



Messaging middleware audits and remediation for Data Security Compliance: The new frontier of PCI, SOX, HIPAA and Regulatory Governance

An Evans Resource Group Whitepaper for the Chief Risk Officer, Chief Security Officer, and Internal Auditor

Published August 2009

© Evans Resource Group, 2009

Evans Resource Group 575 Madison Avenue, Suite 1006, New York, N.Y. 10022

Tel 212.937.8443

Fax 212.937.8422

info@evansrg.com

www.evansresourcegroup.com

Why This White Paper Will Be Worth Your Time

We've all heard "an email is like a postcard", now let us introduce you to messaging middleware... messaging middleware is like industrial strength email and WebSphere MQ is the de facto standard for messaging middleware, with over 10,000 installations globally. Other vendor systems include MSMQ, Tibco and SonicMQ. When IBM created MQSeries in the early '90s, MQ administrators were not concerned about security regulations. Instead, they focused on installing and configuring the product in the fastest and most straightforward way that would provide the required connectivity for applications, without placing any constraints on usage. As use of MQSeries grew, it was still common to implement it in an "out of the box" configuration, which is designed for ease of implementation, and does not have any security constraints. Even though MQSeries provides for a secure configuration, this has not traditionally been the primary focus of MQ Administrators.

However, the result of not implementing security within messaging middleware can lead to severe consequences. As security concerns are now paramount in the marketplace, messaging middleware networks are coming under increased scrutiny, and companies will now fail PCI audits due to non-configuration or mis-configuration of security constructs, and these issues will also apply to the failure of SOX and HIPAA audits as auditors are now being educated about messaging middleware. Security regulations such as HIPAA were enacted in 1996, SOX in 2002 and PCI in 2006, all later than the initial growth of messaging middleware.

The default configuration of WMQ allows anonymous administrative access to the WMQ Command Server (console), permitting arbitrary remote code execution abilities to anonymous users across the network. The implication of this is serious and transparent to anyone in the security community.

It has been determined that most current messaging middleware installations (over 90%) are not configured to properly utilize built-in product functionality that reduces and/or eliminates security threats. In addition, messaging middleware requires additional programming of security exits to ensure that all PCI compliance requirements are met.

...BUT THE PROBLEM IS NOT A THEORETICAL ONE

Card companies may impose fines on their member banking institutions when merchants are found to be non-compliant with PCI DSS. Acquiring banks may in turn contractually oblige merchants to indemnify and reimburse them for such fines. Fines could go up to \$500,000 per incident if data is compromised and merchants are found to be non-compliant. In the worst case scenario, merchants could also risk losing the ability to process customers' credit card transactions.

Businesses from which cardholder data has been compromised are obliged to notify legal authorities and are expected to offer credit-protection services to those potentially affected. There may be other consequences besides the fines. Cardholder data loss, whether accidental or through theft, may also lead to legal action being taken by cardholders. Such a step will result in bad publicity, which may in turn lead to loss of business.

Compliance with the U.S. Public Company Accounting Reform and Investor Protection Act of 2002 (the Sarbanes-Oxley Act) requires that IS organizations work with other departments — mostly finance and legal — in almost-unprecedented ways. However, a recent Gartner survey conducted with 75 senior regulatory compliance managers at U.S.-based public companies revealed that this is not happening. Thirty-seven percent of the companies represented have no involvement from their IS departments in Sarbanes-Oxley compliance processes. Finance is the area contributing the most people toward compliance activities, with an average of six people participating in Sarbanes-Oxley-related activities. Operations groups contribute an average of three. So where is the IS group?

This apparent lack of IS department involvement is alarming for several reasons. Financial documentation and controls are heavily dependent on IT systems. If IT systems are not being included in the audit process, there is a risk that companies will not be Sarbanes-Oxley-compliant. The Big Four auditing firms, on which most large companies are relying to help them through the compliance process, are recommending the use of the Committee of Sponsoring Organizations of the Treadway Commission (COSO) Risk Management Framework. This document is designed to instill "risk and control consciousness," and is a model for discussing and comparing risk management and internal controls. The COSO Framework contains specific IT-related advice. It is therefore surprising that the financial personnel who are driving the Sarbanes-Oxley efforts within their companies have not, in every case, involved their IS departments in the compliance process.

Our conclusions regarding messaging middleware relate to the education of the IS departments and their duties as they are not now involved in Sarbanes-Oxley, specifically:

- Messaging middleware is now a strategic target of hackers
- Trillions of dollars of transactions flow through messaging middleware each week
- Administrative access to messaging middleware is a SEV-1 finding on a data security audit.

Although we do not deny that Sarbanes-Oxley is, and must be, a finance-driven initiative, IS departments also have a major role in compliance.

In this part of the Sarbanes-Oxley Spotlight, we outline the reasons and methods that your IS group should become familiar with to ensure full Sarbanes-Oxley compliance. Established systems must be audited and tested. CIOs, if not already involved with Sarbanes-Oxley committees at their companies, should gain representation there immediately. Your company cannot be compliant without you.

ABOUT THIS WHITEPAPER

This white paper discusses key issues around WebSphere MQ, its security requirements for PCI compliance, SOX and HIPAA, and outlines a program of assessment and remediation resulting in a PCI, SOX and HIPAA compliant MQ network. This white paper also briefly discusses Evans Resource Group, the group behind this white paper, and their relevant offerings.

INTRODUCTION

Messaging middleware audits and remediation for Data Security Compliance: The new frontier of PCI, SOX, HIPAA and Regulatory Governance

An Evans Resource Group Whitepaper for the Chief Risk Officer, Chief Security Officer, and Internal Auditor

This whitepaper explains the history, relationship and future of messaging middleware and regulation, where the gaps are, and how to remediate them to secure your business.

Why do you care? Many auditors focus on what they know. They may be comfortable with routers and network segmentation. However, enterprise messaging middleware is a complex component of the network. Not only that, the control over the enterprise that can be directed by messaging middleware that is not properly secured is profound. Your auditors may have declared you secure; however if your messaging middleware has not even been reviewed we would ask you to reconsider that assessment.

