




US 20080228839A1

(19) **United States**(12) **Patent Application Publication**  
**Fung et al.**(10) **Pub. No.: US 2008/0228839 A1**(43) **Pub. Date: Sep. 18, 2008**(54) **APPARATUS FOR HIGHLY AVAILABLE  
TRANSACTION RECOVERY FOR  
TRANSACTION PROCESSING SYSTEMS**(60) Provisional application No. 60/359,226, filed on Feb.  
22, 2002.**Publication Classification**(75) Inventors: **Priscilla C. Fung**, Union City, CA  
(US); **Alexander J. Somogyi**,  
Basking Ridge, NJ (US)(51) **Int. Cl.**  
**G06F 17/30** (2006.01)(52) **U.S. Cl.** ..... **707/202; 707/E17.007**(57) **ABSTRACT**

Correspondence Address:

**FLIESLER MEYER LLP**  
**650 CALIFORNIA STREET, 14TH FLOOR**  
**SAN FRANCISCO, CA 94108 (US)**

A highly available transaction recovery service migration system in accordance with one embodiment of the present invention implements a server=s Transaction Recovery Service (TRS) as a migratable service. In one embodiment of the present invention, the TRS is a server instance or software module implemented in JAVA. The TRS migrates to an available server that resides in the same cluster as the failed server. The migrated TRS obtains the TLOG of the failed server, reads the transaction log, and performs transaction recovery on behalf of the failed server. The migration may occur manually or automatically on a migratable services framework. The TRS of the failed server migrates back in a fail back operation once the failed primary server is restarted. Failback operation may occur whether recovery is completed or not. This expedites recovery and improves availability of the failed server thereby preserving the efficiency of the network and other servers.

(73) Assignee: **BEA SYSTEMS, INC.**, San Jose,  
CA (US)(21) Appl. No.: **12/127,082**(22) Filed: **May 27, 2008****Related U.S. Application Data**(63) Continuation of application No. 11/325,028, filed on  
Jan. 4, 2006, now Pat. No. 7,406,618, which is a con-  
tinuation of application No. 10/341,041, filed on Jan.  
13, 2003, now Pat. No. 7,178,050.




