will be no slower than connecting to RDMA (remote direct memory access) but there is definitely a difference in cost. Memory is cheaper and prices will continue to fall. Also, on any given server, the memory residing in local DMA is a few orders of magnitude faster to access than storage arrays because of context switching. How much faster? Latency for RDMA over IB is about 250 µs (billionths of a second) versus about 85 µs for local DMA or just a couple of µs (billionths of a second) for cache. Intel reports that a Linux context switch on an i5 core is about 3000 µs (see blog.tsunanet.net/2010/11/how-long-does-it-take-to-make-context.html); and so applications will be optimized to house the memory in the server that will access it most frequently.

The key area of opportunity that supports both MPP (Massively Parallel Processing) and fault tolerance that very few – if any - products have leveraged to date is the ability to leverage RDMA (Remote Direct Memory Access). Any server that implements an InfiniBand HCA can be engineered to access RDMA but only NonStop is engineered to provide a fault-tolerant framework for RDMA. But is there a need for such a capability? Let's dig deeper.

HP Nonstop has been the designated database and hub-system platform of choice for mission-critical core processing architectures in many Fortune 500 companies. The reasons customers select the HP NonStop platform are usually three-fold; mixed-workload processing, scalability and high availability concerns.

The basic premise of NonStop is continuous application availability. Availability, as defined by NonStop, is more than just up-time; it presumes applications performing at acceptable service levels with appropriate response times. NonStop systems are in environments where on-line transaction processing capabilities, response time, security and accuracy are paramount. The NonStop hub may provide stand-in processing at times when the back-end ERP and legacy systems may be experiencing an outage or planned downtime. HP's own IT call centers and websites (HP eats its own cooking) interrogate the NonStop enterprise data store (iHub) to determine order information, ship dates and whether or not a cross-sell opportunity exists for the current customer. These opportunities involve a different style of processing which requires a robust mixed-workload capability (i.e. consistent response times to users and applications while processing batch jobs and large queries), which is a known capability of HP NonStop. When future processing requirements are uncertain, the MPP design allows scalability in a graceful and predictable manner. Adding processors adds additional capability up to a theoretical limit of 4,080 logical processors. Currently with the Quad-core Integrity Blade (Itanium and x86) systems, this is a physical limit of 16,320 cores. These, of course, are known capabilities of NonStop and they fit extremely well with the requirements we are seeing with IoT and stream processing systems. Let's take a look at what's on the horizon.

Why Latency is Important

While NonStop has long been recognized as the database and hub-system platform of choice for mission-critical core processing architectures in many Fortune 500 companies, scalability and fault tolerance has come at the expense of price and performance compared to SMP architecture. There is something very intriguing about IB and how it relates to secondary storage, shared-memory and transaction coordination in a distributed computing environment that will have a profound impact on future application architectures. It all boils down to throughput capabilities and how they are expanding logarithmically; so as to actually invalidate past architectures. Competent systems architects are aware of the following basic rules of thumb about I/O latency:

Operation	Latency (µs)
Disk	7500000
Disk with RAID 10	15000000
XP 20000 Storage array (worse case)	9000
Fiber SAN	5000
Linux/Intel E5-2620 Context switch	3000
10 Gb Ethernet	1200
HP StorServ 7450 (3PAR flash storage)	700
InfiniBand	250
RAM	83
CPU Cache	3

Figure 1 - Relative latency of computing infrastructure components

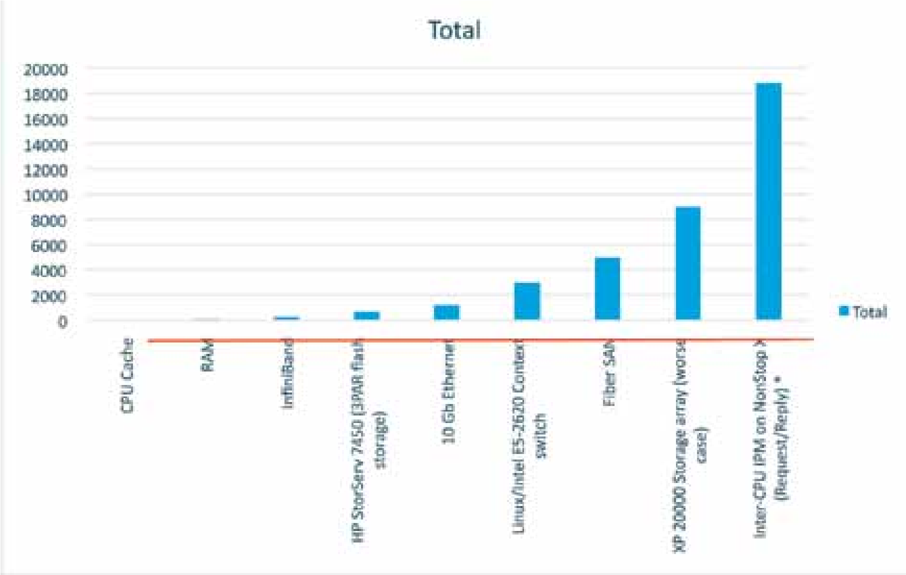

Figure 2 - Relative latencies (µs)

Latency is the critical metric to consider when it comes to performance because this is how long a task must wait before starting to process what has been sent or retrieved. The following graph puts all of this into perspective in nanoseconds with *typical* latencies for each operation of interest. We could find no metrics for the cost of a context switch on NonStop (although we did observe that a Guardian IPM request/reply across processors on a NonStop X server took 18,800 µs) so we included the cost of an Intel i5 on Linux as a reasonable approximation. All the values represented were found on various web sites and are only represented as reasonable approximations:

Platform	Writes/ Sec	Updates/ Sec	Reads/ Sec	Deletes/ Sec
HP Model 30 NIKE RAID 1 raw file access across a single mirrored drive with no indexing	246	159	116	159
Synchronics 1000 with RAID 1 Solid State	61	61	72	61
Synchronics 1000 with RAID 1 conventional disk drives	44	44	45	44
Informix database access on our K-460 with HP model 20 NIKE storage array.	16	50	200	113

Figure 3 - SQL I/O rates versus other file systems

Conventional disk I/O is such a huge drag that we had to factor it out to get a clear sense of the relative latency of the other elements.

Bandwidth and channel overhead become the next critical component to measure. We are ignoring this for now because ultimately, all devices can be engineered to transfer at the IB channel rate. With InfiniBand, all of these interfaces can be leveled to 250 µs latency (i.e. the red line of Figure 2) with a theoretical sustained aggregate transfer throughput of 100 Gb/sec on the latest (e.g. Mellanox SB7790 100 Gb FDR InfiniBand switches, 3PAR solid state storage arrays from HP that are connected as InfiniBand TCA devices, etc.) available technologies. This means that a process can directly access the memory of a remote server at this incredible I/O rate (i.e. 12.5 billion bytes per second) and it can retrieve data from secondary storage at this same rate. Computer processors typically can't keep up with sustained throughput like this today. Processors and their operating systems are the new bottleneck.

From Partitioned Disks to Partitioned Memory

How might new system architectures take advantage of this? If you are building a shared memory solution, you will want the shared memory to reside on the server that will access and update it most frequently to leverage extra 170 µs less latency (i.e. DMA versus RDMA) but you will still provide RDMA access at µs (billionths of a second) speeds in the same order of magnitude so that the data can now realistically remain at rest for remote accessing processes too. That is a key premise behind The Machine (www.hpl.hp.com/research/systems-research/themachine/). If you are doing a database access, the whole access path can be radically accelerated with the use of IB; but there is still considerable processing overhead involved in executing a SQL query. Here is an admittedly dated (i.e. over a decade old) comparison of SQL versus other forms of I/O that one of the authors actually measured to illustrate the point:

Ideally architectures should only do database operations when needed and should wherever possible, accelerate applications I/O through the use of shared memory. Shared memory can now be distributed across an IB fiber network.

Here is another thought to ponder. We stated at the outset that IoT is "OLTP-like" but there are some interesting differences. When there is a deluge of digital and analog data, it is unlikely that businesses will wish to save all of it to permanent storage. It may well only be interested in aggregating 'normal' data or in the values that are outside the expected mean. That said, there will likely be a requirement to save the exception data and in fact, it may need to be fault tolerant. If this is the case, it will not be desirable to have a transaction monitor involved in the update because that will be a drag on performance. How will conventional architectures be able to ensure updates while surviving any single point of failure? Only NonStop is presently engineered to meet this need. How? By delegating such operations to a NonStop process pair whose primary holds the critical data and checkpoints changes to the backup process residing in another CPU. If the primary fails, the backup will elegantly and seamlessly provide the single-point-of-failure recovery that only a NonStop can achieve. That is entirely the point behind IDC's AL4 capabilities (IDC report on availability levels – 4 being the highest).

Are the Days of Map/Reduce Architectures Numbered?

Today cloud-based systems are utilizing various flavors of map/reduce technology (e.g. Hadoop, PIG, Simple Messaging Service, etc.) by replicating data across multiple servers and coordinating their updates with sophisticated monitoring frameworks. However there is very little to date that has been built which meets or exceeds IB aggregate throughput. If you consider, for example, the ZooKeeper framework for managing semaphores in a distributed computing environment, a given semaphore's performance deteriorates very rapidly if the rate of updates of all semaphores in aggregate exceeds 10% of the overall pool of resources (see zookeeper.apache.org/doc/current/zookeeperOver.html#Performance); not particularly scalable. If instead, that semaphore resides in a particular processor that can be accessed across an IB fabric, billions of operations per second are theoretically possible with no need to replicate (replication with map/reduce being the cloud methodology for achieving scalability and resilience).

New Core Capability for NonStop

What this really means is that shared-nothing multiprocessing – what is also generally known as MPP – architecture will soon eclipse symmetrical multiprocessor (SMP) architectures that are currently in vogue. Why? Because SMP is constrained by Von Neumann architectures (sequential processing see en.wikipedia.org/wiki/Von_Neumann_architecture) whereas MPP architectures deal with parallel processing naturally by providing the kinds of synchronization and failover mechanisms we take for granted on NonStop - but that are lacking in other operating systems. To date, NonStop has been at a disadvantage to the SMP shared-memory applications of competing platforms primarily owing to the high-latency IPM requirements mandated by the shared-nothing environment. With IB and the natural complimentary semantics it shares with WRITEREADX, READUPDATEX, AWAITIOX, etc. the disadvantage is about to be turned to advantage. The requirements of the new order include parallelism, scale and fault tolerance which will be combined with the speed advantages of IB. SMP-only systems are already hitting a wall. The increasing number of cores already have incredibly complex programming to optimize the use of threads. When core counts start reaching 64, 128, 256 and beyond; threading becomes untenable. MPP is a far more scalable architecture, as everyone will eventually come to appreciate. Emerging exascale standards such as MVAPICH2 (mvapich.cse.ohio-state.edu/overview/) are predicated on it.

NonStop has many of the new requirements embedded such as scalability, reliability, security and of course fault tolerance. It was mentioned earlier that not all elements of the new stream processing applications need to be fault tolerant. In this second installment, we have demonstrated how NonStop is uniquely positioned to meet the fault tolerance challenge with the correct parallel processing architecture to meet the rigorous demands of the most demanding indestructible computing environments at the extreme velocities Web 3.0 is expected to bring – all over the radically increased velocity made possible by InfiniBand networks. What about hybrid, converged architectures? How can HP's NonStop with these new IB capabilities participate in a hybrid architecture to leverage the economies of lower-priced platforms? We'll explore that in Part III.

SIDEBAR FOR NONSTOP FUNDAMENTALS

Massively Parallel Processing versus Symmetrical Multi-Processor

Massively Parallel Processing (MPP) is a system architecture that presumes each processor has its own RAM and that other processors cannot access it. Each processor is presumed to operate as an autonomous system that has mechanisms for coordinating work with other processors – typically a message bus. Symmetrical Multi-processor (SMP) is a system architecture that presumes two or more processors access the same shared RAM. To do so, synchronization mechanisms (i.e. typically semaphores) are used to coordinate access between processes running in competing CPUs.