First, with Hannaford Brothers, and now with Heartland, the “trusted internal network” is the new frontier of data theft. Enabling SSL is great for protecting messages on the wire but if administrative access is left exposed, the attackers can disable SSL or skip sniffing traffic entirely and instead just browse the messages passing through the message queues. The answer to this is not redoubling security at the perimeter. The answer is to apply meaningful controls at the messaging layer. An auditor familiar with your messaging technology would seem to be a valuable asset if the goal is to actually assess security and not merely to pass the audit.

Hannaford was reportedly the first breach of data in transit. Heartland was the biggest card data breach ever. If the bad guys are only up the H’s, what is in store for firms in the I – Z range? We prefer to think that it’s strict auditing of the messaging layer and not massive name changes to monikers starting with A – G! One of these two alternatives actually could make a difference. The other is about as effective as what we have in place today.

In light of the information in this whitepaper you will be provided with data points to the understand the significance of securing ALL aspects of the network, especially the messaging middleware, and that there is a solution for IBM’s WebSphere MQ, and that this is the only PCI, SOX and HIPAA compliant messaging middleware solution available at this time.

Table of Contents

- Introduction
- The Audience for this white paper
- What is PCI?
- What is SOX?
- What is HIPAA?
- PCI Requirements
- History of PCI
- PCI Controversies and Exposures
- What does this tell us?
- What is messaging middleware?
- What is IBM's WebSphere MQ?
- WMQ Features
- WMQ Communication
- Who uses WMQ?
- What is the issue?
- The Importance of Understanding the Requirements
- How Assessment Works
- The OOMA Report
- What is being done?
- Conclusion
- About Evans Resource Group

THE AUDIENCE FOR THIS WHITEPAPER

Messaging middleware is the glue that holds applications together and WebSphere MQ is the de facto standard for messaging middleware, with over 10,000 installations globally. When IBM created MQSeries in the early '90s MQ administrators were not concerned about security regulations. They focused on installing and configuring the product in the fastest and most straightforward way that would provide the required connectivity for applications, without placing any constraints on usage. As use of MQSeries grew, it was still common to implement it in an "out of the box" configuration, which is designed for ease of implementation, and does not have any security constraints. Even though MQSeries provides for a secure configuration, this has not traditionally been the primary focus of MQ Administrators.

The security people implemented for their MQ networks fifteen years ago was based on assumptions that are not true any longer. Back then we had obscurity on our side. MQ was new and it was just plumbing. Not only did nobody outside the WMQ community know anything about it, but at the time nobody was attacking "network plumbing". The systems of record and the user-facing systems were, and still are, the primary attack targets. But there is evidence that things are changing on both counts.

Messaging in general has gone mainstream and WMQ holds around 90% of the market. Many more people now know about WMQ, including the bad guys. And even though MQ is still considered plumbing, the plumbing itself is now understood to be a valuable target.

However, the result of not implementing security within WMQ can lead to severe consequences. As security concerns are now paramount in the marketplace, MQ networks are coming under increased scrutiny, and companies are now failing PCI audits due to non-configuration or mis-configuration of MQ security, and these issues also apply to failing SOX and HIPAA audits as well.

The default configuration of MQ allows anonymous administrative access to WMQ, permitting arbitrary remote code execution abilities to anonymous users. The implication of this is serious and transparent to anyone in the security community.

Security regulations such as HIPAA were not enacted until 1996, SOX until 2002 and PCI was not launched until 2006. This newly realized threat to data security has led to the launch of this program of audit and remediation for messaging middleware. It is noteworthy to highlight that not only is IBM's WebSphere MQ exposed by this but also the likes of Tibco, SonicMQ, and MSMQ, etc. IBM's WebSphere MQ will provide the only PCI, SOX and HIPAA Compliant messaging middleware solution on the market.

It has been determined that most current MQ installations (over 90%) are not configured to properly utilize built-in product functionality that reduces and/or eliminates security

threats. In addition, WMQ requires additional programming of security exits to completely ensure that all PCI compliance requirements are met.

Many companies have not invested in securing the messaging network for many years and yet have still passed audits up to now. They thought their security was in good shape. They had just never considered it in the context of PCI requirements up to this point.

From a PCI perspective, if your clients have payment applications running on an undifferentiated network, the audit scope covers the entirety of that network. The purpose of segmenting the network is to isolate the PCI applications onto a much smaller network with strict controls on ingress and egress. In addition to providing better security, this greatly reduces the scope of the audit. If you think about what it would mean to audit every single device and user on your intranet versus auditing a small and self-contained network segment, it's easy to see why this is such a big deal. The audit would be impractical with a large undifferentiated network. For a SOX or HIPAA audit this scope does not apply, it is the entire network that is audited.

Most of our clients will wait until they failed the audit to remediate their messaging middleware security. Admittedly, they were unaware that their security would be a problem because they had done what, in any situation other than regulatory compliance, would be considered a *very* good job of securing messaging middleware. Problem is that in all the prior audits, the messaging middleware network pretty much flew under the radar. For the last fifteen years, messaging middleware has just been plumbing. If you were a messaging middleware administrator during that time, chances are good your company could have had an audit every year and you knew nothing about it because the messaging middleware network was always out of scope. No auditor would have come knocking at your door, no emails from the enterprise security team would have showed up in your mailbox. If you are a messaging middleware admin and you knew about the audits at all, chances are it was because you had lunch with someone who was directly impacted.

When we ask clients why they have not enabled security we hear the following:

“Well it’s on our trusted network...”

“We are PCI Compliant...”

“Nobody does this...”

Our response to this is: no one would enable a web server without security and we find the exact opposite to be true with messaging middleware. For example, no one enables SSL for administrators to remotely administer the MQ queue managers, even though this is a fundamental security exposure.

In light of the information in this whitepaper there is an increasingly good chance the next several audits your firm performs will include the messaging middleware network in scope, and you are likely to conclude Sev-1 findings against it. What that means is that many installations, even those with relatively good security, that have previously passed their audits, are going to fail an audit. You will also receive baffled responses from your

© Evans Resource Group, 2009

MQ team as to why they passed previously, but not this time. We will provide you with the answers you need to educate and train them and prevent this from happening again.

Our purpose is to educate you, and help you to become PCI, SOX and HIPAA Compliant with messaging middleware depending on your requirement. This paper serves as background as to how the problem was created over time, how regulations impact it and how to solve it.

