WebSphere MQ V7.5

Morag Hughson
hughson@uk.ibm.com

IBM

How WebSphere MQ meets your Connectivity needs

- *Dynamic network that delivers the **data** you require from wherever it resides to wherever you want it in whatever way you want it at whatever time you want it*

IBM

WebSphere MQ V7.1: Feature Summary

WebSphere MQ V7.1
Announced: 4 October 2011
Availability: 11 November 2011

Simplification

New Feature	*Benefits*	
Multi-Version Install capability on Distributed platforms	Makes it easier to deploy and upgrade systems and stage version to version migration	Unix and Windows support for multiple versions of MQ V7.x (AND one copy of MQ V7.0.1) down to fixpack levels. Relocatable installation support. Applications can connect to any Qmgr
Enhanced Security	Simplified configuration Enhanced Authorisation and Auditing	IP address Authorisation capability Additional crypto algorithms More granular authorisation for non-local queues Application Activity Reports
Cloud Support	Simplifies and support Cloud deployments	Official HVE images
Enhanced Clustering	Improves ease-of-use	Authorisation on Cluster Q rather than XMIT Q on Dist. Platforms Bind-on-Group Support
Multicast capability	New messaging QoS provides low latency with high fan-out capability	MQ Pub/Sub Topics specify can map to multicast Group Addresses Provides direct interoperability with 29West LLM
Improved scalability and availability on z/OS	Further exploitation of z196 Customer control over CF storage use CF Connectivity Loss improvements	Code contention reduced to improve multi-processor linear scaling Use of MQ Datasets rather than DB2 significantly improves "large" message capability Structure rebuild capability for CF Connectivity Loss scenarios
Improved Performance on Dist platforms	Improved multiprocessor exploitation	Various code improvements

IBM

WebSphere MQ V7.1: Feature Summary

Enterprise Messaging

N O T E S

- This page shows the highlights of the new release in one chart. The rest of this presentation goes into the details.
- A one-word summary of this summary is "simplification": making it easier to own, run and work with MQ.

- One part of the MQ V7.0.1 rationale was to deliver new function via the service stream, without requiring a full new release and migration cycle. Lessons learned from that have fed into V7.1, which has been designed to be more capable and more flexible when adding function through this channel.
- These new functions can be optionally enabled. The default is that new function requires adminstrative action to enable it, so that there is no unasked-for change in behaviour when you install a fixpack

IBM

© 2014 IBM Corporation

WebSphere MQ V7.5: Content Summary

Enterprise Messaging

- For Windows, Unix and Linux

WebSphere MQ V7.5
Announced: 24 April 2012
Availability: 20 June 2012

Simplification

New Feature	*Benefits*	*Details*
Integrated Installation	Makes it easier to deploy systems Simpler licensing	Combines several products into a single package Common experience
Enhanced Clustering	Improves ease-of-use Improves application isolation	Split Cluster transmission queue
Java Application Identification	Makes it easier to distinguish applications	Applications no longer to all have the same name
AMS channel interception	Provides a level of message protection even when application environment cannot run AMS	Interception in the SVRCONN still protects messages before hitting queues
FTE Logger Options	Can write FTE audit records to flat file	No longer a requirement for an enterprise database Easier to read data immediately

IBM

WebSphere MQ V7.5: Content Summary

Enterprise Messaging

N O T E S

- Note that z/OS and System i are not included in this announcement
  - These V7.5 platforms were selected based on research on which ones would benefit most from the integration. The lack of z/OS and i **at this point** reflects the slower rate at which new MQ versions are, typically, consumed on these platforms; it should not be taken as suggesting anything more about the value and eventual delivery of new functions across all primary platforms.
- Improved application isolation:
  - WebSphere MQ V7.5 includes improved ability to scale for differing workload environments by the ability to configure multiple transmission queues in a WebSphere MQ clustered environment. This enables applications with different workloads and performance requirements to operate at their own rate without impacting other applications.
- Enhancements for the managed file transfer capabilities:
  - In addition to being tightly integrated as a part of WebSphere MQ V7.5, which enhances the runtime control, the managed file transfer capabilities are also enhanced. There are additional choices for storing file transfer audit information with the addition of the file system as an option. There is also be greater customization for the content and format, as well as more options for logging.
- Enhancements for WebSphere MQ Security:
  - With WebSphere MQ V7.5 the Advanced Message Security feature built into all MQ clients, end-to-end encryption is allowed by updates to WebSphere MQ objects for customers who are licensed to use that function.

IBM

© 2014 IBM Corporation

WebSphere MQ V7.5

- Integrated Messaging Offering
  - Single install, packaging & tooling for all Messaging options
  - Reduce time to value, simplify usage
- What's being delivered?
  - Integration of MQ with MQ FTE, MQ AMS and MQ Telemetry
  - Single install, common integrated tooling and management, simplified licensing and entitlements
  - Updated MQ Explorer tooling for all platforms
  - More complete, easy to use messaging infrastructure, enabling you to gain full range of messaging, swiftly & easily
- All messaging functions & capabilities available to all customers, new and existing with rich choice of qualities of service
  - Removal of charge for MQ XA client
  - Reduced pricing metric for standard MQ Telemetry client
    - Lower cost for larger numbers of clients

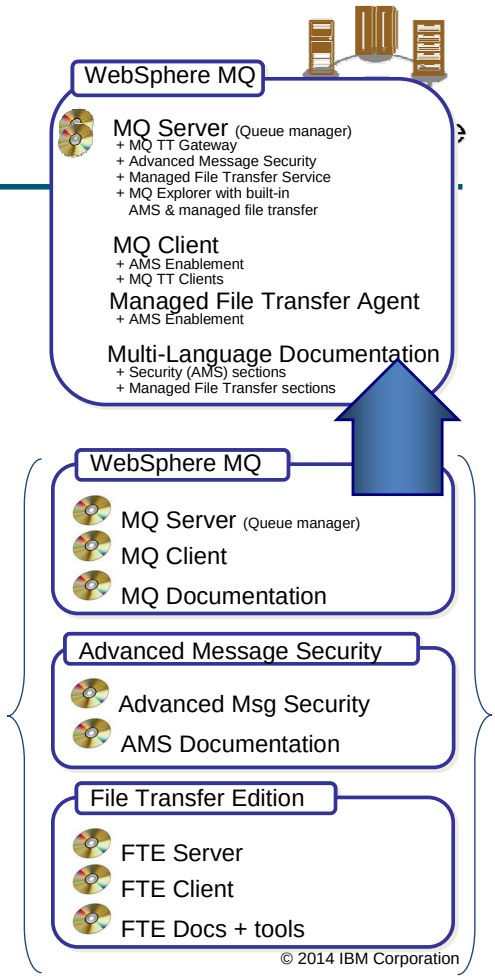

WebSphere MQ V7.5

Enterprise Messaging

N O T E S

- Integrated managed file transfer:
  - WebSphere MQ File Transfer Edition has been available as a separate product offering for a number of years. WebSphere MQ V7.5 extends this capability from being a separate offering, to being an integrated optional feature of the WebSphere MQ server component as the WebSphere MQ Managed File Transfer Service. Now all WebSphere MQ servers can install, subject to entitlement, this server-based capability. Additional separately installable endpoints, WebSphere MQ Managed File Transfer Agents, are included in the package, but entitled separately to allow the managed file transfer infrastructure to be extended to any point in the enterprise connection to the WebSphere MQ server deployments.
- Integrated advanced message security:
  - WebSphere MQ Advanced Message Security, also available for a number of years, allows customers to protect the security of their messages from application to application, without the need to change the application code itself. With WebSphere MQ V7.5, this capability is included as a part of the install, making it simpler for customers to see the function, and to have it installed should they wish to buy license entitlement to use it.
- Wider access to the Extended Transactional Client for all customers:
  - With the availability of WebSphere MQ V7.5, the Extended Transactional Client is available for use in all client deployments without additional licensing.
- Using WebSphere MQ Telemetry Standard Client:
  - WebSphere MQ Telemetry was included as part of the WebSphere MQ V7.1 offering, providing wider access to this capability for customers who wanted to deploy the WebSphere MQ Telemetry clients on suitable endpoints and connect them to their WebSphere MQ servers. Connecting these clients to WebSphere MQ servers required purchasing entitlements based on the number of Telemetry clients connecting to a WebSphere MQ Queue Manager at any one time. With the availability of WebSphere MQ V7.5, use of the WebSphere MQ Telemetry Client requires purchasing an entitlement for each WebSphere MQ Server installed that will have WebSphere MQ Telemetry clients connected to it, with no limit to the number of clients connected. The WebSphere MQ Telemetry Advanced Client entitlements are still based on the number of clients connected at any one time.
- Enhancements for the use of WebSphere MQ as a Trial:
  - Customers wanting to make a rapid start connecting applications with WebSphere MQ can take advantage of the availability of WebSphere MQ as a free download in a Trial version and with WebSphere MQ V7.1 this feature was enhanced to enable trial versions of WebSphere MQ to be upgraded to a full production license. With the integration of additional entitled capabilities in WebSphere MQ V7.5, a Trial download is still available, but where entitlement has not been purchased at the end of the trial for capabilities that have been installed, then those capabilities without entitlement must be removed from the system.

