

The Connection

A Journal for the Hewlett Packard Enterprise Business Technology Community



New Solutions on NonStop to Protect Your Sensitive Data

- ▶ Implementing Tokenization & Access Control
- ▶ Where Tokenization Fits Within Your NonStop Security Plan
- ▶ NonStop File Integrity: Check It! Protect It!



Security Solutions for your HPE NonStop Environment



Exceeding your HPE NonStop security,
compliance & encryption needs for over 30 years

Learn more at
www.xypro.com



Monitor 100% of Your Processes. In 100% Real Time.



ALTERNATIVE THINKING ABOUT PERFORMANCE MONITORING:

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

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

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

Technology for better business outcomes.

HP NONSTOP RPM

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

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



NonStop

Technical Boot Camp

**SAVE THE
DATE**

November 13-16, 2016

Fairmont San Jose Hotel | San Jose, CA

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

REGISTER ONLINE TODAY!

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

Table of Contents



14

14 Security? A Never-Ending Story!

Greg Swedosh



20

20 Privacy Throughout the World

Tim Roake



27

27 You Don't Know What You Don't Know (But Neither Does Anyone Else)

Luther Martin, Stacia Topping & Amy Vosters



42

42 What is Identity-Based Encryption (IBE)?

Josh Lubliner

- 10 Breaches are from Mars Security is from Venus Steve Tcherchian
- 13 Welcome to BITUG Matthew Whiteman
- 18 NonStop Family Legacy Mandi Nulph
- 24 New Solutions on NonStop to Protect Your Sensitive Data Prashanth Kamath
- 29 NonStop File Integrity: Check It! Protect It! Callum Barclay
- 31 Leveraging a Big Data Analytics Engine for Meaningful Insights Keith B. Evans & Paul J. Holenstein
- 35 Implementing Tokenization & Access Control Andrew Price & Scott Uroff
- 38 Where Tokenization Fits Within Your NonStop Security Plan Thomas Burg
- 40 Format Preserving Encryption Karen Martin

Columns...

05 A Note from Connect Leadership

Rob Lesan

06 News from HPE's NonStop Enterprise Division

Andrew Bergholz

08 ADVOCACY Migrating IBM Power Systems to HPE Open-Standards Platforms

Dr. Bill Highleyman

16

NonStop Innovations Deep Dive Tributary Systems Plans for Continued Growth with New CRO and Product Development

Mandi Nulph

44

Back for More...

Richard Buckle



OmniPayments

Financial Transaction Switch

Migrate to OmniPayments

Pain-Free Transition, No Disruption to Customer Services

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

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

Affordable OmniCloudX

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

OmniPayments Fraud Blocker

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

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

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



A Note from Connect Leadership

I had a proud moment the other day. Kind of like when you get to brag about your children (or grandchildren):

I was enjoying a beverage with a group of Mainframe support team members. We had just completed a long day of meetings on HP NonStop topics. Only one of the Mainframe team had attended. At some point, the question was asked "What is a NonStop server and why do we have them?" The company has had their servers for decades. Everyone knows about the machines, but like so many other places, no one ever asks. To my utter amazement, the answer straight from the Mainframe guy was: "It is a mainframe system. HP makes them. They run application x".

Unprompted, without influence, an old school IBM Mainframe systems person tossed this out. The look on my face must have been an odd mixture of happiness and confusion as I have NEVER heard anyone who wasn't raised on NonStop calling my systems a mainframe. We spent the next few hours educating each other on the benefits (and difficulties) of managing our chosen systems.

What I wish I could reproduce at any gathering of platform advocates, whether it be Windows, Mainframe, Linux or NonStop is that spark of understanding. We all enjoy what we do (at least I hope we do) and there is a certain devotion/dedication and pride that goes along with it. At the end of this discussion, we all laughed and really appreciated what the others go through every day.

The next time you run into someone who isn't as savvy as you are on the subject of our favorite platform, take the time to bring them up to speed. While you are at it, try and build a bridge by trying to see why they love what they do as much as you. You never know where the next convert may be.

Thanks,
Rob Lesan

Rob Lesan
XYPRO

2016 Connect Board of Directors



PRESIDENT
Rob Lesan
XYPRO



PAST PRESIDENT
Henk Pomper
Plusine ICT



VICE PRESIDENT
Michael Scroggins
Washington St. Community College



DIRECTOR
Trevor Jackson
SOCAN



DIRECTOR
Jamil Rashdi
wasl Asset Management Group



CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide



HPE LIAISON
Janice Zdankus
Enterprise Group, HPE

The Connection

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

Kristi ElizondoCEO
Stacie Neall Managing Editor
Kelly Luna Event Marketing Mgr.
Janice Reeder-Highleyman..... Editor at Large

Technical Review Board

Dr. Bill Highleyman
Karen Copeland
Thomas Burg
Bill Honaker
Justin Simonds

We welcome article submissions to *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.

© 2016 by Connect
All company and product names are trademarks of their respective companies.



News from HPE's NonStop Enterprise Division

The NonStop business continues to shine in the summer months as the NonStop X product ramps up and the NonStop Vision invigorates our partner and customer base on the bright future of the NonStop platform. HPE Discover 2016 in Las Vegas had a really cool HPE-IT demo on their future state with the vNonStop solution as their transactional engine. Meg also gave mention to NonStop in her keynote speech at Discover. The NonStop business is gaining momentum as we pick up the pace of innovation across the NonStop portfolio.

Innovation has greatly accelerated, leading off with the introduction of the NonStop X family and, in L16.05, the NonStop Application Direct Interface (previously known by its internal name, Yuma). There's much more in the works, including the virtualized NonStop (vNonStop) and other cool new features up and down the software stack. The development teams are in startup mode, and they're responding to my frequent "can you do this?" challenges with "yes, we can." Some of the architectural philosophies introduced with the NonStop X platform have enabled our development team to unleash their creative potential. We're very excited by the technological advances we're making and the new solution spaces that they will open up. We're also excited by how well we've been able to contain the changes in support of vNonStop within the lowest levels of the system, allowing most of our privileged code to just work on either the converged, physical platform or the virtual platform without recompilation. We have several internal partners running on the vNonStop POC system, and "It Just Works" for them too.

It takes much more than a great platform and software stack from HPE to let NonStop systems reach their potential in the composition of robust, mission-critical solutions. In May we hosted a one-day Partner Technical Symposium in our Palo Alto headquarters, where we were able to share information on our future plans and get on-the-spot feedback from our partners in interactive sessions. We also met with many partners individually for more in-depth discussions on where they and we saw new opportunities in an x86-based, virtualized world. The feedback from partners has been very positive. I'd like to thank Karen Copeland and Phyllis Longbons for planning and carrying out the event, as it really helps all of us align our respective products with the NonStop Vision.

This issue is focused on security, which is at or near the top in importance for almost every customer that I talk to (and I talk to a lot of you!). Two big, relatively recent security innovations have been the introduction of HPE Format Preserving Encryption (FPE) and HPE Secure Stateless Tokenization (SST). The win for HPE FPE is that classic symmetric encryption using cipher suites such as AES produces an encrypted result that is larger, and sometimes much larger, than the original plaintext value. This wreaks havoc on database column definitions, message formats, and anything else that assumes the original size. HPE FPE is able to provide strong encryption while preserving the data element's format. HPE SST is a form of tokenization that does not require the maintenance and protection of a token vault. HPE Enterprise Security recognized the criticality of these data protection options as part of a comprehensive security suite, and brought them

into HPE by acquiring Voltage Technologies. The Voltage products, including the Secure Stateless Key Manager, are now collectively known as HPE SecureData Enterprise. As their name suggests, the products were designed from the perspective of data-centric security to provide end-to-end protection for sensitive data throughout your enterprise, not just while the data is resident on a particular platform.

Adding classic encryption to an existing application can be a big challenge, partly due to the changes in data formats and partly due to the need to open up the application. There are certain situations where you might want to make explicit code changes to use SecureData – for example, when a customer wants to extract every ounce of performance. In most cases, an application need not know that its sensitive data has been masked. HPE FPE and HPE SST make it possible to encrypt/decrypt and tokenize/detokenize data without the application being aware of it. comForte's HPE cF Data Security and XYPRO's HPE XYGATE Data Protection products complement SecureData and are able to seamlessly protect Enscribe and SQL/MP data, and work is in progress on SQL/MX.

This issue is chock full of interesting articles about security. You will learn more about HPE SecureData Enterprise and the partner security products that were recently made available through HPE in the article by Prashanth Kamath U, HPE NonStop Product Manager for security. The HPE Data Security team has provided articles on HPE SecureData introducing Hyper FPE and Hyper SSL, as well as articles on how Identity-Based Encryption works and how privacy is looked at both in the US and around the world. Our partners have contributed related articles on how to upgrade your security plan to include encryption and tokenization by Thomas Burg, CTO of comForte and the intersection of tokenization and access control by Andrew Price and Scott Uroff, XYPRO's VP Technology and Chief Architect. Other aspects of security are featured in articles on File Integrity Monitoring by Callum Barclay, CTO of Computer Security Products, Inc., and the ongoing arms race between attackers and defenders by Steve Tcherchian, XYPRO's CISO. Don't miss out on Richard Buckle's latest column, information on BITUG and how to analyze data and metrics for your Big Data environment from Gravic.

Customers have more choices today than ever before on how to deploy NonStop. The NonStop Vision will continue to deliver even more choices with innovative products for the cloud and digital core. We have an opportunity here to drive NonStop into new markets and deliver our unique value proposition to more customers and applications worldwide. It's an exciting time for the HPE NonStop business and for the customers who rely on us. We are glad to have you along on this journey.

Andy Bergholz

Andrew Bergholz

Senior Director of Development of HPE NonStop



Everyone has a moment when business continuity becomes real.

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

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

For more information, please see the Gravic white paper:

Choosing a Business Continuity Solution to Match Your Business Availability Requirements



ShadowbaseSoftware.com

Business Partner



**Hewlett Packard
Enterprise**

Hewlett Packard Enterprise has many years of experience in migrating mission-critical applications from IBM Power Systems running IBM's UNIX operating system AIX to HPE (previously HP) open-standards platforms. As HPE has demonstrated, most of these migrations create significantly less expensive operating environments, often cutting costs by more than 50%. At the same time, the HPE open-standards environments meet or exceed the performance and availability attributes of the original Power Systems.

Compelling Reasons for Migrating from Power Systems to HPE Open Systems

Several pain points are motivating organizations to consider migrating their mission-critical applications from IBM Power Systems to open systems.

Dwindling ISV Support

More ISVs are dropping support for Power Systems. In fact, the Gartner Group predicts that by the year 2020, 70% fewer applications will run on UNIX. As time goes on, vendors may stop supporting their applications on UNIX. They may drop the applications entirely or migrate them to Linux.

IBM claims that you can run Linux applications on Power. However, Power Linux implementations are not binary compatible with mainstream Linux distributions on x86 platforms. Linux applications must be certified by the ISV before they can be run on Power Linux, and it remains to be seen whether ISVs will make this commitment with the narrow market represented by Linux on Power Systems.

High Cost

By offering x86 platforms that deliver the highest levels of uptime, HPE allows customers to maintain their mission-critical service-level agreements (SLAs) with vastly lower costs for software licensing, hardware support, and power consumption.

A significant consideration is the cost of the Oracle database-management system, as many of the applications being migrated use Oracle databases. Oracle charges twice as much per processor core for Power Systems than it does for x86 platforms. Furthermore, Oracle RAC (Real Application Cluster) costs \$11,500 per core on an x86 system compared to \$23,000 on an IBM Power System.

Support for Cloud Computing

Moving workloads to a cloud environment requires defining a virtualized, standardized platform to deploy applications onto a wide range of public cloud-service providers. If workloads remain on Power Systems, the only cloud-deployment option is a high-cost cloud from IBM. Since cloud computing requires common software across all platforms, Power Systems cannot be brought readily into this flexible environment. Tools for publicly available cloud environments such as Microsoft's Hyper-V, ESX from VMware, and Xen from Citrix commonly run on Linux, Windows, or both but not on Power Systems.

HPE's Approach to Migration

Applications must continue to provide uninterrupted services during and after the migration. To aid this fundamental requirement,

workloads for large applications often can be migrated using a phased approach, migrating one at a time, until all target workloads are running on the new platform.

Who Does the Migration?

Is the migration performed by HPE, by the customer, or through cooperation between the two? This decision may be different for each migration. Often, several migrations must be planned and executed; and the mix of participants may vary with each.

Clients often leverage an experienced migration-services vendor to perform all aspects of a migration. This approach can reduce risk and can provide a single point of ownership for the migration project. Migrating from Power Systems to open systems requires knowledge of many application environments, including online transaction processing, batch processing, and enterprise resource planning (ERP), all of which need to be accessed on-demand and cannot be down.

Migration becomes more complicated if the client assigns some of its IT personnel to the project. HPE's approach allows the client to determine the level of staff involvement in the migration.

Managing Application Upgrades During a Migration

In many cases, in addition to migrating from Power Systems to HPE open-standards platforms, clients may want to upgrade their applications. This process can involve changing third-party packages, replacing a custom application with a third-party application, or modifying a custom application.

In general, the application should be migrated first before it is upgraded in order to minimize risk. Unless an application no longer is suitable for production, it is better to leave further changes until after the migration in order to minimize the migration time frame.

In-House Developed Code

A common high-risk area is the migration of in-house-developed code that is older and mission-critical. HPE has developed a portfolio of software tools that automates the migration of in-house code, including C, Java, and scripts.

Clients often will decide to migrate in-house code if they have only a small number of applications or if the applications are not mission-critical. If the customer decides to perform its own migrations from Power Systems to HPE open-standards platforms, it must recompile the applications and run functional and system tests to identify errors and other issues. HPE has found that recompiling Power C code on Linux will identify only about 25% of the needed code changes. The rest are discovered during testing or, even worse, in production.

HPE has developed code-analysis tools that reduce the time and effort of a manual review process. These tools identify the dependencies and required code changes of an application prior to migrating from Power Systems to open systems and can reduce testing time by as much as 70%. The tool set is especially appropriate for analyzing mission-critical code since it reduces the likelihood of bugs showing up in a production environment.

ISV Applications

For packaged applications from vendors such as SAP, Oracle, and dozens of others, the vendor typically offers services to migrate the application to Linux on x86. However, a major application from a vendor typically uses many other

services such as WebSphere, MQ, or SQL databases. The vendor doesn't necessarily offer services for these migrations. All of these services must be migrated, and they must still work with each other afterwards.

HPE can manage the migration of the entire ecosystem - applications, application servers, middleware, databases, and other components. HPE offers robust and comprehensive tools to support the migration of Power Systems DB2 databases to other databases.

HPE Servers with a Spotlight on HPE Integrity Superdome X

The industry-leading HPE ProLiant x86 portfolio delivers comprehensive, versatile compute offerings for datacenter efficiency across diverse workloads and applications. In HPE's mission-critical portfolio, let's take a closer look at the x86-based mission-critical platform, the HPE Integrity Superdome X server.

The x86-based Superdome X supports industry-standard operating environments like Linux and Windows but draws upon decades of HPE's UNIX server experience, delivering levels of availability, processing power, and serviceability typically found only on UNIX platforms with proprietary processors. For example, a Superdome X can be divided into electronically isolated hard partitions called HPE nPars. Each hard partition runs its own copy of the operating system and applications in isolation from the other partitions, making it an ideal environment for migrated workloads.

A recent comparison of the costs of an IBM Power System and an equivalent HPE system shows a TCO (total cost of ownership) savings with Superdome X of 41%, including a 75% reduction in hardware costs, a 38% reduction in software costs, and a 30% reduction in software support costs. In addition, Oracle licensing costs offer substantial savings, as described earlier.

HPE's Methodology for Migration

HPE provides a set of four core migration services to ensure fast, predictable results for most mission-critical systems. They include the following:¹

Transformation Workshop and Platform Advisory Services

- What happens during a UNIX migration?
- What is the process to reduce risk?
- What are the migration options?
- Which is the best platform for the application workloads?

Migration Business-Case Service

- Building a case to migrate for a specific application environment
- Is the migration financially viable?
- Is the migration technically viable?
- Is the timeline valid?
- What are the risks and mitigation strategies?

Migration Design and Planning Service

- How to ensure a successful migration
- Full scoping
- Environment, application, code, and data analysis
- Migration planning and timelines
- Detailed implementation proposal and statement of work

Migration Implementation Service

- Execute the migration plan.
- Migrate applications, code, and data.
- Product replacement
- New infrastructure
- Testing, rollout, and follow-on support

Case Studies

Pharmacy Chain

A chain of pharmacies has separate operations for wholesale distribution and retail functions. It had long relied on SAP for ERP and customer relationship management (CRM) running on IBM DB2 under AIX on Power Systems.

The company moved to SAP HANA and off its legacy databases to reduce costs and to increase scalability and flexibility. It moved to a Superdome X with two HPE nPars, one for the wholesale operations and one for the retail operations. Each has different requirements in terms of concurrent users and database size. Separating the environments made them easier to manage.


Manufacturing

A manufacturing customer migrated its SAP retail system with DB2 to Integrity Superdome X. The company uses one HPE nPar as the production environment for its SAP ERP application and the other nPar for SAP HANA.

RI-Solution

RI-Solution, located in Germany, has deployed two HPE Integrity Superdome X servers running Linux with three HPE nPars per server to deploy SAP applications. This configuration allows RI-Solutions to consolidate and standardize its hardware infrastructure, contain costs, increase availability, simplify business processes, and improve the performance of its mission-critical SAP applications.

Summary

Hewlett Packard Enterprise has over three decades of experience migrating many types of complex workloads for enterprise customers. Through its diverse experience in delivering successful migrations, including IBM Power Systems to open-standard platforms, HPE has learned what it takes to implement a successful migration and to manage the inherent risks. The performance, availability, and scalability of HPE's broad compute portfolio deliver the power and capacity needed to ensure migrated applications meet service-level agreements so that customers can drive increased productivity and business growth. Through industry leadership and innovation, HPE's server portfolio offers a comprehensive array of industry-standard platforms designed to help customers confidently modernize their data centers and take their businesses to the next level. HPE employs proven processes and unique migration tools for IBM Power System migrations to open systems, and it has developed proven approaches to maximize the ability of the target environment to deliver better results for the line of business while reducing costs, often by more than 50%. 

.....

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



Breaches are from Mars

Security is from Venus



Steve Tcherchian, CISSP > CISO > XYPRO Technology

There is quite a large disconnect in the way security breaches are evolving versus how security solutions and resources are keeping up to address them, much like the book from John Gray covering relationships and the different motivations, of men and women. Unlike the book though, we're not trying to come to a happy medium – we're trying to keep the war like Mars at bay. As a security strategist, I'm constantly evaluating what is possible to help identify gaps and opportunities. The one thing I have learned over the course of my career is:

The only thing constant in cybersecurity is that attackers' methods will continue to evolve. They get smarter, more resourceful and are impressively ever patient.

The HPE Integrity NonStop server is not only a foundation of the HPE Server business, it is also central to countless mission-critical environments globally. For the longest time, security of these powerful systems and the "Mission Critical" applications they run remained mostly static and under the radar while high profile attacks on other platforms have taken the spotlight. That hasn't lessened the risk and exposure of the NonStop server. It's actually created a gap. With globalization and introduction of new technologies for the NonStop server, this security gap will only increase if not addressed.

Interestingly enough, the NonStop server isn't the only mission critical enterprise solution in this situation. There are some colorful parallels that can be drawn between applications running on the NonStop server and those running in SAP environments. Both are in highly mission-critical environments and vital to the revenue generation of an organization, and they frequently run payments applications like ACI's BASE24 and other homegrown applications. This creates some interesting security challenges. In a recent The Connection magazine article, Jason Kazarian, Senior Architect at HPE described legacy systems as "complex information systems initially developed well in the past that remain critical to the business in spite of being more difficult or expensive to maintain than modern systems". His article went on to point out the security challenges of legacy applications. In summary some of these types of applications can tend to be unsupported, security patches aren't readily available and if they are, they aren't applied in a timely fashion because of fear of disruption, and they don't have a lot of the security features modern applications would have. This makes detecting and addressing security risk and anomalies a greater challenge than it already is.