Single Point of Failure

A single point of failure is any hardware or software component that should it fail, will bring down the entire system. Such single points of failure typically include CPU, RAM, controller, bus, critical device driver, LAN, communication line, power supply, etc.

Checkpoint

A checkpoint is a NonStop-specific term relevant only to NonStop Process Pairs. The purpose of a checkpoint is to ensure that the two processes have identical process state so that if there is a single point of failure, the backup process can take over the completion of processing without the requesting process needing to do any exception handling. The purpose of checkpoints are to make primary process failures transparent. Checkpoints occur on critical region boundaries to ensure all-or-nothing processing process semantics. The primary process sends checkpoints on critical region boundaries to the backup and waits for acknowledgement before proceeding to the next critical region step. It is up to application architects to determine what the critical regions of a process are.

Multi-core Processors

Multi-core processors are a more recent evolution of CPU architectures whereby sophisticated chip logic and compilers allow processing to be broken into threads of critical regions and submitted to multiple cores simultaneously. Instead of a processor needing to run at ever faster clock speeds, cross-section computing power can be aggregated across multiple processors to achieve the same effect.

NonStop Process Pair

A NonStop process pair refers to a primary process residing in a particular CPU and a designated backup process residing in another CPU. They share the same process name but each have a different CPU:PIN (i.e. process id) pair. They can be configured to be amnesia backup processes (i.e. they know nothing about the state of the other process) or they communicate with each other to preserve state using checkpoints.

IPM

An IPM (Inter-process Message) is a type of message that is sent across a message bus to tie the processors of an MPP system together. In the context of NonStop, this is a message that is sent between any two processes within a given NonStop node or across an EXPAND network of NonStop nodes. These messages are unsolicited requests that are sent to the $RECEIVE message queue of a specified process by another NonStop process. On NonStop X, InfiniBand is the bus fabric.

Transaction Monitor (TMF)

A Transaction Monitor is any daemon or mechanism that can ensure ACID properties of a given database update or transactional all-or-none execution consistency respectively. On NonStop systems, the transaction monitoring is achieved with the Transaction Monitoring Facility (TMF) subsystem. It is fully integrated with the file system and IPM.

RDMA

RDMA (Remote Direct Memory Access) is a capability specific to ServerNet and InfiniBand architectures whereby the memory of a given CPU can be directly referenced by a process in another CPU without the CPU that owns the memory being involved in the I/O operation. There is no context switching, interrupt or trap handling needed to service the I/O. Everything happens in user mode during the process's execution time slice for maximum efficiency.

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

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

YOUR NONSTOP AUDIENCE JUST GOT BIGGER!

Read The Digital Version of The Connection Now

The Connection
Where In The World Is My Credit Card Data?
The Connection
Managing the Media Life Cycle with HP StoreOnce & ETI-NET BackBox
The Connection

http://www.connect-community.org/general/custom.asp?page=theConnectionArchive

Why an Active/Passive Business Continuity Solution is Not Good Enough

Keith B. Evans >> **Shadowbase Product Management** >> Gravic, Inc.

The costs of prolonged downtime of critical business IT systems are significant (potentially to the point of shuttering the company). These potential costs are compounded by the fact that the many events which can lead to such outages are not rare; it is a case of when, not if. This likelihood of outage events is only acceptable if you have a complete, documented, and well-tested business continuity plan in place. Maybe you think that you do, but the data does not support this idea. Many companies are operating with the mistaken belief that their business continuity plan will work when the time comes, or even if it does work, that the plan is good enough to prevent significant consequences to the company. Read further to find out whether this false sense of security applies to you.

Business Continuity Architectures: Pros and Cons

The chosen availability architecture is the primary factor in determining how effective your business continuity plan will be when the time comes. To discuss the typical ones, we first need to understand a couple of terms:

- *Recovery Point Objective* (RPO) is the maximum acceptable amount of data loss arising from an outage of an active system. In practice, it is the data updated in the period between the last time the data was saved to (remote) recoverable media, and the point of failure.
- *Recovery Time Objective* (RTO) is the maximum acceptable time for recovery from an outage. In practice, it is the period between the time of failure and the point at which services are restored to an acceptable level.

Different business continuity architectures have different attributes with respect to RPO and RTO (Figure 1). Let us briefly review each of these major architectures.[1]

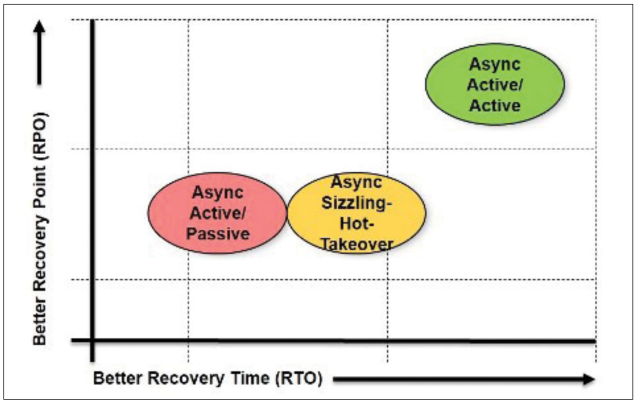

Figure 1 – RPO and RTO for the Various Business Continuity Replication Architectures

Active/Passive – Classic Disaster Recovery

In this architecture, all transactions are executed on a single system (the active node), and the database updates are replicated by Shadowbase software[2] to a backup system (the passive or standby node). In the event of a failure of the active node, a failover to the backup node is executed, the applications are brought up with the local (synchronized) database opened for read/write access, users are switched to the backup node, and processing resumes. This architecture and failover sequence are by far the most common, but are also the most flawed.

The key issue with this architecture is that it is *very difficult to test the backup node and failover procedures*. Proper testing requires an outage of the primary node and may take a long time. Therefore, failover testing is very often not performed at all or not to completion when it is attempted (because it may take longer to fully test than the available outage window).

It is also possible that restarting the production system after the test has completed may not work, which is another reason why testing may be avoided. Because of this lack of testing and the resulting uncertainties surrounding the state of the backup system and the takeover procedures, when a real outage occurs, management is often slow to initiate a failover in the first place, further delaying recovery. Hence, this architecture is risky, the state of the backup system (and procedures to failover) are not really known, failover faults are likely to occur causing the failover to be unsuccessful, or at least take a long time. For all these reasons, this architecture has the probability of a high RTO, often several hours or even days. While a basic active/passive architecture offers some protection, it is by no means the best solution. It should really only be considered as a starting point, or used for non-mission-critical applications.

Active/Almost-Active – The Sizzling-Hot-Takeover (SZT) Architecture

While it looks almost the same as a classic active/passive architecture, *sizzling-hot-takeover* (SZT) has one major difference which makes it a much improved solution. The difference is that while all transactions are still routed to a single active node, the backup node has the applications already up-and-running, with the local database open for read/write access.[3] The key benefit of this ability (versus an active/passive architecture) is to ensure the backup system is ready to go when

[1] Note that each of these architectures use asynchronous replication, where there is a slight delay between when the data is updated on one system, and is safely replicated/stored on another system. This delay accounts for the data loss in the event of an outage.

[2] For a much more detailed description of the various business continuity architectures and their total cost of ownership, see the Gravic Shadowbase white papers, Choosing a Business Continuity Architecture to Meet Your Availability Requirements and Fingers Crossed? Or What is Your Business Continuity Plan for the Inevitable?

[3] Not all data replication products allow the backup database to be open for application read/write access during replication, but Shadowbase software (from www.gravic.com/shadowbase) has no such restriction.

you actually need it. Since the applications are up-and-running on the backup node with the local database open read/write, it is easy to send test/verification transactions against test/verification accounts to validate the backup system at any time, with no impact to the active system. Hence, the backup system can be regularly validated, and becomes a known-working system. (For all intents and purposes, it is a fully active system, with the exception that it is not processing online transactions.)

When an outage occurs of the primary node, the decision to fail over can be made immediately, with confidence that failover faults will not arise, and the failover will succeed quickly. As a matter of fact, as a best practice, failovers should be performed regularly (e.g., weekly or monthly) to test the process and build the confidence of the staff in performing them. Businesses running active/passive architectures simply do not have the same confidence level as those running SZT architectures. Therefore, this architecture gives a much better and repeatable RTO versus classic active/passive architectures. SZT is the minimum level of business continuity solution which should be employed for mission-critical applications.

Active/Active – Partitioned

In a partitioned active/active architecture, the applications are active on all nodes, transactions are routed to all nodes, and each node has a copy of the database, which is kept synchronized by bi-directional data replication. To avoid data collisions[4], for example, the data (or requests) are partitioned so that transactions are routed to a specific node based on some key in the data, or from which user the transaction originated. The database may be split by customer name, and all transactions for customers A-M are executed on one node, and customers N-Z on the other, with their changes being replicated to the other node to keep the databases synchronized. This architecture provides the key benefits of active/active while avoiding data collisions.

The benefits of this architecture compared to classic active/passive and SZT are:

- On failure/outage, only half the users (in a two node configuration, fewer if more nodes are used) are affected and have to be switched. The other half of users see no outage at all, i.e., better RTO.
- There is about half as much data loss (in a two node configuration, fewer if more nodes are used), i.e., better RPO, because only the updates in the replication stream on the failed node are lost. The updates in the replication stream on the remaining node(s) are unaffected, and will be replayed once the failed node is recovered.
- There are little to no testing costs/issues, and no failover faults. All systems in the configuration are known to be working at all times (which is also true for SZT).
- Better system capacity utilization as all nodes are performing productive work.

Active/Active – Route Anywhere

This architecture is the same as the active/active partitioned model described above except that the partitioning aspect is removed. Any transaction can be executed by any node (hence the name, "route anywhere"). This architecture has all the benefits of the active/active–partitioned model, but in addition, eliminates two of the issues with that model. It does not require partitioning (which may not be possible in all cases), and since transaction routing is unconstrained, workload can be evenly load-balanced across nodes.

There is always a price to pay, and in this case it is the possibility of data collisions. For some applications, data collisions may be practically impossible. For example, it is highly unlikely the same credit/debit/ATM card would be used simultaneously for multiple transactions. But if collisions are possible, they must be identified and dealt with immediately. Data replication should include functionality to automatically detect, report, and resolve data collisions. User exits may also be provided to enable more sophisticated processing of data collisions if necessary.

All business continuity architectures are not created equal. Figure 1 helps to visualize the differences between these various architectures with respect to the parameters of RPO and RTO.[5] As far as RPO is concerned, active/passive and SZT solutions are similar. For RTO, an SZT architecture trumps an active/passive architecture. But an active/active implementation beats active/passive and SZT architecture on all counts.

Business Continuity Architectures: Total Cost of Ownership

There is another way of looking at the various business continuity architectures, which provides an even more striking view of the differences between them and their relative benefits, and that is to look at the *total cost of ownership* (TCO). Active/passive configurations are cheaper and less complex to implement, however, when looked at through the lens of TCO, they have a (very) false economy.