The following pages will address a broad audience intended for Chief Risk and Security Officers, Internal Auditors, CIOs, QSAs, and laypersons. We have taken the liberty to include background referential data on regulatory, messaging middleware history and other pertinent details to provide the complete story for each interested party.

What is PCI DSS?

The **Payment Card Industry Data Security Standard** is a worldwide information security standard assembled by the Payment Card Industry Security Standards Council (PCI SSC). The standard was created to help organizations that process card payments prevent credit card fraud through increased controls around data and its exposure to compromise. The standard applies to all organizations which hold, process, or pass cardholder information from any card branded with the logo of one of the card brands.

The standard is maintained by the Payment Card Industry Security Standards Council, which maintains both the PCI DSS and a number of other standards, such as the Payment Card Industry PIN Entry Device security requirements (PCI PED) and the Payment Application Data Security Standard (PA-DSS).

Validation of compliance can be performed either internally or externally, depending on the volume of card transactions the organization is handling, but regardless of the size of the organization, compliance must be assessed annually. Organizations handling large volumes of transactions must have their compliance assessed by an independent assessor known as a Qualified Security Assessor (QSA), while companies handling smaller volumes have the option of self-certification via a Self-Assessment Questionnaire (SAQ). In some regions these SAQs still require signoff by a QSA for submission.

Enforcement of compliance is done by the bodies holding relationships with the in-scope organizations. Thus, for organizations processing Visa or Mastercard transactions, compliance is enforced by the organization's acquirer, while organizations handling American Express transactions will deal directly with American Express for the purposes of compliance. In the case of third party suppliers such as hosting companies who have business relationships with in-scope organizations, enforcement of compliance falls to the in-scope company, as neither the acquirers nor the card brands will have appropriate contractual relationships in place to mandate compliance. Non-compliant companies who maintain a relationship with one or more of the card brands, either directly or through an acquirer, risk losing their ability to process credit card payments, and being audited and/or fined.

PCI Requirements

The current version of the standard (1.2) specifies 12 requirements for compliance, organized into six logically related groups, which are called "control objectives."

Control Objectives	PCI DSS Requirements
Build and Maintain a Secure	1. Install and maintain a firewall configuration to protect

© Evans Resource Group, 2009

Network	cardholder data
	2. Do not use vendor-supplied defaults for system passwords and other security parameters
Protect Cardholder Data	3. Protect stored cardholder data
	4. Encrypt transmission of cardholder data across open, public networks
Maintain a Vulnerability Management Program	5. Use and regularly update anti-virus software on all systems commonly affected by malware
	6. Develop and maintain secure systems and applications
Implement Strong Access Control Measures	7. Restrict access to cardholder data by business need-to-know
	8. Assign a unique ID to each person with computer access
	9. Restrict physical access to cardholder data
Regularly Monitor and Test Networks	10. Track and monitor all access to network resources and cardholder data
	11. Regularly test security systems and processes
Maintain an Information Security Policy	12. Maintain a policy that addresses information security

History of PCI and Data Security Standards

PCI DSS originally began as five different programs: Visa Card Information Security Program, MasterCard Site Data Protection, American Express Data Security Operating Policy, Discover Information and Compliance, and the JCB Data Security Program. Each company's intentions were roughly similar: to create an additional level of protection for customers by ensuring that merchants meet minimum levels of security when they store, process and transmit cardholder data. The Payment Card Industry Security Standards Council (PCI SSC) was formed, and on December 15th, 2004, these companies aligned their individual policies and released the Payment Card Industry Data Security Standard (PCI DSS).

In September 2006, the PCI standard was updated to version 1.1 to provide clarification and minor revisions to version 1.0.

PCI is one of the multiple data security standards that have emerged over the past decade; BS7799, ISF Standards, Basel II, Gramm-Leach-Bliley Act (GLBA), Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act of 2002.

Version 1.2 was released on October 1, 2008. Version 1.1 "sunsetted" on December 31, 2008. V1.2 did not change requirements, only enhanced clarity, improved flexibility, and addressed evolving risks/threats.

The exposure what we will address in this white paper applies to all the data security standards listed above.

PCI Controversies and Exposures

It is suggested by some IT security professionals that the PCI DSS does little more than provide a minimal baseline for security.

"The fact is you can be PCI-compliant and still be insecure. Look at online application vulnerabilities. They're arguably the fastest growing area of security, and for good reason — exposures in customer-facing applications pose a real danger of a security breach." - Greg Reber

Still others believe that PCI DSS is a step toward making all businesses pay more attention to IT security, even if minimum standards are not enough to completely eradicate security problems.

"Regulation--SOX, HIPAA, GLB, the credit-card industry's PCI, the various disclosure laws, the European Data Protection Act, whatever--has been the best stick the industry has found to beat companies over the head with. And it works. Regulation forces

companies to take security more seriously, and sells more products and services." - Bruce Schneier

Companies have had security breaches while being registered as PCI DSS compliant. In 2008 one of the largest payment service providers, Heartland Payment Processing Systems, suffered a data breach which has been estimated by some as exceeding one hundred million card numbers. Other notables include the Hannaford Brothers and the Okemo Mountain Resort, each of which was PCI compliant. It has been noted that this may be an indication of the limits of a snapshot certification; the evaluation cannot ensure that the target company will maintain the good practices seen in an audit. This explanation does not, however seem to explain the compromise of merchants such as Hannaford Bros Co., which received its PCI DSS compliance certification one day after it had been made aware of a two-month long compromise of its internal systems.

The definition of *compliant* has also been open to interpretation, especially regarding how temporary such a declaration might be. Declaring a company *compliant* appears to have some temporal persistence, yet the PCI Standards Council General Manager, Bob Russo, indicates that liabilities could change depending on the state of a given organization at the point in time when an actual breach occurs.

Similar to other industries, a secure state could be more costly to some organizations than accepting and managing the risk of confidentiality breaches. However, many studies have shown that this cost is justifiable.

WHAT DOES THIS TELL US AS AUDITORS?

We are responsible for data security. We must understand the technology behind the regulations that lead to compliance, and as that technology matures and changes we must reach out and bring security to the data network. It is not enough to say that we secure the perimeter. It has to really be secure to the letter of the standard. To that intent we would like to focus your attention on messaging middleware, the glue that holds together applications and databases and allows for message delivery across the network.