Multi-Version Installation

- MQ on Unix and Windows can install multiple levels on a system
  - Relocatable to user-chosen directories
  - Can have multiple copies even at the same fixpack level
- Simplifies migration
  - Can move applications as needed, not all at once
  - No need for parallel hardware
- Easier for ISVs to imbed MQ in solutions
  - Can install in “private” locations without worrying about other copies
  - Reduces support concerns
- Permits a single copy of V7.0.1 to remain on system
  - So existing systems can be migrated
  - Must be 7.0.1.6 or later

Multi-Version Installation

NOTES

- With this release, you can install more than one copy of MQ on Unix and Windows platforms.
  - It is not available on System i
  - z/OS already has this capability, in a different form
- This will simplify migration strategies, as you can continue to use one version of MQ and only gradually migrate applications to a new version, without needing parallel hardware.
- When installing MQ you can choose the directory into which it will be installed. There is no longer a requirement to use /opt/mqm (Linux, most Unix) or /usr/mqm (AIX).
- Third party applications can imbed MQ under their own private directory if they wish, and can choose which versions of MQ they support, without worrying about the “visible” version of MQ that a user might be exploiting.
- You can leave an existing copy of MQ V7.0.1.6 (or later) on your systems, and this new feature will work alongside it. So you do not need to move to V7.1 before starting to exploit the multiple installation capabilities.

Multi-Version Installation: Concepts

- Main concept is an installation
  - Refers to the directory containing the binaries from a particular version of MQ
  - Can have a descriptive name
- One installation can be designated as primary
  - Strongly recommended on Windows where some OS-specific elements have to be registered
  - Optional on Unix, creates symlinks to commands and libraries in /usr
    - Not created by default so your PATH will not always find MQ commands
- Queue Managers are owned by a specific installation
  - Governs the level of function available when the queue manager is running
  - Ownership can be changed to a newer installation for migration

Multi-Version Installation: Concepts

N O T E S

- An installation is the collection of binaries that make up the MQ runtime, its commands a libraries. The precise details vary by platform, but conceptually you first define an installation, optionally giving it a name and description. You then run the install program, which will copy the code from media onto the disk in your chosen directories.
- Generally, there are no “default” paths created to an installation. However a primary installation does insert itself into default locations. You can only have a single primary installation on a system.
- On Windows, one installation is always recommended to be denoted as the primary. This is necessary because some OS functions can only deal with a single location, for example how DLLs are registered for MSCS.
- On Unix, you do not require a primary installation unless you want to recreate links from /usr/bin and /usr/lib to the MQ commands and libraries in that installation.
- If you still have V7.0.1 on your system, that will always be the primary on all platforms.
- Each queue manager is considered as being owned by an installation. This governs the level of function available when running a queue manager. When you install a newer level of code, the queue manager can be moved to being owned by that newer installation, and hence new function can be used. Ownership is not changed automatically; you must execute the new setmqm command. Queue manager control commands such as strmqm must be issued from the directory associated with that owning installation

Multi-Version Installation: Application Impacts

- Existing applications “know” where the MQ libraries are
  - Embedded path or PATH/LIBPATH/LD_LIBRARY_PATH
  - Has always been a fixed location on Unix
- When MQ libraries move, apps will need to know where the new location is
  - /usr cannot be assumed
- New application libraries able to connect to any version of queue manager
  - Libraries such as libmqm, libmqic etc redesigned
  - Dynamically loading dependent libraries associated with the corresponding qmgr
  - If your app can find one V7.1 libmqm, it can connect to any qmgr, including future versions
- **MIGRATION NOTE:** Exits that invoke the MQI will need to be updated
  - Such as API Exits
  - Do not want exits to pull in different libraries than main application
  - Extended interface provides pointers instead for invoking MQI

Multi-Version Installation: Administration Impacts

- **MIGRATION NOTE:** Commands not in default PATH for Unix
  - Unless you have a primary install, all the control commands need explicit path
  - “ksh: crtmqm: not found” will be typical error message
- The dspmq, dspmqver & dspmqinst commands work for all installations
  - Shows installation details; which queue managers exist and owning installations
  - Other administration operations require PATH to point at install-specific commands
    - Get an error if you try to administer a queue manager the command does not own
- Installation details held box-wide
  - /etc/opt/mqm/mqinst.ini or Windows registry
- Default data paths unchanged
  - Still have /var/mqm/mqs.ini on Unix
  - On Windows, many other registry items are changed or removed
    - Now have qm.ini text files with the same content

Administration Examples

```
$ /usr/mqm/bin/dspmqver -i
Name:        WebSphere MQ
Version:     7.1.0.0
Level:       p000-L110915
BuildType:   IKAP - (Production)
Platform:    WebSphere MQ for AIX
Mode:        64-bit
O/S:         AIX 6.1
InstName:    Installation1
InstPath:    /usr/mqm
InstDesc:    My default installation
DataPath:    /var/mqm
Primary:     Yes
MaxCmdLevel: 710

Name:        WebSphere MQ
Version:     7.1.0.0
InstName:    Installation2
InstPath:    /usr/mqm2/usr/mqm
InstDesc:    A second installation
Primary:     No
```

```
$ dspmq -o installation
QMNAME(V71A)
       INSTNAME(Installation1)
       INSTPATH(/usr/mqm)
       INSTVER(7.1.0.0)
QMNAME(V71B)
       INSTNAME(Installation1)
       INSTPATH(/usr/mqm)
       INSTVER(7.1.0.0)
QMNAME(INST2QM)
       INSTNAME(Installation2)
       INSTPATH(/usr/mqm2/usr/mqm)
       INSTVER(7.1.0.0)
```

```
$ /usr/mqm/bin/endmqm INST2QM
AMQ5691: Queue manager 'INST2QM' is
associated with a different
installation.
```


Administration Examples

NOTES

- Here we see some of the commands that assist with installation management.
- The dspmq and dspmqver commands know enough about the system to show all of the installations, their directories, and the associated queue managers
- However queue manager control commands must be issued from the correct directory; they will return an error if you try to run them against a queue manager that is not associated with that installation.
- Tools that inspect the Windows registry or update it explicitly may need to be changed as most of the information that was previously in there now moves to flat-text ini files like the Unix platforms.

Security: Channel Authentication Records

- Set rules to control how inbound connections are treated
  - Inbound Clients
  - Inbound QMgr to QMgr channels
  - Other rogue connections causing FDCs
- Rules can be set to
  - Allow a connection
  - Allow a connection and assign an MCAUSER
  - Block a connection
  - Ban privileged access
  - Provide multiple positive or negative SSL Peer Name matching
- Rules can use any of the following identifying characteristics of the inbound connection
  - IP Address
  - SSL/TLS Subject's Distinguished Name
  - Client asserted user ID
  - Remote queue manager name

Channel Authentication Records – Notes

N O T E S

- Channel Authentication records allow you to define rules about how inbound connections into the queue manager should be treated. Inbound connections might be client channels or queue manager to queue manager channels. These rules can specify whether connections are allowed or blocked. If the connection in question is allowed, the rules can provide a user ID that the channel should run with or indicate that the user ID provided by the channel (flowed from the client or defined on the channel definition) is to be used.
- These rules can therefore be used to
  - Set up appropriate identities for channels to use when they run against the queue manager
  - Block unwanted connections
  - Ban privileged users
- Which users are considered privileged users is slightly different depending on which platform you are running your queue manager on. There is a special value '*MQADMIN' which has been defined to mean "any user that would be privileged on this platform". This special value can be used in the rules that check against the final user ID to be used by the channel – TYPE(USERLIST) rules – to ban any connection that is about to run as a privileged user. This catches any blank user IDs flowed from clients for example.