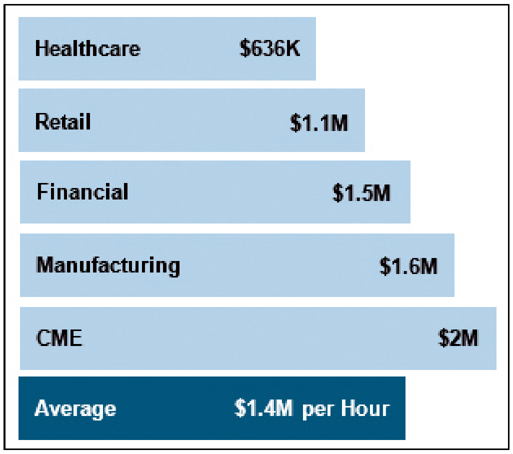

Figure 2 – Average costs per hour of downtime across various industries

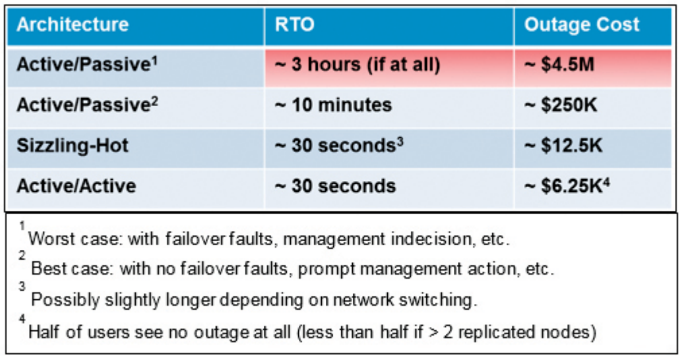

Architecture	RTO	Outage Cost
Active/Passive[1]	~ 3 hours (if at all)	~ $4.5M
Active/Passive[2]	~ 10 minutes	~ $250K
Sizzling-Hot	~ 30 seconds[3]	~ $12.5K
Active/Active	~ 30 seconds	~ $6.25K[4]

[1] Worst case: with failover faults, management indecision, etc.
[2] Best case: with no failover faults, prompt management action, etc.
[3] Possibly slightly longer depending on network switching.
[4] Half of users see no outage at all (less than half if > 2 replicated nodes)

Figure 3 – Estimated Service Unavailability Costs for a Financial Application

[4] A data collision occurs when the same data record is updated simultaneously on two active systems, which after replication to the other system results in both copies of the data record being incorrect.
[5] Note that with respect to RTO and RPO, there is no difference between asynchronous active/active – partitioned and asynchronous active/active – route anywhere.
[6] Sources: Network Computing, the Meta Group, Contingency Planning Research.

Winning with HP Shadowbase

Real-Time Data Replication and Integration for a Nonstop World

HP Shadowbase – Data Replication Software for Continuous Business

- **Make Your Business 'Nonstop'** HP offers the Shadowbase product suite running on HP Integrity NonStop and other server platforms, comprised of software solutions that address business continuity, system upgrades without downtime, real-time business intelligence, and master data management to deliver a true 24x7 "nonstop" enterprise.
- **HP Shadowbase Data Replication** software enables active/passive, sizzling-hot-takeover (SZT), and fully active/active business continuity architectures to suit any application needs, providing rapid recovery from unplanned outages in times ranging from minutes to immediate, from disaster recovery to disaster-tolerant continuous availability.
- **HP Shadowbase Data and Application Integration** software enables low-latency, real-time data and event distribution between heterogeneous systems, databases, and applications, providing data warehouse feeds and enabling rapid development of new business services to achieve competitive advantage.
- **HP Shadowbase Zero Downtime Migration (ZDM)** software provides the means to eliminate planned downtime, keeping your business services online while routine system maintenance or complex and disruptive upgrades and migrations are performed.
- **HP NonStop Shadowbase Compare** software compares a target Enscribe file or NonStop SQL table to its source, and reports any discrepancies found, which is helpful for validating that a target database matches its source, and for satisfying regulatory/auditing requirements.
- **HP Shadowbase Data Management Utilities** provide the tools to monitor and, if necessary, correct data in order to detect and resolve anomalous behavior, ensure continuation of proper business operations, and satisfy audit compliance.

Find out how **HP Shadowbase** can help you be ready for anything.
Visit www.hp.com/nonstopcontinuity and www.shadowbasesoftware.com.

GRAVIC®
Shadowbase

Business
Partner

©2015 Gravic, Inc. All product names mentioned are trademarks of their respective owners. Specifications subject to change without notice.

First, to emphasize this point, let us put some dollar values on the cost of downtime (Figure 2).[6] As shown, these monetary costs are non-trivial, to say the least. Now, using the average cost per hour of downtime for a financial application of $1.5M/hour as an example, and using reported industry averages for typical periods of recovery time (RTO), we can estimate actual outage costs for the various business continuity architectures (Figure 3).

Obviously, basic active/passive architectures are very expensive when looked at in terms of TCO. They may be easier and cheaper to implement, but when outages do occur, they are likely to cost you much, much, more in the long run. Even in the best case scenario, with a well-tested system and a trouble-free failover, a basic active/passive configuration is still going to be about 20 times higher in outage costs when compared to an SZT configuration. For a worst case scenario (much more likely given the difficulties of testing and probability of failover faults as previously discussed), it is about 36 times more costly at about $4.5 million *per outage* (assuming the recovery *only* takes three hours, but it could take much longer).

The cost differences become even more apparent when viewed graphically (Figure 4). Given the marginal incremental cost and complexity, coupled with the significant decrease in potential outage costs, there is really no reason why anyone would run the risk and not move immediately from an active/passive to at least an SZT architecture.

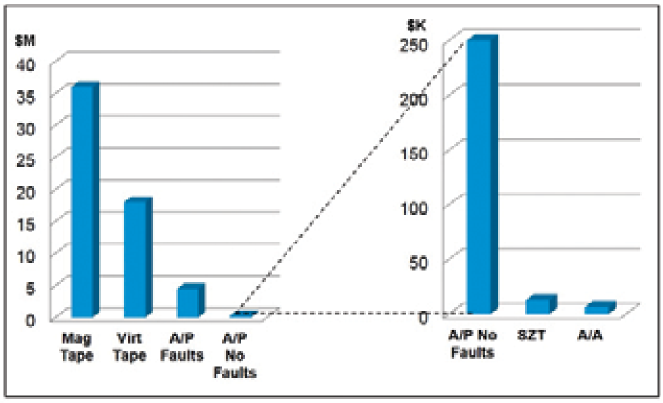

Figure 4 – Estimated Service Unavailability Costs for a Financial Application

As well as considering the cost of downtime, the cost of lost data needs to be considered. Again, using industry averages for the amount of data lost in the event of an outage (the amount of changed data which has not yet been safe-stored on a backup system), we can estimate the cost of lost data for the various business continuity architectures, across various industries (Figure 5).

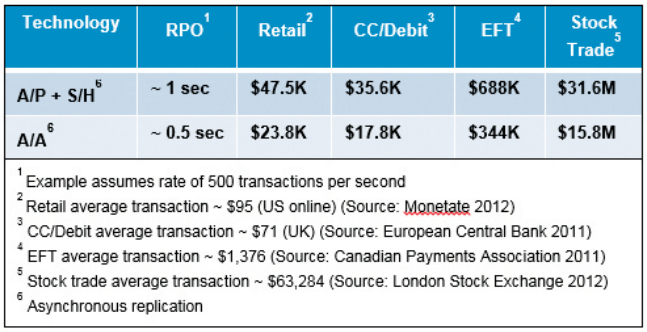

Technology	RPO[1]	Retail[2]	CC/Debit[3]	EFT[4]	Stock Trade[5]
A/P + S/H[6]	~ 1 sec	$47.5K	$35.6K	$688K	$31.6M
A/A[6]	~ 0.5 sec	$23.8K	$17.8K	$344K	$15.8M

[1] Example assumes rate of 500 transactions per second
[2] Retail average transaction ~ $95 (US online) (Source: Monetate 2012)
[3] CC/Debit average transaction ~ $71 (UK) (Source: European Central Bank 2011)
[4] EFT average transaction ~ $1,376 (Source: Canadian Payments Association 2011)
[5] Stock trade average transaction ~ $63,284 (Source: London Stock Exchange 2012)
[6] Asynchronous replication

Figure 5 – Estimated Costs of Lost Data Across Various Industries

In this case, active/passive and SZT are the same since they both lose the same amount of data, but active/active is much better since it only loses half as much data. Again, this difference is more dramatically illustrated graphically (Figure 6).

But even if data loss (RPO) goals based on the average value of a transaction may appear acceptable, some data transactions are much more valuable than others and absolutely cannot be lost:

- Healthcare – lost dosage records can result in patient overdose on medication
- Manufacturing – car manufacturer can tolerate short production line outage, but cannot lose data regarding bolt torque settings, etc., in case of lawsuits from accidents
- Electronic Funds Transfer (EFT) – some transactions are worth millions, even if the average transaction is much lower
- Stock Trades – like EFT, some transactions are worth millions, and stock price is based on previous trades, so none can be lost

Therefore, RPO goals must be set based not on the value of an average transaction, but on the value of the most expensive/critical transaction. If the cost of losing the most valuable/critical data is very high, then an active/active configuration is the only solution, since it has the best RPO characteristics (least data loss).[7]

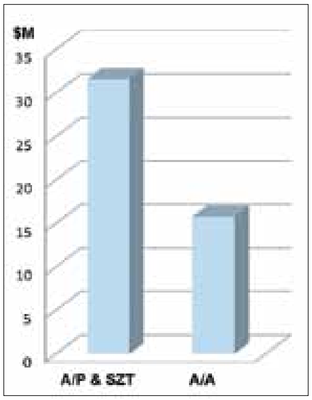

Figure 6 – Estimated Costs of Lost Data for Average Stock Trade Transaction

To summarize, overall TCO decreases by orders of magnitude more than the cost that the business continuity solution increases, as illustrated by Figure 7:

- The better the availability, the greater the complexity and implementation cost
- The better the availability, the lower the outage cost (by orders of magnitude)
- Net result, as implementation cost increases, overall TCO decreases

By this measure, the cost and complexity of an active/active solution is clearly more than outweighed by its superior overall TCO. It also illustrates how much better an SZT solution is in terms of TCO compared with a basic active/passive architecture.

Conclusion

To implement a business continuity plan, the IT architecture to be employed in order to maintain services in the event of an outage (planned or unplanned) must be selected. Many users select, and never get beyond, a basic active/passive architecture, but it has many issues, which can prevent a successful and timely failover. This model is reactive, risky, and provides a false sense of security. The likelihood of an extended outage is high; consequently, the likelihood of a very expensive outage is high. ***Active/passive architectures are simply not good enough for mission-critical applications.***

Though the more sophisticated business continuity solutions (SZT and active/active) are more complex and somewhat more expensive to implement, they are in fact far more cost-effective when looked at in terms of TCO. If you are running an active/

[7] If your application absolutely cannot tolerate any data loss, then contact Gravic, Inc. for more information about a new Shadowbase technology, synchronous replication, which will eliminate data loss entirely.

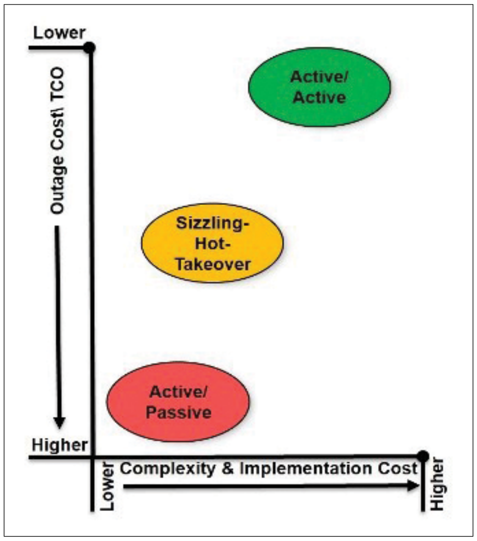

Figure 7 – TCO Versus Complexity/Cost of Implementation

passive architecture, it will probably only take one outage during peak processing hours to realize this fact the hard way. Because SZT is only marginally more complex than active/passive to implement, yet the benefits are significant, ***SZT should be the absolute minimum architecture chosen for applications which must remain available.***

SZT itself should only be seen as a stepping stone to a fully active/active architecture. Active/active cuts in half outage and data loss costs, and significantly improves the utilization of system capacity (i.e., there is no idle backup system). ***An active/active architecture provides the only acceptable solution for applications with high-value transactions where data loss must be minimized, and/or applications which must be continuously available.***

The solution is in your hands. The attention-grabbing outage headlines and long meetings with senior management explaining what happened need not be applicable to your company. With Shadowbase data replication, the solutions are available today to make extended outages a thing of the past. The potential outage costs described above are sobering; you do not want to validate them the hard way. If you are currently using an active/passive configuration, move as quickly as you can to an SZT architecture. If you are already running in SZT mode, congratulations, but also consider whether you should be taking the next step and moving to a fully active/active implementation. You do not want to make the right decision in hindsight, after it is too late to protect the availability of your mission-critical applications and data.

