ID	Location	Comment
ak-1	p21	Chapter 1 "Recovering from Failures" "You can use various techniques to manually recover individual subcomponents" Although this is implied, I'd suggest the above change to "You can use various techiniques to manually recover individual subcomponents after a hardware failure, for example disk failure" **Response**: done.
ak-2	p149	"prirmary" – typo **Response**: done.
ak-3	p153	"Connection Failover" This section needs to clarify these properties in case of enhanced cluster (see separate emails on this from Nigel and me) **Response**: done as part of ND-06 and ND-06.ak below.
ak-4	p155	"which is normally the case" COMMENT: The above needs clarification - does the user needs to do anything in order to fall into "is normally the case" ? or does it mean this is taken cared by the jmsra resource adapter ? - I think it's the later but please check with Nigel. Fair question. This is about MDBs that use shared non-durable topic subscriptions when clientID is not set in the activation spec. In this case JMSRA sets clientID using the ActivationSpec properties groupName and mdbName which are passed in by GlassFish. These values should be the same in each clustered GlassFish instance. Satish: can you please advise what these are set to? Particularly mdbName? We can then replace "If the MDB is configured to receive messages from a non-durable subscription on a topic, then only one MDB instance will receive each message. This feature relies on the MDBs having the same bean name and application name in each instance, which is normally the case." with something more specific. Nigel **Response**: not done yet, pending response from Satish.
SK-1	Page 150	The caution note should move to the next section and should be added after each - Before you begin section **Response**: done.
SK-2	Page 150	"Caution – After configuring a broker cluster for a GlassFish cluster using Embedded or Local JMS hosts, you must follow special procedures to migrate to another type of broker cluster, as described in “To Migrate Between Types of Embedded or Local Conventional Broker Clusters” on page 152." This note is only if you are moving from an existing setup where some amount of JMS related activity has happened. So, the note should be reworded to state this. Something like - If the need arises to convert an existing cluster where JMS related activity has occurred, you must follow special procedures to migrate to another type of broker cluster, as described in “To Migrate Between Types of Embedded or Local Conventional Broker Clusters” on page 152 before running the configure-jms-cluster command. Failing to do so could lead to data corruption and could render you setup un-usable. I think "JMS related activity has occurred" can probably not be relied on user's judgement at all times.

ID	Location	Comment
		amy **Response**: given that SK-1 moves this caution to each “To Configure a GlassFish Cluster to Use …” section, I have rewritten it into a “Do not …” statement.
ND-01	Page 18	"High Availability Java Message Service" This has two subheadings, "Connection Failover" and "Message Queue Broker Clusters", but doesn't make clear that connection failover is dependent on the brokers being clustered and is not a self-contained feature. I suggest swapping the order here, putting the subsection "Message Queue Broker Clusters" first and the subsection "Connection Failover" second. We can then fine-tune the text for "Connection Failover" as follows: Change: "If the primary JMS host (Message Queue broker) in use by a GlassFish instance fails, Message Queue can automatically fail over connections to the failed JMS host to other hosts in the JMS host list, maintaining JMS messaging semantics." "For more information about JMS connection failover, see “Connection Failover” on page 153." To: "The use of Message Queue broker clusters allows connection failover in the event of broker failure. If the primary JMS host (Message Queue broker) in use by a GlassFish instance fails, connections to the failed JMS host will automatically fail over to other hosts in the JMS host list, allowing messaging operations to continue" "For more information about JMS connection failover, see “Connection Failover” on page 153." (I've added an introductory sentence which explains this is a feature of clustering, and simplified the syntax of the following sentence. I omit the bit about JMS messaging semantics as being unnecessarily technical here, since it would need qualifying and explanation). **Response**: done, but with slight rewrite to keep “maintaining JMS messaging semantics”, which Amy specifically asked for in an earlier review.
ND-02	Page 22	"Recovering Message Queue" I note the discussion with Ed and Amy about this and hesitate to intevene, but I don't particularly like the current wording. The first sentence "JMS service configurations, including JMS host configurations for Embedded and Local JMS hosts, are stored in the Domain and are recovered when the Domain Administration Server (DAS) is recovered." doesn't mean anything to me. What does "recovered" mean here? What kind of failure are we recovering from? As for the remaining paragraphs. I think any discussion of recovery should list the possible kinds of failure and describe how to recover from each. I'm not proposing text here as this needs to be agreed with Ed and Amy, but this is the structure I think we should be following. Note that in several cases, there is no need to do anything and we should say this. If Ed and Amy approve I can help work this out into more polished text. *** Conventional cluster; Recovering from broker process failure: If the broker was "embedded" then restart the Glassfish instance. If the broker was "local", and the GlassFish instance is still running then shut it down first. Recovering from broker machine failure: Restart the machine and the GlassFish instance. Recovering from failure of the disk hosting the broker data directory: Replace disk, restart the machine, restore data directory from

