

ID	Location	Comments
ak-01	p22 Recovering Message Queue	".., including JMS hostconfigurations, are stored in the Domain", change to ".., including JMS host configurations for Embedded or Local JMS hosts, are stored in DAS" Response: done.
ak-02	p23 Recovering Message Queue	"according using" typo Response: done.
ak-03	p157 Connection Pooling and Failover	"When the Address List Behavior attribute is random (the default), .." somewhat contradicts with the later sentence "The default value for the Address List Behavior attribute is priority, if the JMS type is of type LOCAL.", Please check with Satish to clarify what is the default Address List Behavior for Embedded, Local and Remote Response: done (extensively rewrote this section, covering this comment).
ak-04	p157 Connection Pooling and Failover	"When failover occurs, MQ transparently transfers the load to another broker ..", change to "When failover occurs, MQ automatically transfers the load to another broker .." Response: done (extensively rewrote this section, covering this comment).
ak-05	p157 Connection Pooling and Failover	"To specify whether the GlassFish Server tries to reconnect to the primary broker ..", need to clarify what does "primary broker" referring to here Response: done (extensively rewrote this section, covering this comment).
ak-06	p157 Connection Pooling and Failover	"select the Reconnect checkbox", need to clarify on which screen Response: done (extensively rewrote this section, covering this comment).
ak-07	p158 Load-Balanced Message Inflow	"set the maxNumActiveConsumers property of the physical destination to a large value", change "a large value" to "a value > 0" Response: done.
ak-08	p158 Load-Balanced Message Inflow	"To ensure that local delivery is preferred", change to "To ensure a message-driven bean first try to connect to the local colocated broker" Response: done.
ak-09	p158 Load-Balanced Message Inflow	"This setting is the default for GlassFish Server instances that belong to a cluster", see comment ak-3 Response: done.
ak-09'	p159	8 links under "The following topics are addressed here" and the link titles look crowded, for clarity maybe combine close-related topics to a higher level title ? Response: formatting of this list is beyond my control (it's done by the pdf generator). However, changes to titles as suggested by Nigel make it less crowded.
ak-10	p159 About Message Queue Broker Clusters	"When a broker or a connection fails, ...reconnect to another broker in the cluster", change to "When a broker fails, ...reconnect to another broker in the cluster", or to be worded to include the case of auto-reconnect to possibly the same broker if only connection failed but broker running Response: done.
ak-11	p159 About Message Queue Broker Clusters	"based on where information about the cluster configuration is stored", change to "based on where the cluster configuration change record is stored" Response: done.

ID	Location	Comments
ak-12	p159 About Message Queue Broker Clusters	change all occurrences of "the information about the cluster's configuration" to "the cluster configuration change record" Response: done.
ak-13	p159 About Message Queue Broker Clusters	"to keep abreast of changes to the configuration", change to "to keep abreast of changes to the cluster configuration change record" Response: kept as is. Flow is still clear though, due to use of "cluster configuration change record" in the preceding sentence.
ak-14	p159 About Message Queue Broker Clusters	"Thus, brokers can access the cluster configuration information", change to "Thus, brokers can access the cluster configuration change record" Response: kept as is. Flow is still clear though, due to use of "cluster configuration change record" in the preceding sentence.
ak-15	p159 About Message Queue Broker Clusters	Enhanced broker clusters: some comment as ak-10 Response: done.
ak-16	p159 About Message Queue Broker Clusters	This section should refer to MQ Technical Overview on "Broker Clusters" so that the same content are not repeated here. I think here it can simply give a brief summarization of MQ cluster types and service availabilities provided Response: done, with the reference appearing right at the start.
ak-17	p160	"it as its primary broker and has that broker listed first in its JMS host list.", if "primary broker" == "listed first in its JMS host list", then rewording this Response: done.
ak-18	p161	"subcommand in remote mode", confusing "remote mode" here, maybe "subcommand in asadmin remote mode", similar changes in later contents where "remote" is used for the same context Response: changed "Remote subcommands require a running server." to "Remote asadmin subcommands require a running server." and "in remote mode" to "in asadmin remote mode".
ak-19	p162	"Use the change-master-broker subcommand .. JMS hosts", add following at the end "for example, before you remove the GlassFish instance that is associated with the current master broker from the cluster" Response: done.
ND-01	P18 "High Availability Java Message Service"	Change "Connection Pooling and Failover The JMS service pools JMS connections automatically." To: "Failover" Connection pooling is nothing to do with high availability. It's a performance/resource management feature. Suggest you delete all references to connection pooling from this section.) Response: done.
ND-02	P18 "High Availability Java Message Service"	Change: "For more information about JMS connection pooling and failover, see "Connection Pooling and Failover" on page xxx." To: "For more information about JMS high availability, see (name of chapter 10) on page xxx." I make more comments about chapter 10 later. Response: done.