SOX AND HOW IT APPLIES TO MESSAGING MIDDLEWARE DATA SECURITY

One key requirement of SOX is that execs are required to certify under penalty of law that their financial data are true and accurate. For example, most corporations don't let just anybody in the company update the financial records. There are processes in place to limit access to these records, and checks and balances to limit the ability of any one person to make fraudulent entries. If these controls were not present, there would be no accountability and therefore no basis on which to claim that the data are indeed true and accurate.

It's one thing to certify that the aggregate financial data are true and accurate, but what about the actual transactions from which revenue flows? Suppose that it was common practice to let anybody in the company, from the CEO on down to the maintenance and housekeeping staff, access, modify and update customer transactions? Would it be accurate then to say that there were sufficient controls and accountability to satisfy SOX? We would say definitely not. Yet, this is the situation that we encounter in the vast majority of cases in our roles as WebSphere MQ security consultants.

Although WebSphere MQ can be configured to restrict administrative access, in most cases, it is not. The result is that anyone with an IP route to the queue manager (a queue manager is a running instance of WebSphere MQ) can anonymously connect and access any message flowing through the system, delete messages, update them or inject rogue messages into the system. This situation is what we usually find when performing security assessments. The queue manager is completely exposed to the entire intranet. Occasionally we find an installation where anonymous connections have been restricted. In nearly all of those cases though, legitimate users and applications are granted administrative access. This still falls far short of what would be considered effective internal controls.

What is at Risk and Why?

You might be wondering what exactly the exposure is. Message queuing is the glue that connects applications and is ubiquitous in corporate IT shops. Messages may be ATM transactions, patient health data, insurance claims, card payment transactions, travel bookings, fleet dispatch orders or just about anything else. The risk is in either stealing the data or in executing rogue transactions. In addition, WebSphere MQ provides its administrators a function to remotely execute commands on the MQ host server. When ordinary users and applications are inappropriately granted administrative access, they inherit this remote code execution capability on the servers which host the company's most critical business applications.

There are a number of factors which give rise to and perpetuate this problem. Primary among these is that WebSphere MQ security is not well understood. By default a remote connection will pass the user's local ID to the queue manager for authorization. This is

© Evans Resource Group, 2009

widely used as the basis for very granular and often complex security models. What is not commonly known however is that the user can optionally choose to pass in any arbitrary ID, including an administrative one. WebSphere MQ can be configured to authenticate these IDs that are passed in but usually it is not.

Because the ID passed in appears to have been authenticated, there is a general misconception that it has been and can be trusted. In fact, the exact opposite is true. The result is that the companies most exposed are usually the ones that are most confident in their messaging security. They certify the veracity of their financial data simply because they don't know any better.

How to Tell if You Are Impacted

Wondering how good your MQ security is? Try this thought exercise. Ask around to see what the opinion of the MQ security is among the administrators, project managers, architects and other stakeholders. If you find that the company is highly confident in the MQ security, chances are there's trouble brewing.

If the company knows the network is wide open, chalk up some points for self-awareness. The company has either accepted the risk or is in the process of mitigating it.

On the other hand, if the company has deep security skills the confidence will probably not be as high, especially with a large network. The larger the network, the more likely something is missed. There is also the notion that the more you know about information security, the less confident you are. It's an arms race between the white hats and the black hats and you never know how where you stand until you are breached. A shop with deep security skills will have a healthy skepticism and a certain amount of doubt about the security of all their systems, including the messaging middleware. They will have a program in place to subscribe to security updates for all their products, automated scanning, incident response plans and continuing education. When someone tells me they think the messaging network is "fairly secure" or otherwise qualifies their opinion, I have some hope.

But if you find a high level of confidence in the security of the MQ network there are only two possibilities. One is that the company has invested in training the administrators, funded the necessary security hardening, performs regular scans and has some reason to believe the job was done correctly. The other is that the network is wide open and the company does not know how to assess it. Since the first possibility is something that rarely happens, there is a very strong correlation between high confidence and poor security.

If you are one of the executives who signs off on the SOX certification for your company, it's time to check your messaging network. If security has not been enabled correctly, the SOX certification may not be worth the paper it's written on.

WHAT IS HIPAA?

The Health Insurance Portability and Accountability Act (HIPAA) was enacted by the U.S. Congress in 1996. According to the Centers for Medicare and Medicaid Services (CMS) website, Title I of HIPAA protects health insurance coverage for workers and their families when they change or lose their jobs. Title II of HIPAA, known as the Administrative Simplification (AS) provisions, requires the establishment of national standards for electronic health care transactions and national identifiers for providers, health insurance plans, and employers. This is intended to help people keep their information private, though in practice it is normal for providers and health insurance plans to require the waiver of HIPAA rights as a condition of service.

The Administrative Simplification provisions also address the security and privacy of health data. The standards are meant to improve the efficiency and effectiveness of the nation's health care system by encouraging the widespread use of electronic data interchange in the U.S. health care system.

HIPAA's Security Rule and its application to Messaging Middleware

The Final Rule on Security Standards was issued on February 20, 2003. It took effect on April 21, 2003 with a compliance date of April 21, 2005 for most covered entities and April 21, 2006 for "small plans". The Security Rule complements the Privacy Rule. While the Privacy Rule pertains to all Protected Health Information (PHI) including paper and electronic, the Security Rule deals specifically with Electronic Protected Health Information (EPHI). It lays out three types of security safeguards required for compliance: administrative, physical, and technical. For each of these types, the Rule identifies various security standards, and for each standard, it names both required and addressable implementation specifications. Required specifications must be adopted and administered as dictated by the Rule. Addressable specifications are more flexible. Individual covered entities can evaluate their own situation and determine the best way to implement addressable specifications. The Health Information Technology for Economic and Clinical Health Act (HITECH Act), enacted as part of the American Recovery and Reinvestment Act of 2009, imposes notification requirements on covered entities, business associates, vendors of personal health records (PHR) and related entities in the event of certain security breaches relating to protected health information (PHI). The U.S. Department of Health and Human Services (HHS) issued guidance on the subject; HHS and the Federal Trade Commission (FTC) are working to harmonize their respective regulations and are seeking public comment with a view to issuing interim final regulations by August 17, 2009, the deadline imposed by the HITECH Act. The standards and specifications are as follows:

- ***Administrative Safeguards*** – policies and procedures designed to clearly show how the entity will comply with the act

- Covered entities (entities that must comply with HIPAA requirements) must adopt a written set of privacy procedures and designate a privacy officer to be responsible for developing and implementing all required policies and procedures.
- The policies and procedures must reference management oversight and organizational buy-in to compliance with the documented security controls.
- Procedures should clearly identify employees or classes of employees who will have access to electronic protected health information (EPHI). Access to EPHI must be restricted to only those employees who have a need for it to complete their job function.