MIND THE GAP

How can this problem be addressed? Protect what you can. As a first step, be it system, application or data - push the risk down the stack to an area that is more controllable by typical security. For example, tokenizing data used by a legacy application will send an attacker to go search for that data through another method, preferably one better suited for detection.

Have a risk based, layered approach. This will swing the odds in your favor. OK, maybe not completely in your favor, but this approach will provide you with the arsenal you previously did not have: It will create those choke points, provide the visibility needed and help reduce mean time to detection and response.

With the way threats are evolving, those of us responsible for security need to constantly evaluate and assess our capabilities. Let's take a dive into each layer to explore the benefits they provide in an overall security strategy.

Protect

Protection/prevention is the first and most critical layer of any security framework. Without a proper protection layer in place, none of the other layers can be relied upon. Think of the protection layer as the traditional defensive strategy - "the wall built around assets". This includes defining and implementing a security policy as well as hardening of the network, the system and applications. The protection layer is also where users, roles, access control and audit are set up. Key fundamental concepts to consider as part of the protection layer.

- Authentication – Allows a system to verify that someone is who they claim to be. In a HPE NonStop server environment, this can be done using Safeguard, XYGATE User Authentication, or through application authentication.
- Authorization – Determines what a user can and cannot do on a system. Authorization defines roles and access to resources.
- Access Control – Enforces the required security for a resource or object.
- Logging and Auditing – Ensures that all security events are captured for analysis, reporting and forensics
- Encryption and Tokenization – Secures communication and data both in flight and at rest. Examples of products which protect data include VLE, TLS, SSH, Tokenization and more.
- Vulnerability and Patch Management – Ensure timely installation of all RVUs, SPRs and application updates. Prioritize and take recommended action on HPE Hotstuff notices.
- These types of preventative controls are necessary and intended to prevent unauthorized access to resources and data, but they cannot

solely be relied on as a long term sustainable security strategy. Attackers' motivations and sophistication are changing, therefore when prevention fails, detection should kick in while there is still time to respond and prevent damage.

Detect

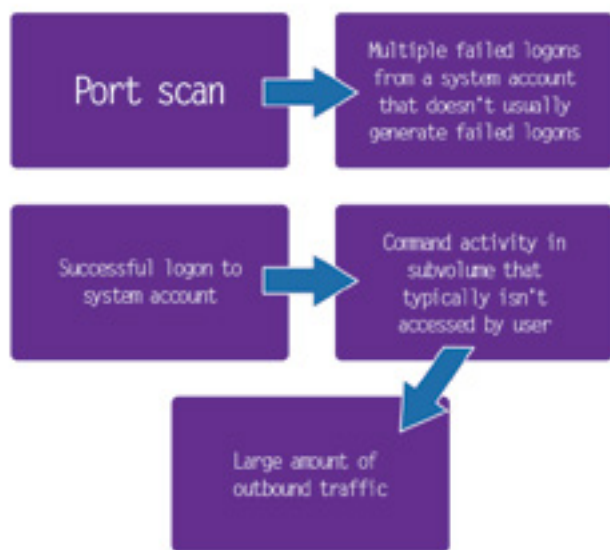
In testimony given before the Senate Subcommittee on Science, Technology and Space, famed cryptographer and cyber security specialist Bruce Schneier said:

"Prevention systems are never perfect. No bank ever says: "Our safe is so good, we don't need an alarm system." No museum ever says: "Our door and window locks are so good, we don't need night watchmen. Detection and response are how we get security in the real world..." "

Schneier gave this testimony back in July of 2001, yet in 2016 where organizations are getting hit by incidents they can't detect, this premise is still valid and critical. In the previous section we discussed hardening systems and building a wall around assets as the first layer of security strategy. I'm surprised by the number of conversations I have with IT and Security folks who still carry the mindset that this degree of protection and compliance is good enough. No matter what level of protection a system has, given enough time, an attacker will find a way through. The faster you can detect, the faster you can respond, preventing or limiting the amount of damage a security breach can cause.

Detection is not a simple task. The traditional method of detection is through setting up distinct rules or thresholds. For example, if a user fails 3 logons in a span of 5 minutes, detect it and send an alert. In most cases that rule is explicit. If the failed logon events spanned 20 minutes, or worse yet, 10 days, it would not be detected. The limitation with relying on rules for detection is they will not alert on what they don't know about. Those low and slow incidents and unknown unknowns – activity not normal on a given system -will fly under the radar and no one would be the wiser until you get a call from the FBI.

The other challenge is correlating events from multiple data sources. Let's look at the incident pattern below.



In this incident pattern, we have events from EMS, Safeguard and XYGATE. The NonStop server could send each individual data



source to an enterprise SIEM, but the SIEM would not have any context to detect the incident pattern as suspicious behavior. A security analyst could create rules to detect the incident pattern, but that's just one use case. The traditional method is to scour through event audit records, try to put the pieces together and then create a rule to detect that pattern in the future. The weakness in that thinking is the incident has already occurred. You're putting a rule together on the off chance it will happen again. However, it's not reasonable or possible to anticipate and define every possible incident pattern before it happens.

A third area of concern is profiling a system and its behavior to understand what is normal behavior for users, applications and the system to be able to recognize when activity is not normal. This can be accomplished through evaluating the system and its configuration, profiling the system over a period of time, profiling user behavior, highlighting risk management and a variety of other intelligence methods. This is where machine learning has a significant advantage. No human could possibly evaluate the volume of data needed to make these types of determinations at the speed required by today's standards. Machine learning is a type of artificial intelligence that enables the system to teach itself. Explicit rules are no longer the lone method of detection. Machine learning can profile a system or network over a given amount of time to determine what is normal to isolate what is not normal. Inserting machine learning as part of a solution process significantly increases abilities to stay on top of what is going on with a given system, user, network or enterprise.

Alert

The third layer relies on alerting. The challenge most environments have as they grow and their infrastructure becomes more chaotic with more tools, more users, more data and more events is that they alert too much or too little. How does one know what to act on and what is just noise? There are solutions that position themselves as being able to do security analytics, but that ends up generating more data from existing data. Now someone needs to determine if the newly formed alert is actionable or just noise.



Going back to our previous failed logon example, if we were to receive 15 different alerts for the same rule, how can one know which alert to pay attention to and which to safely ignore? If you've ever been responsible for responding to security alerts, you know this creates alert fatigue. Back in my early days, mass deleting emails of similar types of alerts was one of my favorite things to do.

Contextualization allows the system itself to determine what is actionable and what is just noise. A solution like XYGATE SecurityOne can evaluate each potential alert and, based on activity that happened previously for that user, IP, system etc..., determine whether the reported activity is business as usual or a serious issue that needs to be paid attention to. Creating new data and new alerts from existing data doesn't solve the problem. Applying context to the new incidents

generated helps focus efforts on those incidents that truly need attention.

Contextualization is key.

Respond

For any of the first three layers to produce value, there needs to be a proper incident response plan. Responding will allow you to deploy your countermeasures, cut off access, send the attacker to a mousetrap or other actions that will assist in minimizing the impacts and recovery of a breach. Containing the breach and quickly recovering from it are the most important steps of this layer. Response and containment comprise of a number of simultaneous activities to assist in minimizing the impact of a breach. These may include but not limited to:

- Disabling accounts
- Blocking IPs and Ports
- Stopping applications or services
- Changing administrator credentials
- Additional firewalling or null routing
- Isolating systems.

This is necessary to slow down or stop an attack as well as the preservation of evidence. Evidence of the attack is generally gathered from audit logs, but coupled with detection and analytics tools can provide access to information in a much quicker and more granular fashion. Being able to preserve evidence is key is forensic investigations of the breach as well as important for prosecution.

Once all the pieces fall into place and there is an incident alert that requires response, how will your organization deal with the issue? Breach incidents are hardly ever the same. There needs to be a level of categorization and prioritization on how to deal with specific incidents. In some cases, you may want to slowly stalk your attacker, where in others, the sledgehammer approach may be the only thing that can preserve data. Does everyone understand their assigned roles and responsibilities? Is there someone in charge? Is there a documented plan? All of these are considerations that need to be accounted for as part of response.

This can be summarized in two words – BE PREPARED.

Resources

On the HPE NonStop server – the protection layer can be addressed with properly configuring Safeguard, implementing protection of data in flight and data at rest and deploying third party security tools available for the system. For alerting and detection, XYGATE Merged Audit with HPE Arcsight can provide the tripwires and alarms necessary for proper detection. For further detail on how to properly protect a NonStop server, HPE has published the HPE NonStop Security Hardening Guide. XYPRO has also published a 10 part blog series on how to properly protect a NonStop server (<http://bit.ly/21nmQiY>).

For the next generation of detection and alerting, XYPRO's newest offering, XYGATE SecurityOne (XS1), bringing risk management and visibility into real time. XS1 correlates data from multiple HPE Integrity NonStop server sources, detects anomalies using intelligence and analytics algorithms to recognize event patterns that are deemed out of the ordinary and suspicious for users, the system and environment. Coupled with HPE ArcSight, the solution can provide a constant, real time and intelligent view of actionable data in a way that was never been seen before.


Strong technology and process is important, but people are paramount to any successful security strategy. Constant security training and development on industry best practices, security trends and attack evolution should be factored into any security program. Without

ongoing training and reinforcement of people, the gap only has an opportunity to widen. An organizations most valuable resource are the people hired to provide security and close the gap. Use them wisely and ensure they have the tools and training to provide the layers of defense required.

En Finale

Cyber criminals don't sit around waiting for solutions to catch up. Security complacency ends up being the Achilles Heel of most organizations. Because of its unique attributes, security on the NonStop server needs to be addressed in a layered approach and Risk Management is a big part of the process. Putting the layers in place to allow us to highlight risk as early as possible to address it is key in dealing with upcoming challenges. This will hopefully help bridge the gap between attacks and security.

We need to recognize the paradigm shift and the change in mindset in how we approach security, and attackers' ability to stay one step ahead of most defenses is central to their strategy. As the NonStop platform evolves and becomes more interconnected, what was put in place previously to address security will not be sustainable going forward. No matter how vendors position their solutions, security is hard, doing the right thing is hard, but that doesn't mean security professionals need to work harder.

From a security professional's perspective, cyber criminals will always be viewed as Mars – warlike. Relentlessly driving to break into systems, get to data, wreak havoc and cause disruption to fulfill their malicious objectives. Meanwhile, cyber security staff need to act more like Venus – clouded in mystery and deliberately avoid being seen while following the enemy. If Mars knows our tactics, Mars can avoid them. Mars is at war. Mars is patient. Mars will continue to attack, low and slow. With the proper security layers in place, Mars will be thwarted by deliberate masking, redirection and detection that hides where the data really is and alerts when the enemy is near. We continue to get smarter by blocking, hiding and redirecting things away in response to attacks. But unlike men and women, Venus in the security world has a goal is to keep Mars at bay forever...or longer... 

Steve Tcherchian, CISSP

Chief Information Security Officer

XYPRO Technology

steve@xypro.com



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



Welcome to BITUG

Matthew Whiteman

Founded in 1984, BITUG is one of the largest and most active Tandem (HPE NonStop) User Groups outside of the USA. We maintain complete independence from HPE (Tandem) whilst continuing a close working relationship, as well as being an affiliated member group of Connect (formerly ITUG) - the International Tandem user group. BITUG is run by its members, for the benefit of its members. Membership of BITUG is free for end users of Tandem (HP NonStop) servers in the British Isles.

Run by a Committee of volunteer BITUG members, BITUG is here to provide the following services to users:

- Co-ordinate user response to developments affecting HPE NonStop users.
- Provide a focus for the representation of British NonStop user interests internationally.
- Convene meetings for members to share NonStop related technical information and NonStop strategic information.
- Hosts special interest groups (SIGs) to provide a forum for specialists in topics of shared interest, for example: systems management, security, contingency, etc.



2014 Big SIG Venue

Institute of Chartered Engineers - One Great George Street

Each year we try to hold at least two one-day special interest group meetings (SIGs) covering various technical issues. We also hold two education days per year. Previous Big and Little SIG events have been held at the following venues:

- Trinity House (opposite The Tower of London)
- HMS Belfast
- Institute of Chartered Engineers (a two minute walk from The Houses of Parliament and Westminster Abbey)
- The Bank of England
- Institute of Directors (on Pall Mall)

The European Tandem User Group event (codenamed eBITUG) will be coming to London over 9th and 10th May 2017. Planning is under way and more information will be released as it becomes available. We look forward to seeing you there.



Big SIG 2014



Big SIG 2016 Passport to Prizes Presentation

For more information, our website is www.bitug.com. Bi-annual newsletters, committee details, Big and Little SIG presentation slides and more can be found there.



SECURITY? A NEVER-ENDING STORY!

Greg Swedosh > Security Consultant > Knightcraft Technology

In today's world we are constantly bombarded with reminders to look at something. The need for check-ups, check-outs and in general, evaluations has turned into an art form. All the latest motor vehicles now tell us as drivers when they need to be brought to the dealers for a service. Our doctors and dentists are quick to inform us that a visit is required to ensure baseline data is updated to reflect our current status. With the changes of seasons, the letter box is filled with reminders to have our windows cleaned and our trees trimmed and then the dreaded "check engine" light appears as we are trying to wrap up a day of shopping.

As a society, whether at our behest or the results of the latest legislation, it seems we just have to endure constant oversight and when it comes to our computers, this too has become the norm. We have always had to endure routine hardware maintenance and set aside time to refresh operating systems and key middleware subsystems but increasingly there's yet another pursuit that is demanding even more of our attention. Security! In a war with parties determined to penetrate our security defenses there's no alternative to checking everything. At least twice!

As a platform that has been depended upon for decades by Financial Institutions (FIs) in support of ATMs and POSs, as well as providing switching capabilities interconnecting retailers, merchants, card issuers, banks and credit unions, etc., NonStop is not immune to needing its security check-up. Increasingly, standards bodies along with government agencies are mandating a level of compliance that requires "return visits to your doctor!" It is recognized that, with the level of sophistication being demonstrated by cybercriminals today, and the much heralded success they have had stealing personal information from the private and public sectors alike, ensuring that your security is reviewed regularly should be a high priority, whether this be by using resources internal to your organization or by bringing in outside expertise. No business or government agency wants to headline the evening news.

When it comes to cybercrime and HPE NonStop systems, hacking a system is about finding a way to access privileged userids such as super.super, or the application owner, which enables a person to then perform functions on the system as that user. Having performed security reviews for a variety of international organizations over the years, I have discovered that a regular and methodical approach to reviews is the only way to ensure that your job is thorough and accurate. There are typically five key steps to perform if you want to review your systems to the required depth.

1. Discuss the aims of the review upfront with management and ensure that expectations are appropriately set.
2. Meet with representatives of all of the various stakeholders of the system, being the system administrators, application support team, developers, database administrators, security folk, operators and so on to get an understanding of what they access the machines for, when they need and how they get access to privileged userids and so on. Don't make assumptions on what you think they need access for. This evolves and changes over time.
3. Collect the relevant information from the systems. This can predominantly be done using non-intrusive methods and no special software. Standard commands from TACL, Safecom, Pathcom and so on, can be run from a non-privileged userid, with the output collected in edit files for later analysis. The only information for which a privileged userid is typically required is to gather information from Safeguard pertaining to userids, aliases and groups.
4. There is a lot of data produced when fully reviewing a system, so a methodical way of analyzing it is required. Having been doing this for some time now, I have built up a suite of macros and utilities that put the data into a format whereby it can be easily analyzed within Excel

and other tools. The more you can automate the better. It is likely that during this analysis phase, more questions will emerge and you will likely need to go back to the various system stakeholders for further information.

5. The final step is preparing some kind of report that should include all of your findings and some recommendations as to how these need to be remediated. There should of course be a section, or even an additional presentation, that gives the key findings and recommendations to the relevant management team.

Often the security on a NonStop system has been set up some time ago, by somebody who is no longer with the institution. In that case, the security tends to keep rolling forward exactly as it is. When customers upgrade their systems, they often just duplicate all existing security settings across from the previous systems. This is usually because nobody is really sure what is set up, what it does and what will happen if they change it. It is also frequently because NonStop shops are typically understaffed when it comes to system management personnel. With tight migration deadlines and a truck load of work, there is often not the resource available to actually make sure that everything that is migrated across actually fits the current requirement and is in line with current best practices.

“Do not be afraid to challenge the security settings that you find and see if there is truly a sound rationale for them being as they are. If something doesn’t appear to make sense to you, it is quite possible that it is no longer a valid setting.”

Don’t point the finger or apportion blame for current settings. Have an awareness that much of what you find will be historical.


When organizations think about security of the NonStop machines, they often think of Safeguard, OSS file security (if relevant) and perhaps encryption. They often forget about the security aspects of the many other areas of the system, that if not set up properly, can allow a non-privileged user to gain super.super powers. This includes subsystems such as Pathway, Netbatch, SCF configurations, the TACL environment, and so on. Every time I approach a security review within a new organization, I wonder if this time will be the time when there are no security configuration vulnerabilities that I can find. So far that has never happened. Typically you will find a number of different ways that users with access to the system could become super.super and then do pretty much anything on the system. This could be something like finding a pathway owned by super.super and secured ‘N’, meaning that any user could come along and add their own server running a program such as SCF, with input and output directed to their paused terminal, so

that when the server is started, it will run a SCF session as super.super, from which a user can run any other program as super.super. Or a sensitive program, such as SQLCI that is owned by super.super and with the ProgID flag set, meaning that any user who executes it is running as super.super. Or something even more basic such as Safeguard secured in a way that would allow unauthorized users to add safeguard records providing them access to sensitive data.

NonStop systems deployed within FIs typically run up against the Payment Card Industry Data Security Standard (PCI DSS). That standard has forced a number of NonStop customers to focus on their security through the prism of compliance. That is, every year they have a PCI assessment of their systems, to establish if they comply with the PCI DSS. Many organizations that are being told they are compliant, feel that this implies that their system is secure. Unfortunately, this is often not the case. The majority of QSAs (Qualified Security Auditors - PCI compliance auditors) who perform the assessments typically know very little about NonStop systems and their security. Mostly they have Unix or Windows backgrounds. When they check for compliance, they work off a checklist and rely on the NonStop system people to provide them with the information. However they often don’t know what information they should actually be receiving and exactly what it should look like. Compounding this is the fact that most organizations approach audits with the ‘don’t tell them more than you have to’ kind of approach. So, at the end of the day, you often have somebody who doesn’t really understand the system, basing their assessment on the minimum amount of information that is provided to them by the customer. This is clearly not a very strong basis for assuming that your system is secure. A number of high profile data breaches of organizations considered to be PCI compliant underlines this point quite emphatically. Compliance does not equal security. You need to address both.

To fully review your systems will typically take four to eight weeks, depending on various factors such as number of systems, number of different types of applications, complexity of environment, etc. It should be approached as a project and resourced appropriately so that the person performing the review doesn’t face continual workload interruptions and can focus exclusively on performing the review. If the person doing this is also responsible for managing the system and responding continually to other work requests, the review will never get done.

To ensure that you feel equipped to handle the task of reviewing your systems, you will need to stay current with any security standards that are relevant to your organization, such as PCI DSS, as well as current best practices for securing your machines. The Security Hardening Guide published by HPE (available in the NonStop Technical Library) is a great starting point for working out the areas that need attention. The PCI DSS Compliance for HPE NonStop Servers technical white paper is available for download on the Knightcraft website for assistance with this standard.

Nobody enjoys taking time for check-ups. However, when it comes to cybercrime and the ingenuity of hackers bent on gaining a financial upper hand and compromising our customer information, there’s no issue about enjoyment. It just has to be done. Regularly checking the security of your systems should not be an afterthought. It should be an annual occurrence. 

Greg Swedosh is the director and senior security consultant of Knightcraft Technology. He is a regular presenter on NonStop security at the NonStop advanced technical boot camp. Knightcraft security and PCI compliance services for the HPE NonStop can be procured directly (www.knightcraft.com), through your HPE account team or through comForte.

NonStop Innovations Deep Dive

Tributary Systems Plans for Continued Growth with New CRO and Product Development

Mandi Nulph >> Marketing Coordinator >> NuWave Technologies

I recently had the opportunity to chat with Tributary Systems Chairman and CEO, Shawn Sabanayagam; VP of Software Development, Glenn Grundstrom; and their new chief revenue officer (CRO), Antonio Rajan. We talked about the company, their new products, and how they intend to position the company for future growth.