Security: Channel Authentication Records

- Simplifying configuration for channel access
  - Clients and queue managers
- SET CHLAUTH definitions control who can use channels
  - Name mapping
  - Access blocking
- Easy to test rules that you define
  - DISPLAY CHLAUTH MATCH(RUNCHECK) can “execute” rules
- Rules can be applied in WARNING mode
  - Not actually blocked, but errors generated
- **MIGRATION NOTE:** Standard rules block clients on new queue managers
  - “Secure by default”
  - Migrated queue managers behave as before until you enable the rules
  - Queue manager attribute CHLAUTH(ENABLED|DISABLED) provides overall control

Security: Channel Authentication Records

N O T E S

- Over the years there have been many requirements raised to make it simpler to block unwanted access to a queue manager. For example, only defined IP addresses should be allowed through.
- With V7.1 many of these rules can now be defined directly to the queue manager and channels.
- A standard set of rules are created when you create a new queue manager or migrate an existing one to run under V7.1. However, the rules only start to be used when you ENABLE them – a migrated queue manager has them DISABLED by default, so as to not disrupt any existing configuration and applications. The default rules block most client access; don’t be surprised to get authorisation failures until you have reviewed the rules. The default rules were chosen to make new queue managers automatically more secure, simplifying the process of securing a system.
- Channel Auth records define the rules that are applied when a queue manager or client attempts to connect through a channel. A number of elements about the partner can be checked, and choices made about whether or not to allow the connection to proceed.
- A pseudo-userid (*MQADMIN) can be used in these rules, which covers the use of any id that would otherwise gain automatic administrative rights over a queue manager. For example, is the asserted identity in the mqm group or blank (so would inherit the channel’s userid). Having a pseudo-userid makes it easier to have the same rules on all platforms, where the actual definition of who is an administrator might vary.
- The DISPLAY CHLAUTH command can be used with the MATCH(RUNCHECK) option to verify a simulated connection (pretending to have an address or id). This means you can test rules from the console without making a real connection.
- Rules can also be defined as “WARN”, causing authorisation events to be generated, but not actually blocking the connection. This may assist in migrating to a secure environment, by not turning off connections immediately.
- To further simplify setting up these rules, the MQ Explorer has a Wizard to take you through the steps.

© 2014 IBM Corporation

Security: Transport Layer Security (TLS/SSL)

- More crypto algorithms supported for TLS
  - Stronger algorithms are now available and recommended
  - MQ V7.0.1 (Dist) added some SHA-2
  - MQ V7.1 (Dist) adds more, with support for the NSA “Suite B” standard which includes Elliptic Curve cryptography
- Some older algorithms (eg SHA-1) should be considered deprecated
  - No plans to withdraw older algorithms immediately
  - But expect them to be removed in a future version of MQ
- Newer algorithms supported by GSKit8 on Distributed platforms and System SSL on z/OS V1R13
  - APAR OA39422 needed on z/OS V1R13 for System SSL
  - APAR PM77341 for MQ on z/OS V7.1
  - IBM i 7.1 Technology Refresh 6 (TR6) and PTF SI48659
  - Fix Pac V7.1.0.3 for MQ on IBM i
- See technote http://www.ibm.com/support/docview.wss?uid=swg21639606
- The GSKit toolkit is now provided inside the MQ installation
  - Will not clash with alternative levels from other MQ installations or other products

Security: SSL

N

O

T

E

S

- MQ V7.1 extends the support for newer cryptographic algorithms, including from the SHA-2 family, and the NSA Suite B set.
- NIST has stated that SHA-1 should be considered deprecated. While MQ has not removed these algorithms, it may be done in future versions.
- Like the earlier FIPS-140 option within MQ, you can choose to enforce Suite B compliance on your channels.
- As the various external toolkits MQ is dependant on; GSKit, JSSE (for Java) and System SSL on z/OS and System i, have made these new algorithms available, MQ has taken advantage of them and provided the ability to use these algorithms in MQ. The technote at http://www.ibm.com/support/docview.wss?uid=swg21639606 summarises all the different environments.

Security: Authorisations for Clustered Queues

- Distributed platforms now have authorisations for non-local clustered queues
  - Making it consistent with z/OS
  - Also consistent with Topic authorisations

- So there is no longer a need to give access to the cluster transmit queue

```
setmqaut –m QM1 –t queue –n SYSTEM.CLUSTER.TRANSMIT.QUEUE
-p mquser +put
                          BECOMES
setmqaut –m QM1 –t queue –n CLUSTER.Q1 –p mquser +put
```

- For fully-qualified puts to remote cluster queues, grant authorisation to the remote queue manager instead
  - A new pseudo-object known to the OAM

```
setmqaut –m QM1 –t rqmname –n QM2 –p mquser +put
```


Security: Authorisations for Clustered Queues

N O T E S

- Access to non-local clustered queues can be authorised at a more granular level than previously. This new function matches something that was already available on z/OS and for Topics on all platforms.
- An object does not need to exist in order to associate ACLs with it, thus authorising a user to open a cluster queue called “CLUSTER.Q1” will check the same profile, a queue profile called “CLUSTER.Q1” whether the object is hosted locally or remotely.
- In addition, to cover the case where an application is putting to a fully qualified queue, a common pattern for server applications using ReplyToQ and ReplyToQMgr fields, a new pseudo-object, the “remote queue manager” is known by the OAM, and authorities are applied to it instead of the transmission queue that will be used to send the message to the remote reply queue.
- The combination of these two changes; local profiles for non-local objects; and ‘remote queue manager’ profiles for replies, means that the cluster transmission queue no longer has to be accessible to anyone wanting to use clustered queues.
- This does not remove the need for controlling access at the receiving queue manager as well (for example by using PUTAUT(CTX)) but it makes it easier to manage the sending side.
- This function can be disabled by a switch in the ini file. It is not a queue manager attribute as we expect customers to migrate to the new model once and then not revert, so it does not need to be as dynamic as real attributes can be.

Application Activity Reports

- New set of events to report on MQI operations by applications
  - One PCF event may contain multiple MQI operations
- Configurable in granularity
  - Amount of data
  - Which applications
- Enables scenarios such as
  - Application audit trail
  - Message duplication
  - Resource usage: which queues or topics are actually being used
  - Problem Determination: most recent MQI calls by applications
  - Application Coding Standards: does everyone use the MQI in the recommended way
  - And more …
- On all Distributed platforms

Extract from Report

```
MonitoringType: MQI Activity Trace
QueueManager: 'V71'
Host Name: 'rockall.hursley.ibm.com'
CommandLevel: 710
ApplicationName: 'WebSphere MQ Client for Java'
ApplicationPid: 18612354
UserId: 'mquser'
ConnName: '9.20.95.106'
Channel Type: MQCHT_SVRCONN
Platform: MQPL_UNIX
================================================================================
Time     Operation  CompCode MQRC  HObj (ObjName)

10:04:09 MQXF_INQ   MQCC_OK  0000  2
10:04:09 MQXF_CLOSE MQCC_OK  0000  2
10:04:09 MQXF_OPEN  MQCC_OK  0000  4 ()
10:04:09 MQXF_INQ   MQCC_OK  0000  4
10:04:09 MQXF_CLOSE MQCC_OK  0000  4
10:04:09 MQXF_OPEN  MQCC_OK  0000  4 (SYSTEM.DEFAULT.LOCAL.QUEUE)
10:04:09 MQXF_INQ   MQCC_OK  0000  4
```