Sales and Support for Shadowbase data replication technology is available from select resellers in certain regions and from HP globally. Contact either Gravic or your HP account team for more information.

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 contact us at sbproductmanagement@gravic.com

Connect to the *Future!*

Connect Worldwide, Hewlett Packard's Enterprise Technology Community, invites you to join a global network of **70,000 colleagues, 39 Chapters** and **19 Special Interest Groups** collectively producing over **50 annual events.**

JOIN TODAY to receive the best value and return on your HP business technology investments!

Communities Served

- » Big Data
- » Converged Systems
- » Cloud
- » HP-UX
- » HAVEn
- » Enterprise Networking, Security, Servers and Storage
- » Infrastructure Software
- » Linux
- » NonStop
- » OpenVMS

CONNECT
Hewlett Packard's Enterprise Technology Community

Why Connect?

- We facilitate communication to ensure you have a voice to HP
- We make learning affordable to help you do more and better business
- We produce events that lead to meaningful business relationships and increased sales
- We foster communities that ensure your questions don't go unanswered
- We relay the latest technology news and product announcements to keep you on top or your game

Membership Dues:

- Individual $50
- Corporate $500

Data Centric Protection

A paradigm shift from perimeter based protection to data centric protection to prevent data breaches

Kong Yew >> Product Engineering Manager >> HP Atalla

New and upcoming threats are always in the horizon. Recent public data breaches clearly indicate that existing solutions, particularly data prevention tools, are not very effective in preventing data loss. In many cases simply using encryption can prevent these data breaches from bearing fruit for those that somehow get a hold of the data. For instance, the Sony Pictures hack involves the release of unencrypted data such as social security numbers, emails and private health information. Since it is now widely known that the impact of leaked confidential information is real and the consequences are impactful, organizations must constantly be on guard to protect their confidential information. The responsibility for what happens with data that is lost falls directly on the company that is the victim of the attack. To prevent data breaches, most large enterprise customers have deployed data leakage prevention tools to detect data exfiltration via multiple channels including USB storage, personal email and SaaS applications. These solutions identify files for classified and sensitive information by rigorously inspecting the content, delivery mechanism and perform some pattern matching algorithm. The fact is no such solution can guarantee with 100% accuracy that all their sensitive information and files have been identified, and classified which is critical to the business. Similarly, it is a fallacy to assume that perimeter protection is sufficient to protect data. Thus, we need think differently on how to protect the data and prevent data breaches.

Challenges of protecting unstructured information

First, most organizations have limited visibility about classified and sensitive data across the company's environment. Without data visibility, most organizations approach data security by protecting the logical and physical security of the network and data egress points. Protecting the organizational perimeter requires serious defense and an in-depth approach utilizing multiple data leakage prevention solutions to continuously detect data exfiltration while minimizing false positives. In fact, finding data exfiltration from insider threats is similar to finding a needle in a hay stack.

Consider the alternative approach of protecting data by understanding the data lifecycle for unstructured data.

Unstructured data such as word documents, excel spreadsheets and emails are continuously created and distributed across the enterprise organization and might potentially be shared externally with third parties. Most likely these documents are not protected and encrypted since a majority of encryption technologies are intrusive and require manual intervention that causes friction with the end users. Even though encryption technologies such as Information rights management are matured and well-established, these technologies face some opposition. Users resist any technology solution that requires disruptive user behavior, even if it's as simple as right clicking on a file to make sure it's protected. Therefore, more and more companies will require a solution that protects the data while ensuring the security policy is enforced and minimizing user friction.

When customers are addressing the challenges of protecting unstructured data, they have to address some of these questions.

- How do we identify sensitive and confidential data such as PII (Personal Identifiable Information), Social Security numbers and credit numbers?
- How do we protect files from accidental leakage?
- How do we detect insider threats? How do we protect data in a borderless enterprise environment that includes Cloud, SaaS applications and BYOD (Bring Your Own Device)?

The consequences of data breaches can be damaging for any organization. It can lead to the loss of confidential documents and intellectual property that can cause real economic damage to the company and erode the value of their information assets. Breaches also degrade the company's reputation and brand value as a result of customer mistrust in your handling of sensitive information. You also face a large potential for market share loss and increased shareholder scrutiny.

Case Study of a large Telecom - ACME

ACME is a large telecom company with global operations in multiple countries and partnerships with many other companies. ACME has regulatory requirements to segregate their customer data within their systems. Customer data is collected, processed and generated within multiple applications

Challenge:

Due to market and regulatory requirements, ACME needs to protect sensitive customer and business data. ACME has developed a well-established corporate-wide security policy. However, ACME also needed a solution to enable consistent and efficient data security policy enforcement without negatively impacting employee workflow and business processes.

To address the challenge ACME wanted to leverage their Information Rights Management (IRM) capability and enable document restrictions to prevent unauthorized access of sensitive documents. They also wanted to proactively monitor data access to restricted customer data and enable data classification and encryption of all sensitive data connected to their users.

Solution:

HP approaches data-centric protection with a platform that embeds data security at the point of data creation or data access with a product called HP Atalla Information Protection and Control (IPC). Atalla IPC enables automatic data classification for unstructured data such as word docs, spreadsheets and other files during a files creation. It also applies protection at that point of data creation and makes that protection

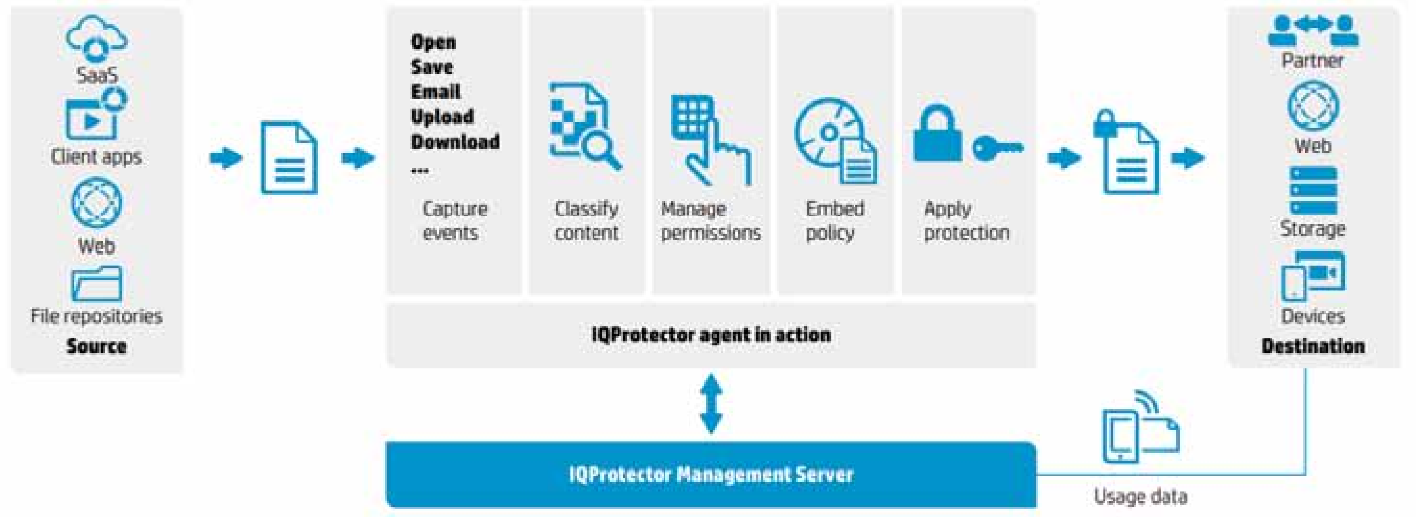

Figure 1 – HP Atalla IPC architecture overview

persistent with the file. This approach ensures that customer sensitive data will always be protected regardless where it resides.

Unlike traditional solutions that attempt to control users, channels, or storage, HP Atalla IPC protects data by uniquely embedding protection within data itself at the moment of creation or initial organizational access in unstructured form. Atalla IPC's IQProtector engine has agents on enterprise hosts instantly identify and classify all new, modified, or accessed sensitive data from any origin. This data, identified with extremely high accuracy, is persistently tagged, enabling comprehensive control over access and usage.

Atalla IPC leverages Microsoft® Active Directory Rights Management Services (AD RMS) to provide information rights management features while providing enterprise features such as automatic classification and protection to automate the process of protecting information. This approach enforces enterprise policies of embedding security at the point of data creation and provides the option for manual user classification to educate users about security awareness.

Benefits of using HP Atalla IPC:

Some of the top level reasons you would want to implement Atalla IPC is that it provides seamless collaboration for both internal and external organizations while persistently protecting the information. Atalla IPC also provides Enterprise level control for information protection while minimizing the user friction when adopting this solution. Atalla IPC is a key way to reduce data breaches by employing data classification and information protection to secure sensitive data.

The core HP Atalla IPC data classification and protection suite includes the IQProtector management server and IQProtector agents that provide persistent multi-format file protection and persistent email protection.

The IQProtector agent is deployed to the endpoint machine and automatically detects events such as file creation, file access and web downloads then classifies and protects the data based on the context of the Information, file content and other meta data. With contextual data classification and protection, sensitive data can be persistently and automatically protected. Alternatively, Atalla IPC also provides the ability for manual data classification based on a user's input and further increases security awareness among employees.

Module options that users will want to consider for their solution include:

- **HP Atalla IPC Scanner Classification and Protection** is essentially a document crawler that scans, classifies and protects existing data on multiple repositories. This solution addresses the use cases to discover existing data that is sensitive and confidential.
- **HP Atalla IPC Bridge for Content Inspection Service** is deployed on enterprise services such as anti-virus, data loss prevention, enterprise content management tool and archiving tool to enable indexing, searching and accessing encrypted content. This solution addresses interoperability among the existing business applications with the encrypted content.

In summary, HP Atalla Information Protection and Control addresses the data centric protection use case by embedding data security at the point of creation. More importantly, Atalla IPC effectively reduces the risk of data breaches. For more information visit www.hp.com/go/AtallaIPC.

Kong Yew, Chan, is the New & Emerging Product Engineering Manager for HP Atalla. He has worked in the security space for more than 10 years. Previously, he has worked in various engineering and management roles from startups to large organizations such as Trend Micro. Kong-yew has a bachelor degree in Computer Engineering from Nanyang Technological University, Singapore and has his MBA from Babson College.

Did You Know?

HP Atalla won the Readers Trust Award for Best Database Security Solution at the RSA awards dinner this year from SC Magazine –with our HP Enterprise Secure Key Manager with HP Secure Encryption.

http://awards.scmagazine.com/
http://media.scmagazine.com/documents/118/botn2015sm_29485.pdf

The Fraud Blocker: Catching the Wrongdoers

Yash Kapadia >> CEO >> OmniPayments Inc.

Steve Anderson had just finished a pleasant meal with his wife at their favorite restaurant. He reached over for the bill and laid his credit card on top of it. The bill and the card were taken by the waitress, and she returned shortly with his credit card and a chit for him to sign.

Three days later, Steve stopped for gas. To his chagrin, his credit card was not in his wallet. Racking his brain for when he last used his card, Steve's best recollection was the restaurant. He paid cash for his gas and called the restaurant. The staff said that his card was not being held by them.

He next called his credit-card company and notified them about the lost card. The company said that it would block the card and issue him a new one. They then confirmed with Steve that within the last two days, he had bought gas three times, had charged two dinners, and had purchased $500 worth of power tools. Steve had made none of these transactions! His card had been stolen and was being used fraudulently.

Fortunately for Steve, he was not responsible for the fraudulent purchases. Depending upon the circumstances, either the merchants involved or the bank that issued the card had to cover the losses. It is for this reason that banks and retailers are intensely interested in payment-card fraud detection.

