



IBM Software Group

Keeping Channels Up and Running

Morag Hughson and Paul Clarke

hughson@uk.ibm.com , paulg_clarke@uk.ibm.com

WebSphere. software

@business on demand software

© 2003 IBM Corporation

V1.0.1

IBM Software Group | WebSphere software



Preface

N

This document is provided by IBM as recommendations based on IBM's expert knowledge of WebSphere MQ and MQSeries and experience with customers using these products and their environments and goals. Acceptance and implementation of these suggestions is at the reader's discretion. You may hear different opinions from other experts. There are several areas within WebSphere MQ that are subject to varying opinions and may vary according to the specific environment in which WebSphere MQ is implemented.

O

After you, the reader, have become familiar with WebSphere MQ you will be able to determine which suggestions you want to follow or you may arrive at opinions that are appropriate for your specific environment.

T

This document is intended for use by Systems Administrators and any others that will support WebSphere MQ networks. Hopefully it will provide the following benefits:

E


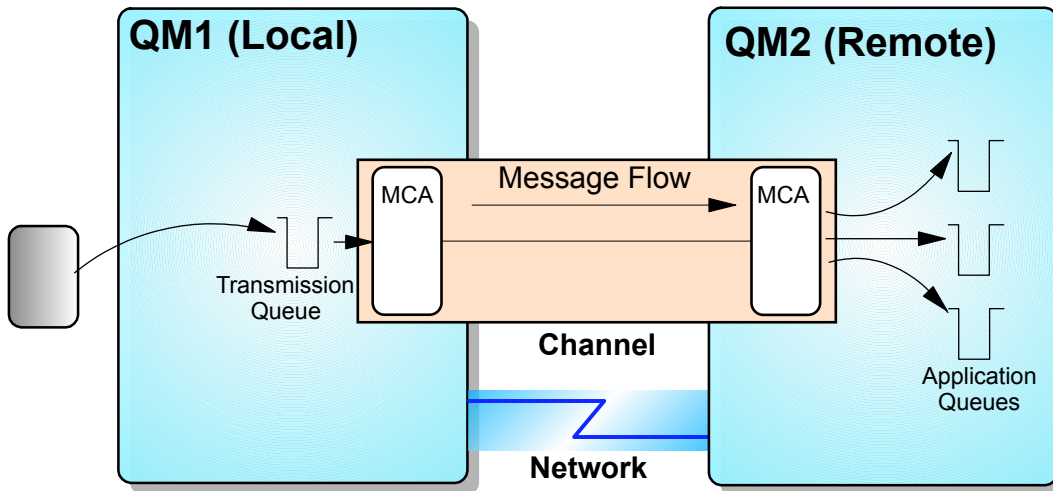
- Give consistency to administration processes
- Provide maximum availability of applications
- Help in avoiding common mistakes that beginners make
- Assist personnel in the early stages of their becoming WebSphere MQ experts
- In general assure a smooth start and success for WebSphere MQ projects

S

Most of the information presented here relates to versions of both MQSeries and WebSphere MQ. Check the documentation for your particular version to ensure that the parameters are supported.

© 2003 IBM Corporation

Channel Architecture



Channel Architecture - Notes

This diagram represents the WebSphere MQ channel architecture. We are going to use this diagram throughout the presentation to illustrate the different channel topics.

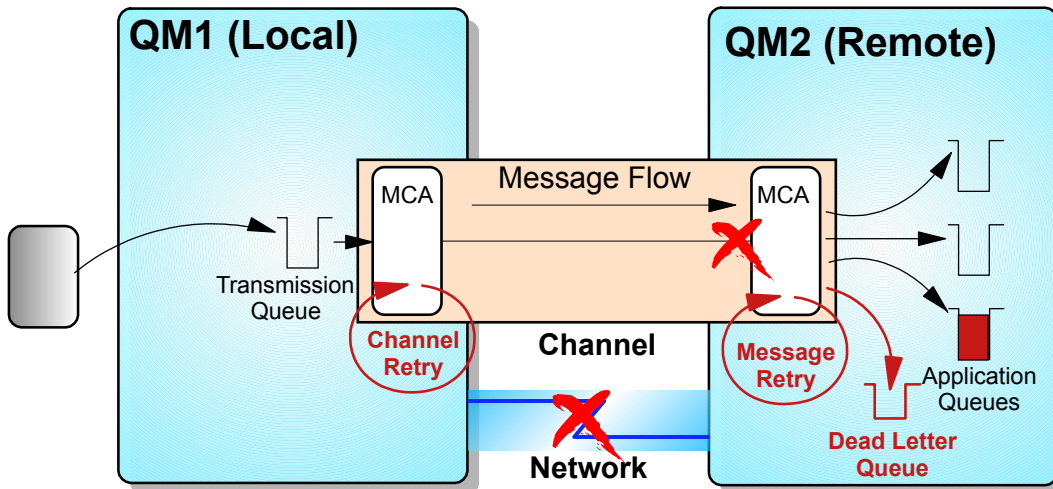
**N
O
T
E
S**

The main part of the diagram shows that applications can put messages (either via remote queue definitions or by explicit addressing) which will land on a transmission queue to be delivered to the remote system. QM1 is our local system in this diagram and QM2 is our remote system. The sending MCA servicing our transmission queue gets the messages off the queue and sends them to the remote system via, the communications protocol specified on the channel definition, where they are received by the receiving MCA and placed on the appropriate target application queue. This pair of MCAs is known as the channel.

That's all very well, but what can go wrong?



What can go wrong?



- Channel can end - remote QMgr ending
- Message delivery not possible - Queue invalid or full
- Network failure
- Performance

What can go wrong? - Notes

This foil acts rather like an agenda of what we are going to discuss today. We will look briefly at several different things that can go wrong and then we will go through each one in more detail on later foils.

N The channel can end. For example, because of the remote Queue Manager terminating. To recover from this type of problem, you can make use of the channel initiator to retry the channels.