ID	Location	Comments
ND-03	Page 19: "How GlassFish Server Provides High Availability"	<p>In the list that follows "GlassFish Server provides high availability through the following subcomponents and features:", Add a link to "Using Message Queue Broker Clusters with GlassFish Server" on page xxx" (or perhaps to the whole of chapter 10)</p> <p>I know we mentioned it earlier, but if we're listing here the ways in which GlassFish provides HA then we need to mention MQ high availability.</p> <p>Response: not done. The structure of this chapter is wonky as is, but I'm not the owner of it; I just contribute the JMS/MQ material in the existing structure. I will confer with the owner about why the structure is the way it is.</p>
ND-04	Page 19: "Message Queue Broker Clusters"	<p>Change "This automatically created broker cluster is configurable to take advantage of the different types of broker clusters supported by Message Queue."</p> <p>To: "This automatically created broker cluster is configurable to take advantage of the two types of broker clusters (conventional and enhanced) supported by Message Queue."</p> <p>Response: done.</p>
ND-05	Page 19: "Message Queue Broker Clusters"	<p>Change: "Additionally, Message Queue broker clusters created and managed using Message Queue itself can be used as external, or remote, JMS hosts to provide JMS messaging high availability to both GlassFish standalone instances and clusters."</p> <p>To: "Additionally, Message Queue broker clusters created and managed using Message Queue itself can be used as external, or remote, JMS hosts. This provides additional deployment options, such as deploying a Message Queue broker on a different machine from the GlassFish instance, or deploying different numbers of GlassFish instances and Message Queue brokers."</p> <p>Response: done.</p>
ND-06	Page 20: "Highly Available Clusters"	<p>Is it worth spelling out that in this context we're not talking about JMS HA?</p> <p>Response: not done. This is out of the JMS/MQ scope. I'll bring it up with the owner of this section.</p>
ND-07	Page 22: "Recovering MessageQueue"	<p>Change "hostconfigurations" to "host configurations"</p> <p>Response: done.</p>
ND-08	Page 22: "Recovering MessageQueue"	<p>Change "maintain by Message Queue brokers" to "maintained by Message Queue brokers"</p> <p>This statement that private data is stored with the GlassFish instances does not apply in the case of enhanced clusters. In that case the message data will be stored in a centralised HA database. Add that in that case you may need to follow any backup/restore procedure defined for that database.</p> <p>Response: done.</p>

ID	Location	Comments
ND-09	Page 157: Chapter 10: "Java Message Service Load Balancing and Failover"	<p>I think the name of this chapter, and its introductory sentence, should refer to "High Availability", for consistency with the previous chapter.</p> <p>I suggest changing the name of the chapter to "Configuring Java Message Service High Availability, Load Balancing and Failover", or, if this is too long, leave out the mention of load balancing since that's a relatively short part of the chapter.</p> <p>Then change "This chapter describes how to configure load balancing and failover of the Java Message Service (JMS) for use with the GlassFish Server."</p> <p>to: "This chapter describes how to configure high availability, load balancing and failover of the Java Message Service (JMS) for use with the GlassFish Server."</p> <p>Response: done.</p>
ND-10	Page 157	<p>Under "The following topics are addressed here:" add a reference to Load-Balanced Message Inflow on p158.</p> <p>Response: done.</p>
ND-11	Page 157	<p>"For more information about JMS connection pooling and failover, see "Connection Pooling and Failover" on page xxx."</p> <p>Response: done, if I understood the comment correctly.</p>
ND-12	Page 157: "Connection Pooling and Failover"	<p>Connection pooling is nothing to do with HA (you get connection pooling in standalone instances as well). Please remove all references to connection pooling.</p> <p>Change: "Connection Pooling and Failover. GlassFish Server supports JMS connection pooling and failover. The Oracle GlassFish Server pools JMS connections automatically."</p> <p>To: "Connection Failover. GlassFish Server supports JMS connection failover."</p> <p>Response: done.</p>
ND-13	Page 159: "Conventional cluster with master broker."	<p>You could add that this is the simplest to configure and is the default for a GlassFish cluster.</p> <p>Response: done.</p>
ND-14	Page 159: "Conventional cluster of peer brokers."	<p>I haven't seen this term before. I like it. Do the updated MQ docs (admin guide and technical overview) use the same term?</p> <p>Response: yes.</p>
ND-15	Page 160: Embedded (at bottom of page).	<p>Add the comment that there is one broker for each GlassFish instance, and that each pair is colocated in the same JVM.</p> <p>Response: done; added above Embedded/Local/Remote list.</p>
ND-16	Page 161: Local (at top of page).	<p>Add the comment that there is one broker for each GlassFish instance, and that each pair is colocated on the same machine.</p> <p>Response: done; added above Embedded/Local/Remote list.</p>
ND-17	Page 161: Remote (near top of page).	<p>I suggest replacing "Any" with the same list as for "Local", since the use of the word "Any" suggests that there are other possibilities.</p> <p>Response: done.</p>