There are two primary methods for fraud detection, *Fraud Blockers and Fraud Monitors*. Fraud Blockers detect potentially fraudulent transactions in real time before they are sent to the issuing bank for authorization. If the Fraud Blocker determines that a transaction may be fraudulent, it rejects the transaction outright; and the issuing bank never sees it. Therefore, a Fraud Blocker is *proactive* in detecting fraud.

Fraud Monitors evaluate transactions as the transactions are being authorized by the issuing bank. They will report suspicious transactions to the issuing bank but typically not in time to reject the transactions. Therefore, a Fraud Monitor is reactive in detecting fraud.

Fraud Blockers and Fraud Monitors are systems that usually are external to the issuing banks. Alternatively, the issuing banks may perform this function; but fraud protection is a computationally complex process that issuing banks would often like to move to external systems.

Fraudulent Card Transactions

Credit cards can be abused if they are lost or stolen. Even worse, recent hacks such as those against Target and Home Depot have snatched the credit-card information of millions of cards. The card numbers are sold on the Darknet (a private part of the Internet open only to trusted accomplices). The stolen credit-card numbers of magnetic-stripe cards can be utilized to create card copies that then can be used as easily as the original cards themselves until the fraudulent card is detected and blocked.

It is for this reason that acquiring banks, issuing banks, and retailers are investing heavily in fraud-detection facilities to deny suspicious purchases or to block the cards themselves. In some cases, the purchases are authorized; but a temporary block may be applied to the card until the cardholder is contacted to see if the purchases were legitimate. Cardholders traveling overseas often receive these calls when they first make a purchase in a foreign country if they have not informed their credit-card companies of their travel plans. In other cases, the card is permanently blocked; and the cardholder is informed that a new card will be sent to him.

The Financial Transaction Switch

Before getting into the detection of fraudulent transactions, let us first understand the path the transactions follow from the retailer to the issuing bank for authorization. As shown in Figure 1, the retailer is provided point-of-sale (POS) terminals by its acquiring bank (or by a company with a relationship to the acquiring bank). The POS terminals are effectively intelligent cash registers, and they create the payment transactions, whether customers pay with cash or with a payment card.

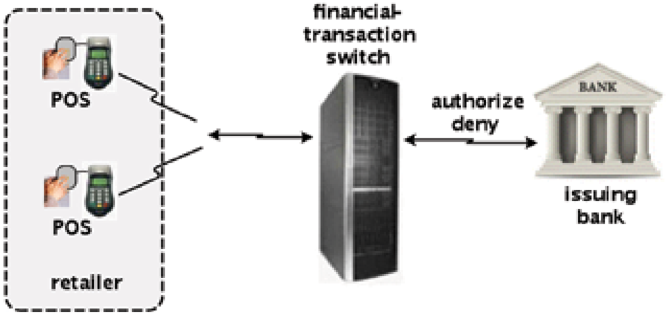

The FinancialTransaction Switch
Figure1

For payment-card transactions, the POS terminals send the transactions to the retailer's financial transaction switch, which forwards each transaction to the bank that issued the payment card (the issuing bank). The financial-transaction switch has access to all of the payment-card networks and routes each transaction over the appropriate network. The issuing bank decides whether or not to allow the transaction and returns an authorize or deny directive to the POS terminal via the financial-transaction switch. If the issuing bank authorizes the transaction, it is completed at the POS terminal. If the issuing bank denies the transaction, the POS terminal rejects the transaction.

Many large retailers can afford to operate their own financial-transaction switches in order to reduce the processing fees charged to them by acquiring banks. The POS terminals of smaller retailers are handled by the acquiring bank, which operates its own financial transaction switch. Alternatively, virtual financial-transaction switch instances are available in clouds, thereby allowing small and mid-size retailers to each have their own virtual switches but pay only for their individual uses of cloud resources.

Fraud Blocker versus Fraud Monitor

Essentially, there are three methods in use today to detect fraud – fraud blocking, fraud monitoring, and issuing-bank fraud detection. In all cases, a pending transaction is matched against a set of rules by a Transaction Screening engine.

The Rules Set

The rule set can be massive and may contain rules from retailers, acquiring banks, and issuing banks. If a transaction violates any of the rules, it is rejected (or in the case of fraud monitoring, it is reported).

These rules allow the establishment of authorization criteria. Examples include:

"Accept no transactions from a particular retailer over the weekend or between 11 PM and 5 AM."

"Decline all AMEX transactions."

"If a customer goes to the same store three times in three days, decline the fourth transaction."

"Accept no transactions from this customer."

"For a U.S. card, do not allow gas purchases in Mexico."

Fraud Blocker

A Fraud Blocker prevents fraudulent transactions from being sent to the issuing bank. A Fraud Blocker may take one of two forms. It may be embedded in the financial-transaction switch, or it may be a separate system that augments the financial-transaction switch.

Embedded Fraud Blocker

Figure 2 shows a Fraud Blocker that is part of the financial-transaction switch. The Fraud Blocker includes a set of fraud rules that transactions must pass in order to be forwarded to the issuing bank for authorization. Augmenting the rule set is a database of past payment-card history. This is needed because some fraud tests depend upon the prior use of the card. For instance, the card may be limited to only a certain number of transactions per day or a maximum amount per day.

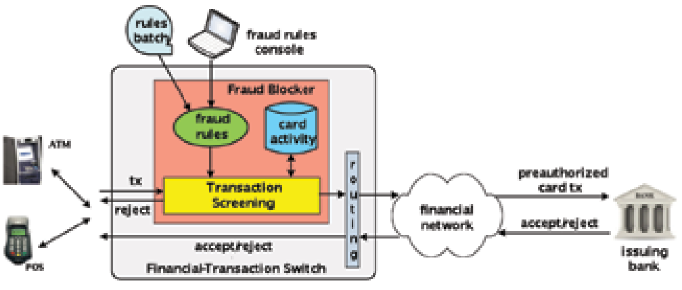

Embedded Fraud Blocker
Figure2

The financial-transaction switch can receive transactions from POS terminals, ATMs, and other devices either via the terminals' acquiring bank or directly in the case of a financial-transaction switch operated by a retailer. When the financial-transaction switch receives a transaction, it invokes its Fraud Blocker. The Fraud Blocker uses its Transaction Screening engine to evaluate the transaction against the fraud rule set. This is a process called preauthorization.

If the transaction fails the fraud check, it is rejected directly by the Fraud Blocker so that the issuing bank does not even see it. This provides a valuable service to the issuing bank as the Fraud Blocker can reduce the number of transactions that the issuing bank must process multifold (in some cases, as high as a factor of five). The issuing bank is not bothered with the "noise" of fraudulent transactions, which can exceed significantly the number of valid transactions.

However, if the transaction passes the fraud test, the preauthorized transaction is routed to the appropriate issuing bank for authorization. The issuing bank's decision – accept or reject – is returned via the financial-transaction switch to the POS terminal.

The fraud rule set can be quite large and is continually changing. It may include rules from retailers, acquiring banks, and issuing banks. Rules are typically updated nightly via a batch run, but they also can be modified online via a fraud-rules console.

Augmented Fraud Blocker

In many cases, it may be advantageous to augment an already existing financial-transaction switch with a Fraud Blocker. Figure 3 shows the technique for doing this.

As with the embedded Fraud Blocker, the augmented Fraud Blocker system uses a rules set that is established by the retailers, acquiring bank, and issuing banks. It maintains a database of past card transaction history. The rules may be updated nightly via a batch run, or they may be updated online via a rules console.

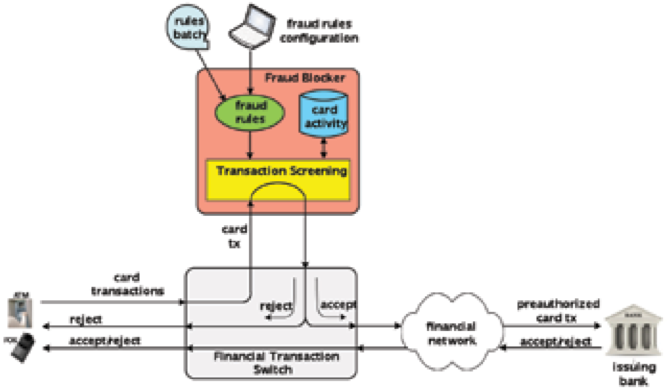

Augmented Fraud Blocker
Figure3

The financial-transaction switch is modified slightly to pass incoming transactions to the external Fraud Blocker, which is running on another system. The external Fraud Blocker checks the transaction against its rule set and returns an accept or a reject indication to the financial-transaction switch.

If the Fraud Blocker rejects the transaction, the financial-transaction switch sends a reject directive back to the terminal creating the transaction, and the terminal will reject the transaction. If the transaction passes the Fraud Blocker's test, the preauthorized transaction is sent to the appropriate issuing bank for authorization. The issuing bank's response – accept or reject – is returned to the financial-transaction switch, which passes it to the transaction terminal.

Fraud Monitor

As opposed to a Fraud Blocker, which is inserted into the transaction flow, a Fraud Monitor observes the transaction flow but is not in the flow path. A typical Fraud Monitor system is shown in Figure 4.

The Fraud Monitor system is much like a Fraud Blocker system. It comprises a Transaction Screening module that has access to a fraud rules set and to a record of card activity. The rules set can be updated by batch runs, or rules may be added via a rules console.

In the case of the Fraud Monitor, card transactions that are sent to the issuing bank by the financial-transaction switch for authorization are also sent in parallel to the Fraud Monitor system. It will check the transaction against its rule set to determine if the transaction is suspicious. If so, it sends to the issuing bank an indication that the transaction may be fraudulent.

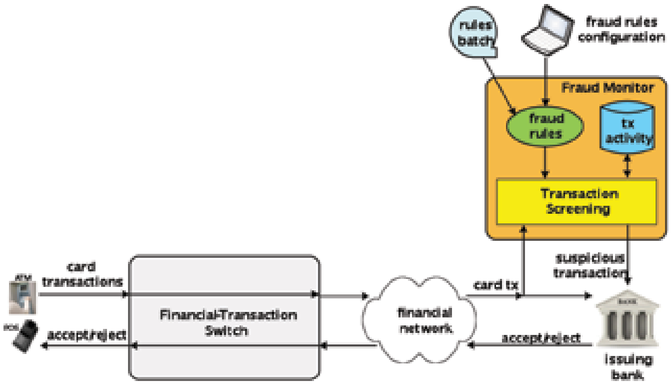

Fraud Monitor
Figure4

During this same period, the issuing bank is processing the transaction. The bank may or may not receive the suspicious indication from the Fraud Monitor before issuing its response to the transaction terminal, but the bank does not wait until the Fraud Monitor reports its findings.

If the Fraud Monitor detects a suspicious transaction and so informs the issuing bank's system, the issuing bank may apply more stringent checks to future transactions from that card; or it may block the card.

Issuing-Bank Fraud Detection

The most common technique in use today is for the issuing bank to perform its own fraud detection. However, the fraud rules set is massive and must be continually updated and maintained. Fraud detection against such a massive rules set is computationally intensive and consumes a lot of the issuing bank's processing resources. Furthermore, maintaining the fraud rules set imposes a significant administrative load on the bank's personnel.

Therefore, there is a move among issuing banks to offload this task to external systems, be they Fraud Blockers or Fraud Monitors. Since Fraud Blockers are reactive and protect the issuing system from having to process fraudulent transactions, they are often the preferred solution.

The Fraud Rules Set

We have mentioned many times that the rules set is massive. Let us take a high-level look at a typical rules set to establish this fact.

Two groups of rules can be configured for preauthorization. One group of rules is used by the fraud detection system, whether it be a Fraud Blocker or a Fraud Monitor. This can be considered Transaction Screening Level 1 (TS1).