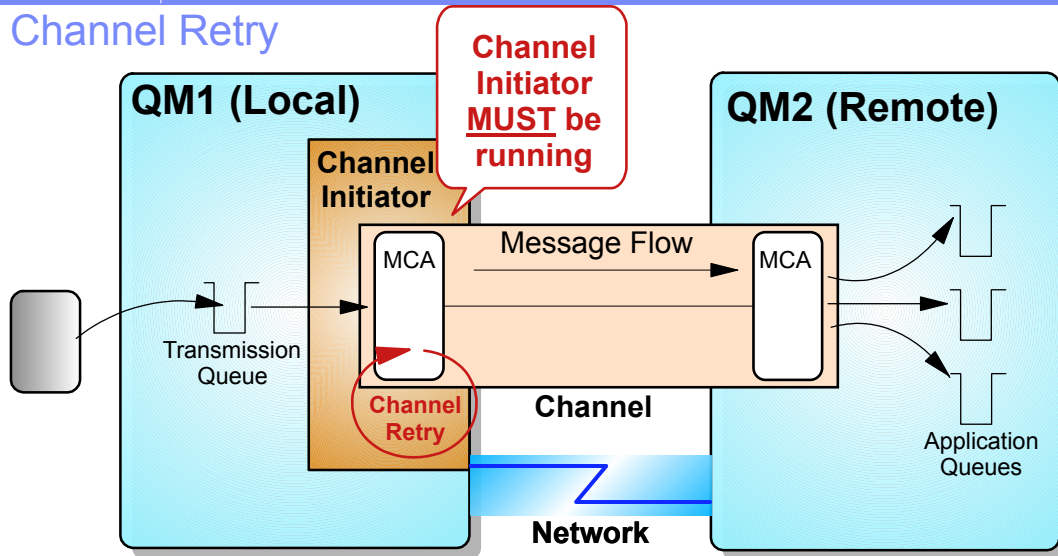
O Message delivery might not be possible. This might be because of an application error, for example the specified queue is invalid, or because of a system problem, for example the queue is full. To continue after this type of problem, you can make use of the Dead-Letter Queue or you can use the Message Retry facility.

T You might suffer a network failure, and depending on the type of failure, there are a variety of ways for your channel to detect this and recover.

E Performance might be an issue, so we will also look at some performance considerations.

S

Channel Retry



```
DEF CHL ....
SHORTRTY(10) SHORTTMR(60)
LONGRTY(999 999 999) LONGTMR(1200)
```

Channel Retry - Notes

The 'catchall' mechanism for ensuring that channels are operating is for WebSphere MQ to restart the channel after a failure. This is known as channel retry. Channel retry can be used on sending types of channels, that is senders, servers or cluster-senders.

N All sender channels should be configured with suitable retry values and the Channel Initiator must be running in order to use the retry mechanism.

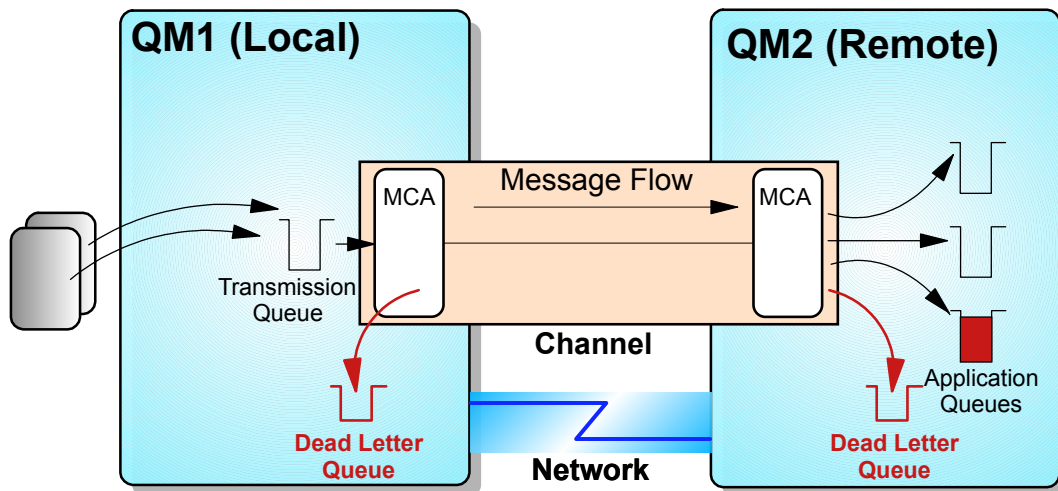
O Retry values should ensure sufficient time for network problems to be resolved, so for example, retrying for 2 minutes is unlikely to be sufficient! There are two different sets of retry values available for your use, a short retry and a long retry. The short retry values tend to be used for a small number of frequent retries, in other words, to attempt to reconnect to the partner after a temporary blip of some sort. The long retry values tend to be used for a larger number of less frequent retries, in other words, if the reason for the channel ending was a problem that is not simply a small blip.

T Be aware that the numbers specified in the long retry values must be enough to cover any outage because if the retry sequence is completed without a successful transfer of messages the channel will require manual intervention to restart it.

E

S

Dead-Letter Queue



- **Storage area for messages which are, for example :-**
 - ▶ **Wrongly addressed**
 - ▶ **Unauthorised**
 - ▶ **Too large**

Dead-Letter Queue - Notes

N It is strongly recommended that Queue Managers have a Dead-Letter Queue. Without a Dead-Letter Queue a simple programming error in the application, by specifying an invalid Queue Name, can be sufficient to bring the Channel down. If several different applications are utilising the same channel, this is clearly unacceptable.

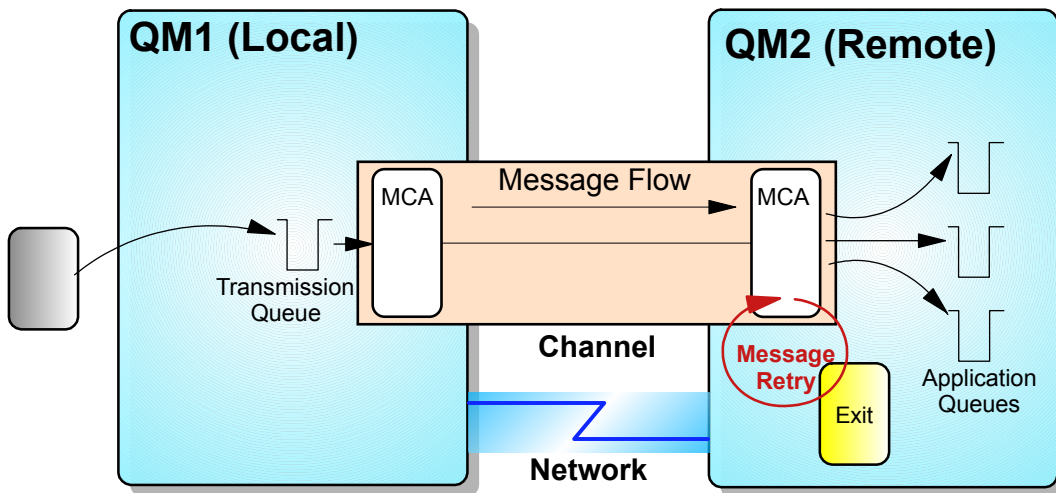
O Note that the DLQ might be used at either end of the channel. At the sending end, the DLQ will be used if a particular message is too large to be sent down the channel. This occurs if the message size exceeds the size of message defined for the channel even though it can be accommodated by the transmission queue. At the receiving end, the DLQ will be used if the message cannot be placed (PUT) onto the target queue. This might occur for any number of reasons, some of which may be transitory. The message retry parameters (defined for a channel) and the Message Retry exit are designed to handle these cases. We will look at Message Retry in a moment.

