Security Requirements Questionnaire

IoPT Consulting, LLC

Author: T.Rob Wyatt
Owner: T.Rob Wyatt
Customer: TBD

Confidential when comppleted

Document History

Document Location

This is a snapshot of an on-line document. Paper copies are valid only on the day they are printed. Refer to the author if you are in any doubt about the currency of this document.

This is a template document. The most current version of this document may be found at the following locations:

Public download – https://t-rob.net/Links

Revision History

Date of this revision: 09 June 201	Date of next revision *N/A*

Revision Number	Revision Date	Summary of Changes	Changes marked
(#)	(-)	(Describe change)	(N)
3.0	20150609	Significant updates, published to web	

Contents

1. Preface 4
2. Intrusion prevention 5
2.1 Authentication 5
2.2 Authorization 7
2.3 Data integrity 8
2.4 Data privacy 8
2.5 Data segregation 8
2.6 Network topology 8
3. Intrusion detection 10
4. Analysis and recovery 11
5. Environment 12

1. Preface

This questionnaire is designed to drive out the IBM MQ security requirements for a given infrastructure or enterprise. Very often the term 'security' is construed narrowly to mean only prevention of unauthorized access. A more comprehensive definition includes not only intrusion prevention but also intrusion detection, the ability to perform forensic analysis, recovery after an incident and an ongoing commitment to staying current on security news such as patches, vulnerabilities and new products or capabilities.

This document asks probing, open-ended questions in all of these areas. Within each category is an itemized list. It is not expected that answers to all questions will be readily available. In some cases analysis will reveal that certain risks can be accepted without the expense of mitigating controls. In other cases the potential impact of a remote risk is so far out of proportion to the cost of mitigating controls that there is no reasonable choice but to make the investment. The important thing is that in each case the risk acceptance or mitigation is a conscious, deliberate decision.

2. Intrusion prevention

2.1 Authentication

IBM MQ leverages the local operating system to authenticate local 'bindings mode' connections. This section will therefore focus on channel connections.

Interactive user to queue manager connections

- What type of tooling is approved for interactive users? WMQ Explorer or other desktop tools? Agent based tool such as Tivoli Omegamon? Central browser-based tool?
- Which categories of interactive user are present and are they required to authenticate in order to connect? Administrators? Developers? Testers? QA? Application Support? Operations? Others?
- Do the authentication requirements differ by environment (i.e. Dev, QA, Prod, etc.)?
- Normally the population of MQ stakeholders with some level of access to MQ includes project team members, integration testers, QA testers, business analysts, application support, network ops, instrumentation, monitoring, and of course the administrators. Is the provisioning of user access so resource intensive that it restricts access from MQ stakeholders?
- Do regular users and administrators have separate access paths to connect to the queue managers?

Queue manager to queue manager connections (internal)

The following questions apply only to the internal network. External (B2B) connections are covered in a later section.

- How do internal queue managers authenticate their connections? TLS? Security exits?
- Are IBM MQ clusters present? If so, how do the cluster channels authenticate connection requests? TLS? CHAD/Security exits?
- Do queue managers accept administrative commends from adjacent queue managers, and if so is this a deliberate design choice?

Application to queue manager connections

- Are client-based applications supported (i.e. those that connect using SVRCONN channels)? If so, are they required to authenticate? How… User ID? TLS? Security exit? Some combination of these?
- In the case of shared environments such as application servers, ESB or brokers, what level of application isolation is provided in the environment and is the same level of granularity required at the queue manger?
- Are there DataPower connections to the queue manager?

External (B2B or extranet) connections

- Where are external connections terminated? In the DMZ at a IBM MQ gateway queue manager? An internal gateway? At the application hosting QMgr? An appliance?
- Is Internet Pass-thru in use? How are these connections authenticated at the IPT node? How are they authenticated at the queue manager?
- Are external connections allowed from IBM MQ clients?
- Are 3rd parties allowed to participate in a local cluster?
- Is DataPower allowed to connect to IBM MQ from an external connection?
- How are external connections authenticated?
- If multiple external business partners are connecting, what level of isolation is required between them?

Accounts management

- Is the 'mqm' **account** (or platform equivalent) considered privileged? What controls are in place to validate requests for access to this account?
- Is the 'mqm' **group** (or platform equivalent) considered privileged? What controls are in place to validate requests for access to this group?
- Are there low-privileged accounts and groups defined for secondary IBM MQ functions such as channel MCAUSER values? What are the controls in place to validate requests for access to these accounts and groups?
- Is there an owner-of-record for the administrative accounts and groups used by IBM MQ? Is the owner required to periodically validate group membership and account entitlements? Can this information be obtained on demand by the account owner of record?
- Are accounts managed centrally or per-host?
- Are service accounts shared by multiple QMgrs?
- Are service accounts shared across multiple staging levels (e.g. Prod and non-Prod)?
- Is there an authoritative and current document mapping accounts and groups to security roles in the MQ network?

Certificate management

- Are certificates self-signed or CA-signed?
- How are certificates provisioned to queue managers, apps and users? If they generated centrally other than in-place, how are they distributed and managed?
- What password policies exist for certificates?
- What policies exist for certificate expiry?
- What policies exist for certificate usage fields?
- What policies exist for Distinguished Name fields?
- Where certificate policies exist, do these differ for human users, applications or queue managers?

- How is certificate expiry and renewal managed? Is there an authoritative source tracking this information?
- Do administrators know how to check certificates for the IsCA , Path Length and other attributes, and the expected values?