Additional rules can be applied via a Card Database that contains additional cardholder information available to the issuing bank (Transaction Screening Level 2, or TS2). TS2 rules are applied by the issuing bank to transactions that have passed the TS1 level of screening in the fraud detection system.

TS1 Rules

At the TS1 level, Transaction Screening validates the card prefix and institution (the issuing bank is identified by the first numbers of the card – the card prefix). It ensures that the credit card or debit card is still current, that the card has not been lost or stolen, and that the transaction is allowed. It also verifies that the retailer and terminal are supported and are not blocked. If currency conversion is required, Transaction Screening performs the conversion.

Certain country codes, ZIP codes, merchant categories, issuing institutions, and card types can be restricted. In addition, credit-card transactions, debit-card transactions, and Not-On-Us transactions can be controlled for specified merchant categories.

The day, date, and time during which the transaction is made can be limited by institution, merchant, terminal, card prefix, and card account number. PIN-based transactions, signature-based transactions, ATM transactions, and POS transactions all can be controlled.

Manually entered transactions can be restricted, and duplicate transactions are rejected.

If the transaction passes the above tests, the transaction amount is checked against the minimum and maximum values for the institution, merchant, and terminal. The transaction is rejected if it exceeds the daily, weekly, bi-weekly, or monthly maximum amount or number of transactions allowed by the rules.

TS1 will determine the routing of valid transactions to the primary authorizer based on the issuing institution, the transaction source (ATM, POS), the account type, and the merchant.

TS2 Rules

The TS2 rules make use of a Card Database if one is available to the Transaction Screening module. The following rules are applied if the corresponding data is in the Card Database:

The card number exists.

The card is active and has not expired.

The account exists and is active.

The card address is valid.

The transaction amount is within the card limits.

There are sufficient funds in the account to cover the transaction.

Establishing the Rules

The rules set involves a number of databases for Prefixes, Institutions, Negative Cards, Positive Cards, Transaction Codes, Retailers, and Terminals. Each transaction type can be configured to use its own set of rules that can vary between issuing institutions, merchants, and even terminals.

Some of these databases are loaded into the fraud protection mechanism via batch flat files. They can be modified from the rules console.

The fraud-detection rules are very dynamic. They can be augmented in real time to protect the bank from evolving fraud schemes. The fraud rules console is used to make instant adjustment to the fraud-detection rules when needed.

Do Smart Chip Cards Defeat Fraud?

Magnetic-stripe payment cards have been substantially replaced with smart cards throughout the world except for the U.S., where the change is just now happening. A smart card, also known as a chip card, contains a small computer chip that encrypts data being transferred between the card and the POS terminal as well as between the card and the issuing bank.[1] The chip card contains many features to prevent it from being cloned. One example is a transaction counter that is sent with each transaction. If a card is cloned, the issuing bank will receive unsynchronized transaction counters and will block the card. Also, sensitive data stored on the card, such as the issuer's private encryption key, is instantly erased if an attempt is made to access it.

In the rest of the world, smart cards are chip-and-PIN cards. The cardholder needs to enter a PIN in order to complete the transaction. This is a strong protection against the use of lost or stolen cards. However, in the U.S., most smart cards that are being issued are chip-and-signature cards. They do not require a PIN. They only require that the cardholder sign a chit for each transaction. Since the signature can be forged easily, and since retail clerks typically do not ask for identification, chip-and-signature cards provide little protection for lost or stolen cards. Why are the U.S. card companies not using chip-and-PIN? It is a competitive issue. If a cardholder has two cards, one a chip-and-PIN and one a chip-and-signature, he is more likely to use the latter because it is more convenient.

If Steve Anderson had been using a smart chip card, would that have prevented fraudulent transactions? Not if he lived in the U.S. His card would have been a chip-and-signature card, and anyone could use it by forging his signature (or by depending upon the sales clerk to not ask for identification).

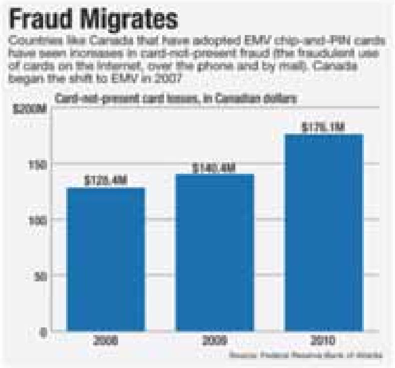

No matter what protection smart cards provide, the protections are good only for card-present transactions. Fraudsters are rapidly moving to use lost and stolen smart cards for card-not-present transactions, such as Internet and mobile purchases. Already in Europe, the rate of online fraudulent transactions is increasing due to card-not-present transactions.

Consequently, the need for fraud protection has not been diminished by the prevalent use of smart cards.

Comparing the Fraud Protection Approaches

We have discussed several approaches to fraud protections – Fraud Blockers, Fraud Monitors, and issuing-bank detection. We have shown that a Fraud Blocker is proactive in that it takes direct action if it detects a fraudulent transaction. It directly rejects the transaction and does not send it to the issuing bank.

On the other hand, a Fraud Monitor is reactive. It simply sends suspicious indications to the issuing bank for transactions that do not pass its rule set. It takes no direct action. A suspicious transaction may still be authorized before the issuing bank receives the suspicious indication from the Fraud Monitor, but the issuing bank is now alerted to the possibility of a card being fraudulently used.

Issuing bank detection is as effective as the use of a Fraud Blocker. However, this activity places undesirable processing and administrative loads on the issuing bank's IT resources.

Therefore, the use of a Fraud Blocker is often the preferred method for fraud protection. The majority of U.S. fraud-detection systems are moving to fraud blocking.

The OmniPayments Fraud Blocker

What Is OmniPayments?

The OmniPayments financial-transaction switch (www.omnipayments.com) is a product of Opsol Integrators Inc. (www.opsol.com). The switch's architecture is illustrated in Figure 5. OmniPayments is a layered architecture and is built upon the fault-tolerant HP NonStop server. All processes are persistent and are automatically restarted should they abort. All database functions such as logging, card parameters, and preauthorization rules are maintained by the NonStop SQL/MX relational database. OmniPayments supports NonStop Enscribe file systems as well as SQL/MP and SQL/MX relational databases.

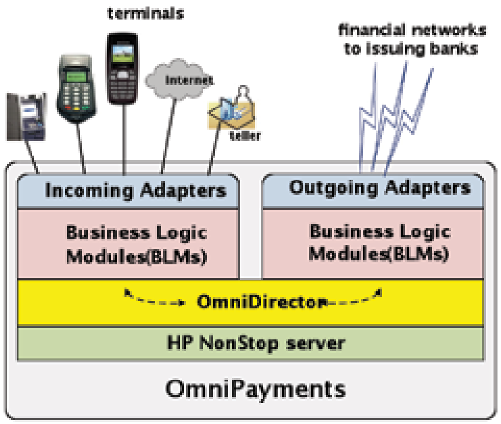

OmniPayments FinancialTransaction Switch
Figure5

The core layer of OmniPayments is Opsol's OmniDirector Enterprise Service Bus. OmniDirector services include data transformation, encryption, intelligent routing, and communication-failure recovery. OmniPayments also provides adapters to support the protocols required to communicate with terminals and with other card-interchange networks.

Business logic modules, or BLMs, supply the business functions of OmniPayments. Each BLM is tasked with providing a specific service. For instance, BLMs were created to support the rules functionality needed for preauthorization.

OmniPayments Fraud Blocking

Embedded Fraud Blocker

The OmniPayments financial-transaction switch supports a Fraud Blocker that can be incorporated into the switch. The structure and operation of the OmniPayments Fraud Blocker is essentially the same as that described earlier in this article for an embedded Fraud Blocker.

An OmniConsole browser interface is provided for the acquiring bank to configure rules for transactions received from its own devices and from financial-transaction networks. The acquiring bank not only can configure preauthorization rules for On-Us transactions that it will process, but it also can configure rules for Not-On-Us transactions that will be sent over financial-transaction networks to other issuing banks.

On-Us transactions that pass the prescreening tests are forwarded to the acquiring bank. If the acquiring bank's authorization system is unavailable, OmniPayments can optionally provide stand-in authorization. Based on a negative card file or other information, OmniPayments can make the decision as to whether or not to authorize a transaction. OmniPayments queues all stand-in transactions and sends them to the bank's authorization system as soon as it becomes available.

Transaction screening can be configured for an OmniPayments switch that serves multiple acquirers and issuers. Each acquirer can configure different rules for different issuers. The rules governing transaction screening are therefore based on the combination of acquirer and issuer.

The OmniConsole rules-configuration browser is protected by ACLs (access-control lists) so that only an authorized administrator can view and configure the rules for his specific acquiring institution.

Augmented Fraud Blocker

The OmniPayments Fraud Blocker also can be configured to provide fraud-blocking services for other financial-transaction switches. For instance, OmniPayments has developed a BASE24-eps Custom Support Module (CSM) to route all transactions received by a BASE24 financial-transaction switch to an OmniPayment's Fraud Blocker running on a separate NonStop system. Transactions received by the BASE24 system are routed to the OmniPayments Fraud Blocker, which returns an accept/reject directive to the BASE24 system. The BASE24 system will forward to the issuing bank for authorization only those transactions that have been approved by the OmniPayments Fraud Blocker.

Fraud Blocking in the Cloud

Fraud blocking is now available to small and mid-size retailers who would like to have their own financial-transaction switches. OmniPayments provides virtual financial-transaction switches in its OmniCloudX, a redundant cloud service that runs in multiple locations (for outage recovery) on HP's NonStop X servers. The new NonStop X servers are based on Intel x86 technology and have improved price/performance compared to the NonStop Itanium series. The cloud-based OmniPayments switches are made available to retailers who only have to pay for the cloud resources they use.

The OmniPayments cloud-based financial-transaction switches in OmniCloudX support fraud protection via an embedded Fraud Blocker. In addition, virtual Fraud Blockers can be provided by OmniCloudX independently of a financial transaction switch.

Individual Preauthorization

OmniPayments has extended its preauthorization services to include individual preauthorization. In addition to rules established for retailers, acquiring banks, and issuing banks, OmniPayments allows companies and even families to establish rules through their issuing banks for card use.

For instance, a small company may issue credit cards to its employees. It can establish various purchase limits on each card to support individual employee activities. Likewise, a family can issue credit cards to its children with the restriction that they cannot be used for alcohol or cigarettes. Multiple credit cards with different restrictions can be issued to a single person.

Opsol Integrators

With successful implementations at many customer sites, OmniPayments is just one member of the Opsol family of solutions for the financial, telco, and other industries. Opsol Integrators specializes in NonStop mission-critical applications and is HP NonStop's largest system integrator.

Summary

Until recently, fraudulent payment-card transactions had to be covered by the merchants. With the advent of smart chip cards, this liability now is being shifted to the banks. If a merchant does not process at least 75% of its transactions through an EMV-enabled POS terminal (whether via chip cards or magnetic-stripe cards) and accepts a disputed or fraudulent card payment, the merchant will be liable for the transaction rather than the issuer. However, even the intelligence built into the new chip cards is insufficient to make a major dent in payment-card fraud.

Fraudulent transactions represent a major cost to banks and retailers. There is a pressing need for technology to significantly reduce the frequency of payment-card fraud. Putting this responsibility on the issuing bank represents a major commitment by the bank for extensive computing and administrative resources. The Fraud Blocker is a proven solution to this dilemma. The OmniPayments Fraud Blocker is available to the issuing banks as well as to retailers and acquiring banks to provide this service, whether or not an OmniPayments financial-transaction switch is currently being used.

Yash Kapadia is the founder and CEO of OmniPayments Inc. and Opsol Integrators, Inc., leading HP NonStop systems integrators for Telco and Financial Services. Opsol's OmniPayments solution is used by banks and retailers for BASE24 replacement. Yash and his team provide several products and remote managed services for NonStop customers. He can be reached at Yash@OmniPayments.com and at +1-408-446-9274.

WHAT'S NEW IN OmniCloudX?