Mandi: Before we get started on your new announcements, could you all give those who might not know you some insight into your background in the NonStop space and how you came to work at Tributary?

Shawn: I have managed Tributary Systems on a day-to-day basis since 2004, but have been involved with the company since I acquired it from its founders with a group of investors in August of 1999.

Antonio: Unlike my colleagues, I do not have a background in the NonStop arena. My specialty is to implement sales and marketing strategies for technology companies that cause rapid revenue growth, which is the impact I hope to make at Tributary Systems.

Glenn: In 2008, I was hired at Tributary to run the software development organization. I manage product development and quality assurance, as well as the installation and support teams. We now support all open and most proprietary operating systems and host server platforms, but we obviously have a strong affinity for NonStop as it is the heritage of the company. Tandem was the foundation on which we built our technology and solutions, including our patented flagship software-defined backup solution, Storage Director®.

Mandi: Could you give us some background on the history of the company as well?

Shawn: Sure. Tributary Systems was founded in 1990 as a solution provider for Tandem systems. We sold used Tandem systems, as well as developed new solutions to add value to Tandem customers all around the world. The company evolved very quickly to being a provider of tape storage and data protection solutions for Tandem customers.

In 1996, I believe, the first official Tandem tape drive using digital linear tape (DLT) technology was introduced, and for a brief while, Tributary Systems competed with Tandem themselves. In 1997, when Tandem was acquired by Compaq, the Compaq NonStop Division (Compaq NSD) decided to outsource key development projects, including tape storage. In 1997, Compaq asked us to integrate and help them with the development of tape products that they could sell as their own. The relationship that we've had with Tandem/Compaq/HPE NonStop has been longstanding, and to this day we have an evergreen OEM supply contract with HPE. We are coming up on 19 years of having a continuous working relationship with NonStop. As the market and need for different products evolved, our relationship with HPE did as well, but Tributary's products have always been consistent with the platform's fundamentals of data integrity, scalability and high availability or fault tolerance.

After 2002, the concept of virtualization on backup storage made its way to the NonStop market, and in 2005, we introduced the product that we now call Storage Director. Storage Director is patented backup virtualization software that was conceived and developed long before "software-defined storage" was an IT trend. Storage Director's architecture made it hardware-agnostic, as well as tape-, disk-, NAS- and cloud-compatible. The any-to-any capability makes Storage

Director unique in the marketplace.

Mandi: Tell me a little bit more about Storage Director and how it has evolved.

Shawn: Storage Director was introduced to the NonStop market in 2006 and it has continuously evolved since then. Storage Director 5.0 was released late last year, and the continuous development of it has made Storage Director a leading-edge multi-platform, software-defined data backup, replication and DR solution. We are very excited to announce that it is the industry's first and highest performance cloud backup solution for NonStop customers. Storage Director 5.0 is completely cloud-compatible and eliminates the need for a second data center for disaster recovery (DR). Instead, we are able to offer customers a secure cloud backup, which provides a second copy of their mission-critical data, vaulted to any Amazon S3-compatible cloud architecture. Storage Director safely replicates the encrypted data to the cloud and is able to efficiently restore it while adhering to RPO and RTO objectives the customer specifies at a fraction of the cost of having a separate DR datacenter.

Storage Director was already unique when it was first brought to market. It is the only solution in the NonStop space that is a multi-platform, any host server to any storage medium solution for data backup, and it has been that way from day one. It was never just a NonStop-only solution, and that is very important. Most recently, NonStop introduced their NonStop X server platforms, based on the Intel x86 chip architecture that further opens up the platform and allows customers to have more applications that run on their NonStops. HPE has successfully made the NonStop platform less proprietary, more open and more versatile.

We have focused on mirroring that value proposition for our NonStop customers by offering a backup solution that not only backs up, protects, replicates and manages mission critical data throughout its lifecycle for NonStop data, but we also can back up all other HPE platforms including SuperdomeX, HP-UX, Open VMS and all IBM platforms, including mainframe z/OS, AS/400/IBM i and AIX, as well as all open platforms and operating systems.

Storage Director was conceived as a fairly versatile backup, replication and DR solution with strong data migration capabilities in customers' heterogeneous environments. As we were the first to come to market with a broader-than-NonStop solution, we were also first to market with the cloud solution that allows NonStop customers to securely replicate their data to the cloud.

Mandi: Some people, especially in NonStop, seem to be very wary of cloud backup. What do you think about this?

Shawn: Cloud means different things to different people. Storage Director does not run as an application in the cloud yet. What is available today is a way for customers to eliminate the need for a second data center and implement Storage Director in a private cloud or a hybrid cloud. The hybrid cloud implementation would involve the ability for NonStop customers to vault a second copy of their data securely to any Amazon S3-compatible cloud. Storage Director uses AES 256-bit encryption with checksum to securely replicate data to any S3-compatible cloud, thereby ensuring both security and data integrity.

Data from NonStop applications is not one size fits all. Many solutions do not allow data to be separated into pools and protected in

different locations, on different media, with different retention policies. Storage Director is a policy-based data backup solution that has, from day one, enabled the separation of data into different categories, or pools, working with the host server's backup application and protecting different sets of data with different policies. This increases the efficiency and reduces the cost and footprint of the hardware needed to protect NonStop data.

Mandi: Do you have any new partners or customers that you're excited about working with?

Shawn: We acquired a large NonStop customer late in 2015 when AT&T switched to our solution. That was a big win. We had multiple smaller wins in terms of new customers last year, and we are on track to acquire many new customers and partners this year.

Part of the strategy that Antonio is implementing is to acquire partners that would sell our data backup and protection solutions in both enterprise environments and medium-sized businesses. We have always focused on expanding our business beyond NonStop: several years ago, we established a partnership with IBM, and now we are in multiple IBM environments providing backup solutions to customers who have both HPE NonStop, and IBM mainframe or IBM i/AS400. Now we are going into an even broader market space, opening up the product line with a purpose-built solution that is targeting the SMB market.

Antonio: The strategy is for us to create a value proposition and a buzz around the channel partnership methodology. Salespeople will always want to offer a product that is a) easy for them to sell, and b) highly differentiated. I have asked product engineering to create something that is unique in the market, and to that end, we are going to be releasing a new, all-in-one, affordable appliance this year for the SMB market. It's something that we feel is a unique and, more importantly, carries a price point that no one else can offer.


Mandi: Aside from that, what else are you planning for the future?

Shawn: Antonio is also putting together a business model for backup-as-a-service (BaaS), which will be a very cost-effective, all-in-one backup solution. In addition, customers will also have the option to simply subscribe to a Tributary Systems backup-as-a-service offering. This requires a significant rethinking of our business model, as well as setting up new partnerships. Antonio is in the middle of developing that plan and partnerships so we can start offering that this year, both to enterprise and SMB customers through our partners.

Mandi: That's all very exciting!

Shawn: Yes, it is. HPE NonStop, its customer base and the ecosystem of partners are all going through some significant changes at this time with the introduction of the NonStop X systems. But we are all excited about the future of the platform. Tributary has served the NonStop market for over 26 years; and we continue to invest, reinvent ourselves and bring new and relevant solutions to market for our customers.

Mandi: Thank you for your time.

Shawn: It was our pleasure. 



.....

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



**SPEAK NOW.
BE HEARD.**

Talk. To. Us. Submit your
Advocacy request **HERE!**

<http://ow.ly/4nnJtF>

When Ernie Guerrero founded NuWave Technologies in 1999, he hoped that the company would one day become a family legacy, but he had no way of knowing how that would pan out. Now, as Ernie steps down as the company's president and assumes the role of Chief Revenue Officer (CRO), his daughter, Gabrielle, takes on the leadership role that Ernie has groomed her for over the past seven years.

"This has been the plan for a long while. It has always been my intention for Gabrielle to move into this position," Ernie said of his daughter. "Last year we made a shift in our priorities to focus our resources on completing the LightWave products. Now that we have made significant progress on LightWave Server™ and LightWave Client™, I will focus on managing and supporting our growing sales team, which is a role I am very familiar and comfortable with. This change in structure is timely and extremely positive. As we grow, we need a more resilient structure and we also need to be more formal about making major decisions and separating responsibility and functions. Gabrielle will bring this discipline."

NuWave Technologies was originally founded during the .com era and Y2K explosion as a consulting group focused on providing IT services, including custom development and project management for large corporations and government organizations. Soon after the .com implosion and uneventful Y2K conclusion, Ernie repositioned the firm, which is why the company and its employees have been able to enjoy long-term success. Today, NuWave is still recognized for its consulting services, including HPE NonStop modernization, migration, and custom development; but it is even more well-known for providing exceptional middleware solutions for customers in the NonStop space.

NuWave's flagship products, SOAPam® Server and SOAPam® Client, were launched in the early 2000s, and at that point, the company became known for connecting the NonStop server to other platforms, applications, and Web services. The company's latest solution, LightWave Server™, uses JSON messages and RESTful APIs to send NonStop data to modern clients like mobile apps and browser-based applications.

As the company began to grow, Ernie began to recruit an exceptional team to support it. What better place to start than with your own family? Gabrielle joined NuWave as the company's marketing coordinator after she completed her bachelor's degree in business administration from Boston University in 2008. Over the years, she gained valuable insight into the industry, not only from tapping into her father's decades of experience in the HPE NonStop space, but also from taking a hands-on approach to figuring out how to market to the NonStop community. Over the years, Gabrielle spoke to and surveyed dozens of NonStop users to learn about their needs and wants. She also started a NonStop-focused blog called NonStop Innovations and interviewed numerous vendors and service providers in the space to better understand the technology available to the community.

In the past several years at NuWave, she had the opportunity to learn about many different aspects of the business by performing tasks related to administration, sales, marketing, and business development. In preparation for her transition into the role of CEO, Gabrielle has also begun a master's of business administration (MBA) program at world-renowned Babson College in Massachusetts.



Gabrielle Guerrero takes on the role of CEO at NuWave Technologies


"My father has laid the foundation for NuWave over the past 17 years," said Gabrielle. "After working in the NonStop space for over 10 years in development, project management, and sales roles, he started NuWave and built the company from the ground up. Now it is my job to take the company to the next level. The hardest part will be filling his shoes, since he has such a magnetic personality. It seems like most people in the industry not only know him, but call him a friend—he makes running a company look easy. Fortunately, he'll still be attending some of the larger trade shows, so I won't have to hear 'Where's Ernie?' all the time. People always seek him out at conferences because he is so easy-going and positive."

For the last 16 years, Ernie has been in charge of sales, pricing, customer relations and revenue management. As CRO, he'll be able to focus more on these specific aspects of the organization. He will fill NuWave's increasing need for additional leadership in customer relations, and sales. He will also regularly measure and analyze productivity and effectiveness, with a goal of continually improving.

"With the company in a state of growth, we need someone to make sure all of our revenue-generating departments stay on track," Gabrielle explained. "My father has the technical NonStop experience, project management expertise, and knowledge of customer needs to be able to provide guidance in these areas."


The change in leadership could not come at a more pivotal time for NuWave Technologies. With new trends and technology advancements coming to the forefront of the NonStop sector, and with NuWave growing in headcount and launching two new products this year, the company is primed to make great strides in the space. Gabrielle's vision for NuWave is to be an innovator in the NonStop space.

"My goal is to focus the company solely on NonStop, and to launch innovative, intuitive products quickly. In the past, we've tried to do too many things, and now, as a growing company, it's critical that we hone in on what we're best at and start being first-to-market. We have exceptional talent, great products, and excellent relationships with all of our NonStop customers, so we're going to leverage those assets as we move forward."

With Gabrielle's vision for how to propel NuWave Technologies forward, the company can only continue its forward progress in the HPE NonStop space. 



Ernie Guerrero, founder and owner, now fulfills NuWave's need for a chief revenue officer



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

Everyone is lining up to get LightWave Server



Don't get left behind.

LightWave Server™ uses JSON and REST technology to send data to modern clients like mobile apps and browser-based applications that run on virtually any platform.



Learn more at
www.nuwavetech.com/lwsconnection

NuWave
TECHNOLOGIES



Privacy Throughout the World

Tim Roake

Senior Software Engineer at HPE Security – Data Security

Imagine you took a great photo of your family that wasn't quite perfect, so you send it to a friend who is expert at touching up images. The friend does a great job, however you later find this person is displaying the image in a portfolio of work. Your privacy seems to have been invaded, but what law could possibly protect you? To further complicate the issue, you live in Milwaukee and your friend lives in London.

Fortunately English common law comes to the rescue, determining that there was an intrinsic confidentiality agreement between the parties that had been broken. In America the privacy tort law of misappropriation could have been used, but what protections exist in other countries? Here we will examine the growth and differences in privacy laws, and the efforts to create a more common privacy framework throughout the world.

Privacy Law through Organic Growth

The old expression “an Englishman's home is his castle” describes a fundamental right of privacy. This was first cited in the late 16th century by a number of authors and was enacted into English law in 1628 as part of the “The Institutes of the Lawes of England” treatises written by Sir Edward Coke.

This common law enactment was exported to United States where the rules were explicitly declared in the Fourth Amendment's search and seizure protections, which interestingly were influenced by the perceived intrusions of the British when investigating purported tax evasion by the colonists.

But the word “privacy” never found its way into the US Constitution and Amendments. It was left to the courts to determine that the constitution protected two types of privacy: the freedom from government intrusion into a person's home, property and self (the Fourth Amendment), and the right to make decisions privately without government interference (First, Third, Fourth and Fifth Amendments). Personal information privacy is a third type, but courts have never determined that the constitution guarantees this. The seminal 1890 paper “The Right To Privacy” by Warren and Brandeis was the first to examine the meaning of privacy and named the increasing intrusions of technology such as the instantaneous photograph and the newspaper into one's privacy. Their prescience was quite remarkable as they believed the time had come for a common law right to privacy, particularly with regards to new technology and an ever changing society.

From this came the four privacy tort laws still valid in America today:

1. Intrusion into one's private life and affairs (intrusion).
2. Public disclosure of embarrassing private facts (private facts).
3. Unwanted publicity of private individuals (false light).
4. Misappropriation of a name or likeness for financial advantage (appropriation).

Warren and Brandeis cogently formulated their findings to a large degree based upon the confidentiality case from the English Courts of Prince Albert v. Strange. This case centered on Strange's publishing of private etchings made by the royal family, and was found to have breached confidence such that the author had the right to protect the etchings for his “private use and pleasure”. Subsequent to the publishing of “The Rights to Privacy” the English Courts repeatedly considered adopting American privacy tort laws but have never done so, even though the torts were derived from one of their own court cases. Instead they developed different understandings of confidentiality and privacy from the very same case.

The four American privacy tort laws established in 1960 reduced the significance of confidentiality and focused on intellectual property rights, whereas English law made confidentiality central to their laws of privacy

based on Prince Albert v. Strange. English law protected personal and commercial information under laws of confidentiality and nondisclosure through trust and reliance in relationships; American information law relied on rules of nondisclosure to protect the “inviolate personalities” as defined by Warren and Brandeis, and has to date remained defined by the privacy tort laws, with the addition of some sectoral laws such as the Health Insurance Portability and Accountability Act (HIPAA), Fair and Accurate Credit Transaction Act (FACTA), the Children's Online Privacy Protection Act (COPPA) and the currently proposed Email Privacy Act.

Privacy through International Guidelines

In contrast to the individual and somewhat haphazard American laws, international privacy laws have developed largely through enactments of international agreements developed by the UN, the Organization for Economic Co-operation and Development (OECD), the European Court of Human Rights, the Council of Europe and the Asia-Pacific Economic Cooperation (APEC).

The OECD initially created seven principles for the protection of personal data: notice, purpose, consent, security, disclosure, access and accountability. Later the OECD established 11 key guidelines in 1990 and 2013 for the protection of personal information:

1. Data collection should be lawful and with consent (**Consent**).
2. Data must be relevant, accurate, complete and kept up-to-date (**Accuracy**).
3. The purpose for data collection must be clearly stated (**Purpose**).
4. Personal data can only be used for purpose except by consent or by authority of law.
5. Personal data must be kept secure from unauthorized access, modification, destruction or disclosure.
6. Data collection policies, purpose and controlling entity should be available.
7. An individual has the right to find their data and rectify errors (**Access**).
8. The data controller should be accountable to these principles.
9. Government at the highest levels should implement national data privacy strategies.
10. Organizations must have available core operational mechanisms for privacy protection.
11. Notification of data security breaches must be made to the affected individuals and to an authority.

Consent, accuracy, purpose and access principles are the most commonly enacted principles worldwide. The Asia Pacific Economic Co-operation (APEC) also created principles based on OECD guidelines.

Development of Regional Privacy Laws Europe

The growth of privacy in most nations often led to confusion between the fundamentals rights of privacy for the individual versus the evolving privacy rights attached to personal information (PI), as occurred when The European Union began formalizing privacy rights in the 1953 European Convention of Human Rights (ECHR), Article 8 stating “Everyone has the right to respect for his private and family life, his home and his correspondence”. All Council of Europe member states subsequently ratified the convention.

Some clarity was created with the initial seven OECD personal information privacy guidelines established in 1980 but they were nonbinding and European data protection laws remained varied leading to impediments in transferring personal information data between member states. To overcome this the Data Protection Directive was created by the European Parliament and Council in 1995. This codified the principles of data privacy and data transfer within Europe, and established minimal requirements for personal data information transfers outside Europe.

Very few countries have met the requirements (Canada, New Zealand, Israel and Argentina being notable exceptions). America, relying on tort laws and a few sectoral laws, clearly did not meet requirements so the Safe Harbor Principles were set up whereby complying U.S. organizations were able to transfer personal information data with Europe. Currently a new set of Safe Harbor Principles are being developed due to a ruling from the European Court of Justice that existing guidelines are invalid [ECJ C-362/14].

Further strengthening and unification of data protection of personal information within Europe is occurring under the General Data Protection Regulation (GDPR). This regulation was adopted in April 2016 and will fully take effect by 25th May 2018. As this is a regulation by the European Commission member nations will automatically enact enabling legislation.

But it must be remembered that Europe contains many nations and cultures each with their own history and understandings of privacy. It was largely the aftermath of the Second World War that led to the current unified and codified privacy laws. However individual nations still enact their own laws as exemplified by France demanding Google apply the legal decision of the “Right to be Forgotten Ruling” [ECJ 2014] across ALL their domains, and Germany forcing Google Maps to provide street view opt-out capability.

Canada

Canada enacted the Charter of Rights and Freedoms in 1982 with no mention of privacy but it did provide rights against unreasonable search and seizure. This was tested in 1984 *Hunter v. Southam Inc.* and found to only protect against violations of reasonable expectations of privacy. In *R. v. Vu*, 2013, it was determined that computers were to be treated as a separate place and required a specific warrant.

The Canadian Privacy Act (1985) specifically limited collection, use and disclosure of personal information by government organizations with a Privacy Commissioner having power to receive and investigate complaints and make findings and recommendations.

Canada enacted the Personal Information Protection and Electronic Documents Act (PIPEDA) in 2001 based upon the OECD Privacy Guidelines and the Canadian Standards Association (CSA) Model Code for Protection of Personal Information. The EU Commission found PIPEDA provided an adequate level of protection in December 2001. This regulated all private sector entities that acquired personal information from commercial activity. The main principles are:

1. Consent.
2. Purpose.
3. Accuracy.
4. Access.
5. Data collection and use policies must be transparent and understandable.

Provincial privacy laws can displace PIPEDA if the laws are considered “substantially similar” or superior. This determination is made by the national privacy commissioner.

The Canadian Anti-Spam Law, CASL, was established in 2010 and provides the following controls:

1. Prohibits electronic spam and installation of spam/intrusive programs.
2. Prohibits false/misleading email and unauthorized collection of personal information.
3. Regulates spyware, botnets and malware.
4. Necessitates sender identification and unsubscribe capabilities.
5. Consent via opt-in mechanisms (stricter than U.S. CAN-SPAM act).

Canadian courts have also been investigating invasion of privacy issues for over 100 years and many cases suggested a need for privacy rights, most recently because of gaps in statutory frameworks (PIPEDA does not cover intrusions by individuals). In 2012 the Ontario court accepted a tort of intrusion upon seclusion did exist in Canadian law [Jones v. Tsige, 2012 ONCA 32].