Revocation of access

- What is the process for revoking access to IBM MQ for each of the categories of access described above (QMgr-to-QMgr, application, interactive user)?
- For interactive users, do external revocation processes, such as employee termination, invoke IBM MQ access revocation?
- For TLS channels, what type of revocation is used? CRL? OCSP? Removal of self-signed certificates from the keystore?
- Are IBM MQ TLS channels configured to connect if the revocation responder is unavailable?

Internal Certificate Authorities

- What physical and procedural controls are in place to protect root and intermediate certificates?
- How long is the certificate chain from the queue manager or user certificate to the root certificate?
- What is the process to request, revoke or renew certificates?
- Are there separate root certificates for Production versus other environments? If not, are there separate intermediate certificates?
- What type of revocation responder is provided?

2.2 Authorization

- Is there a requirement to restrict administrative access to IBM MQ? For interactive users? For applications?
- Is administrative access ever required from one queue manager to another - i.e. using one queue manager as a proxy to administer another?
- Describe the level of authorization granularity required for non-administrative interactive users? Should all non-admins have the same level of access? Different authorizations per department or application or function?
- What level of authorization granularity is required for client-connected applications?
- Is Publish/Subscribe functionality required? If so, what level of authorization granularity is required? By base topic? At points within the topic hierarchy?
- Who is in the mqm group (or platform equivalent)? Who is root (or platform equivalent)?
- For Active Directory domains, are there any domain groups nested in the Administrators or mqm group? If so, who is in those groups?
- Do any interactive, non-administrative users require the ability to create a queue, other than by using a model queue?
- Do non-administrators require the ability to set attributes of the queue manager?

2.3 Data integrity

- Are there any requirements for cryptographic assurance of data integrity for any class of data on the network?
- If so, does the requirement apply to data in transit? Data at rest?
- Is there a requirement for data integrity over channels that do not use certificate authentication?

2.4 Data privacy

- Are there any requirements for cryptographic assurance of data privacy (encryption)?
- If so, does the requirement apply to data in transit? Data at rest?
- Is there a requirement for data privacy over channels that do not use certificate authentication?

2.5 Data segregation

- Are there any requirements to segregate classes of data on the network?
- If so, is this for performance reasons?
- Is this for regulatory compliance (PCI, HIPAA, SOX, GLBA, etc.)?
- Does data segregation extend to the queues? Channels? Per queue manager? At the network or subnet level? Per-host?

2.6 Network topology

- Please describe the IBM MQ network topology. Base patterns include point-to-point, clustered, hub-and-spoke, B2B, gateway, client concentrator, and more.
- Is there an ESB, broker or other central component acting as a virtual (or actual) hub?
- Are there gateways or other instances where messages hop through two or more queue managers?
- Are there choke points where connections from many applications or queue managers are funnelled through a much smaller number of concentrator queue managers?
- How is DR handled?

3. Intrusion detection

- What requirements exist for detection of intrusion attempts?
- What types of monitoring are in place for… log files? Directories? IBM MQ event messages? Real-time traffic statistics? Processes? Configurations?
- What is the process for reviewing intrusion detection reports or logs? Is this scheduled or only in response to an event?
- Are there any automated reports of security relevant events? Which events are reported? How are these reports delivered? Is there both negative and positive event reporting (i.e. does lack of a report mean 'everything is OK')?
- Which non-administrative groups or accounts have access to the Dead Letter Queue? Event Queues? What level of access is granted?

4. Analysis and recovery

- Are privileged accounts keystroke logged?
- Is the mqm account (or platform equivalent) considered privileged?
- What is the backup strategy for IBM MQ queue managers? How often are configurations backed up?
- What provision is made to recover in-flight messages orphaned when a node goes down? Can the messages be automatically reconciled and recreated? Does recovery rely on disk or message replication?
- Are IBM MQ error logs saved and archived?
- Is there a poison message standard that is enforced against project teams or software vendors?
- How are Dead Letter messages handled? If they are retried, are they logged first?
- Is there at least one person whose job includes the responsibility to maintain currency on IBM MQ security issues and vendor patches? What proportion of time is dedicated to security currency?
- In the event a security patch is released, what contingency exists in the budget and/or staffing to push the fix through the application development lifecycle? How long would this take to accomplish?
- Do the IBM MQ administrators, the developers and the other stakeholder teams have sufficient capacity to respond to a security relevant event such as prioritizing implementation of a security patch without severe impact to existing timelines?

5. Environment

- Is the company subject to any regulatory compliance such as PCI-DSS, HIPAA, SOX, GLBA, etc.? If so, have the messaging systems previously been in scope for the compliance review or audit?
- Do the messaging systems carry any Personal Health Information (PHI), Personally Identifiable Information (PII), cardholder information, account information or sensitive internal information not otherwise covered under 'compliance'?
- Are there contractual obligations based on messaging system availability or performance?
- Does the messaging network carry mission-critical, high-value messages where loss or compromise would result in catastrophic financial or reputational impact? What is the threshold for acceptable loss or compromise of messages? Zero? Several? Dozens?
- What non-WMQ security controls are in place on the internal network? Is the messaging network hosted on a flat network or is it segmented? Are these controls in place for non-production as well as production systems? Who can get an IP route to the queue managers?
- Are all nodes, including non-Production nodes, hosted in physically secure locations?
- Are any production and non-production WMQ nodes co-resident on the same server?