T It should be noted that the use of the DLQ - at either side of a channel - may affect the order of delivery of messages. If there is some problem with a message which is a member of some sequence then this message will be DLQ'd and the message sequence will be interrupted.

E Some 'lazy' applications require messages to always be delivered in the same sequence as they were sent. These applications require that there is no DLQ. The consequence of not having a DLQ is that a 'bad' message on a channel will terminate the channel and therefore prevent any application from receiving their messages. For this reason it is strongly recommended that installations do have a DLQ and that applications are modified to cope with messages to occasionally be delivered out of sequence. This modification is not difficult and there are various techniques which can be employed such as using the Correlation Id as a counter, having a secondary queue or using message grouping.

S

Message Retry



```
DEF CHL ....
MRRTY(10) MRTMR(1000)
MREXIT(msg-retry-exit-name) MRDATA()
```

Message Retry - Notes

Consider using Message Retry, particularly if you are running without a Dead-Letter Queue.

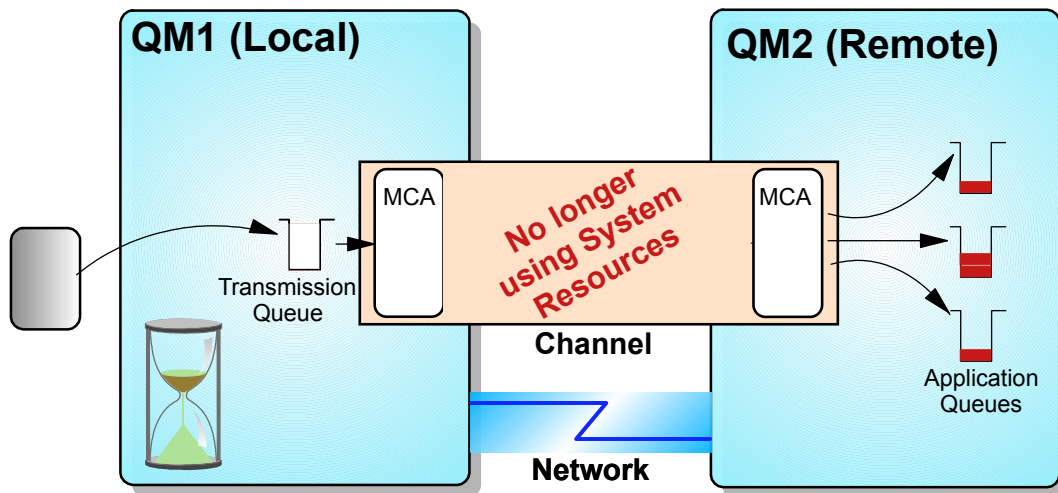
The message retry mechanism can be used with or without an exit. Without an exit the receiving MCA simply retries the failing put on each interval. The values for the time between retries and the number of retries are taken from the channel definition. If used with an exit, the exit is passed these numbers but can change them to be any value. Using an exit with the message retry mechanism allows a more intelligent processing of the failure. You could, for example, redirect messages if they can not be put to target queue.

The use of an exit allows the retry decision to be made on a per message basis.

Message Retry is currently not supported on z/OS.

N
O
T
E
S

Disconnect Interval



DEF CHL DISCINT(6000)

- **Combine with Triggering**

Disconnect Interval - Notes

The Disconnect Interval or DISCINT is the channel attribute on a sending type channel used to control how long a channel should remain running even though there are no messages to send.

N

The value you set depends on several criteria.

- How many channels you have and how critical system resource is
- How reliable your network is
- How expensive the communications session is

O

The ideal value should be longer than the average time between messages to that partner. This means that the channel will disconnect when there is no work to be delivered to that partner.

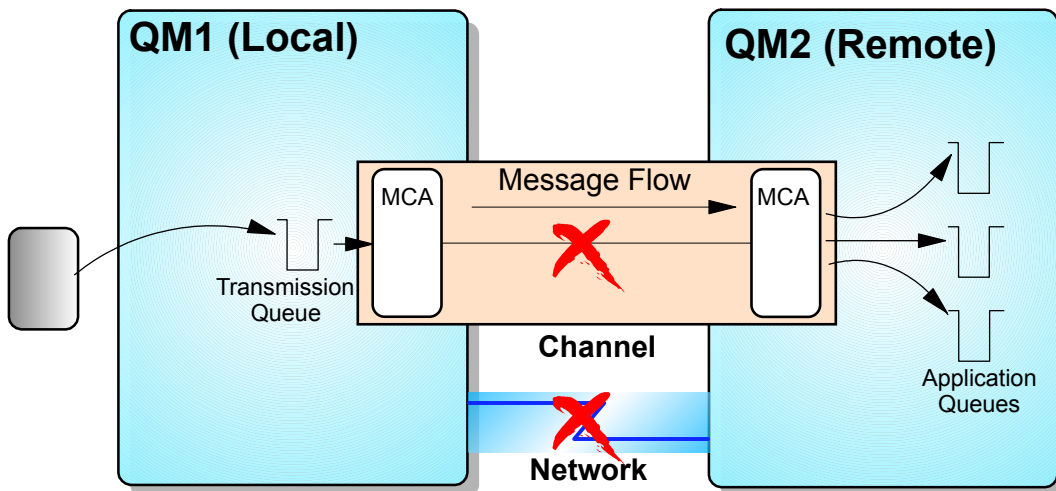
T

You should combine the use of a disconnect interval of your channel with channel triggering to restart the channel. This way your channel is started automatically when there is work for it to do, and ended automatically when there is no more work for it to do.