FIG. 1A




FIG. 1B




FIG. 2

300




FIG. 3

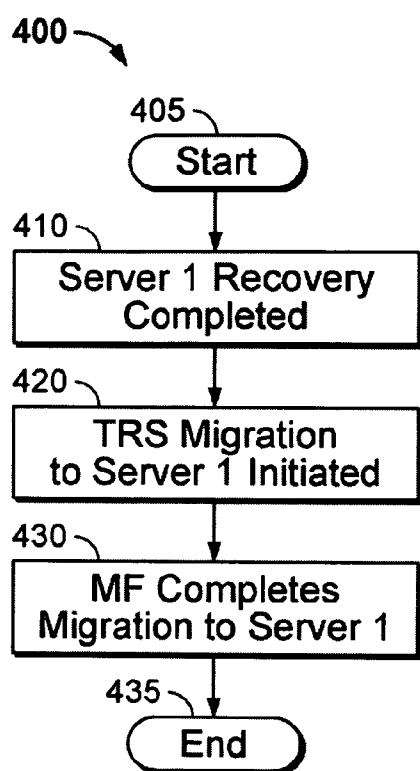


FIG. 4




FIG. 5

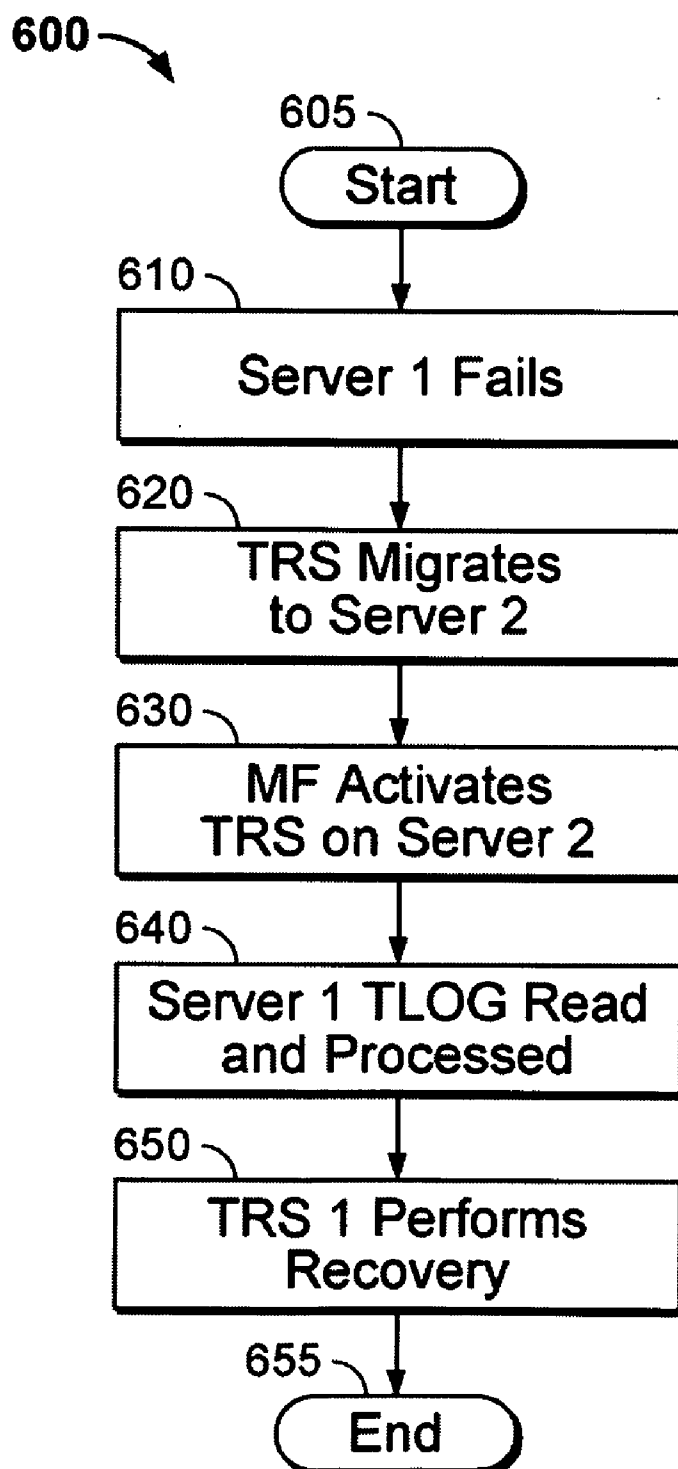


FIG. 6

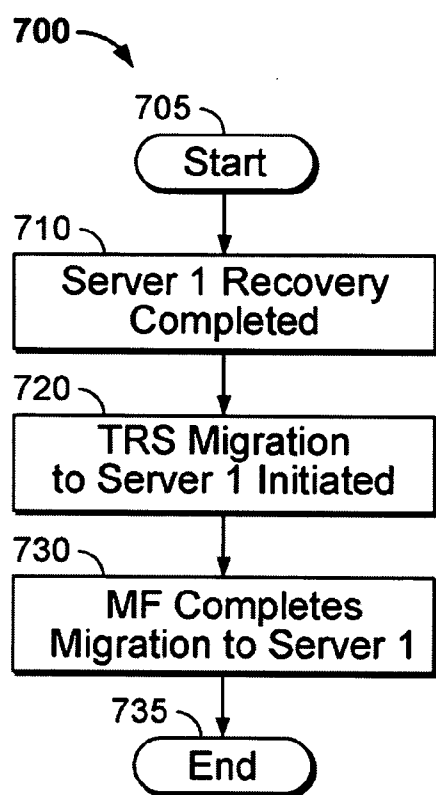


FIG. 7




FIG. 8



## APPARATUS FOR HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS

### CLAIM TO PRIORITY

[0001] This application is a Continuation of U.S. patent application Ser. No. 11/325,028 entitled "APPARATUS FOR HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS" by Priscilla C. Fung et al., filed Jan. 4, 2006 (Attorney Docket No: BEAS-01173US4), which is a continuation of a U.S. patent application Ser. No. 10/341,041, now U.S. Pat. No. 7,178,050, entitled "SYSTEM FOR HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS", by Priscilla C. Fung et al., filed Feb. 22, 2003 (Attorney Docket No: BEAS-01173US2), which claims priority to U.S. Provisional Patent Application No. 60/359,226 entitled "HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS", by Priscilla C. Fung et al., filed Feb. 22, 2002 (Attorney Docket No. BEAS-01173US0), which are incorporated herein by reference.