Mexico

Article 16 of the Mexican Constitution guarantees the right not to be disturbed in person, home or documents without written order by an authority, and from 1996 the constitution also contains the explicit guarantee of privacy of private communications, however it does not have a general information privacy statute.

In 2010 the Federal Data Protection Act was enacted. This was an omnibus law that incorporates the 1995 EU Data Protection Directive the APEC Privacy Framework [Mexico Data Protection] with the following main principles

1. Accuracy.
2. Purpose.
3. Consent.
4. Data must be lawfully collected, processed and disclosed.
5. A data controller must be designated and must provide privacy notices to the person for whom data is collected from.
6. The information is gathered under “habeas data”.

The “habeas data” concept is a judicial measure giving the individual referred to by the information legal ownership of his/her data. They must be told the content and purpose of data holding pertaining to public records.

Argentina

The Argentine Constitution contains an article protecting the home, correspondences and private papers from unreasonable search and seizure, and in 1994 added the “habeas data” cause of action.

In 2000 Argentina adopted the Law for Protection of Personal Data (LPPD) based on the European Union Data Protection Directive. Its main principles are:

1. A controlling body will enforce the regulations and establishes sanctions for violations.
2. Accuracy
3. Consent.
4. Purpose.
5. Access.

In 2003 the EU determined Argentina was EU Data Directive compliant and trans-border flow of information was permitted.

Brazil

The Brazilian Constitution Article 5 explicitly protects privacy, and also contains a constitutional right to habeas data. However as with Mexico it has no omnibus information privacy law.

There is consumer protection (1990) regulating personal data record keeping and federal law (1996) regulating wiretapping.

The proposed Data Protection Bill (2011) has only reached draft status in 2015 and is not yet enacted.

Africa

Surprisingly a number of African countries now include data privacy statutes in their laws. This includes Angola, Benin, Burkina Faso, Cape Verde, Cote D'Ivoire, Gabon, Ghana, Madagascar, Mali, Mauritius, Morocco, Senegal, Seychelles, South Africa and Tunisia.

All countries have data protection authority registration requirements and all but Ghana have cross-border data transfer limitations. **Only Ghana and South Africa have data breach notification requirements.** [Africa Privacy]

South Africa has the most complete information privacy law. The 1996 Constitution includes a right to privacy and protections against unreasonable searches and seizures of person, home and communications.

The Protection of Personal Information (POPI) Act was approved in 2013. POPI is largely based on European legislation and principles set out as:

1. Consent.
2. Purpose.
3. Accuracy.
4. Access.
5. The Responsible Party is accountable for ensuring compliance.
6. The responsible party must inform the Data Subject and the Information Regulator before processing the data.
7. The responsible party must ensure security and integrity of the data.

The POPI Act, although approved, has not yet commenced, and still requires proclamation by the president. It is expected to take effect late 2016 with business having to comply within one year of that date, however the strict consent rules are seen as placing an onerous burden on their many small businesses.

Middle East

There is a perception that privacy within the Middle East is not well developed and that the word “privacy” has no equivalent in Arabic, with the closest meaning being “to be alone” or “loneliness” [Jacqueline Klosek, *The War on Privacy* 62 (2007)]. However the intrinsic right to privacy is certainly provided in the constitutions of Jordan, Saudi Arabia and Egypt in terms of the privacy of the home.

Dubai

Dubai enacted the Electronic Transactions and Commerce Law in 2002 to restrict ISPs from disclosing customer data, and in 2004 implemented a Data Protection Law and improved it in 2007 to include creation of an Independent Office of Commissioner of Data Protection to follow the Data Protection Directive of the European Commission, however the EU Data Protection Act does not consider it as having adequate levels of data protection currently.

Israel

Israel has no constitution but its Basic Law, Article 7 states the following rights to privacy:

1. Everyone has the rights to privacy and intimacy.
2. Consent must be obtained before entry into person's private premises.
3. There can be no searches of a private premises, a person or personal effects.
4. Writings, records and conversations of a person are confidential.

Israel provides the strongest privacy and data protection laws in the Middle East and meets the EU Data Protection Act in terms of protection, thus permitting cross-border transfer of personal information with Europe. Its Protection of Privacy Act (PPA) 1981, updated 2014 applies to public and private entities. It mandates registration of PI databases and consent, accuracy, purpose and access.

Further to the PPA there are infringements of privacy that would please Warren and Brandeis in the case of:

1. Spying/trailing a person.
2. Listening in when prohibited by law.
3. Photographing a person in a private domain.
4. Publishing a photograph that could cause humiliation or contempt.
5. Publishing an identifiable photograph of an injured person.
6. Publishing an identifiable photograph of a deceased person.
7. Using without permission the contents of a letter or other writing that is not intended for publication.
8. Using a person's name, title, picture or voice for profit.
9. Infringing a duty of secrecy or confidentiality.
10. Using or passing a person's private information other than for its original purpose.
11. Publication of a person's personal affairs from his or her private domain.

These infringements of privacy can be seen as a superset of infringements covered by the four American Privacy Torts.

Asia/Pacific Australia

Australia's Constitution has no expression of protecting the right to privacy, however it inherited the tort laws from England regarding privacy based on confidentiality. It inverts the search and seizure powers by determining when such action does not break trespass laws.

Australia enacted the Privacy Act in 1988 and amended it in 2001 to include the private sector as well as the public sector. It had 11 Information National Privacy Principles (NPPs) based on the OECD Privacy Guidelines, however to reduce the burden of implementation and compliance private companies were able to substitute their own policies if they were considered adequate. Small businesses were totally exempted from the NPPs.

In 2012 amendments to the privacy act created 13 new Australian Privacy Principles (APPs), and specifically gives more power to the Office of the Australian Information Commissioner to enforce regulations and also includes new credit reporting principles. It does not however conform to the APEC Cross Border Privacy Enforcement principles.

Japan

Constitution article 13 states a “right to life, liberty and pursuit of happiness”. This was confirmed as a right to privacy by the Supreme Court in 1963.

The Personal Data Protection Act (PDPA) came into effect in 2005. It includes all businesses but with the limitation that it only applies to databases with sizes over 5000 individuals. This reduces the burden for a substantial amount of small enterprises. They also excluded media, religious, political and scholarly organizations.

The resultant omnibus privacy law principles are:

1. Clearly defined purpose provided to the individual.
2. Consent.
3. Access.
4. Certain “competent” government ministers are responsible for supervision and enforcement.

There are no formal mechanism for complaint against breaches by businesses, but rather places responsibility on the business to use best efforts to resolve the complaints. It is the “competent minister's” discretion as to what actions will be taken against such businesses.

China

China has no omnibus privacy laws but its constitution provides protections for privacy, defamation, intrusion into the home and monitoring of communications. However government security officers and prosecutors can issue their own search warrants without judicial oversight.

In January 2013 a wide-ranging set of guidelines came into effect but they are not legally binding.

Hong Kong

Hong Kong's laws are in most part similar to the laws of England (including privacy torts), and a precondition of the 1997 handover to China was that these laws remain in place for 50 years. The enactment of the Personal Data Ordinance Law in 1996 has provided Hong Kong with some of the most highly regarded privacy laws in Asia. The Ordinance was revised in 2013 and contains the following main points:

1. There must be prior consent from data subjects before personal data is used in targeted marketing and there must be opt-out rights.
2. Data subjects have the right to confirm personal data is being held.
3. Data subjects can obtain their personal data and have it corrected.
4. Data subjects can complain the the Privacy Commissioner about suspected breaches and claim compensation.

Singapore

Singapore is well known for its less than perfect privacy policies. Their constitution does not recognize a right to privacy, and the laws on search and seizure in many cases permit searches without warrant, such as retrieval of text messages, emails and web surfing history. And there is evidence of a wide level of government monitoring of the internet [The Internet and Political Control in Singapore, Gary Rodan SingaporeInternet]. But to a

large extent the populace of Singapore does not object to these measures.

The Personal Data Protection Act was passed in 2012 and enacts the basic structure of the EU Data Protection Directive. It only covers the private sector. Main points are:

1. Cross-border data transfers are only permitted to countries with adequate safeguards.
2. There is no provision for safeguard of sensitive data.
3. Does not apply to business data.
4. Organizations must designate a data protection officer.
5. Organizations are not required to register its use of personal information with the data protection authority.

South Korea

South Korea's constitution guarantees:

1. Liberty and secrecy with respect to one's private life.
2. Secrecy of communication.
3. Control over one's personal information and its dissemination.
4. Restrictions on search and seizure.

South Korea enacted the Personal Information Protection Act in 2011 and it came into effect in 2012 in terms of omnibus laws and sectoral laws. The statute applies to private and public sectors and has a Data Protection Commission requiring government and businesses to provide Privacy Compliance officers. It is considered to be the strictest of all privacy laws in Asia and follows the EU model closely.

However the reality is that the government often invokes Constitution Article 37 allowing them to dramatically reduce privacy rights due to "security or public law and order issues", seemingly for political purposes. It has been noted that the South Korean people are so used to this that it is **considered the norm** [SK].

India

India's Constitution Article 21 protects a "right to liberty" and the Supreme Court in 1964 concluded this meant a right to privacy.

India also has sectoral privacy protection for telecommunications and financial transactions, and it enacted the Information Technology Rules in 2011 for the protection of data privacy which provided the following main requirements:

1. Organizations must establish a privacy policy.
2. They must identify any sensitive data collected.
3. They must provide data security.
4. Purpose.
5. Consent.
6. Access.

The rules only apply to Indians, and do not apply to the many call centers servicing overseas customers.

Russia

From a western viewpoint Russia has never been seen as a champion of privacy and fundamental rights, and is still viewed as exercising overt control and surveillance over its populace. It has however attempted to install rights through its constitution. This was formulated in 1993 and establishes:

1. The right to privacy of person and family, and protection honor and name.
2. The right to privacy of correspondence and communications except by order of court.
3. PI gathering, storage or use is forbidden without consent.
4. The home shall be inviolable except by order of a court.
5. Everyone shall have the right to freedom and personal inviolability.

The constitution was largely based on Mikhail Speransky's constitutional project and the French constitution, and in terms of privacy benefits from being relatively contemporary - the modern concepts of privacy are tightly integrated within its framework.

In 2003 the Communications Law was enacted and provided protections against interceptions of electronic communications and The Russian Federal Law on Personal Data was enacted in 2006.

Its omnibus laws contain a number of substantive principles including:

1. Consent.
2. Purpose.
3. Accuracy.
4. Access.
5. Consent can be revoked at any time.
6. In general, processing of sensitive data is prohibited.
7. Personal data must be kept secure.


In 2011 the Cross Border Data Transfer Rules were enacted:

1. Permits transfer to countries following Europe's 1981 convention.
2. Created a list of other approved countries.
3. Permits transfers only on consent by the individual.

In September 2015 a data localization law became effective. This required all organizations to keep copies of all electronic communications for a period of six months at data storage sites within the Russian Federation. In contrast, the European Union struck down a comparable data retention policy in 2014.

Summary

The development of data privacy laws has resulted in a distinct bifurcation of the treatment of personal information from the traditional concerns of privacy throughout much of the world.

The rapid codification of these laws has occurred as a result of the rapid change in information technology and the associated increase in storage and transfer of private data. It remains to be seen if the laws in Europe versus those developed in other nations will ever achieve a reasonable level of compatibility, nor is it known if the costs to businesses of implementing the requirements are even economically feasible. With this in mind it is quite possible that America's organically developed and sectoral privacy laws may in the end provide the most efficient framework if new safe harbor agreements can be formulated that do not impose onerous burdens on business or government organizations. As with most costs in life, it'll only work if the price is right. 

Timothy Roake is a Senior Software Engineer for Hewlett Packard Enterprise - Data Security, and is responsible for providing multi-platform APIs for the HPE FPE technology. His interests include the application of FPE to encoded data series, cryptographic sponge functions and the history of privacy. He has worked in the data communications and telecommunications industries including Skype and Microsoft, military and aerospace industries as well as startups in the mobile domain.

Did You Know?

Built for today, extensible through standards. As one of the founders of the OASIS technical committee for the Key Management Interoperability Protocol (KMIP) with the best server available for interoperability, HPE ESKM can extend your encryption key management support beyond current key management operations to include a wide range of new applications planned for the future. In this way, ESKM improves on total cost of ownership over time by increasing re-use value through a single, central system to maintain that enforces security policy and auditing controls across a wide range of IT solutions.

www.hpe.com/software/datasecurity



New Solutions on NonStop to Protect Your Sensitive Data

Prashanth Kamath > Sr. Product Manager > HPE Mission Critical Servers – NonStop

For the past few decades, HPE Integrity NonStop systems have been the preferred mission critical computing platform in industries such as financial payments. The key reason for this is the mission critical computing features that NonStop systems offer which very few other platforms can match. With a Massively Parallel Processing (MPP) architecture, that provides unmatched scalability, and an integrated software stack that builds high availability right up to the application layer, NonStop offers unique benefits that has made it a platform of choice for the industry.

The IT systems supporting enterprises in industries such as payments, retail and healthcare contain and manage very sensitive information such as customers' credit card numbers and personal information such as medical history. A database containing millions of these customer records is an attractive target for malicious hackers, who try every possible means to steal the data and monetize it in various possible (and often creative) ways. Their success comes at an enormous cost if the target organization becomes a victim of such an attack. Such organizations end up paying dearly in terms of regulatory fines, lost business, loss of reputation, customer compensation cost and so forth.

Various industry and government regulations have been in place and new ones are in the works which aim at protecting consumers from such breaches and guiding the industry towards implementing solutions and practices which mitigate these risks. The Payment Card Industry Data Security Standard (PCI DSS) and General Data Protection Regulations (GDPR) are examples, but there are many others in different geographies and industries. The regulations are fairly comprehensive and cover all aspects of protecting sensitive data. It's rather difficult to explain these regulations in a few lines but the overall philosophy is to:

- Protect the computing environment from external and internal attacks
- Protect the data throughout its life (at creation, as it traverses from one node to another, as it's processed, and on media - whether live or archived)
- Implement strong authentication and access control measures
- Ensure adequate logging mechanisms to enable forensic analysis in case of a breach
- Document and enforce practices and policies; educate employees

Given the importance of data protection, NonStop has been offering security solutions for many years in order to enable you to meet these stringent security requirements using standards-based cryptography. Beginning in early 2016, the HPE NonStop Enterprise Division (NED) launched several new products with an aim to provide additional modern, data protection solutions to NonStop customers. These are essentially two product suites:

- HPE SecureData Enterprise
- HPE SecureData companion products for NonStop

HPE SecureData Enterprise

HPE SecureData Enterprise is a unique, end to end data protection platform used by enterprises in a variety of environments. This product is offered by HPE Software group's Data Security business unit formed with the merger of the erstwhile Atalla products group and Voltage Inc., which was acquired by HPE in early 2015. The NonStop Enterprise Division

(NED) and the HPE Data Security group have teamed up to offer HPE SecureData Enterprise to NonStop customers.

In traditional data protection methods, customers employ different techniques for the different environments that the data passes through. Examples are user access control, file encryption, TLS or SSH protocols for data in transit, disk/volume encryption for secondary storage etc. Each of these may involve separate cryptography, hand-shake, key protection etc. and the data may be in the clear while in between the stages. Overall, this offers a piecemeal approach and not the best protection for your data.

HPE SecureData Enterprise approaches this topic from the perspective of data-centric security that comprehensively protects the sensitive data in an enterprise. Using a data-centric approach to security, the sensitive data is protected right where it is created, as it traverses through the network, while it's processed/stored in different nodes, used in analytics and when it is archived. At all stages, the data is in encrypted or tokenized form such that, even if there is a successful breach, the data is unusable by the cyber-attacker. Figure 1 below compares the data-centric security offered by HPE SecureData to the traditional methods of data protection.

Three key technology elements are at the core of this solution: HPE Format Preserving Encryption (FPE), HPE Secure Stateless Tokenization (SST), and HPE Secure Stateless Key Management.

HPE FPE is a technology used to encrypt data without changing its original format. While it provides the same encryption strength as the

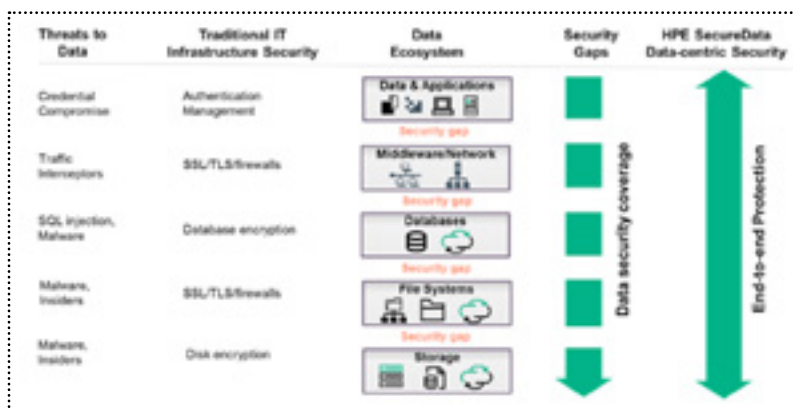


Figure 1: Data centric security versus traditional IT Security

traditional encryption technologies do, the key advantage of FPE is that, because it preserves the data format through the encryption process, the database which stores the data or the applications which process it do not need to be modified, and the majority of applications and processes operate on the data in its protected form—no decryption necessary for use. This drastically brings down the cost and complexity of transforming an existing solution from an unprotected to the protected form, and reduces the exposure of sensitive data to attack. Figure 2 on page 25 compares the results of traditional AES encryption and FPE of plaintext data.

HPE SST is a related technology available in HPE SecureData, and is used to protect sensitive data elements in a file or a database by replacing them with tokens. HPE SST is recommended for use with Primary Account Numbers (PANs) used in payment cards. In this solution, a token table consisting of a static, pre-generated table of random numbers, created using a FIPS-validated random number generator, resides on the platform. A PAN, the data to be tokenized, uniquely maps to a token in that table but has no relationship to it. That token is stored in the system (in files

 Tax ID 934-72-2356		 First Name: Gunther Last Name: Robertson SSN: 934-72-2356 DOB: 20-07-1966	
FPE	345-753-5772	First Name: Uyewjqp Last Name: Muwruwbp SSN: 253- 67 - 2354 DOB: 18-06-1972	
AES	8juYE%Uks&dDFa2345*WFLERG	lJa&3k24kQotugDF2390*32 00WicHk2[*872wsW 0luqwtuwwr*foU0w1@	

Figure 2: Comparison of outputs from a sample data encrypted using FPE and AES

and databases) in place of the PAN (plain data) which is now said to be "tokenized". Only trusted applications are allowed to detokenize and derive the original PAN. Hence, in contrast to the traditional tokenization technologies where a separate "token vault" is maintained, the HPE SecureData solution has no token vault and hence no cost or management complexities associated with it, and no database or vault to be targeted for cyber-attack. Because the system does not store plaintext data in any form, it is outside the scope of PCI audit, thus greatly reducing the costs of PCI compliance. Moreover, token vaults grow in size with the amount of customer data maintained in them, which adds to the management complexities and challenges in scaling and application performance. An SST-based solution, in contrast, does not have these challenges and hence is highly scalable, typically yielding a strong ROI.

HPE Secure Stateless Key Management simplifies the key management

environment such as:

- Modifying the application code to invoke APIs to encrypt or tokenize the sensitive data elements as they are written to the database and decrypt/detokenize them as they are read back
- Optimizing the implementation for the scale-out architecture of the platform (e.g. offload encryption/tokenization to server classes)
- Adding special handling for non-native code (SecureData libraries are native only)
- Integrating with a log management solution such as HPE ArcSight.

While this is all doable and gives full control to the application, it can be complex, time and resource intensive depending on the nature and spread of the overall solution. Moreover, many customers use third party ISV products and modifying that source code may not be a feasible option.

comForte 21 GmbH and XYPRO Technology Corp., NED's software

partners who have been popular among the NonStop customer base for security products for many years, each have developed solutions for the NonStop platform that protect sensitive data using SecureData without modifying the application sources. These vendors have closely partnered with the HPE Data Security group to develop these products, which complement SecureData and greatly simplify its implementation on NonStop. The products are HPE NonStop cF Data Security (from comForte) and HPE NonStop XYGATE Data Protection (from XYPRO). These products have many commonalities in their feature sets but also some significant differences. They also differ in the way the products are structured.