SupportPac MS0P V7.1

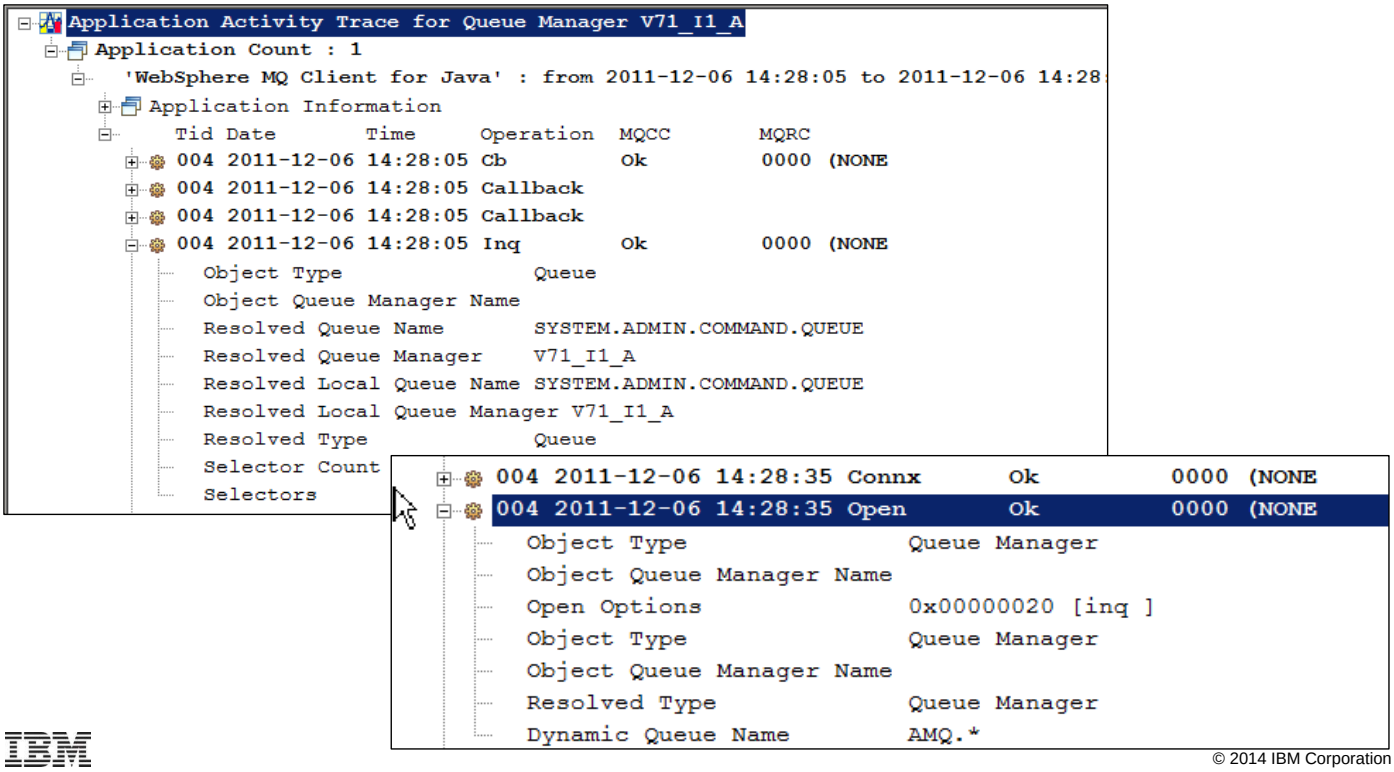

Application Activity Reports

N O T E S

- New for the Distributed platforms is the ability to report on all the MQI operations from an application. This is similar to an MQI trace, but differs in several respects.
  - It is running “inside” the queue manager so has access to more than just the MQI parameters passed by the application. There is some additional information reported, such as the real queues used by the application, not just the name passed to MQOPEN.
  - The output follows the same style as many other reports, in that it is a set of PCF events, where each event holds details about multiple MQI calls
- Applications and their relationships and resources can be analysed without inspecting the application source code.
- Like other events, it is possible to redefine the event queue to be a topic alias so multiple consumers can work with these messages.
- An ini file defines the required granularity – you can have reports of all message data for all applications, but that might be excessive. Changes to the ini file can be made dynamically without restarting the queue manager; just cycle the queue manager attribute that determines whether or not these reports are to be collected.
  - You can also permit applications to disable their own data collection
- Now these reports are generated, many interesting requirements can be met by analysing or using the data. A sample program (source included) formats these events and you can use this as the basis of more sophisticated tools.

Clustering

- "Bind on group"
  - All messages within a logical group are routed to the same queue manager
  - Workload balancing is done for each group
  - Simpler for applications that use message groups
  - Previously would have had to close and reopen the queue
- New option in the MQI and DEFBIND attribute for queues
- Once a group has started its path to a selected queue manager, messages in that group will not be reallocated in the event of a failure
- New sample **amqsclm** to monitor queues and redistribute delivered messages
  - If a queue has no getters, block further deliveries and redistribute existing messages
  - Includes source code, so easy to modify
- Split Transmit Queue Requirements
- Separation of Message Traffic
  - With a single transmission queue there is potential for pending messages for cluster channel 'A' to interfere with messages pending for cluster channel 'B'
- Management of messages
  - Use of queue concepts such as MAXDEPTH not useful when using a single transmission queue for more than one channel
- Monitoring
  - Tracking the number of messages processed by a cluster channel currently difficult
  - Some information available via Channel Status

Clustering – Split Transmit Queue

- With V7.5 a queue manager can automatically define a PERMANENT-DYNAMIC queue for each CLUSSDR channel.
  - Dynamic queues based upon new model queue "SYSTEM.CLUSTER.TRANSMIT.MODEL"
  - Well known queue names: "SYSTEM.CLUSTER.TRANSMIT.<CHANNEL-NAME>"
- Controlled via attribute affecting all cluster-sdr channels on the queue manager

```
ALTER QMGR DEFCLXQ( SCTQ | CHANNEL )
```

- Also have manual definitions
  - Multiple queues can be defined to cover all, or a subset of the cluster channels.

```
DEFINE QLOCAL(APPQMGR.CLUSTER1.XMITQ)
      CHLNAME(CLUSTER1.*) USAGE(XMITQ)
```

- Automatic and Manual are not mutually exclusive
  - They could be used together

MQ Clients

- A client is now available on System i enabling connectivity from C and RPG programs without needing a local queue manager
  - Platform already had a Java client
- MQI libraries like libmqm connect to local and remote queue managers
  - Smart switching for clients, as well as handling multi-version systems
- API Exits available in C clients
  - Same interface as available for local binding applications

MQ Clients

N

O

T

E

S

- The System i platform gets a C-based client, like other Distributed platforms. This complements the Java client that is already available
- Part of the “smart” API switching libraries that are needed to support multi-version installations can now also handle the differences between local and client connections. This makes it simpler to develop new applications as you do not need to compile/link them differently for the different environments.
- API Exits are also now available for the C client libraries, matching the interfaces on the server side. There are some minor differences, such as how XA verbs are handled, but these should not affect existing exits.

Java Application Identification

- Java client applications now fill in APPLTAG field
- No longer appear as “WebSphere MQ Client for Java”
- Application-provided property
- Or the Main class

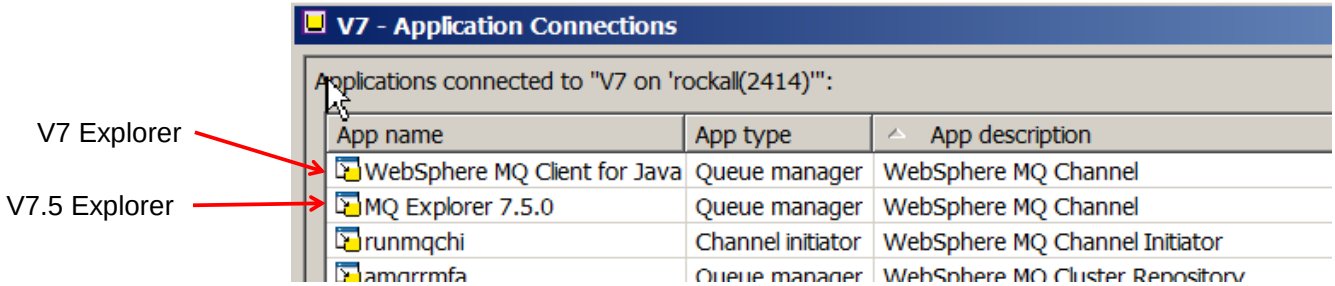

MQ Clients – Multicast

- Publish/Subscribe is enhanced to support multicast communication
  - Uses technology from the MQ Low Latency Messaging product
  - So it is interoperable with LLM
- Provides new Quality of Service
  - Low latency with high fan-out
  - Provides higher speeds for non-persistent messages
  - Provides higher availability as queue manager can be removed without affecting flow
  - Provides “fairness” as all recipients of a message get it at the same time
  - Higher scalability as additional subscribers cause no additional traffic
- Mapping MQ topic space to multicast group addresses
  - Can have mix of multicast and queue-based subscribers
  - Topic objects have associated COMMINFO objects to define addresses and other attributes
- Supports direct communication from publisher to subscriber, bypassing QMgr
- Queue manager maintains status and statistics for monitoring

Multicast Publish/Subscribe

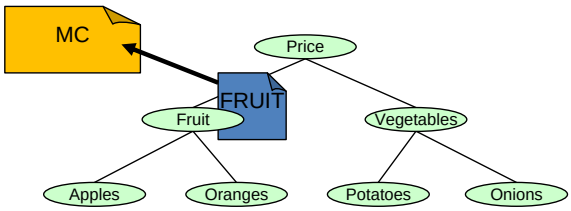

- Normal MQ connection still required
- Messages flow directly between clients
- Subscriber matching done at each client