This requirement calls for the separation of duties within development and production environments within messaging middleware.

- The procedures must address access authorization, establishment, modification, and termination.

This requirement must address the requirement for data at rest.

- Entities must show that an appropriate ongoing training program regarding the handling of PHI is provided to employees performing health plan administrative functions.
- Covered entities that out-source some of their business processes to a third party must ensure that their vendors also have a framework in place to comply with HIPAA requirements. Companies typically gain this assurance through clauses in the contracts stating that the vendor will meet the same data protection requirements that apply to the covered entity. Care must be taken to determine if the vendor further out-sources any data handling functions to other vendors and monitor whether appropriate contracts and controls are in place.
- A contingency plan should be in place for responding to emergencies. Covered entities are responsible for backing up their data and having disaster recovery procedures in place. The plan should document data priority and failure analysis, testing activities, and change control procedures.
- Internal audits play a key role in HIPAA compliance by reviewing operations with the goal of identifying potential security violations. Policies and procedures should specifically document the scope, frequency, and procedures of audits. Audits should be both routine and event-based.
- Procedures should document instructions for addressing and responding to security breaches that are identified either during the audit or the normal course of operations.

- **Technical Safeguards** – controlling access to computer systems and enabling covered entities to protect communications containing PHI transmitted electronically over open networks from being intercepted by anyone other than the intended recipient.
 - Information systems housing PHI must be protected from intrusion. When information flows over open networks, some form of encryption must be utilized. If closed systems/networks are utilized, existing access controls are considered sufficient and encryption is optional.
 - Each covered entity is responsible for ensuring that the data within its systems has not been changed or erased in an unauthorized manner.
 - Data corroboration, including the use of check sum, double-keying, message authentication, and digital signature may be used to ensure data integrity.
 - Covered entities must also authenticate entities it communicates with. Authentication consists of corroborating that an entity is who it claims to be. Examples of corroboration include: password systems, two or three-way handshakes, telephone callback, and token systems.
 - Covered entities must make documentation of their HIPAA practices available to the government to determine compliance.
 - In addition to policies and procedures and access records, information technology documentation should also include a written record of all configuration settings on the components of the network because these components are complex, configurable, and always changing.
 - Documented risk analysis and risk management programs are required. Covered entities must carefully consider the risks of their operations as they implement systems to comply with the act. (The requirement of risk analysis and risk management implies that the act's security requirements are a minimum standard and places responsibility on covered entities to take all reasonable precautions necessary to prevent PHI from being used for non-health purposes.)

Enforcement Rule

On February 16, 2006, HHS issued the Final Rule regarding HIPAA enforcement. It became effective on March 16, 2006. The Enforcement Rule sets civil money penalties for violating HIPAA rules and establishes procedures for investigations and hearings for HIPAA violations; however, its deterrent effects seem to be negligible with few prosecutions for violations.

WHAT IS MESSAGING MIDDLEWARE?

Messaging middleware is computer software that connects software components or applications. The software consists of a set of services that allows multiple processes running on one or more machines to interact across a network via data messages. This technology evolved to provide for interoperability in support of the move to coherent distributed architectures, which are used most often to support and simplify complex, distributed applications. It includes web servers, application servers, and similar tools that support application development and delivery. Messaging middleware is especially integral to modern information technology based on XML, SOAP, Web services, and service-oriented architecture.

Messaging middleware sits "in the middle" between application software working on different operating systems. It is similar to the middle layer of a three-tier single system architecture, except that it is stretched across multiple systems or applications. Examples include database systems, telecommunications software, transaction monitors, and messaging-and-queuing software. For purposes of this white paper we are focusing on the messaging-and-queuing software, which provides the infrastructure to connect distributed platforms.

What is IBM's WebSphere MQ?

IBM WebSphere MQ is a family of network communication software products launched by IBM in March 1992. It was previously known as MQSeries, a trademark that IBM rebranded in 2002 to join the suite of WebSphere products. WebSphere MQ is IBM's Message Oriented Middleware offering. It allows independent and potentially non-concurrent applications on a distributed system to communicate with each other. MQ is available on a large number of platforms (both IBM and non-IBM), including z/OS (mainframe), OS/400 (IBM System i or AS/400), Transaction Processing Facility, UNIX (AIX, HP-UX, Solaris), HP NonStop, OpenVMS, Linux, and Microsoft Windows.

WebSphere MQ, a member of the WebSphere family from IBM, is the most popular system for messaging across multiple platforms, including Windows, Linux, IBM mainframe and midrange, and Unix. WebSphere MQ is often referred to as "MQ" or "MQSeries".

There are two parts to message queuing:

- Messages are collections of binary or character (for instance ASCII or EBCDIC) data that have some meaning to a participating program. As in other communications protocols, storage, routing, and delivery information is added to the message before transmission and stripped from the message prior to delivery to the receiving application.
- Message queues are objects that store messages in an application.

A Queue Manager, although not strictly required for message-oriented middleware, is a Websphere MQ prerequisite and system service that provides a logical container for the message queue and is responsible for transferring data to other queue managers via message channels.

There are several advantages to this technology:

- Messages do not depend on pure packet-based transmissions, such as TCP/IP. This allows the sending and receiving ends to be decoupled and potentially operate asynchronously.
- Messages will be delivered once and once only, irrespective of errors and network problems.

WebSphere MQ is a key component in IBM's SOA strategy providing the universal messaging backbone across 80 different platforms. The growing importance of Service Oriented Architecture (SOA) and the growth of Web Services and other connectivity mechanisms are clearly important developments. Because of the loosely-coupled nature of the message queuing model, a large number of existing MQ customers feel that they are already adopting SOA principles. The MQ Service definition supportpac MA93 allows MQ applications to be catalogued as Software assets which can then be reused and composed as Web Services.