E

S

Network Failure



- Heartbeats
- TCP/IP Keepalive
- Adopt MCA

Network failure - Notes

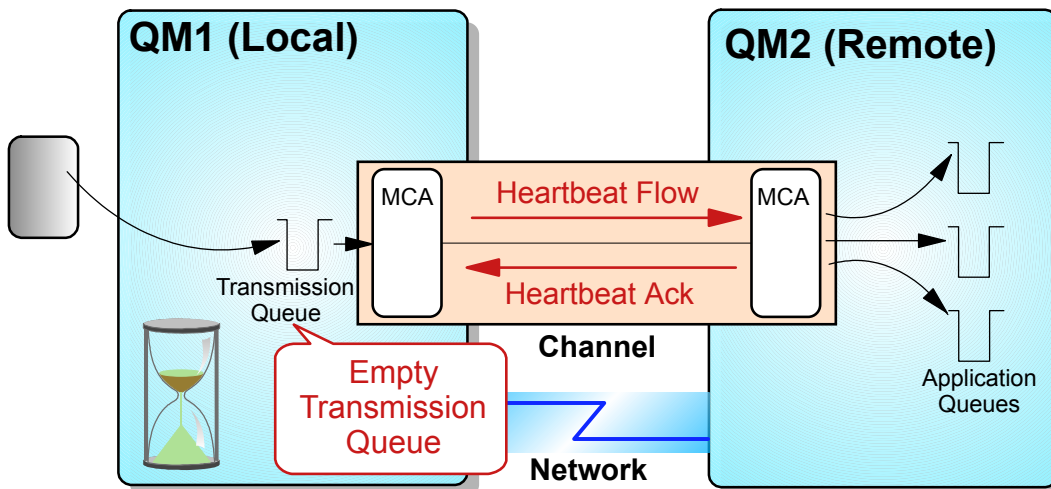
Any network failure reported to the channel will cause the channel to end. If the network failure is not reported to the channel the channel will not end, it may not notice a network failure.

N The Channel is reliant on the lower level communications functions to report any network outages. For example, the resting state of a RECEIVER channel is on a recv call. If the recv call does not return the Channel is unaware that connection to the partner is lost. One of the difficulties with MCAs was recognising in a timely way when the network had actually failed, as the comms protocols didn't always return errors at the right times. Typically, LU6.2 has been much better at this than TCP/IP (it's a more complex protocol).

O Various enhancements to the MCAs have improved the situation considerably. We will look at a few of these over the next few foils.

T
E
S

Heartbeats



DEF CHL HBINT(300)

Free Buffers ✓

Close Queues ✓

Heartbeats - Notes

N

The Heartbeat Interval or HBINT is the channel attribute which controls how often the sending end of the Channel should check that the receiver is ok when there are no messages to send. The value is specified in seconds on both ends of the channel and they negotiate an appropriate value to use at channel startup. This value is the least frequent interval of the two specified. If one or both ends specifies 0 which means no heartbeats, then this value, being the least frequent value possible, is the one chosen.

O

By default channels are defined with a heartbeat interval of 300 seconds (5 minutes). A relatively high value was chosen as the default so that there would be very little impact on the network. In practice a lower value of say 60 seconds is better in a lot of circumstances.

T

This heartbeat is actually a short message sent to the receiving MCA (or CLNTCONN during MQGET) to check that it is still available. It is only sent when there is no other activity on the channel for at least Heartbeat seconds. These messages solve two problems; sender MCAs may now get an immediate error return from the network, instead of sitting on their MQGET; and they also allow receiver MCAs to do some processing such as checking for the termination of the queue manager before sending their response.

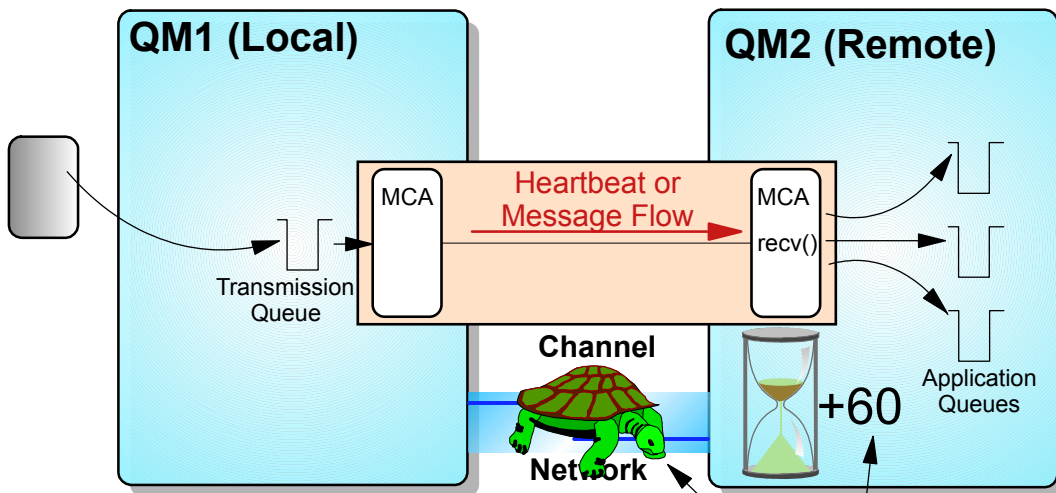
E

Additionally, since the channel is not transferring messages regularly the channel chooses not to maintain its cache of storage and Queues. So, at heartbeat time a channel will also free any unnecessary storage buffers and message areas and close any cached queues.

S

So, for performance, the Heartbeat value should not be too small. But note that low heartbeat values will cause no extra network traffic provided the channel remains busy. Consider an application sending a message every two seconds down a channel. Even a low heartbeat value of 3 seconds will never cause a heartbeat to be sent since there will always be an actual message to send.

Receive Wait Time Out



```
DEF CHL .... HBINT(300)
MQRCVBLKTO: +60
RCVTIME=+60
```

Reflect Network Latency ✓

Receive Wait Time Out - Notes

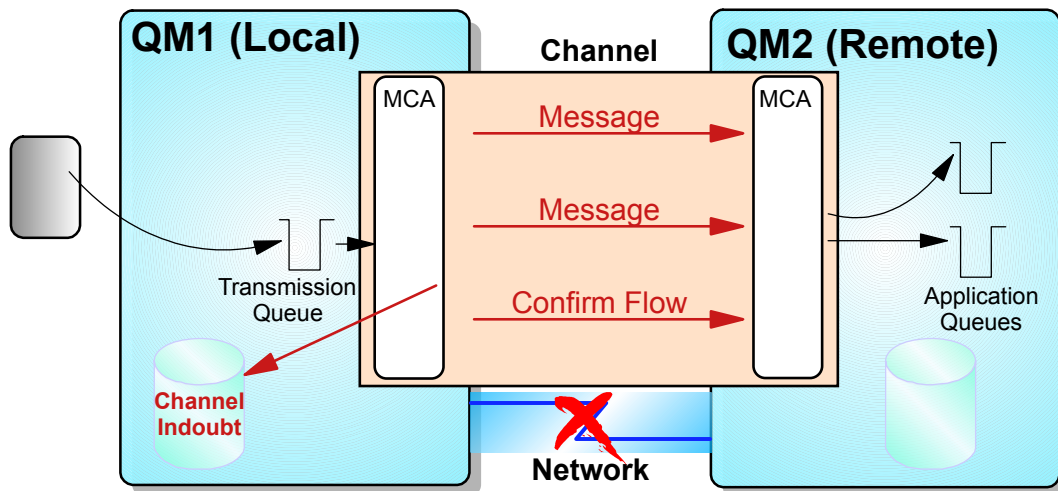
N
O
T
E
S

As we have already mentioned, the heartbeat interval used by the sending end of the channel to periodically send this heartbeat message, is a negotiated value, so both ends know what the value is. Without a heartbeat the receiver end of the channel can not predict when the next piece of data should arrive down the channel since the arrival of messages is non-deterministic. However, with a heartbeat the receiver 'knows' that within the heartbeat interval either a message or a heartbeat should arrive. If it doesn't then the communications session must have broken. This is used as the basis for TCP/IP network failure recognition. For CLNTCONN channels, this time-out is also used while waiting for a response to an MQI command (except when the SVRCONN is on z/OS which does not support the client variant of heartbeats).