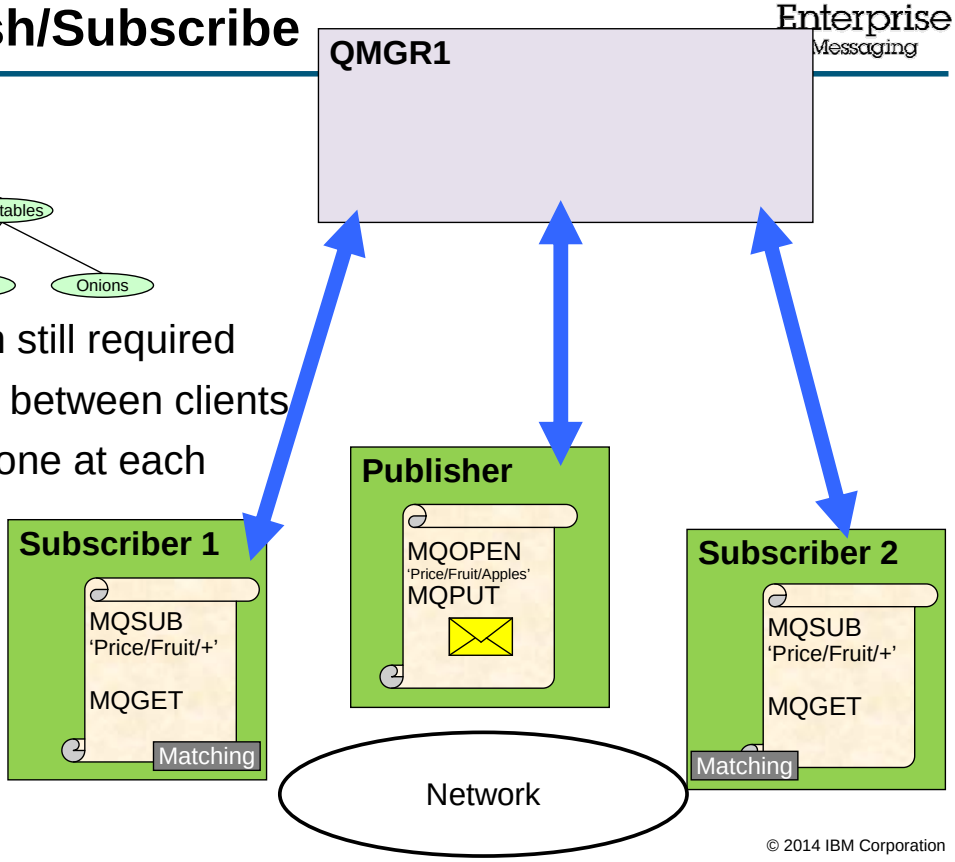

Multicast Monitoring

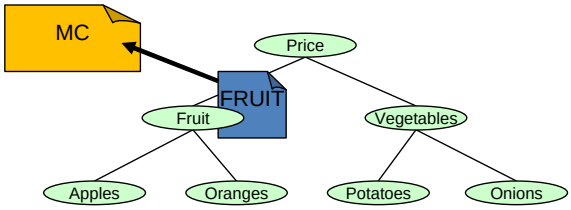

- Monitor interval in COMMINFO
- Monitor messages update TPSTATUS
- Event messages created if required based on COMMEV

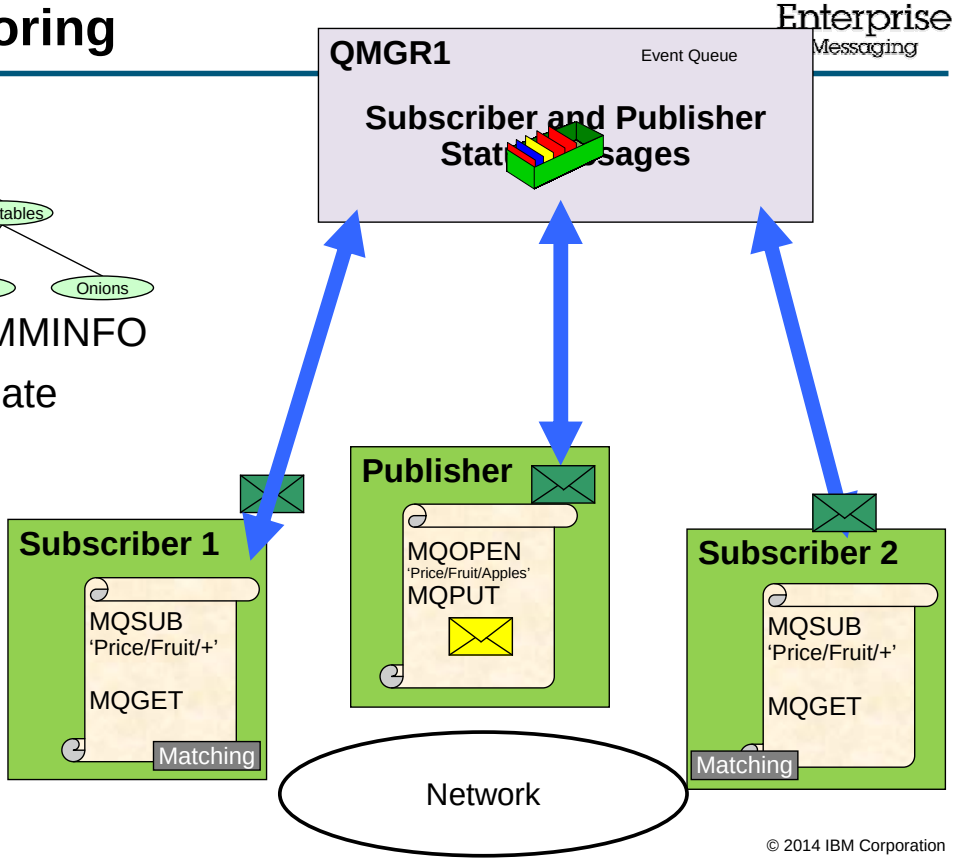

Multicast Bridging and Queueing

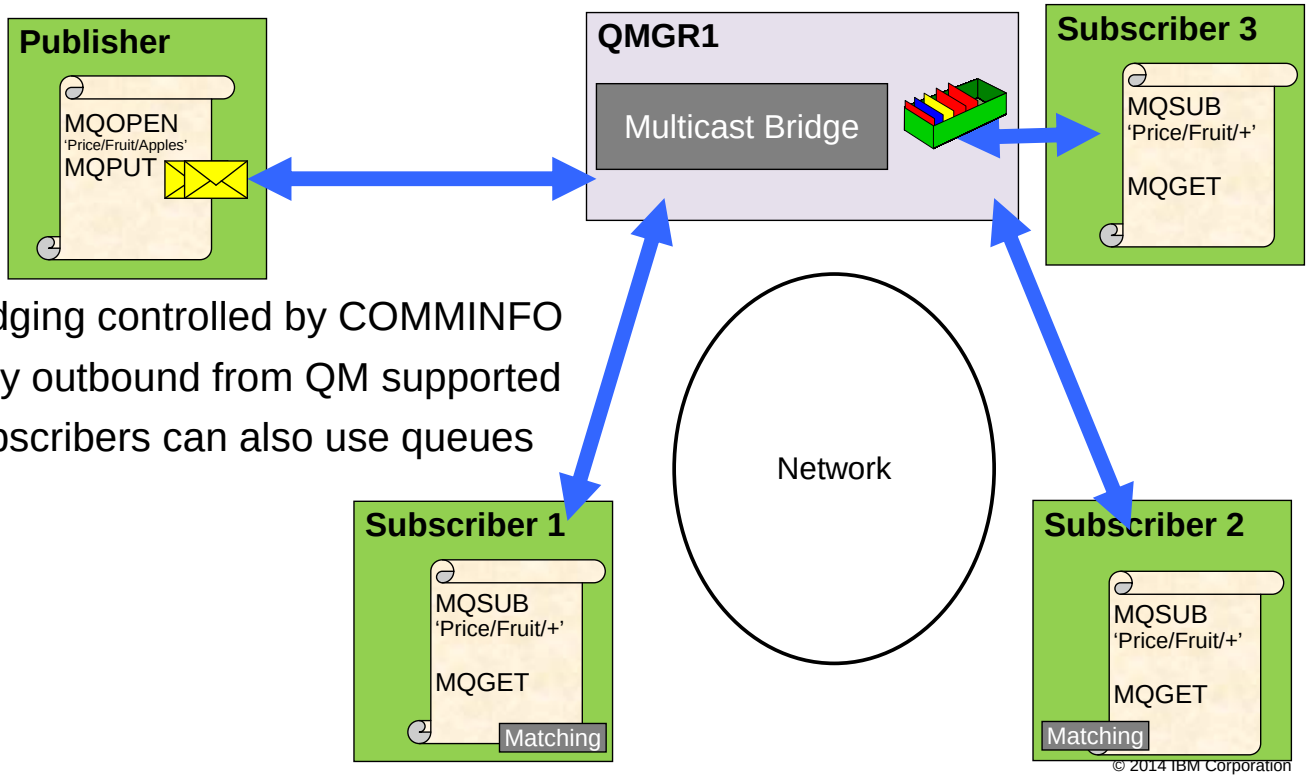

- Bridging controlled by COMMINFO
- Only outbound from QM supported
- Subscribers can also use queues

Multicast MQI Support

NOTES

- Supported
  - Normal MQI
  - MQGET & MQCB
  - Message properties
  - MQSUB with selection by message property
  - Managed subscriptions
  - Data conversion (MQGMO_CONVERT)
- Limitations
  - C client only
  - Pub/Sub only
  - No transactionality or persistence
  - No durable subscriptions
  - No message grouping or segmentation
- Uses LLM technology but does not include all LLM features

Channels

Enterprise Messaging

- See the MQ version of connecting partner
  - Level of clients and queue managers available in channel status
  - For example a V7.0.0.1 client shows as RVERSION(07000001)
  - Can distinguish Java, C, .Net client programs
  - Helps administrator determine whether partner needs upgrading
- Distributed platforms now use DISCINT to disconnect idle clients
  - ClientIdle qm.ini parameter ignored
  - Consistent with z/OS
- Alternative channel batch control based on byte counts
  - BATCHLIM attribute
  - Useful when a transmission queue holds mix of large and small messages
  - Can make batch time (latency) more consistent
  - Batch is ended when first of either bytes or messages transferred reach configured limit
- Per-channel control of Dead Letter Queue
  - New channel attribute USEDLQ(YES|NO)