ID	Location	Comment
		backup and restart GlassFish instance. *** Enhanced cluster Recovering from broker process failure: There's no urgent need to do anything. Other brokers will have taken over client connections and the messages of the failed broker. When convenient, shutdown and restart GlassFish instance. Recovering from broker machine failure: There's no urgent need to do anything. Other brokers will have taken over client connections and the messages of the failed broker. When convenient, restart the machine and the GlassFish instance. Recovering from failure of the disk hosting the broker data directory: There's no urgent need to do anything. Other brokers will have taken over client connections and the messages of the failed broker. Replace disk, restart the machine, restore data directory from backup and restart GlassFish instance. Recovering from the failure of the highly-available database: Consult vendors' documentation. **Response**: not done, as I haven't seen agreement from Ed and Amy on this. Regardless of their opinion, though, the intro sentence of the parent section, "You can use various techniques to manually recover individual subcomponents after hardware failures such as disk crashes", coupled with the content of its peer sections, suggest that what you are proposing is beyond the intended scope. I have spoken with Paul Davies, the "owner" of the HA Admin Guide for this release, and he readily admits that this whole chapter is a patchwork hodgepodge that's simply grown over releases. He's more than willing to look at the whole chapter's structure for the March library update. Personally, I don't think recovery from hardware failure is an HA topic at all, but instead one that should be covered in the Admin Guide.
ND-03	Page 147	The words here are definitely better now. However "configure the high availability features of the Java Message Service (JMS) available..." is a bit of a mouthful in the way that "available" appears to be used twice. **Response**: I've tightened it up. I also think that having two "The following topics are addressed here:" lists on the same page, one at high level and one at lower level, looks a bit odd. Can we really not change the words "The following topics are addressed here:" to something more specific for each list? **Response**: I couldn't change the words, but was able to get rid of it altogether. I grafted the list intro into the first para. The sentence "It also provides instructions for configuring the types of Message Queue broker clusters to support the types of JMS hosts that GlassFish Server clusters can use." doesn't really make sense. Configuring types of MQ clusters to support types of hosts? I think this should be changed to something simpler such as "It describes the different cluster and broker types that are supported and how to configure them". **Response**: done. BTW, much better than what was there.

ID	Location	Comment
ND-04	Page 154	"Address List Behavior" Since we're merging two sections together here a bit of reordering is required, so this whole section now reads: "Specifies how the Java Message Service selects which JMS host in the JMS hosts list to connect to. If this attribute is set to Priority, the Java Message Service tries to connect to the first JMS host specified in the JMS hosts list and uses another one only if the first one is not available. If this attribute is set to Random, the JavaMessage Service selects the JMS host randomly from the JMS hosts list. The default for Embedded and Local JMS host types is Priority, and the default for the Remote JMS host type is Random. For Embedded and Local JMS host types, the Java Message Service ensures that the Message Queue broker servicing a clustered instance appears first in that instance's JMS host list. Thus, having Priority as the default Address List Behavior ensures that an application deployed to a clustered instance is first connected to that instance's co-located broker. If there are many clients attempting a connection using the same connection factory, use the Random setting to prevent them from all attempting to connect to the same JMS host. **Response**: superseded by ND-06 and ND-06.ak below.
ND-05	Page 154	"Load-Balanced Delivery to MDBs" This is much better. Change "one MDB instance will receive each message" to "one MDB instance across the whole GlassFish cluster will receive each message." Note that this change needs to be made in two places. **Response**: done and done.
ND-06	Page 153ff	Replace the content of "Connection Failover" with content provided in 2/23/2011 2:55 AM email (which is an amendment of a 2/22/2011 8:39 AM email) **Response**: done.
ND-06.ak	Page 153ff	Update Nigel's writeup in 2/23/2011 2:55 AM email with comments in 2/23/2011 1:23 PM email. **Response**: done.