

ID	Location	Comment
AK-01	p18	"Message Queue transparently transfers that broker's load to another JMS host in the JMS host list" CHANGE TO "Message Queue can automatically failover connections on the failed JMS host to other JMS host(s) in the JMS host list" <b>Response:</b> done.
AK-02	p23 "Recovering Message Queue"	Please see separate emails from Ed and me on this section. <b>Response:</b> done.
AK-2	p161	"The default value is . 3A value" – typo <b>Response:</b> done.
AK-3	p161-162	The configurable JMS service attributes and their descriptions appear to duplicate what's in p291 in GlassFish Admin Guide, should one of them simply reference to the other instead of repeat the attribute list ? <b>Response:</b> they differ in that this one uses the GUI and GUI labels, and the one in the GlassFish Admin Guide uses the CLI and dotted names.
AK-4	p163	"This first broker is the local co-locatedMessage Queue broker" CHANGE TO "This first broker is the co-located Message Queue broker for Embedded and Local JMS host type" <b>Response:</b> done, but affected by ND-3.
AK-5	p163	"and in messages being delivered out of order" CHANGE TO "and in JMS message order semantics not preserved" <b>Response:</b> done.
AK-6	p164	"a broker failure rarely results in significant delays in message delivery and messages are always delivered in order." CHANGE TO "messages owned by the failed broker are delivered by the failover broker as soon as it takes over, and JMS message order semantics are preserved." <b>Response:</b> done.
AK-7	p164	"Note –Despite the message service availability offered by both conventional and enhanced broker clusters, they do not provide a guarantee against failure and the possibility that certain failures, for example in the middle of a transaction, could require that some operations be repeated. It is the responsibility of the messaging application (both producers and consumers) to handle and respond appropriately to failure notifications from the messaging service." COMMENTS: The above paragraph is toward standalone MQ client applications. For GlassFish application components, there is no interface exposed to allow a Java EE component to listen for MQ "failure notification". I think this paragraph should be reworded to say that "... application component using transactions should respond to JMS exceptions appropriately which including rollback the transaction, retry the operation if applicable. For more information please reference MQ Java Developer's Guide "Handling Exceptions When Failover Occurs" though it's written toward standalone MQ client applications" <b>Response:</b> done.
AK-8	p165	"Note that this configuration, with Embedded brokers, is the default for GlassFish Server clusters." CHANGE TO "Note that this configuration, with Embedded and Local brokers, is the default for GlassFish server clusters". <b>Response:</b> not done. The default for GlassFish clusters <b>is</b> Embedded, <b>not</b> Embedded and Local. (There can only be one default.)

ID	Location	Comment
AK-9	p167	<p>"Follow this procedure, for example, before you remove from a GlassFish cluster the instance associated with the current master broker." ADD to the end of above sentence "However, changing master broker involves data migration and configuration synchronization, this operation should be used with caution, and in general you should avoid to remove the GlassFish instance associated with master broker from the GlassFish cluster"</p> <p>COMMENT: Please check with Satish for recommended best condition or required conditions to do this operation.</p> <p><b>Response:</b> not done. According to Satish's DHQA on 11/8, changing the master broker doesn't involve data migration or configuration synchronization that the user must do.</p>
ND-1	Page 161	<p>You changed the name of this chapter from "Java Message Service Load Balancing and Failover" to "Configuring Java Message Service Failover, Load Balancing and Clustering"</p> <p>This is slightly awkward because load balancing and failover are features available when using clustering, whereas this suggests that clustering is a separate feature. My preference would be a chapter title of "Java Message Service High Availability", since the purpose of this chapter is to draw together all the MQ HA features.</p> <p><b>Response:</b> done.</p> <p>Even if you decide to keep the existing chapter name, I suggest the following chapter intro:          "This chapter describes how to configure the high availability features of the Java Message Service (JMS) available when using GlassFish Server. It covers how to configure message queue broker clusters and how to use them to provide connection failover and load balancing"</p> <p><b>Response:</b> done.</p>
ND-2	Page 161: "Connection Failover"	<p>"GlassFish Server supports JMS connection failover, providing several options that control how connection failures are handled. " Connection failover relies on the use of clustering (of either kind), whereas we don't mention the word cluster anywhere in this section. I suggest moving this section to the end of the chapter, after the discussion of clustering.</p> <p>We can then change          "GlassFish Server supports JMS connection failover, providing several options that control how connection failures are handled." to "The use of message queue broker clusters provides JMS connection failover, providing several options that control how connection failures are handled."</p> <p><b>Response:</b> done.</p>

ID	Location	Comment
ND-3	Page 162 "Load-Balanced Message Inflow"	<p>Oh, gosh. This chapter is a muddle and mixes multiple features together. I'm sorry I didn't spot this before, but I hope it should be simple to remedy.</p> <p>I suggest moving this to the end and renaming. I suggest changing the heading to "Load-balanced delivery to MDBs" and using the following text:</p> <p>"When a message-driven bean (MDB) application is deployed to a GlassFish server cluster, incoming messages will be delivered randomly to MDBs irrespective of which instance they are running in."</p> <p>(Note to Mike: note that this is a feature of GlassFish clusters, not broker clusters)</p> <p><b>Response:</b> done.</p> <p>(note to Mike: the material on ActivationSpec is not specific to HA or clusters and is not appropriate here, so I have omitted it)</p> <p><b>Response:</b> done.</p> <p>"If the MDB is configured to receive messages from a durable subscription on a topic then only one MDB instance will receive each message, irrespective of which GlassFish instance that MDB is running in. This feature relies on the MDBs being configured to have the same Client ID in each instance, which would normally be the case.</p> <p>"If the MDB is configured to receive messages from a non-durable subscription on a topic then, again, only one MDB instance will receive each message, irrespective of which GlassFish instance that MDB is running in. This feature relies on the MDBs having the same bean name and application name in each instance, which again would normally be the case.</p> <p>"For more information, see "About SharedTopic Subscriptions for Clustered Containers" in the Message Queue Administration Guide."</p> <p><b>Response:</b> done.</p> <p>This left me wondering what to do with this paragraph:</p> <p>"To support multiple consumers using the same queue, set the maxNumActiveConsumers property of the physical destination to a large value. If this property is set, the Message Queue software allows up to that number of message-driven beans to consume messages from same queue. The message is delivered randomly to the message-driven beans. If maxNumActiveConsumers is set to -1, there is no limit to the number of consumers."</p> <p>This is actually wrong in respect of MDBs, and in any case the default is unlimited anyway, so it's unnecessary. I suggest deleting this paragraph.</p> <p><b>Response:</b> done.</p> <p>We also have this paragraph to put somewhere:</p> <p>"To ensure that local delivery is preferred, set addresslist-behavior to priority. This setting specifies that the first broker in the AddressList is selected first. This first broker is the local colocatedMessage Queue instance. If this broker is unavailable, connection attempts are made to brokers in the order in which they are listed in the AddressList. This setting is the default for GlassFish Server instances that belong to a cluster."</p> <p>This actually belongs to the paragraph labelled "addressListBehavior" in the section on load balancing. So please move it to there (after "...is Random.").</p> <p><b>Response:</b> done.</p>