IBM

Channels

Enterprise Messaging

N O T E S

- Some small but useful enhancements to channel controls. These are not all of them!
- The RVERSION and RPRODUCT values on channel status can tell an administrator what is connecting to a queue manager. The information has been sent between systems since V7.0, and is now exposed to users. Any client or queue manager that is at V6.0 or older will not send this data, so the lack of detail will indicate old systems.
- Both z/OS and Distributed platforms have ways of forcing a client to be disconnected when it has done no work for a while; with V7.1 those mechanisms are made consistent with use of the DISCINT channel attribute
- Traditionally, channels commit a batch after having sent 50 messages or when they reached an empty transmission queue . The amount of data that might be sent with 50 messages could vary wildly from, for example, 50 * 1K to 50 * 100MB depending on the pattern of messages. This means that there is no way to tell the channel to commit the batch sooner when some of these large messages appear and the problem will appear as a slow channel due to the need to re-transmit a very large quantity of data if there is a network outage. Adding a control based on bytes makes the transmission time more consistent. There is no precedence between BATCHLIM and BATCHSZ; whichever value is reached the first will cause the batch to be ended
- A per-channel control can be set on whether DLQ processing should be followed. A channel with USEDLQ(NO) will stop on error, even if there is a qmgr-defined DLQ

IBM

z/OS Performance and Availability

Enterprise Messaging

- z196 Scaling improvements for both non-shared and shared queues
  - Have successfully processed more than ONE MILLION non-shared messages/sec through a single queue manager
  - Have also successfully processed 150K shared msgs/sec with 3 queue managers
- Improved performance by using SMDS for large messages on shared queues
  - Using DB2 BLOBs to store large (>63KB) messages is expensive
- Shared Message Data Sets (SMDS) removes DB2 for large message storage
- Shared VSAM datasets increase shared queues capacity and performance
  - All queue managers in the QSG can access the datasets
- CF Structure message reference still controls locking, ordering, deletion etc.
  - So every message still has a “pointer” in the CF
- Rules control offload message size and % Structure-full offload trigger
- Structure rebuild when connectivity to CF is lost improves availability of Shared Queues
- GroupUR function from MQ V7.0.1 for Distributed QSG connections available for CICS usage
  - CICS 4.2 can use this to enhance the MQ Group Attach originally provided in CICS 4.1

IBM

Shared Message Data Set Concepts

Enterprise Messaging

N O T E S

- Offloaded message data for shared messages is stored in data sets.
- Each application structure has an associated group of shared message data sets, with one data set per queue manager.
  - Named using DSGROUP parameter on CFSTRUCT definition
- Each queue manager owns a data set for each structure, opened for read/write access, which it uses to write new large messages.
- Each queue manager opens the data sets for the other queue managers for read-only access, so it can read their message data
- When a message with offloaded data needs to be deleted, it is passed back to the queue manager which originally wrote it, so that the queue manager can free the data set space when it deletes the message.
- Messages too large for CF entry (> 63K bytes) are always offloaded
- Other messages may be selectively offloaded using offload rules
  - Each structure has three offload rules, specified on the CFSTRUCT definition
  - Each rule specifies message size in Kbytes and structure usage threshold
  - Data for new messages exceeding the specified size is offloaded (as for a large message) when structure usage exceeds the specified threshold
  - Default rules are provided which should be sufficient for most cases and can be set to dummy values if not required

IBM

Loss of CF Connectivity

NOTES

Admin Structure

- Queue managers will tolerate loss of connectivity to the admin structure without terminating if
  - The QMGR CFCONLOS attribute is set to TOLERATE
  - All the queue managers in the QSG are at V7.1
- All queue managers in the QSG will disconnect from the admin structure, then attempt to reconnect and rebuild their own admin structure data.
- If a queue manager cannot reconnect to the admin structure, for example because there is no CF available with better connectivity, some shared queue operations will remain unavailable until the queue manager can successfully reconnect to the admin structure and rebuild its admin structure data.
- The queue manager will automatically reconnect to the admin structure when a suitable CF becomes available on the system.
- Failure to connect to the admin structure during queue manager startup is not tolerated, regardless of the value of CFCONLOS.

Application Structure

- Queue managers will tolerate loss of connectivity to application structures if
  - They are at CFLEVEL(5)
  - The CFCONLOS attribute is set to TOLERATE
- All queue managers that lose connectivity to an application structure will disconnect from the structure
- The next action depends on whether it is a partial or total loss of connectivity
  - Loss of connectivity is partial if there is at least one system in the sysplex that still has connectivity to the CF that the structure is allocated in.
  - Loss of connectivity is total if all systems in the sysplex have lost connectivity to the CF that the structure is allocated in.
- In the case of total loss of connectivity
  - The structure will (probably) need to be recovered using the RECOVER CFSTRUCT command.
  - Non-persistent messages will be lost.

IBM

Partial Loss of CF Connectivity: V7.1

NOTES

- In the case of partial loss of connectivity, queue managers that lost connectivity to the structure will attempt to initiate a system-managed rebuild in order to move the structure to another CF with better connectivity
- If the rebuild is successful, both persistent and non-persistent messages will be copied to the other CF
- Queue managers that didn't lose connectivity to the structure may experience a slight delay during system-managed rebuild processing, but shared queues will remain available
- If an application structure cannot be reallocated in another CF with better connectivity, queues on the structure will remain unavailable until connectivity is restored to the CF that the structure is currently allocated in
- Queue managers will automatically reconnect to the structure when it becomes available

IBM

CICS and WMQ - status

- Maintenance for CICS TS 3.2 & 4.1 needed for exploitation of new V7 API
  - pub/sub, message properties and async. consume
- CICS 4.1 provides MQ Group Attach with WMQ V6 and upwards
  - This (of course) requires an additional Qmgr in the LPAR
  - CICS chooses the Qmgr to connect to;
  - A Resyncmember(yes) setting overrides group attach if indoubts are outstanding
    - Forces a reconnect to the last Qmgr (which, of course, may still be down)
- CICS 4.2 provides extended MQ Group Attach when used with WMQ V7.1
  - Exploiting the GroupUR feature available from WAS in V7.0.1
  - This also (of course) requires an additional Qmgr in the LPAR
  - WMQ chooses the Qmgr
- WMQ V7.1 also provides 64-bit support for Java Classes for CICS

IBM

IMS and WMQ (1)

- Transaction Expiration (IMS Bridge)
  - A service parameter in V7.0.1 (GA) can be used to propagate the remaining 'time-to-live' for a message from the MQMD expiry field to the OTMA header when being sent to IMS.
    - If a transaction expires in IMS and a response is required, MQ will generate an expiration report – just like a regular MQ expiry report – which will be placed on the Reply-to Queue.
  - WMQ V7.0.1 PTFs required
    - WMQ V7.0.1 – PM47795
  - IMS PTFs required
    - IMS V11 – PM05984
    - IMS V10 – PM05985
    - IMS V12 - TBD
  - WMQ V7.1 features
    - WMQ V7.1 provides the expiration capability provided IMS PTFs are applied
  - Result is that with MQ V7.0.1 and V7.1 we have consistent expiry reporting regardless of whether the transaction expires on the MQ queue or in IMS

IMS and WMQ (2)

- IMS Resource Monitoring (IMS Bridge Flood Prevention)
  - Introduced in IMS v10 via PK70458
- MQ V7.1 will respond to health status reporting from OTMA
  - If OTMA reports a flood warning, WMQ V7.1 will reduce the rate that new transactions are sent to that specific IMS partner
    - Other IMS partners are not affected
  - All TPIPEs will continue to be suspended if a full flood condition occurs
    - TPIPEs will be resumed when the flood is relieved
  - On relief of a flood (or flood warning) WMQ V7.1 will increase the message traffic being sent to IMS traffic in a gradual manner to avoid immediately re-flooding IMS

Scalability & Performance – Distributed platforms

- Performance measured and improved for a range of scenarios
  - Hardware capabilities have evolved over years to have more CPUs, more memory etc
  - MQ topologies have evolved to have more clients and larger/fewer queue managers
- "Fastest MQ ever": better performance than V6 and V7
- Multicast faster than traditional non-persistent
  - Over 5x for one-many publications
- Performance reports now available from SupportPac site
- Design changes to MQ Explorer reduce its footprint and improve performance
- Now does not include full Eclipse development workbench
  - But Explorer can be easily added to other Eclipse installations and products
- Many Explorer installs are supported within the overall multi-version support
  - But each Explorer only fully manages queue managers associated with its own installation
  - Use client connections for other installation queue managers on same machine

	V7.0.1	V7.1
Time to install MS0T	203 seconds	92 seconds
Startup Time	6 seconds	4 seconds
Connect to 100 queue managers	At least 53 seconds	7 seconds
Enable and disable Sets for 100 queue managers	35 seconds	1 second