The receiver will actually time-out if no data is received within twice the Heartbeat interval if the negotiated Heartbeat Interval is less than 60 seconds, or 60 seconds beyond the negotiated heartbeat interval if it is greater than or equal to 60 seconds, by default, before assuming there has been a communications failure. This qualifier can actually be tuned to reflect your network latency, on the distributed platforms, by changing the environment variable MQRCVBLKTO. There is also another environment variable MQRCVBLKMIN which allows you to set a minimum value to be used when calculating the timeout based on the heartbeats. The behaviour can be set on z/OS in V5.3.1 by the use of XPARM attributes RCVTIME to qualify the timeout and RCVTMIN to set the minimum value.

So.. a small value will detect failures more quickly but may also be more expensive in terms of network traffic and CPU resource if your channel is lightly loaded.

Assured Delivery



- **Synchronisation Data written to disc at both ends**
 - ▶ Allows rescynchronisation after a network failure
 - ▶ Messages in an in-doubt batch cannot be reallocated by clustering algorithm

Assured Delivery - Notes

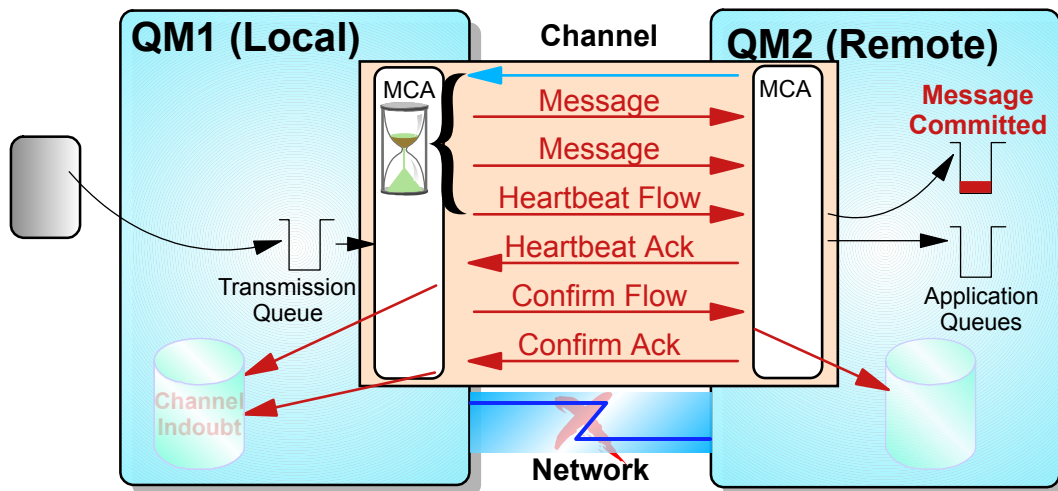
**N
O
T
E
S**

Here we can see the two ends of the channel transferring a batch of messages. Note how the MCA's write data to disk at the end of the batch. This is only done for recoverable batches. The data written contains the transaction id of the transaction used at the sending end to retrieve all the messages. Once the sender end has issued a 'confirm' flow to its partner it is 'indoubt' until it receives a response. In other words, the sender channel is not sure whether the messages have been delivered successfully or not. If there is a communications failure during this phase then the channel will end indoubt. When it reconnects to its partner it will notice from the data store that it was indoubt with its partner and so will ask the other channel whether the last batch of messages were delivered or not. Using the answer the partner sends, the channel can decide whether to commit or rollback the messages on the transmission queue.

This synchronisation data is viewable by issuing a DIS CHSTATUS(*) SAVED command. The values displayed should be the same at both ends of the channel.

Note that if the channel is restarted when it is indoubt it will automatically resolve the indoubt. However, it can only do this if it is talking to the same partner. If the channel attributes are changed or a different Queue Manager takes over the IP address or a different channel serving the same transmission queue is started then the channel will end immediately with message saying that it is still indoubt with a different Queue Manager. The user must start the channel directing it at the correct Queue Manager or resolve the indoubt manually by issuing the RESOLVE CHANNEL command. Note that in this case the user should use the output from DIS CHS(*) SAVED to ensure that the correct action COMMIT or BACKOUT is chosen.

Batch Heartbeats



DEF CHL BATCHHB(500)

- If time since last communication is greater than Interval, heartbeat flow is sent

Reflect Network Stability ✓

Batch Heartbeats - Notes

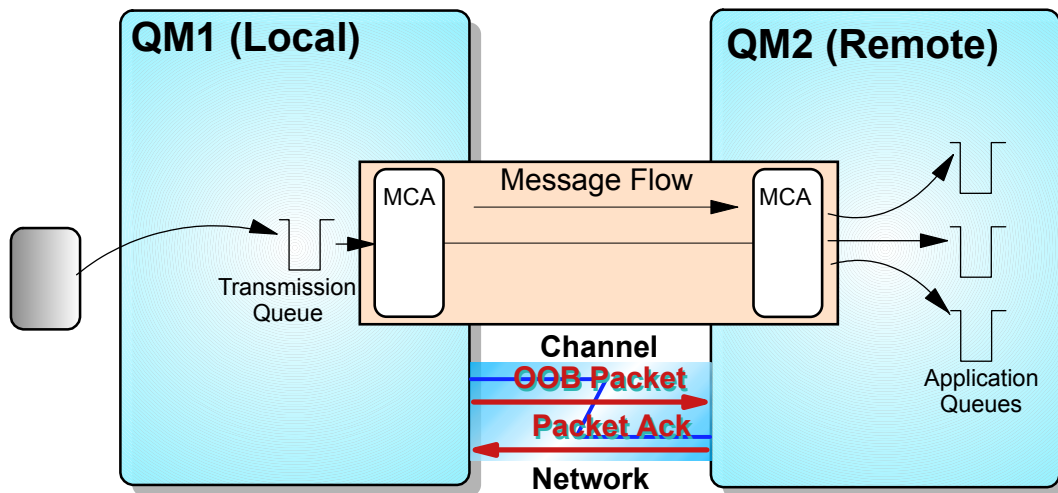
**N
O
T
E
S**

The problem with channels that end indoubt is seen with the clustering algorithm. If a clustering channel ends abnormally, any messages that were to be delivered by this channel are reallocated to another channel and another instance of the target cluster queue. They cannot however be reallocated if they are part of a batch of messages that the sender channel is indoubt about whether the receiver channel has got them.