### CROSS REFERENCE TO RELATED APPLICATIONS

[0002] The present application is related to the following United States patents and patent applications, which patents/applications are assigned to the owner of the present invention, and which patents/applications are incorporated by reference herein in their entirety: U.S. patent application Ser. No. 10/366,075, entitled "SYSTEMS AND METHODS FOR MIGRATABLE SERVICES", by Eric M. Halpern, filed Feb. 13, 2003 (Attorney Docket No. BEAS-01195US2), which claims priority to U.S. Provisional Patent Application No. 60/358,418, entitled "HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS", by Eric M. Halpern, filed Feb. 21, 2001 (Attorney Docket No. BEAS-01195US0) and to U.S. patent application Ser. No. 10/341,207 entitled "METHOD FOR HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS", by Priscilla C. Fung et al., filed Jan. 13, 2003 (Attorney Docket No. BEAS-01173US1), which claims priority to U.S. Provisional Patent Application No. 60/359,226, entitled "HIGHLY AVAILABLE TRANSACTION RECOVERY FOR TRANSACTION PROCESSING SYSTEMS", by Priscilla C. Fung et al., filed Feb. 22, 2002 (Attorney Docket No. BEAS-01173US0).

### COPYRIGHT NOTICE

[0003] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

### BACKGROUND

[0004] Distributed networks are well known to programmers and computer system architects. A distributed network may include multiple nodes, computers, or servers. As used herein, a server is defined as a software process. A node or

cluster is a group of servers that may exist on a single hardware machine or share a physical resource such as a memory disk. Each server in a network usually has applications or objects that perform different functions. An application on a particular server may be initiated by another server or by the server it resides on. Distributed networks are advantageous in that several applications required to accomplish a task or a process may be distributed among several servers. The distributed applications may then be called upon when needed. Processes invoked simultaneously may be run on different servers instead of weighing down a single server and processor. This advantageously distributes processing power and contributes to a more efficient network.

[0005] Distributed transactions can span multiple servers, and servers often host resource managers (e.g. database connection pools or JMS queues) which participate in distributed transactions. As a result of distributed transaction participation, locks or other internal resources can be held up in the resource managers (e.g. databases locks are acquired for database records that are updated in a distributed transaction) on behalf of the distributed transaction until the distributed transaction is completed. For each distributed transaction, a particular server acts as the coordinator, which drives the participating transactional resources to commit atomically and thus the transaction to completion, via the Two Phase Commit (2PC) protocol. In the first phase of the 2PC protocol, the coordinator logs a record of the transaction and its participants persistently in its TLOG files after all participants are prepared successfully. Once prepared, all participants hold on to the acquired locks or other internal resources for the transaction until it is told to commit or rollback by the coordinator. In the second phase of the 2PC protocol, the coordinator commits all the participants, which then make the updates durable and release the locks and other internal resources. After all participants are successfully committed, the coordinator then releases the log record from its TLOG. Thus, if a coordinator fails, all the in-flight transactions that are logged in its TLOG files cannot be driven to completion, and thus all participants cannot release their locks or other internal resources, until the coordinator is restarted. Thus, with systems of the prior art, transaction recovery cannot take place before a failed server restarts. This limits the availability of transaction recovery of the failed server and thus the availability of other XA resources (e.g. JMS backends).

[0006] In addition to unexpected server failure, a server may be brought down intentionally. Application servers are often configured to run on specific machines to service client requests. These machines are brought down for periodic maintenance, machine servicing, and other reasons. As a result, the servers located on the downed machine are not able to service client requests to that machine or perform recovery of in-doubt transactions until the servers are restarted.

[0007] One approach the prior art has taken to address this problem is to migrate servers and their TLOG files to a backup or alternate machine. This allows unfinished transactions in a TLOG to be processed thus improving the availability of the failed server and preserving the operation and efficiency of a network. One such server migration system for use in a distributed network is included in the BEA TUXEDO application. TUXEDO supports migration of multiple servers residing on a machine. The servers must either consist of a group of servers or all the servers that reside on a machine. A group of servers within the TUXEDO application is defined

as a collection of servers or services on a machine often associated with a resource manager.

**[0008]** An administrator manually migrate servers using the TUXEDO application. The administrator specifies a primary machine and a secondary or back-up machine for each group of servers. Once a server group has failed or been deactivated by a user, a user may manually migrate the servers from the primary machine to the secondary machine. The primary then becomes the acting secondary machine, and the secondary becomes the acting primary machine. When the group of servers is to be moved back to the original primary machine, the user shuts-down the back-up machine and then migrate the server group back to the original primary machine.