Management of Distributed platforms

Enterprise Messaging

- New integrated command (dmpmqcfg) to extract configuration
  - Fulfills the role that MS03 (saveqmgr) has done over many years
    - Backup your configuration, change control, rebuild systems etc
  - MAKEDEF already available on z/OS
  - Different syntax than MS03, but similar function
- MQSC commands equivalent to setmqaut/dspmqaut
  - So you don't need to drop out of the command interface to modify security
  - Can simplify scripting of configuration changes
  - No current plans to remove *mqaut commands
- Multi-instance Queue Managers on Windows
  - The need for domain controllers ("domainlets") has been removed
  - New option when creating queue managers to define ownership

IBM

Management of Distributed platforms

Enterprise Messaging

N O T E S

- Probably the most commonly-used SupportPac has been MS03 (saveqmgr). In MQ V7.1, the concept is now built into the product and formally supported. The new dmpmqcfg command has a slightly different syntax, but fulfills the same role of extracting a queue manager's configuration and displaying it in MQSC syntax.
- Dmpmqcfg can connect to local queue managers or use client connections to remote systems.
- New MQSC commands are available that are equivalent to the set/dsp/dmpmqaut commands. These may be more convienent when you are already inside runmqsc, rather than dropping out to the command line, and certainly more convenient when scripting configuration changes. The new dmpmqcfg for example can put authorisations in a single MQSC script for replay, rather than having to run separate commands.
- On Windows, the requirement for multi-instance queue managers to be domain controllers (even if limited in scope eg "domainlets") has been removed. When a queue manager is created, you can now name a group that both machines share knowledge of, and that group gets appropriate ownership of resources such as the files and directories that make up the queue manager.

IBM

© 2014 IBM Corporation

MQ Cloud Support: Pre-Connect Exit

- Supports movement by some to “Utility Compute”, Private Cloud configs, etc.
  - Rapid provision of applications allied with need to further decouple Client/Server connectivity
  - Server applications might move location – new addresses or queue managers
- MQ Client connects to a “service” rather than specific Queue Manager
- Can transparently change location of MQ server-side applications
  - No client code changes needed
  - No configuration files need to be updated at the client machine
  - JMS/XMS applications already do this via JNDI lookup
- Exit run during MQCONN queries a repository to discover real location
  - MQ V7.1 incorporates the LDAP implementation from SupportPac MA98
  - Operates similar to CCDT but without many of the downsides
    - Record based instead of file based
    - No copying of CCDT files around

MQ 7.5: Enhancements to newly-integrated components

- Managed File Transfer
  - Logger can now write to a file
- AMS
  - V7.0.1.2 enhancements
    - Supports SHA-2 Digest algorithms
    - Command and Configuration Events for Policy changes
    - Audit trail of who has changed configuration
  - SVRCONN interception

2013 Updates

IBM

2013 enhancements: Enterprise messaging & social

Enterprise Messaging

Enhancement	*Benefits*	*Details*
MQ V7.1 VUE for z/OS	MQ Value unit edition product introduced for z/OS OTC pricing for zNALC partition workloads	One time charge (OTC) price metric instead of MLC Run new workloads that are qualified to run on System z New Application License Charge (zNALC) LPARs Includes code and entitlement for WebSphere MQ Client Attach, allowing FoC connections from MQ Clients to z/OS
New Twitter ID "@IBMIoT"		Internet of Things focused communications channel on twitter Bringing you news from around the Internet of Things rapidly expanding space
The developerWorks Messaging communities begin their facelift	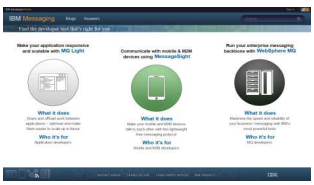	Top level community introduced … **https://www.ibmdw.net/messaging/** Links down to each of the three main communities Transformation continues

IBM

Internet Pass-Thru – SupportPac MS81

Enterprise Messaging

- Previous MQIPT 2.0 release was in 2008
  - Many new operating systems have been released since then
  - New cryptographic algorithms since then
- Redesigned to use the JRE-provided Java Secure Sockets Extension (JSSE)
  - Previously used a separate SSL implementation
- This gives consistent SSL/TLS support
  - Also has consistent certificate DN attributes with MQ V7.1
- US NIST standard SP800-131A recommends:
  - Support for SHA-2 (e.g. SHA-256, SHA-384 and SHA-512)
  - Elliptic Curve encryption

IBM

JMS in Liberty Profile

Enterprise Messaging

- WAS Liberty profile does not include a copy of the MQ JMS client
  - For consistency with other resource providers such as databases
  - Provider specific classes come from the resource provider
- The RA provided with MQ can be easily configured in Liberty
  - Define one variable in server.xml to tell the *wmqJmsClient-1.1* where the client is

```
<variable name="wmqJmsClient.rar.location"
value="${shared.resource.dir}/wmq/wmq.jmsra.rar"/>
```

- Requires the MQ 7.5.0.2 rar or a fix on top of 7.5.0.1
- More information is at
  - https://www.ibmdw.net/wasdev/2013/06/14/using-websphere-mq-with-the-liberty-profile/
  - http://www.ibm.com/support/docview.wss?uid=swg21633761

IBM

New client support for HP NS (aka Tandem)

Enterprise Messaging

- Extended Transactional client support at an MQ V7.1 release level
- Full MQ V7.1 function for NonStop applications accessing MQ on other systems
- Client is free and fully supported
  - C / C++ extended transactional client
  - Cobol
  - pTAL support
  - JMS client
- Supports both Guardian and OSS applications
  - including OSS multi-threaded applications
- Available as SupportPac MAT1

IBM

Now available

Enterprise Messaging

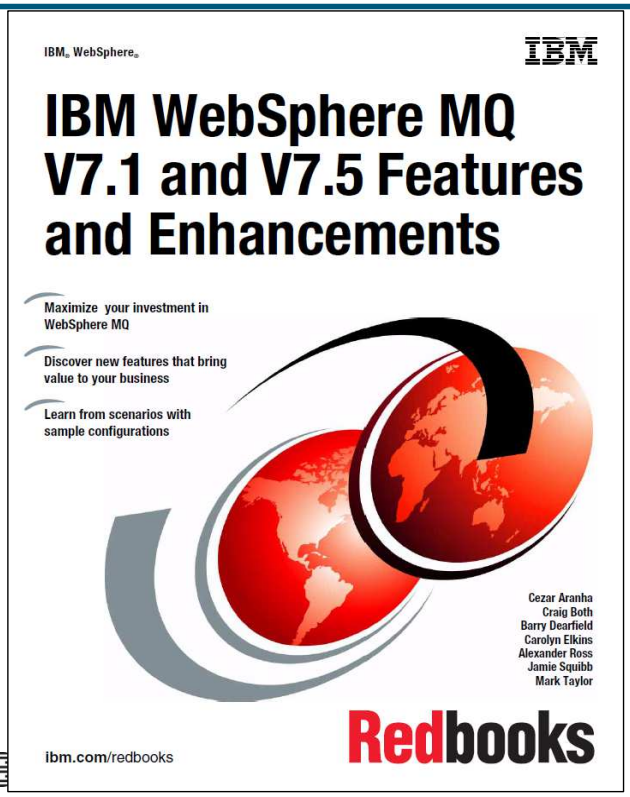

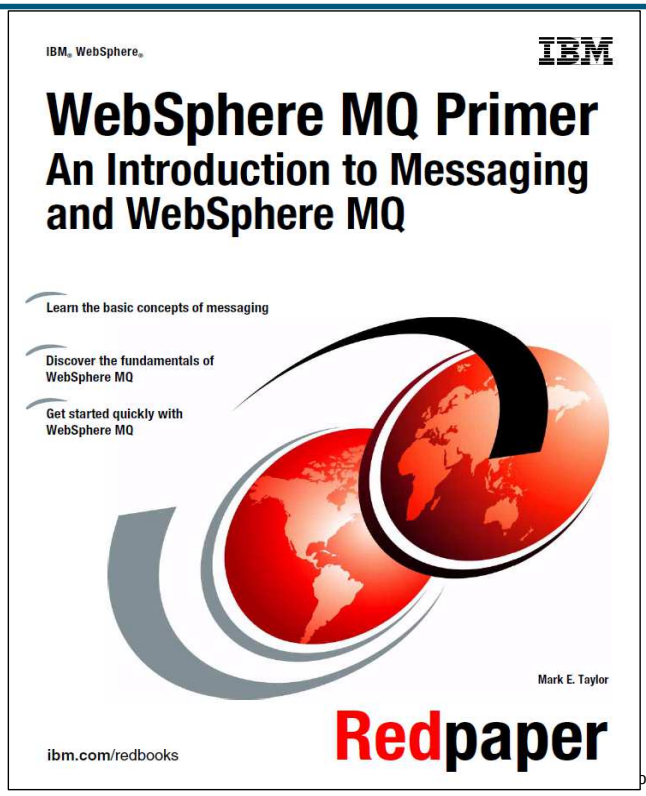

2013 enhancements: Developer

Enhancement	*Benefits*	*Details*
MQ Advanced for Developers – free for devt use!!	Development use only edition Free of charge	Full MQ Advanced V7.5 functionality Available free of charge for development use only Download from developerWorks, MQdev community Best effort community assistance
IBM MessageSight for Developers – free for devt use!!	Development use only edition Free of charge	Virtual MessageSight image, updated to V1.1 Available free of charge for development use only Download from developerWorks, IBM Messaging community Best effort community assistance
Mosquitto - Open source MQTT server	Simple, open access to get on-board with MQTT and MQTT-SN	Proposed Eclipse Mosquitto project for open-source MQTT server for lightweight devt use MQTT-SN (MQTT for sensor networks) specification available
IBM MQ Light – alpha release	Straight forward, easy to use messaging API for writing scalable and responsive applications	Download and feedback via developerWorks community Available free of charge for development use Currently new API available in Java, more languages to come Regular updates planned introducing new features and responding to feedback