To reduce the likelihood of a channel ending in indoubt state, an extra flow has been added to the channel protocol. This flow is known as a batch heartbeat. This flow is sent after the messages for the batch have been sent, but before the channel marks itself to be indoubt and sends the confirmation flow. Having received a response from the receiving end of the channel, we can be fairly sure that we can complete the batch. There is now an extremely small window where the receiver channel can end and leave us stuck in indoubt state. Without the batch heartbeat flow this window is much larger (up to the time since the last flow received from the partner).

To determine whether this batch heartbeat flow is sent, we note the time we last received a flow from our partner. If this time was longer ago than the interval specified in the BatchHeartbeat or BATCHHB parameter on the channel, then we will send the flow, otherwise we will not. The interval specified is therefore a reflection of the stability of your network. This interval is specified in milliseconds.

TCP/IP Keepalive



- QM.INI **KEEPALIVE=YES**

- XPARAM **TCPKEEP=YES**

DEF CHL ... KAIN(360) (z/OS only)

TCP/IP Keepalive - Notes

TCP/IP also has its own heartbeat protocol, called KeepAlive, allowing it to recognise network failures. The use of KeepAlive is recommended.

N On z/OS only, this can be configured as a channel parameter, allowing different channels to have specific time-out values, or simply use the system values. This channel parameter still requires KeepAlive to be switched on in TCP/IP before it can be used.

O On all platforms, this can be configured on a Queue Manager basis in WebSphere MQ in the `qm.ini` file or equivalent. The actual values used for the timers and counters in this case, are set for the entire TCP/IP stack on the machine. Typical values are for network failure to be noticed after a default 2 hours. This is probably too high for a WebSphere MQ environment, but other applications on the same machine may also be wanting to use KeepAlive and so values may need to be a compromise between the requirements of these different applications.

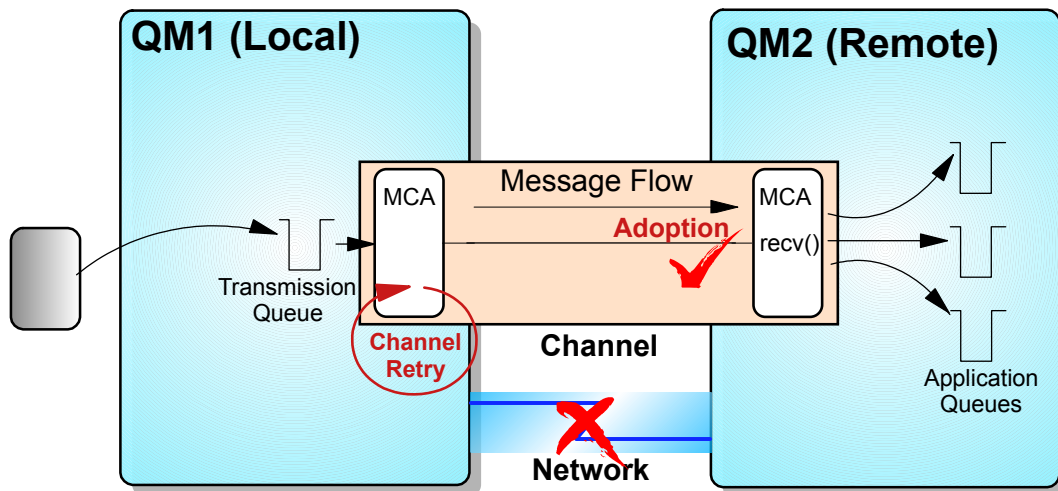
T You need to read your TCP/IP manuals to discover how to configure KeepAlive. For AIX, use the "no" command. For HP-UX, use the "netttune" command. For Windows NT, you need to edit the registry. On z/OS, update your TCP/IP PROFILE dataset and add/change the INTERVAL parameter in the TCPCONFIG section.

Note: On z/OS TCP/IP APAR PQ75195 is recommended

E The use of heartbeats and receive wait time out remove the need for KeepAlive when both ends support both options, however KeepAlive can still be used as well. When communicating with WebSphere MQ implementations which do not support both heartbeats and the non-blocking reads, you should still use KeepAlive.

S Keepalive is strongly recommended for SVRCONN channels since even if client heartbeats are available, they are only used during and MQGETs.

Adopt MCA



- **QM.INI** **AdoptNewMCA, AdoptNewMCATimeout, AdoptNewMCACheck**
- **XPARM** **ADOPTMCA, ADOPTCHK**

AdoptMCA - Notes

If a channel fails, for example a communications error, there is no guarantee that both ends of the channel will detect the failure at the same time. Often the sending end of the channel detects the failure first since it is the end trying to send messages down the socket. What will happen therefore is that the sender will try to reconnect to the target Queue Manager. When the connection comes into the target machine the Queue Manager sees that it already has a channel running with that name and from that location. Consequently the new connection is rejected.

By configuring `AdoptNewMCA` you can tell the Queue Manager that if a new connection arrives and there is already a channel with that name running from the same network address and from the same Queue Manager then the existing channel should be ended and replaced with this new instance. In other words, the channel instance should be 'adopted' by the new connection.

The various configuration parameters in the channels stanza of QM.INI or equivalent (or on z/OS use the XPARMS) allow the behaviour of Adopt to be changed slightly. Set the list of channel types which may be adopted in this manner to turn on adoption

```
AdoptNewMCA=SVR,SDR,RCVR,CLUSRCVR,ALL,FASTPATH
```

```
ADOPTMCA=YES/NO (on z/OS)
```

Set the wait interval between ending the 'adopted' channel nicely and ending it immediate.

```
AdoptNewMCATimeout=60
```

Clearly the channel name must match but additionally you can specify that the network address must be the same and also the remote Queue Manager name. This prevents a rogue connection from a different machine on the network from adopting already running channels. Set which attributes must match between the Adopted and new connections.

```
AdoptNewMCACheck=QM,ADDRESS,NAME,ALL
```

```
ADOPTCHK=NONE/QMNAME/NETADDR/ALL (on z/OS)
```

Performance

- **Some questions to think about**
 - ▶ What is arrival rate on transmission queue?
 - ▶ What are required response times?
 - ▶ How long does a batch take to transmit?
 - ▶ What is the startup time for a channel?
- **Answers depend on your environment**
 - ▶ Application Usage
 - ▶ Operating System/Hardware
 - ▶ Network Bandwidth
 - ▶ Message Sizes