**[0009]** Though a TLOG cannot be migrated by itself in Tuxedo, an administrator may manually migrate a TLOG file to a back-up server as a secondary step to of migrating a server. The TLOG migration is a manual process performed with tmadm commands. To migrate a TLOG in TUXEDO, an Atmadmin@ session is started and all servers that write to the TLOG are manually shut-down by a user. Next, the user dumps the TLOG contents into a text file, copies the name of the TLOG file to the back-up machine, and reads the text file into the existing TLOG for the specified back-up machine. The user then forces a warm start of the TLOG. Though a user may manually migrate a TLOG in this manner, TUXEDO does not support having multiple TLOGs per server.

**[0010]** There are several disadvantages to the prior art such as the TUXEDO application. Tuxedo does not support the migration of anything less than a group of servers. Thus, if a single server has crashed in a system or requires maintenance; multiple servers must be shut-down in order to migrate the server. Tuxedo requires that all servers that write to a particular TLOG file must be shut-down while the TLOG file is migrated. Tuxedo also does not support multiple TLOGs residing on a single server. In Tuxedo, there is only one TLOG for a group of servers. Once servers of a machine or group have migrated, and the corresponding TLOG is migrated thereafter, the secondary machine hosts only the migrated TLOG. Additionally, all migration steps in Tuxedo are done manually, including a complete shut-down of the secondary server when failing back to the original primary or master server. What is needed is a migration system that addresses the deficiencies of existing migration systems such as the one in Tuxedo.

## SUMMARY

**[0011]** A highly available transaction recovery service migration system in accordance with one embodiment of the present invention implements a server=s Transaction Recovery Service as a migratable service. In one embodiment of the present invention, the TRS is a server instance or software module implemented in JAVA. Highly available transaction recovery of a server within a cluster is achieved by migrating the TRS to another available server in the same cluster. This allows the backup server to read the transaction log and perform recovery on the behalf of the failed server. Each server in a cluster has a corresponding TRS, which maintains ownership of the server=s TLOG. When a primary server fails, the failed server=s TRS migrates to an available secondary server that resides in the same cluster as the failed server. The primary server and secondary server share access to the same memory disk. While residing on the secondary server, the migrated TRS obtains access to the TLOG of the failed pri-

mary server, reads the transaction log, and performs transaction recovery on behalf of the failed server. Multiple TRS instances may reside on any server, all of which performing transaction recovery on a single server. The migration may occur manually or automatically on a migratable services framework. The TRS of the failed primary server migrates back to the primary server in a fail back operation once the failed primary server is restarted. Failback operation may occur whether recovery is completed or not. No servers need to be shutdown during TRS failover migration to a secondary server or during TRS failback migration to the primary server. This expedites recovery and improves availability of the failed server thereby preserving the efficiency of the network and other servers.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0012]** FIG. 1a is a block diagram of a transaction recovery service migration system in accordance with one embodiment of the present invention.

**[0013]** FIG. 1b is a block diagram of a transaction recovery service migration system after failover in accordance with one embodiment of the present invention.

**[0014]** FIG. 2 is a diagram of a flow chart showing manual migration failover operation in accordance with one embodiment of the present invention.

**[0015]** FIG. 3 is a diagram of a flow chart showing manual migration failback operation after recovery is complete in accordance with one embodiment of the present invention.

**[0016]** FIG. 4 is a diagram of a flow chart showing manual migration failback operation before recovery is complete in accordance with one embodiment of the present invention.

**[0017]** FIG. 5 is a diagram of a flow chart showing automatic migration failover operation in accordance with one embodiment of the present invention.

**[0018]** FIG. 6 is a diagram of a flow chart showing automatic migration failback operation after recovery is complete in accordance with one embodiment of the present invention.

**[0019]** FIG. 7 is a diagram of a flow chart showing automatic migration failback operation before recovery is done in accordance with one embodiment of the present invention.

**[0020]** FIG. 8 is a diagram of a flow chart showing operation of automatic migration failback before recovery is complete in accordance with one embodiment of the present invention.