IBM MQ Light

- Simplifying Development for developer-led Messaging Applications
- "MQ Light" - a new messaging form factor, easy for developers
  - Straightforward download and unzip – zero install
  - Simple Web UI– no need to be an infrastructure specialist
- Focus on your application – using simplified messaging API
- Support for popular languages and tooling with seamless migration from development into production
- MQ Light runtime not for advanced messaging uses cases, enterprise backbone or highly available/scalable apps

Developer coding in JavaScript, Java, Python, Ruby, C#, PHP

Builds messaging into applications using preferred IDE or CLI and tests in developer sandbox

Download “MQ Light” alpha and try it today

Enterprise Messaging

https://www.ibmdw.net/messaging/mq-light

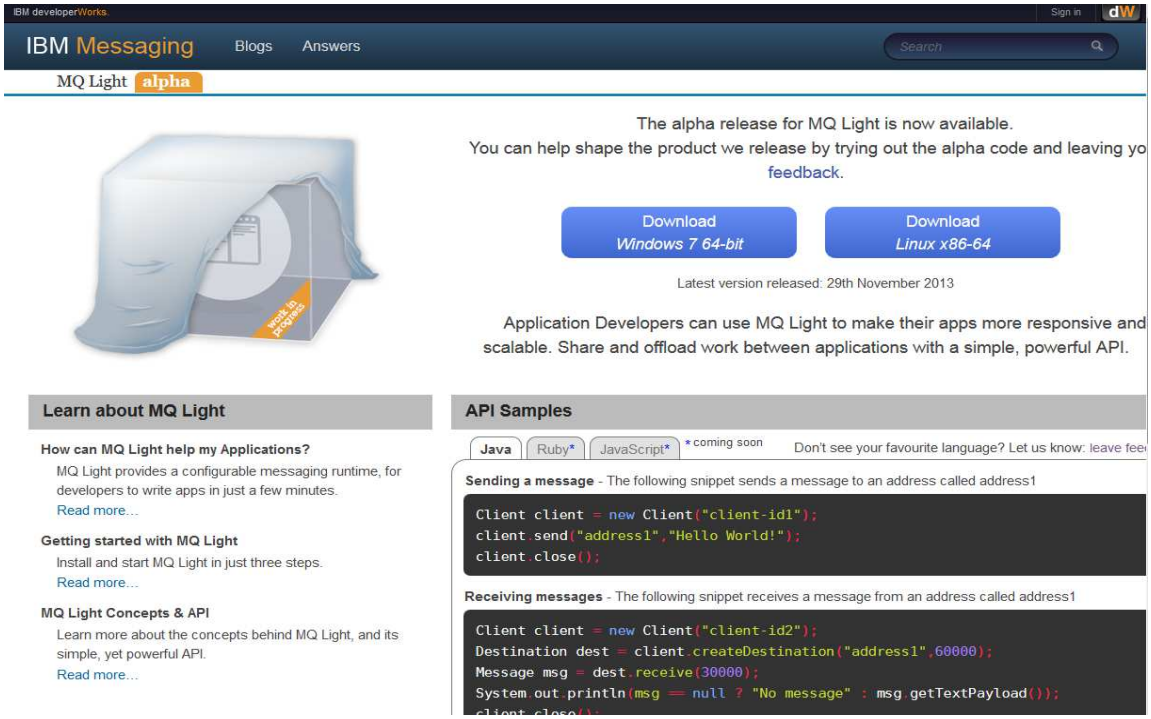

IBM

2013 enhancements: Internet of Things, mobile & m2m

Enterprise Messaging

Enhancement	*Benefits*	*Details*
IBM MessageSight V1.1	High-performance, reliable, and scalable messaging - designed specifically for M2M and Mobile scenarios	Balance mobile workloads across your enterprise application server clusters Single sign-on from mobile devices to your enterprise apps Disconnected app notification that messages are waiting
IBM Internet of Things Cloud – statement of direction	Fully cloud hosted solution for internet of things Freemium, PAYG & custom pricing	Scalable connectivity for small and large numbers of devices Uses industry-standard, MQTT transport Registration of devices for access to the cloud service offering Ability for devices to produce and consume events and messages in near real-time
Mobile Messaging clients	Updated clients for mobile platforms	Updated clients for Android, iOS & hybrid mobile apps

IBM

IBM MessageSight – a messaging Appliance

- Extends the IBM Messaging family with a secure, easy to deploy appliance-based messaging server, optimized to address the massive scale requirements of machine to machine (m2m) and mobile use cases
- A million connections, and millions of messages per second
  - Exploits hardware acceleration for performance

- Designed to sit at the edge of the enterprise
- Can extend your existing messaging infrastructure or used standalone
- Complements MQ - provides an offload/accelerator for edge of enterprise scenarios
- Supports familiar APIs with a mixture of standard and high-speed protocols

IBM MessageSight V1.1 - Major enhancements

- Mobile to Enterprise
  - Balance mobile workloads across your enterprise application server clusters
  - Single sign-on from mobile devices to your enterprise apps
  - Disconnected app notification that messages are waiting
- Security
  - Implements the Federal government recommendations with NIST SP 800-131A strengthened security
  - Uses client certificate identity for authentication and authorization
- Developer Operations
  - New “at-a-glance” system dashboard
  - Flexible work environments supported with remote monitoring of MessageSight from client devices
  - Virtual MessageSight image speeds mobile and machine-to-machine app development

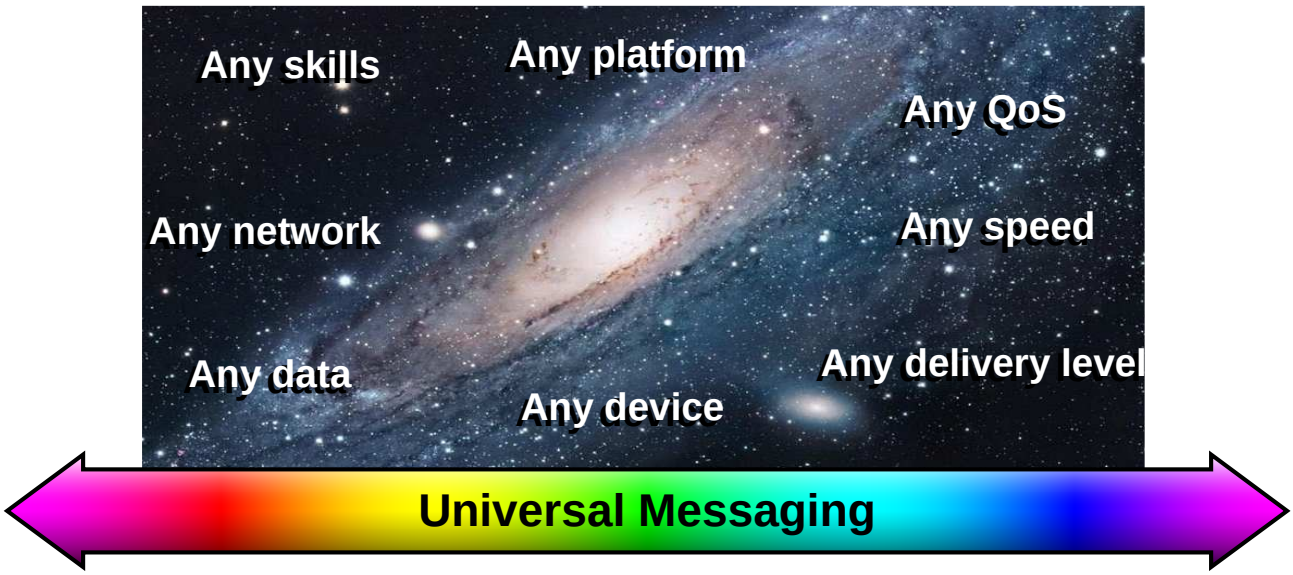

Questions & Answers

Morag Hughson -
hughson@uk.ibm.com