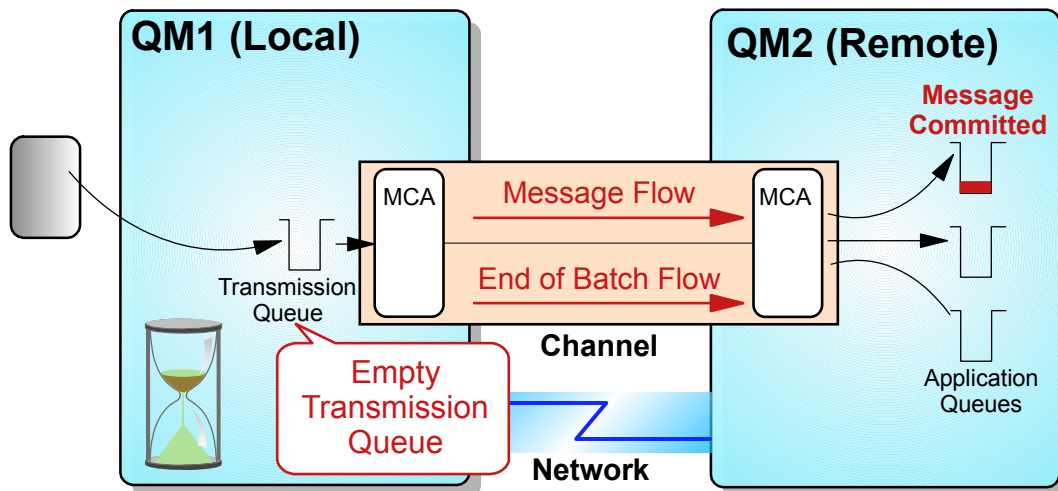


Performance configuration

- **Use BatchInterval to keep batches open**
 - ▶ if arrival rate is slightly slower than sending complete batch
 - ▶ may increase latency, but improve total throughput
- **Set DisconnectInterval high enough so channels kept alive**
 - ▶ Reduce number of times a channel is started/stopped/restarted
 - ▶ Unless network charges by "active sessions" or you are hitting MaxChannels limit
- **Three Scenarios**
- **Using Trusted Bindings**
- **See Capacity Planning Guides**
 - ▶ SupportPac MP00
 - ▶ SupportPac MP16



Batch Interval



- Minimum lifetime of Batch of Messages

DEF CHL BATCHINT(100)

► Recoverable messages will arrive no sooner than BATCHINT

Batch Interval - Notes

N

A recoverable batch of messages is ended as soon as the transmission queue is empty. This can lead to more line-turnarounds and disk activity than necessary. Consider an application putting 3 persistent messages to a transmission queue. There is a good chance that the first message will be retrieved by the channel before the second message is put. Consequently the channel will complete the batch and the second and third message will actually go in a second batch.

O

By adding a small delay to the batch, say, BATCHINT(100) then all 3 messages would be sent in the first batch with only half the disk activity. Note that the messages may well appear to take longer to be delivered though since there is a delay of at least 100ms. BATCHINT has no effect on non-persistent messages NPMSPEED(FAST) since these messages are put outside of syncpoint. Usually applications send database queries as non-persistent which are time critical and database updates as persistent messages which are not so time critical.

T

The net effect and value of BATCHINT depends greatly on the MQ messaging network and messaging profile. The greater the number of channels on the machine the larger the potential savings.

E

One way of helping to determine whether there would be much advantage of BATCHINT is to look at the effective batch size. Issue a DISPLAY CHSTATUS and divide the number of messages by the number of batches. If this number is very low, (say between 1 and 2) then the channel is having to commit for each 1 or 2 messages which is inefficient and may benefit from BATCHINT. Note that this is only an indication since if all the messages are non-persistent BATCHINT won't help at all.


S

Another thing to look for is the disk usage of the channel processes. If you see large disk utilisation then there's a good chance BATCHINT will reduce this cost and improve overall throughput.

Scenario 1

■ Bulk transfer of pre-loaded transmission queue

- ▶ Use a large BatchSize (default of 50 is good),
 - unless using very large messages (cost of retransmission)
- ▶ BatchInt should be 0



Scenario 1 - Notes

Our first performance tuning scenario is a bulk transfer of the messages on a pre-loaded transmission queue.

N You are aiming here to be able to send the messages in large batches therefore reducing the cost of the end of batch processing per message. So a large Batch Size should be used (the default of 50 would be good here) and Batch Interval should be set to 0. There is no need to hold the batch open in this scenario since all the messages are pre-loaded on the transmission queue.

O Your performance balance is that of how many messages you should send in one batch to get a low cost of end of batch processing per message, weighed up against the cost of retransmitting that entire batch should network connectivity be lost part way through. If you are moving very large messages, this balance may need to be achieved with a smaller Batch Size than 50.

T


E

S

Scenario 2

■ Trickle transfer for deferred processing

- ▶ Exact numbers depend on arrival rate
- ▶ Try to achieve high batches
- ▶ Set BatchInt to acceptable delay before transmission e.g. 1 minute
- ▶ But don't make it very large: otherwise you get long-running UoW



Scenario 2 - Notes

N
O
T
E
S

Our second performance tuning scenario is a trickle transfer of message for deferred processing.

You are aiming here to achieve a reasonable batch size of message therefore reducing the cost of end of batch processing per message. This is a classic use of Batch Interval. Messages are arriving at a slightly slower rate than the channel can process them. Holding the batch open for a little longer allows more messages to be batched up together.

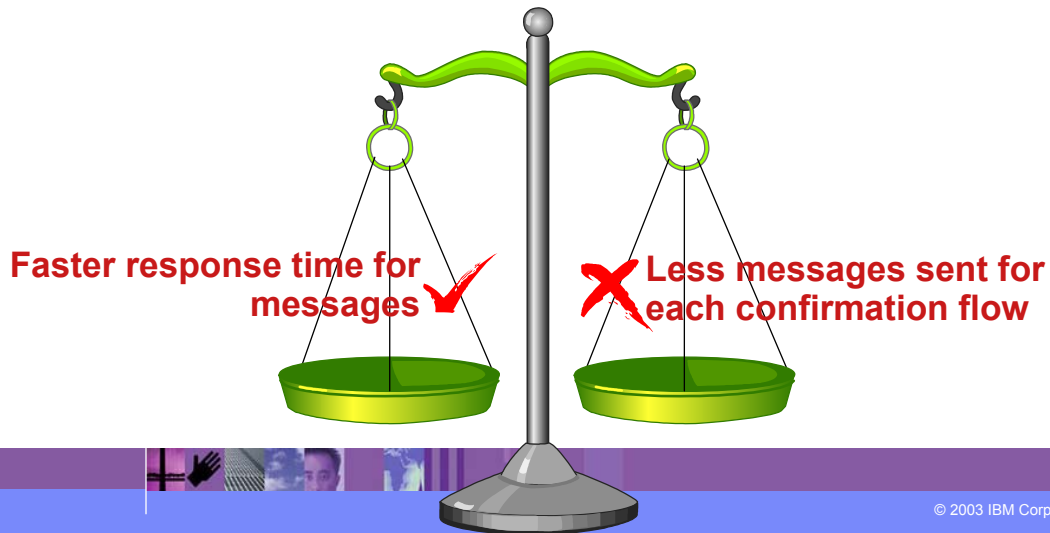
Your performance balance is that of how many messages you can send in one batch, weighed up against the fact that recoverable messages will appear to take longer to be delivered since they are not committed on the target queues until the batch completes, i.e. after at least a Batch Interval.