HTTP-MQ bridge supportpac MA0Y extends the reach of MQ into the Web 2.0 world, MQ provides a Universal Messaging Backbone capable of accessing and delivering business data with a range of Qualities of Service both inside and outside of their enterprises. MA0Y functionality is included in WebSphere MQ V7.0 which contains a major set of JMS support enhancements and an integrated Publish/Subscribe capability. These features enable adoption of event-driven SOA.

MQ also supports Enterprise Service Bus implementations, both proprietary and open source. As for example, the Mule WMQ transport is available with Mule Enterprise Edition version 1.6 or later and, as of 2.2, contains many critical performance and reliability improvements as well as native support for WebSphere MQ-specific features.

There are many ways to access WebSphere MQ's facilities. Some of the APIs supported by IBM are:

- IBM Message Queue Interface for C, COBOL, PL/I, and Java, RPG
- JMS for Java
- Perl interface
- Windows PowerShell
- XMS for C/C++ and .NET
- .NET
- SOAP

There are many other APIs that are unsupported by IBM.

WebSphere MQ provides assured one-time delivery of messages across a wide variety of platforms. The product emphasizes reliability and robustness of message traffic, and ensures that a message should never be lost if MQ is appropriately configured.

It needs to be remembered that a message in the context of MQ has no implication other than a gathering of data. MQ is very generalized and can be used as a robust substitute for many forms of intercommunication. For example, it can be used to implement reliable delivery of large files as a substitute for FTP.

WMQ FEATURES

MQ provides application designers a mechanism to achieve non-time-dependent architecture. Messages can be sent from one application to another, regardless of whether the applications are running at the same time. If a message receiver application is not running when a sender sends it a message, the queue manager will hold the message until the receiver asks for it. Ordering of all messages is preserved, by default this is in FIFO order of receipt at the local queue within priority of the message.

It provides a means for transforming data between different architectures and protocols, such as Big Endian to Little Endian, or EBCDIC to ASCII. WebSphere MQ allows receipt of messages to "trigger" other applications to run, and thus provides the framework for a message driven architecture. In addition, MQ implements the JMS standard API, and also has its own proprietary API, known as the Message Queuing Interface.

Unlike email, MQ itself is responsible for determining the destination of messages by the definition of queues, so processing of sent messages can be moved to a different application at a different destination. MQ provides a robust routing architecture, allowing messages to be routed via alternative paths around a network of MQ managers. MQ can be implemented as a cluster, where multiple MQ implementations share the processing of messages to allow higher performance and load balancing.

WMQ COMMUNICATION

The primary component of a WebSphere MQ installation is the Queue Manager. The queue manager handles storage, timing issues, triggering, and all other functions not directly related to actual movement of data.

Queue managers communicate with the outside world either via a direct software connection, referred to by IBM as a "bindings" connection, or via a network or "client" connection. The bindings connection is limited to programs running on the same physical host as the queue manager, whereas applications using a client connection can connect to a queue manager on any other host in the network.

Bindings connections are generally faster, but client connections allow for a more robust, easily-changeable application design. For instance, with a client connection, the physical location of the queue manager is irrelevant, as long as it is reachable over the network.

Communication between queue managers relies on a channel. Each queue manager uses one or more channels to send and receive data to other queue managers. A channel is uni-directional, a second channel is required to return data. In a TCP/IP based network, a channel will send or receive data on a specific port. A sending channel has a defined destination and is associated with a specific transmission queue, the mechanism by which messages are queued awaiting transmission on the channel; a receiving channel will receive data from any other queue manager with a sending channel of the same name. When a receiving channel receives a message, it is examined to see which queue manager and queue it is destined for. In the event of a communications failure, MQ can automatically re-establish a connection when the problem is resolved.

The "listener" has the function of detecting connections from incoming channels and managing the connection of the sending to the receiving channels. It is the application's network interface to the queue manager. In a TCP/IP network, the listener will "listen" for connections on a specific port.

Local queues represent the location in which data is stored awaiting processing.

Remote queues represent a queue on another queue manager. They define the destination queue, which is one element of the routing mechanism for messages.

To transmit data to a queue on another queue manager, a message is placed on a remote queue. A remote queue is sent via the temporary storage transmission queue associated with a channel. On placing a message on a remote queue, the message will be transmitted across the remote channel. If the transmission is successful the message is removed from the transmit queue. On receiving a message, the receiving queue manager will examine the message to determine whether the message is for itself or is required to forward on to another queue manager. If it is the destination, the required queue will be checked, and if it exists, the message will be placed on this queue, and if not, placed on the dead letter queue. MQ has features to manage efficient transmission of data across a variety of communication mediums, so for example messages can be batched together until a queue reaches a particular depth.

Although the queue is FIFO, it is ordered based on the receipt in the local queue, not the committing of the message from the sender. Messages can be prioritized, and by default the queue is held in priority arrival sequence. Message grouping can be used to ensure a set of messages are in a specific order. Aside from that, if sequence is critical, it is the application's responsibility to place sequence data in the message or implement a handshaking mechanism via a return queue. In reality, ordering will be maintained in straightforward configurations.

The other element of a queue manager is the log, the mechanism used to create the robustness. As a message is placed on a queue or a configuration change is made, the data is also logged. In the event of a failure, the log is used to recreate damaged objects and recreate messages. In the event of a failure (as opposed to clean shutdown) only "persistent" messages will be recreated. "Non-persistent" messages are lost in the event of a failure or forced shutdown. Non-persistent messages can be sent across a channel set to a fast mode, in which delivery is not assured in the event of a channel failure.

MQ is designed to support a wide variety of approaches to application development. Information can be retrieved from queues either by polling the queue to check for available data at suitable intervals, or alternatively MQ can trigger an event, allowing a client application to respond to the delivery of a message.

WHO USES WMQ?

WebSphere MQ is utilized in over 10,000 installations globally and is the de facto messaging standard. In the PCI arena we see it in acquiring banks, merchants, processors, POS and others. Below are the many other sectors WMQ is predominant.