ID	Location	Comments
ND-18	Page 161: Remote (near top of page).	Add the comment that this option allows a different number of brokers to GlassFish instance, and that may be on different machines. Response: done.
ND-19	Page 161: "To Configure a Conventional Broker Cluster With Master Broker as an Embedded or Local JMS Host for a GlassFish Server Cluster"	This heading is a rather convoluted form of words. Can't we drop the words "JMS Host" (which in this context is synonymous with "Broker", and using two terms for the same thing is confusing) and change it to "To Configure a GlassFish Server Cluster to use a conventional broker cluster with master broker" Response: redid all the "To ..." headings to make them briefer.
ND-20	Page 161	Change "Use the configure-jms-cluster subcommand in remote mode to configure a conventional broker cluster with master broker to service a GlassFish Server cluster that uses Embedded or Local JMS hosts." To: "Use the configure-jms-cluster subcommand to configure a conventional broker cluster with master broker to service a GlassFish Server cluster that uses either Embedded or Local brokers" (I don't know why the words "in remote mode" were used previously) Response: it's remote asadmin mode, and has been clarified; changed "that uses Embedded or" to "that uses either Embedded or".
ND-21	Page 161	Change "Note that this configuration, with embedded brokers, is the default for GlassFish Server clusters" Response: done.
ND-22	Page 161: "To Configure a Conventional Broker Cluster of Peer Brokers as an Embedded or Local JMS Host for a GlassFish Server Cluster"	Change heading to "To Configure a GlassFish Server Cluster to use a conventional cluster of peer brokers" Response: redid all the "To ..." headings to make them briefer.
ND-23	Page 161	Change "Use the configure-jms-cluster subcommand in remote mode to configure a conventional broker cluster of peer brokers to service a GlassFish Server cluster that uses Embedded or Local JMS hosts." To: "Use the configure-jms-cluster subcommand to configure a conventional broker cluster of peer brokers to service a GlassFish Server cluster that uses Embedded or Local brokers" Response: it's remote asadmin mode, and has been clarified.
ND-24	Page 163: "To Change the Master Broker in a Broker Cluster Serving as an Embedded or LocalHost"	Change to "To change the master broker in a broker cluster that uses either embedded or local brokers" (N.B. I'm not trying to impose my own capitalisation preferences here) Response: redid all the "To ..." headings to make them briefer.
ND-25	Page 163	Fix incorrect use of bold font in "If the need arises to convert from a conventional broker cluster with master broker to a conventional broker cluster of peer brokers, or the reverse, follow the instructions in "Managing Conventional Clusters" in" Response: This formatting is beyond my control and provided by the pdf generator.
sats_i-1	Connection Pooling and failover (Page 157)	"By default, the JMS service selects the primary JMS host (Message Queue broker) randomly from the specified JMS host list." The above is true only for Remote mode. For EMBEDDED and LOCAL, it should be - "By default, the JMS Service selects the co-located JMS Host (MQ Broker). When failover occurs, MQ transparently transfers the load to another JMS host in the list..." Response: done.

ID	Location	Comments
sats_i-2	Page 164	<p>"Configure the GlassFish Server cluster to use a Message Queue enhanced broker cluster with master broker by using the configure-jms-cluster(1) subcommand: asadmin --passwordfile password-file configure-jms-cluster --clustertype=enhanced --dbvendor database-vendor-name --dbuser database-user-name --dburl database-url --property list-of-database-specific-properties glassfish-cluster-name"</p> <p>This is incorrect there is no master-broker option in enhanced broker clusters. The command will need the --configstore and --messagestore to point to JDBC</p> <p>Response: done.</p>
sats_i-3	To Migrate Between Types of Conventional Broker Clusters (Page 162)	<p>"If the need arises to convert from a conventional broker cluster with master broker to a conventional broker cluster of peer brokers, or the reverse, follow the instructions in 'Managing Conventional Clusters' in Oracle GlassFish Message Queue 4.5 Administration Guide." This should be changed into a note of caution and mentioned right in the beginning of this section.</p> <p>Response: left this section, but added a caution under "Configuring GlassFish Clusters to Use Message Queue Broker Clusters".</p>