 Credit Card 1234 5678 8765 4321		 Tax ID 934-72-2356	
SST	8736 5533 4678 9453	347-98-8309	
Partial SST	1234 5633 4678 4321	347-98-2356	
Obvious SST	1234 56AZ UYTZ 4321	AZS-UX-2356	

Figure 3: Sample PAN and its tokenized forms

needs of the HPE FPE solution through another landmark innovation. It securely derives the key on-the-fly thereby greatly reducing the cost and complexities of key management. It can authenticate key requests using industry-standard identity and access management infrastructure such as LDAP or Active Directory.

HPE SecureData Enterprise provides a comprehensive data security solution across a cross-section of computing platforms commonly used in the industry. It is supported on traditional *NIX environments, IBM mainframe, HPE NonStop, open systems, cloud, mobile and Big Data environments such as Vertica, Hadoop, and Teradata. This breadth of support enables you to standardize on a single solution to address the data protection needs across your enterprise, which may use one or more of these platforms. Apart from protecting the data, HPE SecureData gives you immense benefits in terms of managing the cost and complexity of the solution and eliminating security weaknesses in the enterprise.

HPE SecureData companion products on NonStop

In the earlier section, you read about how HPE SecureData offers data protection solution to enterprises through a data-centric security approach. Customers can implement this solution on NonStop using the software suite offered by HPE SecureData. It consists of two clients, namely SecureData Simple API and SecureData host SDK. Implementing such a solution could require extensive changes to the NonStop application

Both products transparently encrypt/decrypt or tokenize/detokenize sensitive data which an application writes to/reads from the database or files. They accomplish this by intercepting the application IOs and front-ending them with logic that looks for sensitive elements in the data and tokenizes or detokenizes them as they are written to or read from the database. The application is unaware of this operation and continues to function, oblivious of this operation. These products fully encapsulate the actual operations such as SecureData interaction and integration with NonStop architectural elements (e.g. authentication, scalability, logging). Both of these products allow you to easily configure the solution for ACI's Base24 application. They support protection of data stored in OSS and Enscribe file systems as well as SQL/MP databases. Transparent encryption and tokenization for SQL/MX databases is not yet supported. However, you could still protect the data in the SQL/MX database using FPE/SST by modifying the application code as needed.



Figure 4 illustrates conceptually the basic architecture of these products.

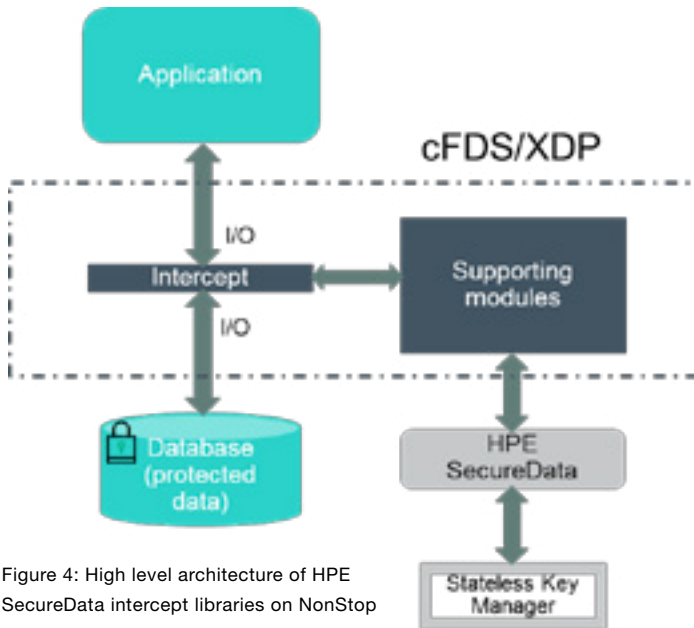


Figure 4: High level architecture of HPE SecureData intercept libraries on NonStop

comForte Data Security (cFDS) comes with a base module and three add-ons. The base module provides the basic functionality described above and protects the data in Enscribe/OSS files. A separate add-on is required to protect the data in SQL/MP databases. cFDS has an advanced module which supports transparent encryption or tokenization of data stored in complex formats such as ISO 8583 (a standard for systems to exchange electronic transactions made by payment cards). It also provides mechanisms to migrate data from an unprotected environment to a SecureData protected environment without requiring down-time. Another add-on called “file protection” can be used to protect large data records stored in unstructured files using traditional encryption instead of FPE or tokenization for optimal performance. It also supports dual control and split knowledge for keys, which can provide an additional layer of security for the keys stored in Hardware Security Modules (HSMs).


XYGATE Data Protection (XDP) comes in three independent modules. The first one is an SDK that provides the basic functionality described above minus the transparent encryption/tokenization feature. It is useful in scenarios where you would want to have a tighter control on implementation but want to use the out of the box integration with platform fundamentals available in XDP. This module also simplifies the integration with SecureData libraries apart from enabling applications written in non-native code to protect the data using tokenization. The second XDP module contains the SDK and also supports transparent encryption or tokenization of sensitive data stored in OSS/Enscribe files. The third XDP module contains the SDK and supports transparent encryption or tokenization of sensitive data stored in OSS/Enscribe files and SQL/MP databases.

Depending on the enterprise environment, implementing data-centric security solution can at times be complex and may require expert advice. In order to help you get the best value of this solution—which protects your customers’ data privacy and neutralizes the effects of data breach—and provide implementation guidance, HPE also offers professional services which can be ordered along with these products.


If you would like to evaluate the solution to help you make the right product choices, HPE can offer you a free trial period during which you can try and implement this solution as a proof of concept and see for yourself the value it provides to your environment. Feel free to contact your account representative should you be interested in exploring this option.

Conclusion

In summary, HPE SecureData and its companion products on NonStop provide comprehensive solutions to help you migrate from an unprotected or semi-protected environment to a fully-protected environment in the shortest period of time. Depending on the nature of the deployment, migrating to a fully protected state in a matter of a few weeks is possible with these solutions.

The need to protect customer’s sensitive data is well understood across many industries today. The question, however, is what is the right method to accomplish this and what solutions and tools are available in the market to help you implement it in an efficient and future-proof manner. With HPE SecureData and the companion products available from NED, you now have arguably the most state of the art solution in the industry for implementing data-centric security for your NonStop or heterogeneous environment. With a single-vendor approach, you also benefit from using a familiar and trusted partner who knows your environment and understands your needs. 

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

Prashanth Kamath U is a Senior Product Manager at HPE NonStop Enterprise Division and manages the product portfolio of the NonStop OS, Security products and Release Version Updates. He is responsible for defining the product strategy, roadmap and life cycle planning. He has over 20 years of experience in the industry. He has Bachelor of Engineering qualification from the National Institute of Technology, Surathkal (India) and Post Graduate Diploma in Software Enterprise Management from Indian Institute of Management, Bangalore.



You Don't Know What You Don't Know (But Neither Does Anyone Else)

How our Brains are Often Deceptive (Without Us Realizing It)

Luther Martin >> Distinguished Technologist >> HPE
Stacia Topping >> Engineering Program Manager >> HPE
Amy Vosters >> Marketing Manager >> SOASTA

You do not know what you do not know, and if you do not consciously compensate for this, you will probably end up making some less than optimal decisions. Fortunately, how to do this is well known. Understanding that we all experience this issue is a good first step to beginning to understand and avoiding the difficulties that it can cause. You can then use various strategies to assist you in making better decisions.

The Dunning-Kruger effect

People with below-average abilities tend to not know that they are below average. This is often summarized as “incompetent people don’t know they’re incompetent,” and it applies to everyone because we all have some areas in which our abilities are below average. This can lead to situations that are either amusing, embarrassing or irritating, depending on your point of view and on the situation. But the more general psychological effect of cognitive biases, situations where our brains work in ways that systematically make non-rational decisions, is that they probably affect how well we perform in our day-to-day jobs including the choices and decisions we make while at work, so understanding how cognitive biases affect our judgment can be very useful.

The fact that incompetent people do not know that they are incompetent has probably been observed informally for thousands of years, but the first dedicated study of this phenomenon was motivated by a puzzling crime which occurred in 1995. That is when an unlucky bank robber named MacArthur Wheeler committed what seemed like an incredible blunder: he used lemon juice to cover his face during a robbery, incorrectly assuming that the same properties of lemon juice that made it viable for use as invisible ink would also render his face invisible to the surveillance cameras in the bank.

Wheeler was apprehended by the local police shortly after committing his robbery, quite puzzled by why his clever scheme has failed. He was even noted to have said in dismay as he was arrested and shown the videotape recording of his robbery, “But I wore the juice.”

David Dunning, a professor of psychology at Cornell University, and Justin Kruger, one of his graduate students were so amused when they read about this incident in the newspaper that they began studying the nature of such blunders.

The subsequent research by Dunning and Kruger¹ uncovered two interesting patterns: unskilled people tend to overestimate their level of skill, while highly skilled people tend to underestimate their level of skill. This turns out to be a special case of a cognitive bias, the manner in which our brains make a predictable error in the way we make decisions. In particular, what is now known as the Dunning-Kruger effect is a specific example of the false consensus effect, where we tend to assume that everyone else thinks in the same way we do.

According to the Dunning-Kruger effect, an unskilled person will tend to assume that everyone else has the same level of skill. As a result, he assumes that his skill level is closer to average than it actually is. Similarly, highly skilled people will tend to assume that everyone else has the same high level of skill, and they conclude that their skills are closer to average than they really are. The net result is that people tend not to have a good understanding of their true abilities or skill level.

Why this matters

No matter how much we might like to think that we are exempt from cognitive biases shown in the Dunning-Kruger effect, the reality is that we are probably not. The decisions that we make throughout the day almost certainly involve more than a few questionable or even bad choices due to both aspects of the Dunning-Kruger effect, as we have both weaknesses and strengths. But there are strategies for dealing with cognitive biases that can help us overcome the limitations imposed by our brains, including the Dunning-Kruger effect.

A good example of the gains that can be realized from working to identify and overcome our cognitive biases is the story of how the Oakland Athletics baseball team managed to improve their standings from a very mediocre record of 65 wins and 97 losses (winning 40.1% of games) in 1997 to a significantly better record of 103 wins and 59 losses (winning 63.6% of games) in 2002. Over this period, the team had the second-best record in American Major League Baseball, while keeping their costs (mostly determined by players’ salaries) extremely low.

This dramatic turnaround, described by Michael Lewis in his 2003 book *Moneyball* and in the 2011 movie by the same name, was largely due to how the Athletics found ways to overcome the cognitive biases that affected the judgment of the managers of other baseball teams. In particular, with the help of economist Paul Podesta, the Athletics’ general manager Billy Bean developed statistical models of player performance and took the bold step of implementing the strategy suggested by the models instead of the one used by industry experts.

The baseball industry now relies extensively on statistical models to optimize the performance of teams while keeping costs as low as possible. But the fact that it took until the late 1990’s for the approach to be considered is interesting. The most reasonable explanation for this is probably not that highly paid managers of baseball teams are incompetent. Instead, it probably occurred simply because managers are human. And because they are human, they are prone to the Dunning-Kruger effect and the assumption that their expert knowledge of baseball was enough for them pick optimal teams. And because they are human, they ended up being incorrect in their assumptions.

If cognitive biases kept the managers of baseball teams from adopting winning strategies for many years, we should not be surprised to learn that they also affect the rest of us too. And just as it was possible for the Athletics to overcome these biases and hire players that a careful analysis of data recommended instead of the players recommended by expert opinions, it is possible for us to do the same.

The easiest way to avoid problems caused by the Dunning-Kruger effect and related cognitive biases is to remove as much of the decision-making as possible from the often flawed process of human process based on experience, opinion and judgment. That is essentially what Bean and Podesta did for the Oakland Athletics. Decision-analysis tools can be used to make optimal decisions and avoid some of the problems that our flawed judgment can cause. But even these tools are not perfect, because the models they are based upon are only as good as we allow them to be, based on our experience and information, and choices and decisions.

Suppose that you want to find optimal pricing for auto insurance. If

¹Kruger, Justin, and David Dunning. “Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments.” *Journal of personality and social psychology* 77, no. 6 (1999): 1121.

your model does not include the gender of the driver then the model will not be as useful as it could be. You will end up overcharging women and undercharging men because historical data has shown that it costs more to insure male drivers than to insure female drivers. Many, but not all, US state governments allow the use of gender in determining the price of auto insurance, which allows the cost of the insurance to more accurately reflect the cost of insuring drivers.

Similarly, there may be factors that contribute significantly to the accuracy of the models that decision-making tools can create, but which users of the tools either overlook (perhaps due to cognitive biases) or are not permitted to use by the regulatory environment in which their organizations operate, such as the limitations imposed on the health insurance industry regarding gender and the price of insurance. Although it is probably extremely unlikely, it might be the case that the height of the wives of baseball players or the astrological signs of the players themselves are strongly correlated with the performance of the players on the field, but unless those particular variables are used in creating a model that tries to predict the performance of players, the model will not capture this information.

And just because there is a statistically significant correlation between two events, does not necessarily indicate a causal relationship between the two events. Finding spurious correlations for which there is probably no causal relationship is something of a hobby for at least a few people. Some such people have mined publicly available data sets, discovered fascinating correlations and posted their discoveries on the internet.

One enterprising person has even published a book showing the most amusing of these. For example, one such spurious conclusions the book purports is that there appears to be a strong correlation between the divorce rate in Maine and the per capita consumption of margarine (having a correlation coefficient $r = 0.992598!$). But basing public policy on this correlation is extremely unlikely to produce useful results. Discouraging people from eating margarine is unlikely to keep couples in Maine either happier or married. Correlation simply does not imply causation.

If one were to apply this to the IT industry, there are probably a lot of interesting correlations waiting to be found in the event logs that your IT systems work so hard at creating. But many of those correlations are probably due simply to chance. After all, even an event that has only a 1 in 1 million chance of occurring probably occurs many times over the course of 100 million trials. And because the quantity of data that event logs for many systems easily capture each month, there are likely to be many actually very rare events that appear in the log files. And they appear there purely by chance, not because they represent that a significant event occurred. So it is important to note that correlation, an event occurring in a log file containing millions of events, does not necessarily imply causation, even over time.

It can also be important to understand that models can change over time. So that what may be a very good model today may end up being relatively useless in a few years. This problem was demonstrated by the dramatic success of Google Flu Trends in predicting flu outbreaks, which was followed a few years later by an equally dramatic failure of the model.

In 2009, researchers from Google and the Centers for Disease Control and Prevention (CDC) published a description of the method they had devised to analyze Google search queries to track influenza-like illnesses in populations and had then used this to create the Google Flu Trends website. The approach was more accurate than the one used by the CDC, which led to Google Flu Trends being touted as an example of the sort of amazing things that are possible when “big data” is correctly analyzed.

Google’s approach to tracking influenza-like illnesses made the news again a few years later, but this time for a different reason. In 2013, a group of researchers noted that the model that had been so successful at predicting the influenza-like illnesses was no longer as good as it once was. The Google Flu Trends model had actually been fairly inaccurate since the 2011 - 2012 flu season, and by 2013, Google Flu Trends was predicting more than double the number of office visits than the CDC’s model was. The model based on internet searches was no longer as good as the one based on actual visits to doctor offices.

However, it was discovered that the changes that Google made in its search algorithms a couple years later greatly decreased the accuracy of the model that was based on the output of searches that were based on the older algorithm. Another important lesson was learned about the capabilities of “big data,” but it was not the lesson that many people were expecting to learn.

From our point of view, an important thing to learn from the Google Flu Trends story is that even the best model may not be accurate over time. So while replacing the judgment of error-prone people with more precise analytic models is probably a very good idea, we also need to understand that even the best model may not be useful forever. In addition to challenging the assumptions that go into making a model, it is also useful to challenge the effectiveness of a model after it has been in use for a while.


Another example of how cognitive biases can affect the decisions that we make may be shown by the reception of the 2000 Stanford doctoral dissertation of Kevin Soo Hoo, in which he did a careful cost-benefit analysis of many information security technologies. His results were somewhat surprising: some technologies that are widely used seem to be hard to justify while other technologies that are not as widely used seem to be easy to justify. And while no one has taken the time and effort to argue that Soo Hoo’s results are inaccurate or incorrect, they are widely ignored by the information security community.

In light of the Dunning-Kruger effect and related cognitive biases that affect our judgment, this should not be terribly surprising. After all, the conventional wisdom that both industry analysts and the marketing departments of security vendors provide for us certainly feels like it comes from reliable sources. But this could be just as reliable as the conventional wisdom that led to the development of sub-optimal baseball teams. The cognitive biases of experts in the information security industry could very well be keeping them from making better decisions, and this could be causing investments in security products that are not as effective as they might be.

A better approach might be to take Soo Hoo’s results more seriously, but a strategy based on them would probably be very hard to actually implement. The same cognitive biases that affect information security professionals also may affect their auditors, and it could be very difficult to convince the auditors that ignoring the conventional wisdom is acceptable just because a mathematical model indicates that the conventional wisdom is less accurate than was commonly assumed.

Summary

The fact is that our brains do not always work as we would like them to. This reality provides an endless source of both entertainment and employment for psychologists. In addition, if we do not understand the unexpected ways in which our brains function, it can lead to problems in our daily life when we make decisions that are not as informed as they could be.

Cognitive biases are fairly common, but it is also fairly easy to account for them. The first step toward avoiding the challenges and oversights that they can cause is to understand that they exist and how they influence our judgment. Once we realize this, it becomes easier to adjust how we make decisions to account for such biases. Removing the error-prone judgment from as much of the decision making as possible seems to be a good way to do this. Replace judgment with impartial analytical analysis and you will be well on your way. 

Luther Martin is a Hewlett Packard Enterprise Distinguished Technologist. You can reach him at luther.martin@hpe.com.

Stacia Topping is an Engineering Program Manager at Hewlett Packard Enterprise. You can reach her at stacia.topping@hpe.com.

Amy Vosters is a Marketing Manager at SOASTA. You can reach her at amy_vosters@yahoo.com.

NonStop File Integrity: Check It! Protect It!

Callum Barclay > CTO & Founding Partner > Computer Security Products, Inc.

File Integrity Monitoring on NonStop

File Integrity Monitoring (FIM) is an important requirement of the PCI data security standard for maintaining confidential (e.g. cardholder) information, and is considered a crucial part of protecting business assets.

NonStop systems are now being used in far more dynamic situations and have more external connections than ever before. The ubiquity of payment cards for personal electronic transactions has changed the security equation in a fundamental way.

Any compromise in security is likely to have far reaching consequences, both for the immediate damage that may be done in terms of financial loss, and for the wider damage done to a merchant's reputation. The security of personal cardholder information has become paramount.

In this context, FIM should be considered an important security necessity, not just for PCI systems, but for all NonStop systems.

FIM and PCI DSS

PCI DSS Requirement 11.5 stipulates that members must "Deploy a change-detection mechanism (for example file-integrity monitoring tools) to alert personnel of unauthorized modification of critical system files, configuration files, or content files; and configure the software to perform the critical file comparisons at least weekly."

In version 2.0 of the security standard, in a clarification to Requirement 11.5.b, it was further specified that it is an audit requirement to "Verify that tools are configured to alert personnel to unauthorized modification of critical files."

PCI DSS Requirement 11.5 version 3.1 further clarifies that unauthorized modifications include changes, additions, and deletions of critical systems files.

It is clear from these excerpts that FIM is a key requirement of PCI DSS, and therefore a FIM solution must be implemented on any system that handles cardholder information.

What is FIM?



FIM includes any technology that monitors files for changes. Assuming that at least some file change is expected on a system, then FIM's primary purpose should be to identify possible "bad" changes so that they can be rolled back or remediated in some way. A "bad change" is any change that is undesirable. This is not the same as an unplanned, unauthorized or suspect change.

An unplanned change is not necessarily a "bad" change. Most system administrators have found it necessary to intervene on occasion to remedy a problem. Their actions might include changing a configuration parameter, or perhaps changing the security of a file due to an oversight.

In both cases, the change is both unplanned and unauthorized. Regardless, the change must be appropriately recorded and reported, then

reviewed and either made permanent or modified.

Of course other changes that may be unplanned and unauthorized can be part of an active security threat, in which case FIM may provide the first notice that the system has been compromised. Accidental change represents no less of an issue and is probably the most likely source of unplanned and unauthorized change.