Commerce

- Retailers and card payment processors
- Retail banks, FRB, clearing houses and other financial institutions
- Healthcare, medical, pharmacy
- Retail, commercial

Infrastructure

- Federal, State and local government
- Military
- Transportation and Administration: Air, bus, train, auto traffic

Outsourcing and B2B

- Customer Relationship Management: Collaboration (SAP, Peopleware)
- Personnel
- Supply Chain Management
- All of the above, and more

THE ISSUES: MIS-CONFIGURATION AND PCI NON-COMPLIANCE

The PCI Data Security Standard enumerates 12 broad principles, each of which is further broken down into several sub-categories. The purpose of the standard is to illustrate categories of risk against which a given classification should be evaluated; the standard lists examples for each item listed. The result is that the standard tends to be interpreted very narrowly against the examples given rather than categorically by the *principles* described.

THE IMPORTANCE OF UNDERSTANDING THE REQUIREMENTS

In one case a forum discussion was debating the intent of a PCI standard that calls for changing vendor security defaults. The standard lists account names and passwords as examples but would apply to any defaults such as, in the case of MQ, channel MCAUSER values. The poster in the forum had inquired whether that PCI requirement meant that their shop needed to set passwords for service accounts that currently had no passwords. In other words, they had provisioned non-login service accounts for their application and were seriously considering adding password entries to the accounts because the PCI standard mentioned a need to change default passwords.

Obviously, the PCI standard was not intended to result in *reduced* security! The example given is there because it is well known that many applications and hardware systems are pre-loaded with default accounts and passwords. If these are not changed, an attacker can just walk right in. But this is just an example of the principle of addressing well-known attack vectors. Any security default that exposes a system would fall under this PCI DSS item. For example, if you install a router that allows administration from internal and external addresses, you might want to disable external administration. Similarly, if you install WebSphere MQ you should disable access to your default client channels, preventing anonymous users from penetrating the WMQ network.

HOW THE MQ SECURITY ASSESSMENT WORKS

Evans Resource Group works to gather the WMQ information and then will provide a report to the Chief Risk/Security Officer, Chief Information Officer or internal auditor committee. This report outlines the findings and maps them to the PCI requirements and how they impact the principles outlined within the framework of compliance. We ask for information on the WMQ environment in terms of the queue managers in each environment, clusters, object authority managers and outside client connections and other MQ object settings. Once this information is provided we analyze it to produce an order of magnitude report that outlines our findings and provides clients with the scope of how “wide open” the WMQ network is based upon this assessment. The sizing of the remediation can then be proposed and a project plan provided to the client to outline roles, responsibilities and phases of remediation.

ORDER OF MAGNITUDE ASSESSMENT REPORT: MAPPINGS OF THE PCI DSS REQUIREMENTS TO WMQ GAPS WITH EXAMPLE FINDINGS

2 Protect the perimeter, internal, and wireless networks. This milestone targets controls for points of access to most compromises – the network or a wireless access point.

- **PCI: 2.1:** Always change vendor-supplied defaults **before** installing a system on the network.

MQ Exposure: No security has been applied to the examined Queue Managers, leaving each MQ Queue Manager in a default “wide open” state, allowing open connectivity to “everyone” as MQ Administrator.

Each Production Queue Manager’s “wide open” state allows any party with an IP route to the Queue Manager to connect, and perform arbitrary remote code execution, either using the MQ Command Server, PCF commands or any third party code.

- **PCI 2.3:** Encrypt all non-console administrative access. Use technologies such as SSH, VPN, or SSL/TLS for web based management and other non-console administrative access.

MQ Exposure: This configuration allows for unencrypted non-console administrative access to all Queue Managers.

3 Secure payment card applications. This milestone targets controls for applications, application processes, and application servers. Weaknesses in these areas allows for compromising of systems and obtaining access to cardholder data.

- **PCI 2.2:** Develop configuration standards for all system components. Assure that these standards address all known security vulnerabilities and are consistent with industry-accepted system hardening standards.

MQ Exposure: The client has not been asked if they have developed MQ configuration standards. However, even if such standards exist, no recognized standard has been visibly applied by the client.

- **PCI 2.4** Shared hosting providers must protect each entity’s hosted environment and cardholder data.
- **PCI 6.1:** Ensure that all system components and software have the latest vendor-supplied security patches installed. Install critical security patches within one month of release.

MQ Exposure: Versions of MQ prior to V6.0.2.2 and V5.3.1.4 have well publicized security vulnerabilities. The version of MQ used by this (example) client has newly publicized vulnerabilities, but as this product version is no longer vendor supported, no security patch will be issued by the vendor. This installation requires an upgrade to a specifically compliant later version of MQ.

- **PCI 6.2:** Establish a process to identify newly discovered security vulnerabilities (for example, subscribe to alert services freely available on the Internet). Update configuration standards as required by PCI DSS Requirement 2.2 to address new vulnerability issues. *Outstanding vulnerabilities have not been addressed. Conclusion: this (example) client either does not have such a process, or they have a process, but are not following it.*

MQ Exposure: as usage and vulnerabilities of MQ are generally “hidden” from the audit process, there is no implementation or enforcement of this requirement. There are specific known vulnerabilities with the version of MQ that this client is using that have not been addressed.

4 Monitor and control access to your systems. Controls for this milestone allow you to detect who, what, when, and how the network and cardholder data is accessed.

- **7.1** Limit access to system components and cardholder data to only those individuals whose job requires such access.

MQ Exposure: The overlapping connectivity between production and development MQ environments proves that this requirement is not met. Role based Access Control requires that these environments be separate, unconnected networks.

- **PCI 7.2** Establish an access control system for systems components with multiple users that restricts access based on a user’s need to know, and is set to “deny all” unless specifically allowed.

MQ Exposure: all MQ components at this installation are set to allow “open access” for all administrative and application functions, without any audit trail for usage or system changes. The production and development environments are interconnected, failing to provide for Role Based Access Control Channels used to provide administrative access are unprotected, therefore no access controls are applied. This client has not enabled MQ’s intrusion detection and logging of detected security violations.

5 Protect stored cardholder data. For those organizations that have analyzed their business processes and determined that they must store Primary Account Numbers, Milestone Five targets key protection mechanisms for that stored data.

- **PCI 3.5** Protect cryptographic keys used for encryption of cardholder data against both disclosure and misuse.

MQ Exposure: by allowing anonymous administrative access, this client also allows access to the cryptographic key store.

PCI 3.6.7 Prevention of unauthorized substitution of cryptographic keys. N/A

MQ Exposure: by allowing anonymous administrative access, this client has opened access to the cryptographic key store, allowing anonymous users to delete, copy and substitute keys / certificates.