OMNIPAYMENTS INC. NOW HOSTS ITUGLIB FOR FREE

FREE PHYSICAL SPACE · FREE MAINTENANCE · FREE POWER · FREE BANDWIDTH

OmniPayments

The financial transaction switch that replaces BASE24

www.omnipayments.com

Above the Cloud

HP Helion OpenStack®

Announcing

Introduction to OpenStack® Self Paced Training

from HP Helion Education Services

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

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

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

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

Why HP

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

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

Learn more at
hp.com/learn/cloud

Available Here

© Copyright 2014 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.
®The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission.

Tokenization - Your Last Line of Defense

Thomas Gloerfeld >> VP Marketing >> comForte 21 GmbH

Thomas Burg, CISSP >> Chief Technical Officer >> comForte 21 GmbH

Introduction to Tokenization products for HP NonStop

The PCI data security standard requires that credit card numbers need to be protected when stored on disk: PCI 3.4 can be paraphrased as "Thou shalt not leave PAN data on disk to be readable in plain". Tokenization has emerged as a favorite technology to implement PCI 3.4 and several articles in this magazine have described both the technology as well as the motivation of using a Tokenization product in detail.

As of today, several vendors offer products on the HP NonStop platform which perform the following basic tasks:

- Provide the mechanics of "Tokenization" by mapping individual credit card numbers (PANs) to a so-called "Token" and doing the reverse operation. Products for HP NonStop are available from comForte (running exclusively on HP NonStop) and HP Security Voltage (running on multiple platforms)
- Provide an "Integration Framework" which enables existing applications to use tokenization for data protection without having to change the source code. Products for HP NonStop are available from XYPRO ("XDP") as well as comForte ("SecurData")

For most HP NonStop customers, only the combination of the tokenization engine and the Integration Framework provides a practical solution as the changes of the underlying application are simply not feasible.

Use cases - Overview

As mentioned in the Introduction, the most typical use case and driver is compliance with PCI 3.4. Following a still growing trend and impact generated by data breaches affecting the card payments industry, auditors are increasingly focusing on the protection of stored data. It is obvious there is already decreasing tolerance for compensating controls in this area making the implementation of a tokenization solution the only option. The PCI Security Standards Council has just released its second document in April 2015 providing guidance for implementing Tokenization of Cardholder details.

In this article we will also look at some other, less obvious, use cases:

- Scope reduction
- Generation of secure test data
- Improving partitioning efficiency
- Encryption of password databases

"Scope reduction" to reduce compliance costs

In its recent “VERIZON 2015 PCI COMPLIANCE REPORT”, available from

http://www.verizonenterprise.com/pcireport/2015/, Verizon states that scope reduction is a key element to bring down compliance costs and to reduce risk:

For all these reasons, it is strongly recommended that organizations look at implementing a sound scope-reduction strategy. This should be done right at the start of your compliance initiative as practically everything else is based on the defined compliance environment. (…) Organizations realized the benefits of moving away from encryption, in particular due to the challenges around cryptographic key management, and the increased frequency of attacks where memory parsing malware is used to extract keys or sensitive data directly from RAM. In 2014, 12% of organizations in our dataset used tokenization.

While your HP NonStop system will never be out of scope as it handles the actual credit card data for live processing, several comForte SecurData customers were able to take systems out of PCI scope by completely replacing the PAN data with tokens.

Creating secure test data

PCI-DSS section 6.4.3 requires that *"production data (live PANs) are not used for testing or development"*. Let us assume you need a huge dataset of "real-looking" PANs for a stress test or functional test - how do you create this data?

A good tokenization solution will allow you to do a "batch copy" where you start with PAN data and end up with tokens. Doing this in production once is a fast and secure way to create test data: When configured properly, the tokens can look like "real" PANs and can even pass the Luhn check.

Thus creating a huge set of test data which is not tied to real PAN data at all becomes an easy exercise.

Improve partitioning efficiency – A very welcome side effect of tokenization

The problem of partitioning large files based on PANs

Load-balancing a large BASE24 installation across available system resources can be a daunting task. One of the problems is to balance random access by card number to the cardholder file (CAF) or transaction log files (PTLF/TLF). To balance the access, these files (or their alternate key files in case of transaction logs) are partitioned by card number ranges across multiple disk volumes.

However, it is very hard (if not impossible) to define the card number ranges for the partitions to be evenly utilized. Furthermore, if cards are added or withdrawn over time, the system tends to get out of balance again. This leads to repeated efforts to “re-partition” the relevant files. The effort includes the regular monitoring and analysis of performance data, the identification of hot spots, the prediction of future trends, definition of new card number ranges for the partitions and finally a full reload (refresh) of the file.

Very often this has to be done on a trial and error basis because it is very hard to predict how card number ranges will contribute to future transaction activity. The need for repartitioning may occur once a year or even more often, each time creating a re-partitioning project which will bind costly technical resources.

The solution: tokenization randomizes the distribution

By its very nature, a Tokenization solution replaces the PAN with a random token. A proper Tokenization algorithm will result in a full randomization of the PAN and thus significantly simplifies the definition of key ranges for partitions of files that are accessed with the PAN as key. After Tokenization the typical "peaks" around certain credit card number ranges (i.e. 37… for Amex) will vanish.

Even if cards are added or withdrawn, there is no need for re-partitioning any more. As an illustration, the following chart shows the effect of the comForte SecurData PAN randomization on CAF partition utilization at one of comForte's customers:

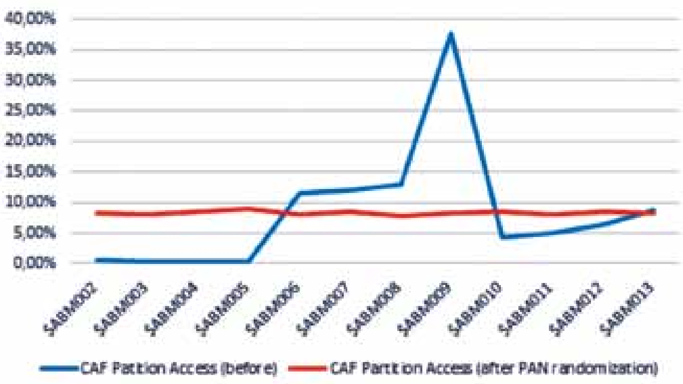

Protect Password databases

Many applications running on HP NonStop have their own user database with associated passwords. Some examples are BASE24, BESS and Atlas - but there are many more, most of them being home-grown. While SafeGuard does a good job of storing passwords of system users (a.k.a. Guardian users), applications are typically left on their own to store passwords.

Historically, simply storing user names and passwords in a database - and securing this database via standard O/S mechanims - has been both "best practice" and "good enough".

With the recent increased security awareness, storing application passwords "in the clear" is no longer an option and comForte has already helped a few customers to move to a secure password storage solution.

This section discusses what options exist to migrate a "legacy application" from non-secure password storage to secure password storage.

What PCI says about safe storage of passwords

PCI puts a clear emphasis on the secure storage and processing of passwords. *"Requirement 8: Identify and authenticate access to system components"* of the PCI-DSS standard version 3.1 contains several specific rules on how passwords need to be processed and stored. The introduction to the section on Requirements 8 states: *The effectiveness of a password is largely determined by the design and implementation of the authentication system—particularly, how frequently password attempts can be made by an attacker, and the security methods to protect user passwords at the point of entry, during transmission, and while in storage.*

Specifically, PCI requires the following: *8.2.1 - Using strong cryptography, render all authentication credentials (such as passwords/phrases) unreadable during transmission and storage on all system components.*

The art and science of storing passwords securely

Protecting passwords sounds simple - but if one starts looking in detail, it is NOT. Why is this the case?

When analyzing the security of any solution to protect a database, the following assumptions are usually made:

- The attacker has gained access to the whole "password database file"
- The attacker has the source code and/or knows the algorithms (if any) used to protect the database

The topic of "protecting a password database" has been an important one for a long time and the industry best practice approaches uses the following key concepts:

- Use one-way hash functions (rather than encryption). After all, there is no need whatsoever to store the actual passwords. One-way hash functions are good enough to verify a user input for "match" versus "no-match"
- Use of additional salt for the hash function. This is mostly used to prevent the so-called "rainbow table" attack against password files[1]
- Use a slow hash function

A good write-up of all this can be found at https://crackstation.net/hashing-security.htm. In summary, it is *not* trivial to add proper protection to an existing application maintaining its own password database.

Options for HP NonStop users

Assuming one has an existing application which implements its own user and password database and to which one wants to add a proper protection layer, the following options are possible:

- Implementation of secure hashing as described in the prior section
- Use of strong cryptography. This is also valid - but also difficult as key management (including regular rotation of keys) needs to be included in the design
- Use Safeguard and Guardian passwords. Safeguard protection of passwords is very good, however this approach requires to add every application user to the Guardian user database
- Leverage other password databases - such as Microsoft Active Directory which can be accessed via various ways. Products such as comForte SecurSSO or XYGATE/UA will make implementing this option relatively painless

All of the above approaches require changes to the application ranging from small to somewhat major.

- as Tokenization is format-preserving, no database changing is required
- no source code changes in the applications are required

[1] See *Security in Computing (5th Edition)*, by Charles P. Pfleeger, February 2015, Chapter 2, Section "Defeating Concealment" for a more detailed explanation

Conclusion

Tokenization is a quickly emerging technology and has already proven to be an effective way to protect critical data stored in databases or files. Tokenization at its best is transparent to the application using the file or database.

Compared with encryption technologies, tokenization offers a number of unique benefits:

- uncompromised system performance and response time as already achieved in high performance payment systems implementations,
- no algorithmic possibilities to revert from tokenized data to real data,
- possibility to take systems out of compliance audit scope thus reducing ongoing costs of compliance

While tokenization alone will by no means reduce the probability of your company being attacked or of the attacker penetrating your first line of defenses, it provides a very effective method to dramatically reduce the potential impact of a cyber-attack. If implemented properly, tokenization makes all the difference between a major data breach and a failed cyber-attack.

Tokenization can be used very selectively to "hide"

- Single data elements (i.e. PAN numbers, PII (personal identifiable information, passwords, social security numbers, etc.) while keeping non-critical data elements in clear text.
- Entire files or databases

For many, tokenization is becoming a best practice for adding an extra layer of data protection to any type of data at rest.

When access control measures fail and intruders manage to get unauthorized access to files and databases, tokenization is your last line of defense making sure that the data will be of no use to the criminal. As a result, the compromise of confidentiality, integrity or availability is averted.

In this article we discussed the current products available and some of their typical and not-so typical use cases. If data stored on your NonStop systems is considered a critical asset within your Enterprise, then the time to look at these products is now.

Thomas Gloerfeld has been associated with the NonStop community for 20 years. Before joining comForte twenty years ago, he held various management positions at ACI Worldwide, both in Germany and the UK. Thomas Gloerfeld can be contacted at t.gloerfeld@comforte.com.

Thomas Burg has an extensive background in systems programming, networking, and security. For more than 30 years, Thomas has worked with a range of computing platforms, including Windows, UNIX, and HP NonStop. At comForte, he has helped guide the company's strategic product direction and orchestrated a range of technology initiatives, such as the company's SSL/SSH encryption suite, which was ultimately adopted by HP within the NonStop OS. Thomas Burg can be reached at t.burg@comforte.com.

2015 NonStop Technical Boot Camp
Register TODAY!
CONNECT
Your Independent HP Business Technology Community
November 15-18, 2015
The Fairmont San Jose Hotel
San Jose, California
#NonStopTBC2015

Discover2015

London | 1–3 December

Join HP at the ExCeL London

Test your Security Knowledge!

Answers from page 9:

1.) C

Explanation: Multi-factor Authentication (MFA) combines two or more independent credentials: what the user knows (password), what the user has (i.e. security token) and what the user is (biometric verification). The goal of MFA is to create a layered defense and make it more difficult for an unauthorized person to access a target such as a physical location, computing device, network or database. If one factor is compromised or broken, the attacker still has at least one more barrier to breach before successfully breaking into the target.