FIM can also be used as part of a change control regime, whereby planned changes are detected and recorded to have occurred as expected.

Components of FIM

Basic FIM functionality should allow the administrator to:

1. Create and store a baseline for specified files and their attributes of interest
2. Update the baseline to take into account planned or allowable change
3. Run periodic checks and report the results
4. Store the results of each check

What to monitor?

A major concern is that FIM generates "noise" about file changes. Too much activity is recorded on too many files. Like excessive audit, excessive FIM can result in a reduction of useful information.

An effective FIM solution must therefore provide flexible integrity check mechanisms able to select files based on name and property. In deciding what files should be monitored, judgment is needed to determine the risk created by a change to a file.

Obvious monitoring choices would include system files. Files which are LICENSE'd or PROGID'd would also be candidates.

Added to these would be key application files, including data, executables and configuration files.

Other files should be added as required. Any files related to cardholder data are especially sensitive. While a file's contents may be dynamic, other file attributes can be monitored to ensure that the correct attributes are in place and remain so.

Types of change

The FIM solution should allow for the specification of different types or categories of change, for example:

- Content, both complete and incremental
- Security settings, including flags for ownership, file permissions and special security bits
- Basic attributes like file type, last modified etc.

Grouping checks by change type simplifies both the scheduling of checks and the review of the results.

Low and high risk changes

As noted, a significant challenge is presented by the potential quantity of change notices that can accumulate. To counter this problem the FIM solution should be configured so that it supports the appropriate specification of risk for each file.

For example, a file may represent a risk if compromised and may also have relatively dynamic content. Monitoring for content change is therefore not useful. It is better to look at the attributes that directly impact the risk-level and monitor for those. These are most likely to be security settings.

Grouping files into filesets based on their risk profile is therefore recommended.

Typical FIM workflow: Change notification

When the FIM has completed its check, there should be a way to quickly identify any changes that have occurred, since trawling through

lengthy detail reports could lead to oversights.

Ideally, an exception report could be created that feeds into a central event management system, in order to ensure that all high-risk changes are considered collectively, and are acted upon promptly.

Real-time alerts

Some files are too sensitive to wait for changes to be analyzed. In this case, event monitoring solutions can be used to immediately invoke both notifications (alerts) and investigative or remedial actions.

Certain changes will also be recorded in the Safeguard audit trail, e.g. security setting modifications. These can also be detected in real-time by event monitoring solutions and made actionable.

Change analysis

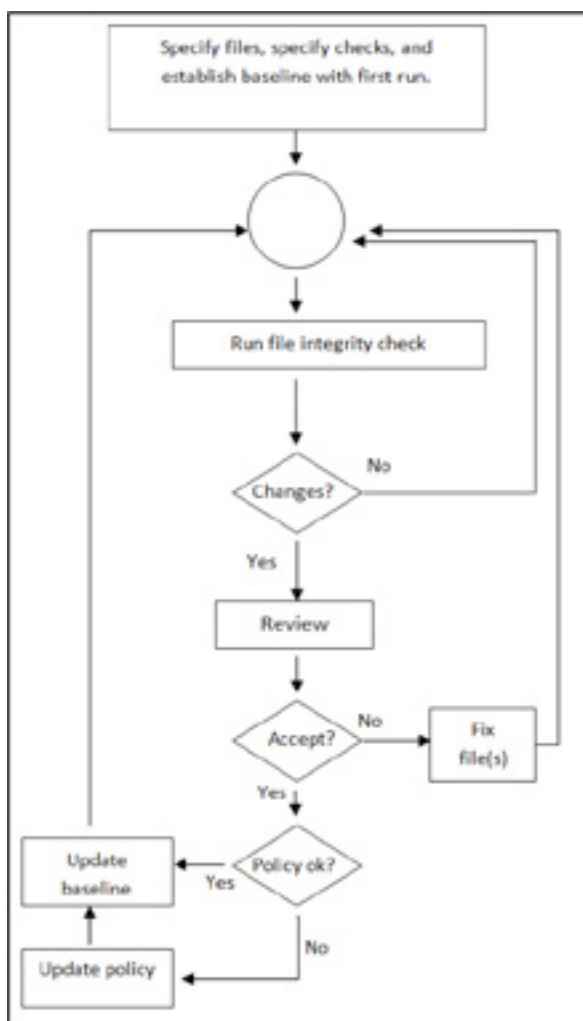
As each monitoring cycle is completed, any changes should be reviewed. The number of changes to be reviewed must be manageable. For some filesets the cycle will occur daily, for others weekly or monthly.

Advanced FIM solutions provide the ability to check multiple filesets, enabling the verification of groups of critical files more often (i.e. daily), and less critical files less often, (i.e. weekly).

Files that have been marked as changed in one of the specified attributes should be reviewed by the individual responsible for the contents or the attributes of the file. If, after review, the change is acceptable, then the baseline for that file is reset immediately. Otherwise, further investigation is required.

Change history

The results of each monitoring cycle should be stored. Information should include the details of the change so that there is a complete record of all changes. Such information can be correlated with audit records to determine who made the change.



Unauthorized, unplanned or suspicious change

If detected changes are not promptly analyzed, the value of FIM rapidly diminishes. The purpose of the analysis should be to establish quickly as possible that the change should either stand or be rolled back.

Unauthorized or unplanned change is not necessarily undesirable change – such as an emergency change made to remedy a production problem. By the same token, change that is authorized may turn out to be bad change. FIM solutions should save all change information for possible future remediation.

Suspicious change is simply change that has not been categorized for acceptance or rejection. Hence, the need to analyze change information as quickly as possible in order to identify it as good or bad.

Once a change is identified as “good”, then a new baseline is set for the file in question.

Change Approval

If a change has been identified as “good” and the baseline must be reset, it is important to require that the reset is properly approved. Otherwise the temptation to let “ok” changes roll through will be hard to resist, particularly in busy environments.

Any update to the baseline should therefore be separately authorized using an additional level of authentication requiring the entry of a special PIN.

Multi-Node File Compare

FIM solutions should also have the ability to compare files on different nodes, e.g. between production and disaster recovery systems. This is to ensure that no unwanted or unauthorized changes occur during the synchronization of files between systems. A record count difference threshold should be used to reduce false positives due to latency across nodes.

Integration with other security products

The significant quantities of data that can be generated by FIM solutions can be more easily managed by integration with other security products, in particular with auditing and event monitoring solutions.

Not all file changes represent equivalent risk. By developing mechanisms to extract high risk file changes and make them part of a broader event monitoring effort for similar security events (attempted breaches, logon attacks etc.), users can make FIM significantly more effective.

Correlating such file changes with other suspicious activity will also become easier – allowing security administrators to provide a more complete view of activity that may threaten to compromise the system.

Conclusion

FIM is a critical requirement for security, and key to PCI DSS compliance. However, the detection of any particular change is just the start of the process.

To be effective, FIM solutions must differentiate low-risk from high-risk change; integrate with other security solutions for log and security event management (including real-time alerts) and support a fully-managed history database of changes.

CSP's File Integrity Checker (FIC) is widely used by financial institutions to deliver FIM in NonStop Guardian and OSS environments, and is tightly integrated with CSP's other solutions for audit, compliance and Safeguard, EMS and Base24 OMF real-time event monitoring. FIC's new “Guardian Fileset Compare” feature permits the attributes of any two file sets on any two systems to be compared against each other. Find out more at www.cspsecurity.com.



As the original founder of CSP in 1987, Callum Barclay leads the technical direction of the company from its headquarters outside Toronto, Canada.

Originally from Edinburgh, Scotland, Callum is now heavily involved in bringing pioneering security and compliance solutions to HPE NonStop customers.



Leveraging a Big Data Analytics Engine for Meaningful Insights

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

The amount of information being generated each year is exploding at an unprecedented rate.[1] It is estimated that 80% of all of the world's data that has ever been created was produced in the last two years, and this rate is increasing. Social media such as Twitter and Facebook, articles and news stories posted online, blogs, emails, YouTube and other videos – they all contribute to what is now called big data.

Big data allows companies to obtain real-time business intelligence (RTBI) that they could never access in the past from their typical internal systems. Think of the customer-sentiment analysis that can be obtained simply from tweets. However, big data is a collection of data sets so large and so complex that it becomes impossible to process with current database-management tools and data-processing applications.

Much (perhaps most) of the content of big data is noise. It has little or no value to an organization. However, buried in this noise are tidbits of invaluable data which may be used to determine what customers are thinking, to plan new products, to find the strengths and weaknesses of competitors, to monitor for fraud and cyber-attacks, to defend against terrorism, and for many other purposes. The challenge is extracting the meaningful data from the noise. This is the task of the big data analytics engine.

A big data analytics engine typically requires a large network of tens, hundreds, or even thousands of heterogeneous, purpose-built servers, each performing its own portion of the task. All of these systems must communicate with each other in real-time. They must be integrated with a high-speed, flexible, and reliable data-distribution and data-sharing backbone.

In this article, we look at several ways to interact with big data to extract valuable business information from the noise.

Big Data

Events are no longer sufficient.

What do we mean by the above statement? After all, business processes and business intelligence are based on events. What did a customer purchase? When was a call made? When was an order delivered? Who logged on to our system and when?

For decades, businesses have captured these events and stored them in transactional databases managed by highly reliable systems. Events drive the business. They determine production schedules, product deliveries, product re-order thresholds, banking, fraud detection, corporate financial statements, and a myriad of other business functions. A business would be paralyzed without its mission-critical online transaction-processing systems.

However, the world has evolved. The amount of available information that can be valuable to a company has rapidly expanded. Tweets, Facebook postings, news articles and newscasts, YouTube videos, the email and customer service calls a business receives – all of them may contain information advantageous to a company for making informed decisions and enhancing competitiveness. This is what we mean by big data—all of this data no matter the source or format.

The data stored in transactional databases represents high-density information. Every element is pre-determined to be important. However, transactional data is a tiny fraction of the total data generated worldwide. The data contained in big data is low-density. Most of the data is noise and has no real value to a company. But some of the data can be extremely important. How is the valuable data identified and extracted from the noise and put to use?

Real-Time Data and Long Data

There are, in fact, two types of big data that have to be managed – real-time data and long data.

Real-Time

Data Real-time data is used for immediate analytics and business decisions. The most immediate data available to a big data analytics engine is data that is streamed (pushed) to it, such as tweets, web clicks, emails, and customer calls. Other real-time data must be pulled from its sources, such as Facebook posts, news stories, and blogs.

Real-time data is characterized primarily by velocity and variety. Real-time data arrives in a variety of formats, and the big data analytics engine must be able to parse and process all of the real-time data that is presented to it with minimal processing delays.

Long Data

Long data is a massive data set that extends back in time over an extended period, such as over the life of an organization, and is important because it places real-time data in its proper perspective. If an organization does not look at events from an historical viewpoint, it will analyze current events as the norm and will be blinded to what has happened previously. It will miss repeated or unusual events and the opportunities (or threats) thereby presented. This perspective is why the addition of long data to an organization's source of information is so important. It provides context for current events.

Consider climate change, for example. Real-time data tells us that our ice caps are melting and that sea levels are rising. Is the culprit our increased carbon emissions, or is it a natural cycle that has gone on for eons? Long data can help answer this question.

The Time Value of Data

The value of information is a function of time. Interestingly, the relationship of value to time is opposite for real-time data and long data, as shown in Figure 1.

Real-time data is typically used for real-time decision making. The older it gets, the less useful it becomes. Some real-time data items may have half-lives of minutes. Others may have half-lives of microseconds (as is the case for algorithmic, high-frequency stock trading).




Figure 1 – The Value of Data Changes Over Time and Type

Long data, on the other hand, provides historical context for real-time data. The more historical data that is collected, the better is the context. Therefore, the value of long data increases over time as more and more data is accumulated.

The Big Data Analytics Engine

The general structure for a big data analytics engine is shown in Figure 2. As mentioned earlier, three types of data sources make up the information that flows into the analytics engine:

- Streamed data is pushed into the analytics engine, including sources such as Twitter and email. Stream processors are provided for each source to parse these streams and to deliver pertinent information to the analytics engine.
- Static data is pulled from other sources, such as Facebook postings and news stories. Fetchers are provided for each static source to fetch new data that has been added to that source and to deliver the fetched data to the analytics engine.
- Transactional data from the organization's transaction-processing systems is sent to the analytics engine for its real-time value as well as for its historical value.

At the heart of the analytics engine are several (typically massive) components. The implementation of each can require tens or even hundreds of commodity servers:

- A batch-storage analytics engine is capable of storing unstructured data of any kind and can search that data rapidly for correlations. It receives all of the pertinent streaming data and all of the data that is being fetched from static sources. It analyzes this real-time data in context with the long data that it has stored to determine patterns of importance. These patterns, or correlations, are sent to other elements of the analytics engine for analysis processing.
- A column-oriented database stores intermediate results. A column-oriented database stores relational tables as columns rather than as rows. Many types of queries deal only with one or a few columns of a row and can be satisfied much more rapidly and efficiently with this architecture.
- An in-memory database typically holds the contents of the column-oriented database. It improves performance by eliminating disk-seek and transfer times and can further significantly speed up queries. Coupled with the column-oriented database, complex analytic queries can be completed in real-time.
- A Complex Event Processor (CEP) combines data from multiple event streams in real-time to create more encompassing events. These latter events are the RTBI generated by the big data analytics engine. They provide in-depth insight into what is happening in the business. The goal of the CEP is to identify meaningful events such as business opportunities or threats and to allow immediate responses through the applications that the CEP feeds.

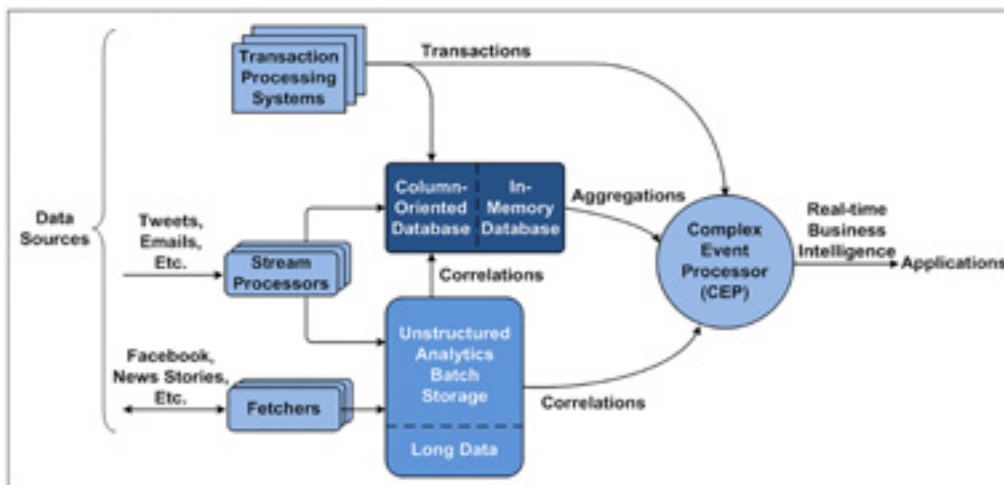


Figure 2 – Big Data Analytics Engine

The Integration “Glue” for the Big Data Analytics Engine

As described previously, a big data analytics engine comprises many different systems with different missions. Each system is implemented on a “best-fit” platform with a “best-fit” database manager. There may be a myriad of heterogeneous platforms, applications, and databases that make up the analytics engine. A powerful, flexible, fast, and reliable data-distribution fabric is required to interconnect these systems.

The data-distribution fabric between the many components in a big data analytics engine must be low-latency and provide high-capacity. It must be fundamentally heterogeneous and be able to deal with any application or database as a source or as a target. It must be able to reformat and restructure data on the fly as it moves data from one source to a totally different target. It must be highly reliable.




Figure 3 – Data-Distribution (DD) Fabric for Big Data Environments

A process-to-process architecture can eliminate disk-queuing points that slow down information delivery. Sub-second replication latency is achieved. An architecture that can be multithreaded, including the communication channels, enables the desired data-transfer capacity to be attained.

In an ideal world, replication should support heterogeneity. It should receive data as it is generated from any supported application or database and deliver it to any other application or database, with support for filtering, redefining, and enriching the information in-flight to satisfy any target environment formatting needs.

For mission-critical software environments, replication must be architected to provide continuous availability. If one of its components fails, it must be automatically restarted. Replication must continue uninterrupted. If the target system fails, the replication engine must queue all events until the target system is restored to service. It should then drain its queue of saved events to bring the target system back into synchronization with the data source and automatically resume replication of real-time events.

Summary


Big data offers the opportunity for businesses to obtain RTBI that they could never reach in the past from their typical internal systems. A big data analytics engine can mine social media, the press, email, blogs, videos, and a variety of other data sources to determine what customers are thinking, to plan new products, to find the strengths and weaknesses of competitors, to monitor fraud and cyber-attacks, to gain competitive advantage, and for many other purposes.

Big Data Protection

Businesses today are driven by data, and the quality of the business depends upon the quality of that data. Consequently, data has become one of a company's most valuable assets, and other people want it. Stealing or corruption of this data can result in significant business losses, pose serious security threats, and result in regulatory violations. As hackers become increasingly sophisticated, protection of data from unauthorized access is a number one priority for any IT department.

Protection of data from unauthorized access within a big data environment becomes much more complicated because the data is being consumed from many sources (trusted and otherwise), and moved between systems for analysis. For example, a data source may be restricted to only a certain set of users; if this data is replicated to a big data repository, measures must be taken to ensure that access to the data remains restricted to only this set of users.

Fortunately, using a data replication engine (such as HPE Shadowbase) for the data-distribution fabric addresses many of these data protection issues. When the data is in motion (being copied between systems) techniques such as IPSec and/or proxy servers (SSL/TLS) can be used to authenticate and encrypt each data packet. For data at rest, as the replication engine applies the data to the big data repository, HPE Shadowbase user exits can be customized to encrypt or obfuscate the data as it is written. Encrypted target file systems can also be used when available. Via these means, the data replication engine ensures that data replicated to a big data repository remains protected, regardless of the source of the data.

HPE Shadowbase replication capabilities can play a significant role in delivering inputs and outputs to key processes for analyzing big data. Wherever there is a need to transfer data from a data source to another target, regardless of the nature of those devices, HPE Shadowbase software solutions can be placed into service to get the job done efficiently and reliably. 

Keith B. Evans works on Shadowbase business development and product management for the Shadowbase product suite, including business continuity, data integration, application integration, zero downtime migration, data utilities, and synchronous replication, a significant and unique differentiating technology. To contact the author, please email: SBProductManagement@gravic.com.

Paul J. Holenstein is Executive Vice President of Gravic, Inc. He is responsible for the Shadowbase suite of products. The Shadowbase replication engine is a high-speed, unidirectional and bidirectional, homogeneous and heterogeneous data replication engine that moves data updates between enterprise systems in fractions of a second. It also provides capabilities to integrate disparate operational application information into real-time business intelligence systems. Shadowbase Total Replication Solutions® provides products to leverage this technology with proven implementations. For further information regarding Shadowbase data integration and application integration capabilities that can assist in solving big data integration problems, please refer to the companion documents [Shadowbase Streams for Data Integration](#) and [Shadowbase Streams for Application Integration](#), or visit www.ShadowbaseSoftware.com for more information.

To contact the author, please email: SBProductManagement@gravic.com.



Did You Know?

Improve your ROI for managing your enterprise encryption keys! Did you know HPE ESKM supports the largest single vendor ecosystem of data center storage and server applications on the market with a large community of third-party IT applications also included? Now you can maintain a single system to manage keys, set policy and audit protection in place across your entire global enterprise of critical HPE infrastructure. www.hpe.com/software/datasecurity

XYGATE® Object Security



A single solution for complete control
of Safeguard and OSS security

Sophistication

RBAC for HPE NonStop

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

"Simplicity is the ultimate sophistication" Leonardo da Vinci



Manage Access
to Information

Learn more at
xypro.com/XOS

 **XYPRO**
Mission Critical Security

Implementing Tokenization & Access Control



Don't Let it all Hang Out!