If delivery time is not an issue, be careful not to set Batch Interval too large, since you then run the risk of generating long running Units of Work.

Scenario 3

■ Synchronous Request/Reply

- ▶ May not want to use BatchInt if response time is key
- ▶ If persistent volume is low consider a smaller BatchSize
- ▶ For non-persistent message flows, use NPMSPEED(FAST) and a non-zero BatchInt



© 2003 IBM Corporation

Scenario 3 - Notes

Our third performance tuning scenario is a synchronous request/reply model where the response time does matter.

N This is a similar situation to scenario 2, except this time the balance is weighted more towards a faster response time rather than more messages in a batch.

O If the persistent message volume is low, you may wish to consider using a smaller batch size and either not using Batch Interval at all, or using a fairly low value in the order of a few seconds to ensure messages are available in a timely fashion.

T If you have non-persistent messages, you should use Non-Persistent Message Speed or NPMSPEED of FAST, and set a non-zero Batch Interval, since non-persistent messages delivered over a fast channel are available immediately. The Batch Size value is still used but does not affect the response time.

To give some example figures:

E With an expected rate of 30 non-persistent msg/sec, set BatchInt=2000 (2secs) and BatchSize=50 for an effective batch size of 50.

S

© 2003 IBM Corporation

Trusted Binding

- **Single process**
 - ▶ For both MCA and Local Queue Manager Agent
 - ▶ No IPC overhead for MCA MQI commands
 - significant for non-persistent messages
- **Configured in queue manager ini file**
 - ▶ Channels stanza

MQIBINDTYPE=FASTPATH

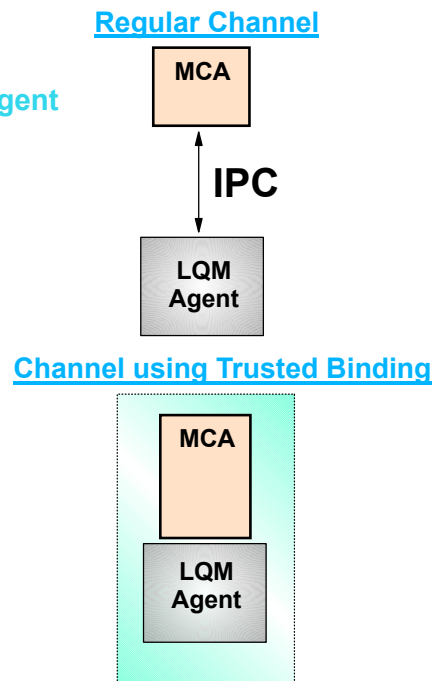
- **Don't use if testing Channel exits**
- **Be careful issuing**

STOP CHL() MODE(FORCE)

(pre V5.3)

STOP CHL() MODE(TERMINATE)

V5.3



Trusted Binding - Notes

**N
O
T
E
S**

By default in the Distributed set of products all Queue Manager operations are actually performed by a process separate from the application known as an agent. Each API call is sent over to the agent process via some IPC mechanism. This is done for reasons of isolation. An application doesn't have direct access to Queue Manager resources or memory. However since each MQI call is a process switch it does have an impact on performance. For a non-persistent message about 60% of the cost of the call is just in the process switch.

All applications can choose to run in 'trusted' mode which means that the Queue Manager functions run directly in-line with the application. We recommend that in general applications should not run in this mode. Channels, however, are product code and therefore suitable to run as fastpath 'trusted' applications. By default they will run as normal applications but by setting MQIBindType=FASTPATH in the Channels stanza in the QM.INI file all channels will run trusted.

Care should be taken with channel exits. Since channel exits run inline with the MCA the exits will now also be in the same memory space as the Queue Manager agent. Channel exits should be tested thoroughly before being run in this mode.

Care should also be taken when using the STOP MODE(FORCE) command on releases prior to V5.3, and on V5.3 care should be taken when using the STOP MODE(TERMINATE) command, either of which could corrupt or hang the queue manager as cleanup may not be completely performed when the channel is killed.

Parameter Reference

Channel	Description
HBINT	Heartbeat Interval
BATCHSZ	Maximum number of Messages in a Batch
BATCHINT	Minimum lifetime of a Batch of Messages
DISCINT	Interval after which a channel will end
SHORTRTY, SHORTTMR LONGRTY, LONGTMR	Channel retry counts and timers
MRRTY, MRTMR MREXIT, MRDATA	Message retry counts and timers and exit (not on z/OS)

QM.INI - Channel Stanza	XPARMS	Description
AdoptNewMCA	ADOPTMCA	Adopt Channel types
AdoptNewMCATimeout		Adopt Quiesce time-out
AdoptNewMCACheck	ADOPTCHK	Adopt parameter check
KeepAlive	TCPKEEP	Use Keepalive
MQIBindType		Channel FASTPATH binding

Environment Variables	XPARMS	Description
MQRCVBLKTO	RCVTIME	Receive Wait Timeout
MQRCVBLKMIN	RCVTMIN	Receive Wait Minimum Timeout



Parameter Reference - Notes

This chart lists those attributes of a channel which affect the performance and recovery capabilities.

Some of the attributes are part of the true channel definition and others are setup in other places such as an INI file.

N

Full details of the parameters on the channel definition can be found in the InterCommunication book.

On the distributed platforms those parameters that are in the QM.INI file are in the Channels stanza or on Windows they are implemented in the Registry

O

Key : HKEY_LOCAL_MACHINE/SOFTWARE/IBM/MQSeries/QueueManager/<QMNAME>
or use MQSeries services

T

Channel Initiator parameters provide some of the same parameters on z/OS that can be found in the QM.INI file on the distributed platforms. The XPARMS are built using the CSQ6CHIP macro.

E

S