2.) A

Explanation: Although an individual may have access to a system, it does not necessarily mean that the user is authorized to access all the programs and files on that system.

3.) B

Explanation: Password rotation is the act of regularly changing passwords. This is standard practice and a requirement in all security compliance programs and regulations like PCI-DSS.

4.) B

Explanation: If at any point in the future you need to recover data you will need to have the key to access the data.

5.) D

Explanation: An idle timeout logs out an interactive session after a set amount of idle time. The idea behind it is if the user is not actively participating in the session, then their session does not need to be kept active. This helps prevent unauthorized access and free up unused system resources.

6.) D

Explanation: Single sign-on (SSO) is a session/user authentication process that permits a user to enter one name and password in order to access multiple applications. The process authenticates the user for all the applications they have been given rights to and eliminates further prompts when they switch applications during a particular session.

7.) A

Explanation: The principle of least privilege (POLP) is the practice of limiting access to the minimal level that will allow normal functioning. Applied to employees, the principle of least privilege translates to giving people the lowest level of user rights that they can have and still do their jobs.

8.) A

Explanation: Tokenization is the term used when PII, PAN or other sensitive data is replaced with a benign substitute. Tokens can be data hashes, surrogate values, or any other identifier that cannot be derived from the original value. The token by itself must have no value to an attacker and purpose of tokenization is to reduce the PCI-DSS footprint.

9.) B

Explanation: POODLE is a vulnerability in the SSL protocol. To ensure that your systems are not vulnerable to POODLE, confirm that you're only running Transport Layer Security (TLS) version 1.0 and above. This will prevent Man-In-The-Middle attacks.

10.) C

Explanation: Spear phishing doesn't rely on volume. Instead, it relies on targeting specific targets in an effort to steal their personal information. Celebrities had their compromising photos revealed that had been stored on Apple's iCloud service not by attacking the technology (although Apple's password practices at the time were a contributing cause), but by using the celebrities fame and publicity to guess their weak passwords. And it's not just celebrities – executives in companies with more than 2,500 employees have a 1 in 2.3 chance of becoming the target of a spear phishing attack.

11.) D

Explanation: Database auditing records all user activity for forensic investigations, unauthorized access, database accountability and deters unauthorized activity.

12.) B

Explanation: Connecting to a unknown WiFi network that provides SSL or other encryption does not ensure that you are connecting to a secure network. Only connect to WiFi networks that you trust, and that you are sure are secure.

13.) B

Explanation: Social engineering is a non-technical method of intrusion hackers use that relies heavily on human interaction and often involves tricking people into breaking normal security procedures and revealing sensitive information like their passwords. It is one of the greatest threats that organizations encounter today. Social Engineering is designed to bypass all of the standard procedures such as Firewalls, Encryption and Antivirus software.

14.) D

Explanation: Here is an example of what a phishing scam in an email message might look like:

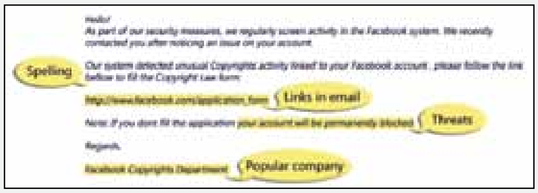

Hello!

As part of our security measures, we regularly screen activity in the Facebook system. We recently contacted you after noticing an issue on your account.

Our system detected unusual Copyrights activity linked to your Facebook account , please follow the link bellow to fill the Copyright Law form:

http://www.facebook.com/application_form

Note: If you dont fill the application your account will be permanently blocked.

Regards,

Facebook Copyrights Department

15.) D

Explanation: Ensure that all hardware and software installed on your computer comes from a trusted source.

Back for More...

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

I have waited until the very last moment to write this issues back page column as I wanted to have time to reflect on the annual America's big tent event for HP –2015 HP Discover. It's an occasion to be "up close and personal" to the HP management and executive team and this year I just happened to run across CEO, Meg Whitman, EVP & CTO, Martin Fink, and Senior VP & GM, HP Enterprise Group, Antonio Neri, inside the HP Labs booth on the exhibition floor. I only mention this as I was taken aback by the casual nature of the run-in as it was very apparent that Whitman was there, with Martin and Antonio, but devoid of the usual group of "minders" I typically assume CEOs have around them at all times.

However, what this also reinforces for me is that HP senior executives today are more willing than at any time in the past to simply mingle with the communities they support and are open to just about any dialogue that may occur – and it was clear to me that the folks they brushed up against were only too willing to engage this trio in conversation. What triggered this occurrence was that the second general session of the event had just wrapped up and as it had featured Martin and what HP Labs had been up to of late, it was an opportunity for Martin to showcase his team and their accomplishments – alongside the lad were stands with different components definitely part of The Machine and so, naturally enough, warranted closer inspection by those most likely to be effected by its introduction.

For the NonStop community, the news coming from this event is that even though there's still several months to go before the split of HP is realized – it's all going to happen as of November 1st and coincide with the beginning of a new financial year – current HP CEO, Whitman, is well into her new role as Hewlett-Packard Enterprise (HPE) CEO and very much focused on all things Enterprise related. Having just unveiled the new logo and color scheme for HPE, it was at the HP Labs where the green rectangles of HPE were on display. So many, in fact, and so green, that many at HP were referring to HP Labs as the "emerald isles".

If you had second thoughts about the future of HP Discover then fret no more – HPE "owns" HP Discover and even with the split still months away, this will remain the showcase, big tent, marketing event where HPE will showcase current and future product offerings and even though it took two attempts, I did find the NonStop X system on the exhibition floor along with Mark Pollans from product management, Mike Rice from development and HP Master Technologist, Justin Simonds never more than a few feet away from the latest offering from the NonStop group, within HP's Enterprise Group. The opportunity to really observe all that constitutes the HP Integrity NonStop X NS7 X1 – pulling open both front and back doors was something I relished doing – proved quite a revelation.

Hybrid computers, featuring NonStop? With 16 slots available for the blades packages, and standard connections to InfiniBand switches, it was very apparent how easy it will become to mix NonStop with Linux (and even Windows) within the current chassis, and I anticipate that middleware and solutions vendors are already working out the kinks of their product offerings that will be optimized in support of such hybrid configurations. And this will only add to the interest many of us within the NonStop community will have as we contemplate what might be announced at NonStop Technical Boot Camp 2015.

If creating two $50 billion companies was a mind-blowing exercise – how do you split all the real estate holdings? – consider all that is involved in creating two IT departments from what already exists. Before CTO Fink was introduced, Whitman brought current EVP & Chief Customer Officer, John Hinshaw, on stage to update us on the magnitude of the task of bringing a new IT organization up and running for all of HPE. The numbers start small, but then escalate in size to something few other companies in the world have ever faced. 6 data centers, 50,000 servers, 2,800 applications (being separated, of course), 75,000 application interfaces (also requiring separation and fine-tuning), with 570 projects overseeing some 172,000 integration tests and with the expectation to start running (in parallel for a time) as of August 1, 2015. Yes, big indeed.

However, what this all comes down to is that at the end of the day, as the new financial year starts, HP will have built expertise with first-hand real-world experience under its belt, capable of talking any potential migration opportunity unlike any other vendor and this is encouraging for all companies, no matter how big or how complex, that may contemplate moving to HP. Like many within the NonStop community up until now I have harbored doubts about HP's participation in the really big projects we often associate with governments as well as the really big global conglomerates, but no more. With NonStop X systems under consideration by some of the companies, as part of a network of distributed systems, perhaps it's not as unrealistic as we may have once thought.

As HP Discover 2015 wound down and the crowds began departing, there was much that reminded me of big user run conferences of the past, even as there was no escaping that this was a premier marketing opportunity for HP. It certainly is not a substitute for the upcoming Boot Camp, but all the same, it does warrant our attention and I encourage all in the NonStop community not to lose sight of just how much value can be derived from attending – yes, your NonStop X system may be hard to find, but on the other hand, don't be surprised by who you might run into as you move around the exhibition floor. You might even be able to strike up a conversation with the very folks who make it all possible and for most of us that can only ever happen at big tent events such as HP Discover.

SQLXPress

Not just another pretty face

An integrated SQL Database Manager for HP NonStop.

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

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

XYPRO®
Mission Critical Security

Learn more at
xypro.com/SQLXPress

©2014 XYPRO Technology Corporation. All rights reserved.
Brands mentioned are trademarks of their respective companies

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

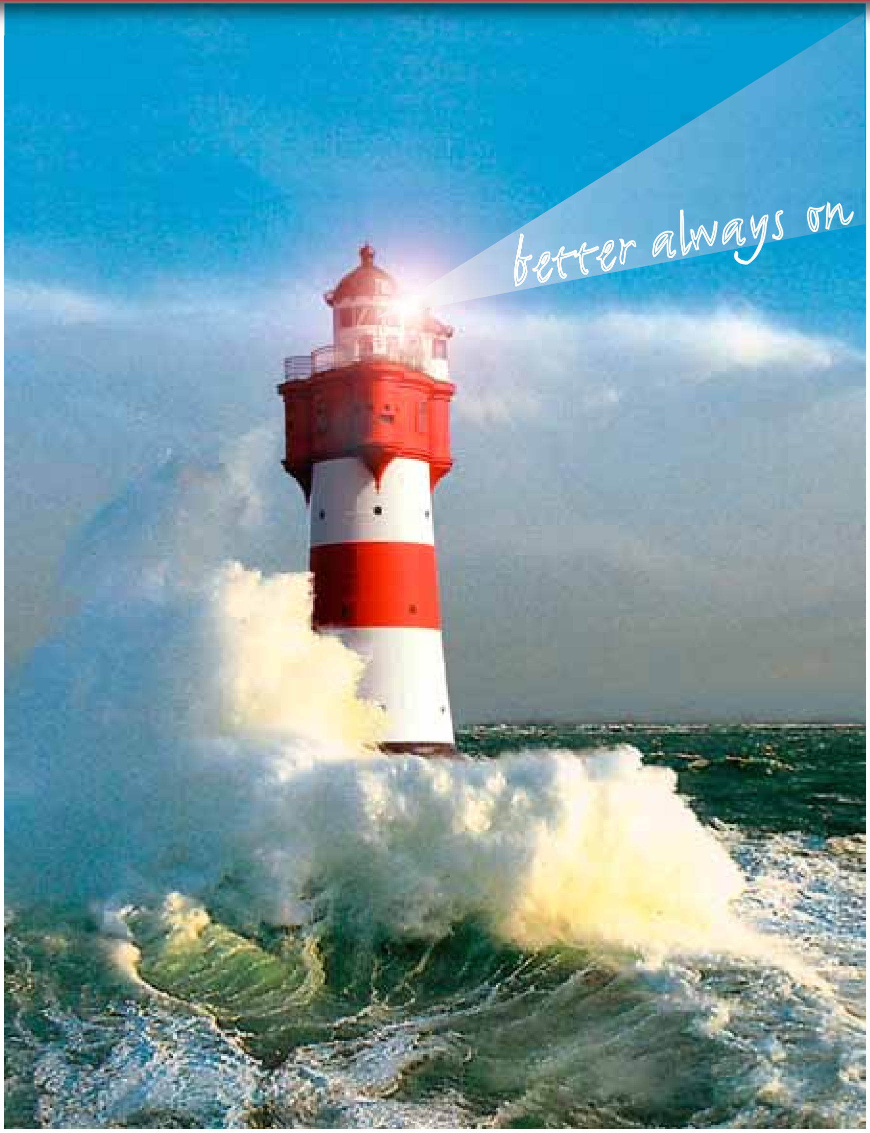

www.comforte.com/better_always_on

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

comForte „better always on" solutions help you gain …

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

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

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

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

comforte®

better always on

NonStop is a trademark of Hewlett-Packard Development Company, LP. All other trademarks are acknowledged. ©2015 comForte 21 GmbH. All rights reserved. Jan12_2015