Andrew Price > VP Technology > XYPRO Technology
Scott Uroff > Chief Architect > XYPRO Technology

In recent years there has been an emergence of several new technologies to protect sensitive data, including Format Preserving Encryption (FPE) and Secure Stateless Tokenization (SST), such as those provided by HPE Security's SecureData product. These products provide excellent capabilities to assist HPE NonStop users in protecting data within their application environments. Both HPE FPE and HPE SST provide strong protection against the exposure of sensitive data but they should not be used alone or to replace traditional access controls. Data protection methods such as FPE and SST need to be carefully considered and planned alongside traditional access controls to ensure all application data is comprehensively protected both from authorized and unauthorized exposure. This article will give a high-level overview of how to implement a best-of-breed HPE NonStop security framework; protecting all sensitive application files and tables using comprehensive access controls, and also selectively protecting the highly sensitive and valuable data those files may contain, such as credit card (PAN) data or personally identifiable information (PII).

Mission critical applications such as those typically found on the HPE NonStop Server are composed of programs and files or tables. There are multiple levels of access requirements for both programs and files. For instance, only certain programs running as certain users should be able to access tables containing application data. This simple access control rule can be challenging to implement on the HPE NonStop Server, as standard HPE Nonstop security controls do not include the granularity features necessary to implement the desired security. For example, using the requesting object file as an attribute that can be used to control file access is not an option. Standard HPE Nonstop security can only control file access by the user running the program. In addition to the emergence of data protection technologies like FPE and SST, XYGATE Object Security (XOS), which uses the Safeguard Authorization Security Event Exit Process (SEEP), can be used to achieve the desired access controls for application security. This solution can use the requesting object file, among others, as an attribute when making access decisions, thus introducing more granularity into the access control matrix. Other partner products, including those from Greenhouse Software, also support the Safeguard Authorization SEEP.

Encryption and Tokenization Options

In addition to controlling the access rights of users and programs to application data, it is often also necessary to encrypt or tokenize sensitive data in tables to prevent its exposure to non-authorized parties. This may

be due to regulations, such as PCI-DSS, industry/corporate regulations, or just a result of the sensitive nature of the data itself. This can create a complex multi-tiered environment, which no single security product can fully address.

Two data protection methods have recently received a lot of focus in the NonStop space: disk (or volume) level encryption and application level encryption/tokenization. As a side note, file encryption is not considered for the purposes of this article, as encrypting entire live application files is generally either impractical, or involves extensive application redesign. Disk level encryption, known as VLE on HPE NonStop, is generally transparent to any logged-on users and therefore only protects against the disk drive being taken off-site and accessed. Due to this constraint, disk level encryption is no longer considered sufficient protection for PAN data, according to the PCI-DSS. Application level encryption also protects against disk drive removal but in addition will also protect the data from being accessed by anything other than authorized users or programs.

There are typically two variations of application level encryption:

- Integrated application level encryption
- Transparent application level encryption

Integrated application level encryption/tokenization is implemented by modifying the application programs to encrypt/tokenize and decrypt/detokenize sensitive columns. This can be a very expensive proposition depending on how many programs need modifying. It may also require the application programmers to have encryption programming knowledge, for instance how to manage keys. Also, this method typically precludes the ability, if required, for operating system utility programs to be able to see unencrypted data, since those programs cannot be easily modified. Using integrated application level encryption can make it difficult to share encrypted data with off box applications because those off-box applications would also have to be modified in the same way. HPE offers the HPE SecureData product for customers which want to use the integrated approach – and companies such as XYPRO, comForte and HPE are able to provide consulting services to assist with implementation if that approach is taken.

Transparent application level encryption/tokenization involves attaching a library to each program that needs to access the protected data. The library intercepts all I/O calls, and, based on its configuration, encrypts and decrypts specified fields or columns for specified programs running as specified users. The library can also be attached to operating system utility programs if required, and then those utilities can see unencrypted data. If this

library uses underlying encryption technologies that are available on multiple platforms, sharing data with off-box applications is relatively easy. HPE offers two transparent application level encryption/tokenization products: XYGATE Data Protection (XDP) and cF Data Security. Both products provide these features and address these needs, using industry-leading HPE SecureData as the underlying 'layer' for encryption and tokenization.

Adding Access Controls into the Mix

When using transparent application level encryption, granular access controls are also important. The encrypting of data has to be combined with the ability to configure which processes, running as which users, running which object files, can access the sensitive data in an unencrypted format. For example, the process running the object file that is used to verify PANs should be granted the authority to see the unencrypted PANs. A process running any other object file should not see unencrypted PANs. An encryption scheme that encrypts and decrypts PAN data for processes running any object file accessing the data provides no better protection than disk level encryption.

Let's look at some examples of how access control to a tokenization system can be implemented in the XDP encryption library.

```
DPGROUP CRD-FILE
FILE $*.*.CRDTBL
FIELD FPE FIELD_POSITION 6:16
REQUESTOR $*.*.CRDFPROG
OPERATION ENCRYPT,DECRYPT
ACL APPL.USER
AUDIT_ACCESS_PASS ON
AUDIT_ACCESS_FAIL ON
```

The above configuration entry says for any file called CRDTBL on the system, there is a 16 digit PAN starting at position 6 in the record that needs to be encrypted with Format Preserving Encryption. This entry only applies to object files named CRDFPROG run by user APPL.USER. All access is audited and can be captured and/or forwarded to a SIEM using XYGATE Merged Audit.


While the above entry controls how the PAN column is encrypted, and which program can see the unencrypted data, it does not control overall access to the file. Access should be controlled to the file so that only those programs that need to access the file are able to. This is important because the file may contain other information that needs to be protected, not just the encrypted/tokenized PAN. This could be implemented with an XOS entry like the following:

```
OSGROUP CRD-TABLE
MASK $*.*.CRDTBL
ACL APP1.USER *
REQUESTOR $*.*.CRDPROG $*.*.ALTPROG
AUDIT_ACCESS_PASS ON
AUDIT_ACCESS_FAIL ON
```

The above entry says that for the file \$*.*.CRDTBL, which is owned by user APP1.USER, the APP1.USER can perform any file system operations on the file when running a \$*.*.CRDPROG or \$*.*.ALTPROG program. Note that in combination with the XDP entry above, while ALTPROG can access the file, it will only see encrypted PAN data.

Two applications, one file?

Being able to control the security of a file based on the requesting object also helps in the situation where two different applications need to share a file when the applications run as different users. Assume that there is a primary application and a secondary application that both need access to one file owned by the primary application. Typically the security to access the shared file would need to be granted to both applications UserIDs. However, this means that any program running as one of the secondary applications UserIDs would be able to access the data. Having a security scheme that includes the object file as one of the access controls means that the one program in the secondary application that needs to access the primary applications file will be the only program that can access it. Any




"When I told management that we needed to increase security around our data, this is NOT what I had in mind..."

other program running as the secondary application's UserID will not be able to access the data. The above scheme could be implemented with an XOS configuration entry like the following:

```
OSGROUP SHARED-APPL-FILE
MASK $DATA.APP1.TXFR
ACL APP1.USER *
APP2.USER (R,W)
REQUESTPR $APP1.OBJ.WRTETXFR $APP2.OBJ.PROCTXFR
AUDIT_ACCESS_PASS ON
AUDIT_ACCESS_FAIL ON
```

The above entry says that for the file \$DATA.APP1.TXFR, which is owned by user APP1.USER, the APP1.USER can perform any file system operations on the file when running the \$APP1.OBJ.WRTETXFR program, and that the APP2.USER, when running the \$APP2.OBJ.PROCTXFR program, can Read or Write the TXFR file.

HPE NonStop servers and most modern computing platforms have always benefited from a layered approach to security – there is no point locking your windows when your front door is wide open. Newer technologies like HPE Format Preserving Encryption and HPE Secure Stateless Tokenization provide another layer in the security administrator's arsenal and can be very powerful when deployed in conjunction with more traditional security mechanisms. Just make sure to plan out your complete implementation so that all users and applications get just the access they need, and nothing more. As an added benefit, you'll also address both PCI-DSS Requirement 7 "Restrict access to cardholder data by business need to know" and Requirement 3 "Protect Cardholder Data". 



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

Scott Uroff is Chief Architect at XYPRO. A Double Master ASE with HPE NonStop Solution Architect and HPE NonStop Systems Support accreditation, Scott has served as XYPRO's senior technical expert since 1992 and has more than 35 years of experience with the HPE NonStop platform. He's the original developer of most of the core elements in XYPRO's product set and a key contributor to the company's published handbooks on HPE NonStop security.

XYGATE® Data Protection



All the benefits of HPE SecureData
No application changes

Simplicity

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



Now available from



Hewlett Packard
Enterprise

Protect your Data
at rest and in transit

Learn more at
xypro.com/XDP

XYPRO®
Mission Critical Security

©2016 XYPRO Technology Corporation. All rights reserved.
Brands mentioned are trademarks of their respective companies

Where Tokenization Fits Within Your NonStop Security Plan

Thomas Burg > CTO > comForte

If you're administering NonStop systems, then chances are your organization already has a security plan. That's because NonStop systems support some of the most important organizations in critical industries, including financial services, telecommunications, and energy, to name a few.

A successful cyber-attack on certain organizations within any of these industries could have widespread effects. That's precisely why they have been staked out for, if not already targeted with, sophisticated attacks by threat actors ranging from criminals and hackers to terrorists and nation-states.

In an ideal world, every NonStop administrator has built upon his or her organization's existing security plan to develop specific NonStop policies and safeguards. However, we know the reality: Security is but one of many hats NonStop administrators wear every day, and there are many hats and only so many hours.

So, if you haven't already, today is a good day to start developing a security plan for the systems you administer. If you already have one, this article can serve as a guide to reviewing and updating your plan.

Developing (or Updating) a Security Plan

This will be challenging because few, if any, people inside your organization will be able to assist you, the NonStop expert, with detailed security planning for your NonStop systems. You know the CEO can't help, and chances are, most CSOs and fellow colleagues can provide only general guidance. You know your systems, the applications they support, and the data they handle better than anyone else — so you must be the one who develops a comprehensive security plan to protect them.

The entire process of developing a security plan is beyond the scope of this article, but I'll provide a general framework, including how tokenization fits within your security plan. And it absolutely should.

First, it's helpful to begin at the highest level, which includes a consideration of people, process, and technology. Diving deeper, you can subdivide the technologies partially or fully under your purview as follows:

- **Networks** -- Both internal and external
- **Hardware** -- Systems and their functions (e.g., application, production, development, disaster recovery)
- **Operating systems** -- Versions
- **Applications** -- Business functionality
- **Data** -- Types and the risk if compromised

You likely know that specific tools have been designed to safeguard these different technologies. Firewalls and intrusion prevention systems protect the network, identity and access management solutions protect your applications and files, and so on. A security information and event management (SIEM) system collects data on potential incidents from security technologies, while analytics help to make sense of the data (e.g., categorizing, prioritizing, visualizing, etc.). Don't forget that a solid security plan will also include incident response, business continuity, and disaster recovery specifically related to the NonStop systems you administer.

Defense-in-Depth for Data: Classic Encryption and Tokenization

Chances are, your organization has considered (at least at the enterprise level) and hopefully implemented some form of data protection. Solutions range from classic encryption and data loss prevention systems (DLP, also known as data leak prevention) to tokenization.

Both tokenization and "classic encryption" protect data effectively if implemented properly, and an ideal payment security solution will use both. One key difference between them is that unlike classic encryption,

tokenization will never change the length of the data it protects, which is extremely attractive since many legacy systems will not allow changes to data field length (e.g., in databases).

There are numerous ways to classify tokens, including single-use and multi-use, reversible and irreversible, cryptographic and non-cryptographic, authenticable and non-authenticable, and various combinations thereof. Payment Card Industry Data Security Standard (PCI DSS) documentation provides detailed guidance.

Unfortunately, many information sources only implicitly reveal the type of token being discussed, and even fewer highlight important distinctions. To make matters worse, the terminology is not fully mature and agreed upon yet. In effect, the audience usually must make assumptions. Let's take a moment to look at the difference between payment tokens (i.e., high-value tokens) and security tokens (i.e., low-value tokens).

High-Value Tokens or Payment Tokens

High-value tokens (HVTs) are values that act as surrogates for actual Primary Account Numbers (PANs) in payment transactions. Importantly, an HVT solution (e.g., Apple Pay) enables the HVT itself to be used as an instrument for completing a payment transaction. To function, HVTs must look like actual PANs.

Here are a few key drivers for using an HVT for actual payment rather than the PAN it is mapped back to:

- **Multiple HVTs can map back to a single PAN** -- In very basic terms, HVTs are helpful because they are created out of "thin air" and multiple HVTs can be and are mapped back to a single physical credit card without the owner being aware of it.
- **Limit range of fraud** -- HVT usage can be limited to certain networks (e.g., Apple Pay) and/or merchants (e.g., Apple, Amazon, etc.) whereas PANs cannot.
- **Bind tokens to devices** -- HVTs can be bound to specific devices. It would be easy to correlate tokens to some physical device identifier (e.g., media access control [MAC] address, International Mobile Subscriber Identity [IMSI], etc.) along with historical location data. Anomalies between token use, physical devices, and geographic locations could then be flagged as potentially fraudulent.

Low-Value Tokens or Security Tokens

Low-value tokens (LVTs) also act as surrogates for actual PANs in payment transactions. However, they exist for a different reason than HVTs. By design, and in contrast to HVTs, LVTs cannot be used in and of themselves to complete a payment transaction. For LVTs to work at all, it must be possible to match them back to the actual PANs they represent, but only in a tightly controlled fashion.

For example, using an LVT solution, a consumer's PAN (e.g., 4485-4269-0687-2380) is tokenized by replacing the actual value with the surrogate value, the token (e.g., 0x3K-9u4L-09e8-03i7). This token obviously cannot be used in place of an actual card number in any transaction. It must always be matched back to the actual PAN (e.g., 4485-4269-0687-2380) to complete a payment transaction. This mapping from LVT to actual PAN is done within a "tokenization system."

A tokenization system converts LVTs to PANs and vice versa, and it can reside both in a separate hardware device, as well as on highly secured servers. The entire point of using security tokens to protect PANs becomes moot if a tokenization system is breached. Hence, securing the tokenization system itself is extremely important.

Where Tokenization Fits Within Your NonStop Security Plan

So now we have distinguished between high- and low-value tokens. But how should you think about data protection, generally, and tokenization, specifically, in your security plan for NonStop systems?

NonStop administrators can think about securing data as it exists in three distinct states:

1. At rest
2. In motion
3. In use at a point of transaction (e.g., a point-of-sale device, ATM terminal, or eCommerce web form)


Classic encryption is an important layer of security for data in all three states, but it has some potential weaknesses. First, classic encryption is only as effective as the scheme for managing the underlying cryptographic systems, which is among the most complex tasks even security specialists undertake. As if this is not complicated enough, there have been instances where the compromise of external organizations (e.g., certificate authorities) have led to another organization's breach. Third, a malicious actor (internal or external) with access (logical or physical) to NonStop systems could also potentially steal the cryptographic keys protecting that data. So if you are relying on classic encryption alone, you could be more vulnerable than you realize.

This is precisely where tokenization — specifically, the use of low-value security tokens for PANs — fits into your security plan, as part of a defense-in-depth strategy. It adds an additional layer to your other data safeguards by rendering PANs indecipherable and unusable — even if an attacker successfully decrypts the (still) tokenized PANs.

Suppose, for instance, a malicious actor gains access to the

be physical or logical by an internal or external party, and the system could be in the merchant, acquirer, or issuer environment. The attacker would encounter one of several scenarios:

1. If the data on the disk drive are not encrypted or tokenized, then nothing prevents the attacker from using the clear text PANs in fraudulent transactions and/or selling them in the online black market.
2. If the PANs are encrypted using classical cryptography but not tokenized, then the attacker could potentially also steal the encryption key (which typically is only a few bytes long) to decrypt the data and then proceed with fraud or resale.
3. However, if the PANs are tokenized via LVTs and the tokenization system itself is secure, then the data would prove useless to the attacker because LVTs alone cannot be used for payment. Without access to the tokenization system, the LVTs cannot be mapped back to the actual PANs. In this way, the stolen LVTs would prove worthless to the criminal without the reference list.

Defense-in-depth of data stored, processed, and/or transmitted by NonStop systems is an important element in your security plan. Whether you are developing a NonStop security plan from scratch or revising an existing plan, make sure you understand the important distinction between classic encryption and tokenization, as well as the necessary role each plays in securing your NonStop systems. 

If you have any question or if you would like to discuss your security challenges please email Thomas Burg at t.burg@comforte.com Or get in touch at info@comforte.com

comForte®

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.



Get Connected!

Sign up today! It's easy and FREE for HPE customers and HPE employees.

Tell a friend!

<http://www.connect-community.org/join-today>

CONNECT
Hewlett Packard Enterprise Technology User Group

SIGN UP TODAY!



Format Preserving Encryption

Karen Martin

The National Institute of Standards and Technology (NIST) recently approved two format-preserving encryption modes for government use.¹ Format preserving encryption (FPE) solves two major problems with encrypting certain types of numeric data: it maintains the format of the data, and it allows the encrypted data to be used as indices in relational data bases. Both of these features will be useful for protecting information such as credit card numbers (CCNs), Social Security Numbers (SSNs), and other personally identifiable information. This article will describe the advantages of using FPE, explain the basic concept of the modes, and discuss why they are secure and effective.

Format Preservation

With most encryption schemes, the ciphertext looks very different from the original plaintext. To encrypt a sequence of decimal digits, like “41111111111111”, the numerals would be encoded as a bit string – a sequence of 1s and 0s – and then encrypted. Our numeral sequence “41111111111111” might be encrypted to “gBRTB1cxAbvMhjFWc+7J9A==” for example.² The first sequence could represent a 16-digit credit card number, but the encrypted version could not.

Many existing IT systems cannot handle format changes without expensive and time consuming modifications. CCNs, for example, are frequently stored in or transmitted through systems that require the data to look like a valid CCN. The payment card processing system is so large and distributed that modifying the systems to accept the encrypted format of the data would be extremely expensive. FPE, however, modifies the encrypted data to fit the existing environment, which may be more cost-effective.

The FPE modes approved by NIST preserve the format by treating the data as numeral strings, rather than bit strings. The data must have a finite set of characters, such as the 10 decimal digits or the 128 ASCII characters. During encryption, the data undergoes a series of transformations. During each transformation, the numeral strings are converted to a bit string, transformed, and then converted back to a numeral string that fits the original format.

In March 2016, NIST approved two FPE modes, designated FF1 and FF3. The “FF” designator signifies that they are “format-preserving, Feistel-based encryption modes”. Feistel networks, created by Horst Feistel in 1970, were a key element in the Data Encryption Standard (DES) encryption algorithm, and have become a very popular building block in cryptographic algorithms. They have been extensively analyzed, and provide a high assurance of security.

Feistel Networks

Feistel networks are a sequence of permutations and encryptions carried out over a number of iterations called “rounds.” They can be built around a variety of encryption algorithms or cryptographic hash functions. The NIST standard requires the Feistel networks to use an approved block cipher with a block size of 128 bits. The 128-bit Advanced Encryption Standard (AES) block cipher is the only cipher

that currently meets this profile.

The FF1 and FF3 encryption algorithms each take three inputs: a numeral string to be encrypted, a key, and an optional tweak. Tweaks, which are described in greater detail in the next section, are used to introduce greater variability to the encryption.

Figure 1 shows one round of a generic Feistel network. The input data is split into two parts – a left side, L_n , and a right side R_n (where “n” is the round number). Each round consists of three steps: 1) a keyed function, F_k , called the round function, is applied to R_n ; 2)




Figure 1 – A single round of a Feistel network

the encrypted R_n is used to modify L_n ; 3) the original, unencrypted R_n and the modified L_n swap positions and are input to the next round. In other words, R_n becomes L_{n+1} , and the modified L_n becomes R_n . The number of rounds used will vary depending on the application.

The NIST specification includes some additional details for the FF1 and FF3 mode. R_n and L_n do not have to be the same length, which allows them to work with character strings containing an odd number of characters. The round function must be a suitable, approved block cipher, which currently means AES. Both modes also take an optional tweak as an additional input for the encryption and decryption algorithms. The tweak and the round number are concatenated to R_n before it is passed to the round function during each Feistel round.