MQ Exposure: MQ supports SSL, however this client has not applied SSL to connections supporting the transmission of credit card payment / order data between the client and the payment processor.

AUDIT FINDINGS

Summary: This client will require remediation of the MQ environment to become PCI compliant. The client will require an upgrade, application of security exits, application of SSL, differentiation of development and production environments and the application of role based access control.

WHAT IS BEING DONE: REMEDIATION OF THE WMQ ENVIRONMENT EVANS RESOURCE GROUP MQSENTRY PROGRAM

Preventing remote code authorization without sacrificing core business operations or disrupting customer interactions can be achieved using a phased approach for remediation. In addition, the need to be flexible, while accommodating future growth and responding to future attack vectors, all reinforce the need for this approach.

There are three primary phases that organizations employ to remediate the WMQ environment without compromising the user experience:

1. Secure at the ADMINISTRATION LEVEL.
2. Secure at the APPLICATION LEVEL
3. Secure at the DATA LEVEL.

SECURING AT THE ADMINISTRATIVE LEVEL

Securing at the administrative level so that customers are guaranteed that all administrative vs. non-administrative access is segmented and that the possibility of remote code execution by anonymous administrative access is eliminated is the first and foremost level of remediation. It is critical to change vendor security defaults in the WMQ environment to prevent unauthorized access to the queue managers. Not changing defaults allows an attacker to walk right in to a well known attack vector. Any security default that exposes a system falls under PCI DSS as an item for remediation. As per our findings approximately 90% of assessments have no WMQ administrative security.

SECURING AT THE APPLICATION LEVEL

We secure data at the application level by applying security exits. This can be done using either a custom approach or a productized solution. The custom approach is less desirable since it is then proprietary to each client. This approach requires the writing of custom code for security exits, and applying the security exits for authentication, along with the application of authorization settings, and the setting of default parameter values that relate to security. Connecting applications will be assigned user Ids and, in some cases, certificates for authentication, and those Ids will be assigned to specific objects within the MQ environment. Selecting the productized approach means the use of Evans Resource Group's MQSentry product, which provides a standardized security framework, as well as fully supported configurable security exits for each platform, and each MQ network type, including multi-platform clusters.

SECURING AT THE DATA LEVEL

The use of SSL based encryption is necessary for enhanced channel data security for sensitive data, including data covered by PCI DSS. In addition, a highly sensitive

© Evans Resource Group, 2009

differentiated network carrying high-volume sensitive data may require the use of WebSphere MQ Extended Security Edition, which can be configured to provide end-to-end protection of data, even within the framework of persistent messaging. Evans Resource Group has the necessary expertise to recommend and apply MQ ESE in such cases.

WebSphere MQ Extended Security Edition features:

- Flexibility to add application-level data protection and remote security policy administration, which includes:
 - Sign each message with a unique private key associated with the sending application.
 - Encrypt individual messages under unique keys, helping to remove the threat of compromising the encryption key through repetitive use.
- Yields a true end-to-end security model by removing the need for you to reengineer and modify your applications to secure message data from within each application.

WebSphere MQ Extended Security Edition also enables remote administration of security policies on queue managers and on individual queues. These include:

- - Put and get access control permissions, including time of day and day of week restriction.
- Data protection options (none, integrity, and privacy).
- Audit options that allow generation of a specific security audit record for each open, put, get, and close operation showing the security policy in place and whether it was successfully enforced.

This can be critical in demonstrating compliance with legislation like HIPAA and similar mandates in other countries.

Remote MQ administration is performed via a Web-based utility, protected via VPN or SSL, with an easy-to-use GUI for setting, viewing, and updating security policies. A delegation capability allows IT organizations to maintain control over the enterprise security infrastructure for WebSphere MQ Extended Security Edition and still grant a specific department or line-of-business the ability to manage its subset of resources. Administration can also be done via scripting using a command line interface, as with base MQ.

- WebSphere Extended Security Edition also includes the ability to:
 - Protect message data while the message is in queue and as it flows across the network
 - Attach digital signatures to individual messages so you can verify if they have been tampered with and identify and trace messages to their point of origin
 - Encrypt messages to shield them from disclosure to unauthorized parties and

- centrally define and enforce security policies via a Web-based interface
- Provide security-specific, fine-grained auditing records to document adherence to security policy, protecting highly sensitive transactions records
 - Help accelerate new application deployment by reducing development time and costs by providing application-level security without requiring complex security coding
 - Provide an immediate ROI upon deployment by securing existing applications without application change

SECURING B2B CONNECTIVITY

MQ is commonly used to maintain B2B connectivity between an enterprise and its partners. A specific auditable condition relates to security constructs around the configuration of a B2B gateway. B2B gateway security must provide for the total separation of partner data and connectivity, and protection from intended, or unintended, access from one partner system to another, and exposure of a partner's proprietary data to another partner. Evans Resource Group has standardized the approach for securing a B2B MQ gateway that provides for full compliance.

CONCLUSION

As with other potential threats to data security, protection of the WebSphere MQ network requires organizations to apply security using a phased approach. When looking at any proposed solution, it is important to understand the following criteria:

- Types of threats, in depth
- Impact on customer experience
- Future growth of the WMQ network

Simply applying only one level of security will not ensure the appropriate response to ensure compliance. Adding risk-based, multi-level authentication and authorization based on role and responsibility provides more robust protection against exposures and provides alignment with required compliance standards.

The application of security exits will ensure an appropriate level of compliance at each segment of the WMQ network, assuring that the administrative, application and data requirements are met. These levels can be triaged by setting appropriate parameters, upgrading to the appropriate versions across platforms and providing authentication with SSL or WMQ ESE where appropriate.

ABOUT EVANS RESOURCE GROUP

Evans Resource Group can secure the WebSphere MQ network for merchants, acquiring banks, processors, healthcare and public enterprises, that are amongst the more than 10,000 organizations in 60 countries using MQ today. Leveraging a program of assessment and remediation constructed in conjunction with IBM's WebSphere Security team, we provide a multi-phased approach to data security that addresses growing risks using our MQSentry program and product. We are also the founders of Reconda International Corporation, the inventors of QN-AppWatch and QN-StatWatch, the industry leaders in browser based WMQ technology. For information, call 212.937.8443 or email info@evansrg.com or visit www.evansresourcegroup.com