Note also that the generic description does not specify how the encrypted R_n is combined with L_n . Many Feistel networks simply XOR R_n and L_n , but FPE needs to combine the streams in a way that preserves the format. If the input data is a 16-digit, base 10 credit card number, for example, the result of every round of the Feistel network must be a 16-digit, base 10 number. When the number is split into two even halves, each half is a string of 8 numerals, which can be interpreted as an integer less than 100,000,000. Both the FF1 and FF3 modes convert R_n and L_n to bit strings, permute the R_n bit string with the round function, add it to the L_n bit string, and reduce the result to a numeral string the same length as the unmodified R_n . At the end of the last round, L_n and R_n are concatenated to give an encrypted numeral string that is the same length as the original plaintext numeral string.

Tweaks

Because FPE is likely to be used in applications where the number of possible numeral strings is relatively small, both modes can take an additional input, called the “tweak” to introduce variability to the encryption. Unlike a cryptographic key, a tweak does not need to be secret.

¹ Dworkin, M. (2016, March). Recommendation for Block Cipher Modes of Operation: Methods for Format-Preserving Encryption. Retrieved April 5, 2016, from <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-38G.pdf>.

² This example shows the Base64 encoding of the ciphertext, which ensures printable characters. Otherwise, the ciphertext would mostly consist of unprintable characters.

The NIST standard uses 16-digit CCNs as an example. The first six digits of a CCN are the Issuer Identification Number (IIN) that identifies the bank that issued the card, the last digit is a checksum, and the rest identify a specific account. Many applications need to leave the first six digits (the IIN) and the last four digits unencrypted. If the middle 6 digits are encrypted, there are only 1 million different plaintexts. If you have a large enough database of CCNs, many of them will share the same six middle digits. When you use the same AES key to encrypt them, they will share the same ciphertext. If one is compromised, they are all compromised. In addition, they will be vulnerable to a dictionary attack. A determined hacker could, in theory, build a table matching the million different plaintexts with their corresponding ciphertexts, and use the table to recover CCNs.

If, however, the unencrypted digits are used to tweak the encryption, dictionary attacks become much more difficult. Even if two CCNs share the same six middle digits, they should not be exactly identical – the IIN or the last four digits must be different. The unencrypted digits can thus be used as a tweak to vary the encryption of the middle six digits, reducing the chance of a successful dictionary attack.

The tweak does not need to be a secret, and many types of information may be used as tweaks. In general, the information should be easily available to anyone with access to the plaintext, statically associated to the plaintext, and, ideally, associated only with that plaintext. A patient's name might be used to tweak a medical record number, for example, or the billing address for a credit card could be used to tweak a CCN.

Because the tweaks are statically associated with the plaintext, the variability they introduce is not random. Every time these modes encrypt the same plaintext/tweak pair, they will produce the same ciphertext. In contrast, many encryption algorithms, including the cipher-block chaining (CBC) mode of AES, use an initialization vector (IV) to provide variability. An IV is a randomly selected value that is changed every time a new plaintext message is encrypted. To decrypt the message, you need the ciphertext, the AES key, and the IV. When you use the AES-CBC mode, the algorithm returns two values: the ciphertext and the IV picked to create that encrypted value.

This random encryption, however, causes a referential integrity problem. Suppose you want to use a CCN to look up a record in a database, and the database uses a random encryption scheme to encrypt the stored CCNs. Every time the CCN is decrypted and re-encrypted, it will be stored as a different value, which means it cannot be used as an index in that database. Of course, you might be able to store and keep track of the IV used to create a particular encrypted value, but that is unlikely to work in legacy systems that require specific formats for data. FPE, a nonrandom encryption system, avoids this problem.

FPE Security

The security of the FPE modes derives from the security of its cryptographic building blocks -- Feistel networks and AES. The security of Feistel networks has been extensively studied since they were introduced in 1979. Michael Luby and Charles Rackoff proved that a Feistel network using a cryptographically secure pseudorandom function as the round function provides strong pseudorandom permutation security. More recently, John Black and Phil Rogaway published a description of an approach to FPE based on Feistel networks and proved

it was secure. Building on Luby and Rackoff's work, they determined that Feistel networks with at least four rounds are as secure as the round function, Fk, used in the algorithm. The FPE modes approved by NIST, which use 8 or 10 round Feistel networks should be as secure as the AES block cipher used as the round function.

Comparison of FF1 and FF3

Although the two modes of operation are generally similar, there are two major differences, as shown in Table 1. FF1 supports a wider range of data and tweak lengths, but FF3 requires fewer Feistel rounds. FF1 will be the obvious choice for encrypting strings longer than 57 decimal digits. It may also be a better choice when the message space is very small, as longer tweaks will provide better security. FF3 may be a better choice, for very small input strings that do not need large tweaks. It only requires eight Feistel rounds to the ten required by FF1, which means it is likely to be slightly faster. In any case, the tweak length will need to be chosen carefully to balance security and performance.

	FF1	FF3
Maximum Input Length	2 ³² decimal digits	57 decimal digits
Tweak Length	0-2 ³² bytes	64 bits
Number of Feistel Rounds	10	8

Table 1. Comparison of FF1 and FF3. Note that the maximum input length for FF3 will depend on the base, which may be anything from binary to base 16. Base 10 was chosen for this comparison.

Conclusion

As NIST's approval of FPE modes for government use indicates, FPE is emerging as a useful cryptographic tool. The primary advantage of FPE is that the encryption maintains the data format. Legacy systems, such as the complex systems that process financial transactions, cannot accommodate new data formats without expensive and potentially error-prone system changes. Classic block cipher encryption modes do not fit this legacy environment, but FPE does. In addition, as a deterministic encryption mode, FPE is suitable for use in relational databases. An SSN encrypted under FPE using an appropriate tweak, will always be encrypted to the same value. This means it can be used as a record index in its encrypted form. Finally, the FPE modes are based on proven, well-understood cryptographic building blocks – Feistel networks and the 128-bit AES block cipher – and therefore, FPE provides a high assurance of security.

NonStop supports two FPE solutions an HPE product from Voltage sold through Xypro and another one from comForte.



Karen Martin is an Information Security consultant and technical writer based in San Jose, CA.

¹ Luby, M. and Rackoff, C. How to construct pseudorandom permutations from pseudorandom functions. SIAM Journal of Computing 17, 2 (Apr. 1988)
² Black, J., & Rogaway, P. (2002, February 1). Ciphers with arbitrary finite domains. Retrieved January 14, 2016, from <http://web.cs.ucdavis.edu/~rogaway/papers/subset.pdf>



What is Identity-Based Encryption (IBE)?

Josh Lubliner > Software Engineering Manager > HPE

Identity-Based Encryption (IBE) is an alternative method of managing Public Key Infrastructure (PKI) for sending secure communications. Ok, but what does that mean? Alternative to what? Why is it a better alternative? Is it better? Should I use it?

Why is it an alternative to traditional PKI management? Let's say I want to send you an encrypted email message. In a traditional PKI system, I would follow this procedure:

1. I would tell you that I want to send you an encrypted message, so please generate a key pair.
2. You would generate a key pair, and send me the public key.
3. I would encrypt my message with the public key, and send you the resulting ciphertext
4. You would decrypt the ciphertext with your private key.

There are several major disadvantages to this procedure. First of all, you have to generate a key pair before I can send you a secure message. Secondly, I have to verify that the public key I received is the same one you sent; if an attacker manages to make me think that his key is actually your key, he can read all the encrypted messages I send to you. Finally, you now have the responsibility of securely storing this public key; if you lose it, you can no longer read any of the encrypted messages I've sent you in the past, or may send you in the future; and if it is stolen, the thief can read all your encrypted messages (at least the ones he can find or intercept).

In traditional PKI, there are a couple of ways to verify that I have the correct public key. One is to confirm with you personally; either we can physically meet and you can hand me your public key, or we can use some kind of out-of-band communication (say, I can call you on the phone) and you can tell me the hash of your public key, which I can verify with what I have. The problem with this solution is that it doesn't scale. If I want to send an encrypted message to one person, it works reasonably well, but if I want to send encrypted messages to a thousand people, or a hundred thousand people, it doesn't work at all.

However, there is one big advantage to this solution: you and I can have a secure conversation, and we don't have to trust anyone (well, ok, if we're using someone else's software to generate the keys, we have to trust them).

The second way traditional PKI deals with the problem of verifying public keys is through certificate authorities. When you generate your key pair, you can generate a Certificate Signing Request (CSR), and send it to some Certificate Authority (CA) that we both trust. They will sign your certificate (which contains your public key), and then I can verify your public key using the CA's certificate. This is how your browser verifies that the web site you're visiting is actually the one you thought you were visiting. Your browser contains certificates from a bunch of CAs that it trusts. The web site's certificate contains the domain name of the website; so if the name in the certificate doesn't match the URL of the website, your browser will warn you.

This solution scales much better (see: the web), but now we have to

trust a third party (the CA). If the CA signs certificates from people who don't own the domains they're requesting certificates for, then the system breaks.

So, how does IBE solve these problems? In an IBE system, if I want to send you an encrypted message, I follow a different procedure:

1. I go to the Key Server (more about this later) and request a public key for your email address.
2. I encrypt my message with this public key and email it to you.
3. You receive the email. If you don't already have the private key, you can go to the Key Server to get it. You will have to authenticate yourself to the key server; for example, the Key Server could email you a code, and you could supply that code to prove you are the person who has access to that email address.
4. You can now decrypt the email.

There are three major differences between this process and the traditional one. First of all, I don't have to wait for you to create a key pair before I can send you an encrypted message! The Key Server derives a key pair based on an internal secret and the email address itself. This guarantees both that for a given email address, the key pair is always the same, and that it isn't possible to guess what the key for a particular email address will be.

So what if I send you an encrypted email and you have no idea how to decrypt it? Well, I can simply send some (unencrypted) instructions along with the encrypted message, explaining how to contact the key server, authenticate yourself, retrieve the private key, and decrypt the message.

The second major difference is that I don't have to verify that the public key I received is the correct one. The Key Server takes care of that. (Of course, I have to be sure that the email address I'm sending to is actually your email address... but that's always true!)

This scales easily. I can send encrypted messages to 100,000 different recipients. I don't have to send requests to 100,000 different people to generate key pairs, and I don't have to meet with 100,000 different people to verify their public keys.

The third difference is that you don't have to worry about securely storing your private key when you're done decrypting the email. You can just discard it, and request it from the Key Server again the next time you need it. No worries about stolen laptops, lost thumbdrives, or crashed hard drives.

The only downside to this process is that we both have to trust the Key Server. After all, the Key Server knows how to derive key pairs for any email address, so anyone who controls the Key Server could decrypt any message. This means that IBE (at least in this form) is only useful in certain environments.

One environment where this works particularly well is in a corporate email system. When we use corporate email, we expect the corporation to be able to read our emails; in fact the corporation may even have retention policies that require it. So IBE is a good choice in this case. And in corporate environments, you probably already have a good way to authenticate users (Active Directory, LDAP, etc.), so it is easier to make sure

users only get their own private keys.

But IBE can do more than just encrypt emails! You probably have applications in your enterprise that encrypt sensitive data; for example, protecting personally identifiable information (PII), such as names and birthdays; or health information; or payment information, such as credit card numbers. Each of these applications has to deal with the cryptographic keys it uses to encrypt and decrypt this information. Those keys have to be stored securely, which can be difficult, error-prone, and costly.

But if your application has an identity, it can simply request the keys it needs from the Key Server (which is already conveniently sitting in your corporate environment so that your users can send encrypted emails to each other). When it needs to encrypt or decrypt some data, it simply sends its identity to the Key Server (along with some authentication information), gets the key, performs the cryptographic operation, and then discards the key, just like you did after you decrypted your email. This means a more secure environment, and possibly even reduced audit scope, depending on the application.

Perhaps your application doesn't use key pairs; perhaps it uses symmetric keys, such as AES keys. This is known as IBSE, or Identity Based Symmetric Encryption. No problem, you just expand the Key Server protocol so you can tell it what kind of key you want. It can generate any kind of key for a given identity, and when we take a closer look at the Key Server, you'll see why.

The Key Server is an application that sits on a server somewhere in your corporate environment, listening for key requests. These requests could be in the form of a web service request, or a proprietary protocol. When a request comes in, it determines what keys are being requested. That is, what the identity is; and whether public, private, or symmetric keys are desired, and what algorithm and key strength they keys should be. It performs authentication on the user making the request, and makes sure that user has authorization to receive the keys it is requesting.

Then the Key Server has to get the requested keys. This is where things are very different from traditional PKI – in a traditional system, a user would, at this point, either retrieve the key from secure storage, or generate a new key if they didn't already have one. But in this case, the IBE Key Server will derive the keys. The Key Server uses a Key Derivation Function (KDF), which uses a one-way hash function to combine the identity itself with an internal secret. A different KDF is needed for each type and size of key users can request. Also, a different secret should be used for each one. The Key Server can then return the derived keys, and then discard them. The Key Server doesn't have to store the keys – it can just regenerate them the next time they are requested. This is why IBE infrastructure is sometimes referred to as stateless key management. The Key Server doesn't have to maintain a database of keys, which would be another threat surface you'd have to deal with.

So what about these internal secrets? Those, of course, are very important and must be treated very carefully. Anyone who has access to the secrets (and knows your KDFs) can generate keys for any identity they want. Luckily, it is just a small number of secrets, so they can be stored, for example, in a Hardware Security Module (HSM) for improved security and compliance. Depending on the environment and security requirements, it could also be stored on a smart card or a Trusted Platform Module (TPM), which is a secure storage subsystem which is now available in many commercial computers.


Let's think for a minute about how much infrastructure we can avoid building and maintaining by using the idea of derived keys, rather than generated keys.

With randomly generated keys, each time you generate a key, you have to store it. Loss of the key would mean that messages or data encrypted

with that key could never be decrypted. This is generally catastrophic; you're either losing valuable customer data, or violating retention policies. So you have to keep a database which associates the identity with the key. That database has to be kept highly secure, while at the same time accessible so that keys can be retrieved by the various applications that need them. The database may have to be replicated to different, geographically separated data centers, and kept synchronized. This is complicated and expensive.

With derived keys, you just have to generate the internal secrets once, and then securely transfer them to your different data centers once. This is a function which is often easily implemented by HSMs, but even if you're not using HSMs, since you only have to do this once (not continuously), it's ok if you have to go through some kind of security ceremony.

By using derived keys, and a central server to derive and provide them, you relieve the users and applications which use cryptographic keys of the burden of generating and storing them. This can simplify application development, help to increase interoperability, and make applications more secure.

IBE, IBSE, stateless key management, and derived keys. Keep these concepts in mind when thinking about how to create PKI in your organization. 

Josh Lubliner is a Software Engineering Manager at HPE in the Data Security group, helping to create tools to secure Enterprise data. In his 25-year career he's worked in fields as diverse as architectural software, online advertising, and fraud detection for online banking. He can be reached at joshua.lubliner@hpe.com

Did You Know?

HPE NonStop Enterprise Division (NED) is pleased to announce the general availability of a new product called NonStop Application Direct Interface (NSADI). NSADI (developed under the program name YUMA) enables low latency and high speed connectivity between applications running on HPE Integrity NonStop X and Linux systems using the Remote Direct Memory Access (RDMA) over InfiniBand technology. Using NSADI you can connect a NonStop X system to multiple Linux servers over an InfiniBand network to establish direct connectivity between applications running on these systems without having to go through the CLIMs.

NSADI opens up new possibilities for architecting modern, mission-critical solutions using closely knit applications running on Linux and NonStop systems. These could be targeted for deployments in financial services, retail, data-analytics, patient care, IoT and several others where the industry needs highly innovative but mission critical solutions.

Back For More...

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

My travels have taken me to gatherings of the NonStop community both small and large. Whether it's been the drizzle of rain in London or the blistering heat of Las Vegas, there has been a presence of NonStop users and vendors in numbers that have impressed me and on each occasion I have run into NonStop stakeholders I haven't met before. When this has happened they have been surprised by the size of the vendor community and the overall mix of product developers, consultants and services providers and yes, indeed, the media.

While it has been participation in the Regional User Group (RUG) meetings that have dominated my schedule for the past couple of months, I have also had the good fortune to attend the Partner Technical Symposium in Palo Alto put on by HPE NonStop Product Management as well as the big-tent marketing event in Las Vegas, HPE Discover 2016. Of them all, the one that really stood out was the Symposium and for two very compelling reasons.

The first being that executing on an idea coming from the folks at HPE isn't always a given no matter how good an idea it may be. There will always be competing priorities within a company the size of HPE so simply talking about it meant little if the event itself never took place. Karen Copeland pulled it off and made true on her promise! The second was just how big a turn out by the NonStop vendor community there was when the day eventually arrived.

I have always been a softie when it comes to meetings of key NonStop stakeholders. While Palo Alto is among the most expensive options when it comes to holding a meeting I gave up a weekend to drive to our niece's place in Half Moon Bay just the other side of Interstate 280, a place that is bathed in fog for much of the year. Seeing anything at all particularly in the mornings or the early evenings is always challenging but this was the only occasion when you could say visibility was poor.

No such comments could be made about the Symposium, of course. Having had to sign confidentiality agreements, the NonStop product managers gave us as much information as they had at the time and while I will not go into specifics here as I am still under that confidentiality agreement, it wasn't so much the new features being described or the dates being floated but rather the openness and insightfulness on display that impressed me the most. No, visibility was terrific and over coffee, the conversations among the vendors quickly turned to just how impressed we all were that HPE NonStop was holding such an event.


At HPE Discover I surprised several parties when I talked about the ecosystem that has grown up around NonStop systems. Apart from the initial shock that there were companies focused on products, consulting, services and yes, even the media, I couldn't help but notice their focus change. With the interest peaked following this particular sound bite, I would then steer the conversation in


the direction of a series of "did you know" observations only to see each party lean in closer to hear more. The NonStop ecosystem is a compelling argument in favor of why NonStop continues to be the premier platform for running mission critical applications needing near real time performance. And with uninterrupted availability! All attributes every member of the NonStop community knows by heart but often isn't understood by everyone interacting with the bigger HPE, including many folks within HPE itself.

It is becoming well known just how big an investment HPE is making in NonStop. At last year's NonStop Technical Boot Camp the rumors were that HPE had invested as much as \$200 million on just porting NonStop to the Intel x86 architecture and I have to believe the project to develop a virtual NonStop option didn't come cheap. However, what is often overlooked is the collective investment being made by the NonStop vendor community.

Of course, such figures aren't tabulated by any independent agency but the numbers must be high and sitting in the room at the Symposium it didn't escape my attention that this occasion represented a sizeable investment for the vendors both in time and yes, money. It was hard to miss just how many CEOs were present in the room and I have to believe that the turn out surprised some at HPE as well. On more than one occasion I saw the head of one HPE executive or another pop through the door even if it was just for a few minutes. I saw vendors present about whom I knew little even as the vendors I did know all seemed to have made it!

Holding the Symposium and having it completely separate from RUGs and big-tent events was very important and was a major contribution to its success. After all, there are times when HPE NonStop folks need to be open and there are times when they have to remain silent. And I believe the NonStop community understands all of this. However, for the past couple of years in my discussions with my clients it's been clear that a judicious amount of second guessing has been at work and for the first time in a very long time, the roadmaps have been fleshed out to a degree that NonStop vendors are no longer lacking the information they need. An ecosystem after all is about shared knowledge and a level of cooperation that leads to sustainable growth and if this Symposium is any indication, there is knowledge being shared and cooperation being fostered so yes, growth will not be far behind.

The traditional fog of Half Moon Bay may have reduced visibility to just a couple of feet but, then again, that is all part of the location's charm. Not having visibility into product plans and timeframes holds no charm whatsoever but with as successful a Symposium as has just been held, few members within the larger NonStop ecosystem can claim ignorance of NonStop any longer and perhaps, that is the right yardstick to be measuring success when it comes to nurturing a community every bit as committed to the success of NonStop as HPE itself. 



XYGATE® SecurityOne™

Security Intelligence and Analytics
for HPE Integrity NonStop™ Servers



Visibility

Faster Threat Detection

Improved Risk Management

Differentiate Noise from Actionable Incidents

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



Reduce Mean
Time To Detection

Learn more at
xypro.com/SecurityOne

XYPRO®
Mission Critical Security

The guiding light for your mission critical business

Improve your NonStop'ness. Better always on!



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

comForte „better always on“ solutions help you gain ...

Better Infrastructure

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

Better Security

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

Better Applications

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

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

www.comforte.com

com.forte®
better always on