## DETAILED DESCRIPTION

**[0021]** A highly available transaction recovery service migration system in accordance with one embodiment of the present invention implements a server=s Transaction Recovery Service (TRS) as a migratable service. In one embodiment of the present invention, the TRS is a server instance implemented in JAVA. The TRS migrates to an available server that resides in the same cluster as the failed server. Highly available transaction recovery of a server within a cluster is achieved by migrating the TRS to another available server in the same cluster. The migrated TRS obtains the TLOG of the failed server, reads the transaction log, and performs transaction recovery on behalf of the failed server. A server may host multiple TRS instances at any time as well as coordinate their corresponding TLOG transactions. In one embodiment of the present invention, though a server may host multiple TRS instances and TLOGS, each TRS and TLOG corresponds to only one server. The migration may

occur manually or automatically on a migratable services framework. The TRS of the failed server migrates back in a fail back operation once the failed primary server is restarted. Failback operation may occur whether recovery is completed or not. No servers need to be shutdown during TRS failover migration to a secondary server or during TRS failback migration to the primary server. This expedites recovery of the failed server and while preserving the efficiency of the network and other servers.

**[0022]** A transaction recovery service migration system **100** in accordance with one embodiment of the present invention is shown in FIG. 1a. System **100** includes servers **110**, **120** and **140**. Each server has a corresponding TRS instance **112**, **122**, and **142**, respectively. Servers **110** and **120** share a common disk **130** while server **140** utilizes a separate disk **150**. Each server has a corresponding transaction log (TLOG) that resides on a disk. Server **110** has TLOG **114** on disk **130**, server **120** has TLOG **124** on disk **130**, and server **140** has TLOG **144** on disk **150**. All servers may reside on a single cluster **160**. Servers within a cluster may reside on the same or different machines (not shown).

**[0023]** In one embodiment of the present invention, each server is associated with only one TLOG. Each TRS has exclusive ownership of the TLOG for its particular server. Thus, TRS **122** has exclusive ownership of the TLOG for server **120**, TLOG **130**. When a particular server fails, the TRS for the failed server may be migrated to an alternate server. The migrated TRS may then perform recovery on the failed server's TLOG while residing on the alternate server. In one embodiment of the present invention, a TRS may only be migrated to a server that has access to the same disk space as the failed server. In particular, the shared disk space must contain TLOG files for the failed server. In another embodiment, an administrator may transfer the TLOG file for the failed server to the disk that the alternate server can access. The shared disk space may be a dual-ported SCSI, a storage area network (SAN), or some other reliable shared disk architecture.

**[0024]** For example, if server **110** fails, the TRS **112** can migrate to server **120** as shown in FIG. 1b. Once at server **120**, TRS **112** performs recovery on TLOG **114** corresponding to server **110**. In this case, server **110** is the primary server and server **120** is the back-up, secondary, or alternate server. A migration of a TRS from a primary server to a secondary server is called failover. In one embodiment of the present invention, a TRS may undergo failover migration to a server that shares access to the memory containing the TLOG of the failed server. In FIG. 1b, TRS **112** could not perform recovery on server **110** if migrated to server **140** because server **140** and server **110** do not share access to memory **130**. If TRS **112** migrates to server **140**, recovery by TRS **112** would require that server **140** obtain access to memory disk **130** or experience another migration to a server with access to disk **130**.

**[0025]** Each TRS is also associated with a migratable target as an alternate server. In one embodiment, administrators can configure a JTAMigratableTarget element for a clustered server. One example of a JTAMigratableTarget configuration is as follows:

```
[0026] <Server Name="server1" Cluster="mycluster"
ListenAddress="campton-1"
ListenPort="7001"><JTAMigratableTarget
Name="server1"
ConstrainedCandidateServers="server1,server2"/></
Server>
```

**[0027]** The runtime information is available from a JTA runtime MBean:

**[0028]** JTARecoveryRuntimeMBean, which can be obtained from a JTARuntimeMBean MBean. In one embodiment, at least two methods of the JTARuntimeMBean MBean may provide access to the JTARecoveryRuntimeMBean. One method is:

**[0029]** JTARecoveryRuntimeMBean[] getRecoveryRuntimeMBeans().

**[0030]** This method returns an array of JTARecoveryRuntimeMBean MBeans that corresponds to the TRS instances that are deployed on the current server. Another method is:

**[0031]** JTARecoveryRuntimeMBean getRecoveryRuntimeMBean(String serverName).

**[0032]** This method returns the JTARecoveryRuntimeMBean MBean that is associated with the specified server. If the corresponding JTARecoveryRuntimeMBean MBean is not deployed on this server, null is returned. The JTARecoveryRuntimeMBean MBean has several methods as well. One method is:

Boolean isActive()

**[0033]** This method returns whether the Transaction Recovery Service is currently activated on the server. Another method is:

**[0034]** int getInitialRecoveredTransactionTotalCount().

**[0035]** This method returns the total number of transactions that are read from the transaction log by the TRS. The administrator may use this information to increase the value of the MaxTransactions attribute of the JTAMBean MBean as appropriate. Another method is:

**[0036]** int getRecoveredTransactionCompletionPercent().

**[0037]** This method returns the percentage of the recovered transactions that are completed by the Transaction Recovery Service. In one embodiment, the name of the JTARecoveryRuntimeMBean MBean is the name of the original server of the Transaction Recovery Service.

**[0038]** Though failover and failback migration usually involve moving a single TRS instance at any time, a server may facilitate multiple TRS instances residing on the server and coordinate multiple transactions for TLOGs corresponding to the multiple TRS instances. In this case, the server performs recovery for multiple TRS instances in parallel. Server **120** in FIG. 1b facilitates recovery for failed server **110** as well as its own recovery and normal processing. In one embodiment, only the primary server may service new transactions. In this embodiment, a back-up server can not service new transactions for a failed primary server. To regain its TRS and service new transactions, the failed primary server must restart and the TRS must migrate back to the primary server. Migration of a TRS from a secondary server back to a primary server is called failback. Failback operation may vary according to whether recovery for the failed server is completed or not before failback occurs. After failback, TRS **112** would again reside in server **110** as shown in FIG. 1a.

**[0039]** In one embodiment of the present invention, manual migration failover is the only migration scenario that requires interaction by a user. An administrator may manually migrate the TRS of a failed server to another available server in the same cluster. The operation of a manual migration failover system in accordance with one embodiment of the present invention is shown in block diagram **200** of FIG. 2. Manual migration failover operation begins at start step **205**. A first server instance (S1) fails in step **210**. This may occur by an act

of an administrator or by server malfunction. In step 220, a user issues a migrate command to trigger the migration of the transaction recovery service for the first server instance (TRS1) from S1 to a second server instance (S2). This is usually done after the user has discovered a failed server or has shut-down a server. In one embodiment, a user may trigger the migration of TRS1 from S1 to S2 using a console implemented as a GUI system. The GUI console may be implemented so as to graphically display different clusters and servers. The user may choose a server having the corresponding TRS to migrate and the back-up server to receive the TRS. In one embodiment, the migration would be performed using a Java Transaction API (JTA). A JTA Recovery tab is provided for each server that allows administrators to specify various attributes of a Migratable Target and perform manual migration of the Transaction Recovery Service associated with the server. The appearance as viewed on a computer screen of a GUI allowing a user to issue a migrate command in accordance with one embodiment of the present invention is shown in FIG. 3. In another embodiment of the present invention, a user or system administrator may trigger a manual migration of a TRS using a command line. A command line administration tool, implemented as a Java program, may allow a user to specify the TRS to be migrated and what server to migrate the TRS to. The command line tool may also require a user to enter a username and password in order to perform the migration. The general format of such a command line command in accordance with one embodiment of the present invention is shown below.

**[0040]** java weblogic.Admin [-url <url>] [-username <username>] [-password <password>] MIGRATE-jta-migratabletarget <server name>-destination <destination server name>

**[0041]** [-sourcedown] [-destinationdown]

In another embodiment of the present invention, manual migration may be triggered by a user programmatically using a JMX MBean. In particular, a JMX MigratableTarget MBean of the TRS may be used to trigger the migration of a TRS from one server to another. An example of the code comprising a MigratableTarget MBean in accordance with one embodiment of the present invention is below.

**[0042]** Though specific code is listed above, an MBean can be configured and implemented in various ways to achieve the result of triggering the migration of a TRS from one server to another. These variations of code are all considered within the scope of the present invention. The present invention is not intended to be limited to the JMX MigratableTarget MBean code example listed above.

**[0043]** Next, the migratable framework detects that the S1 is down in step 230. In one embodiment, the user issued command in step 220 informs the migratable framework that the server is down. When the migratable framework detects the server is down in step 220, the migratable framework moves the TRS to a back-up server. The back-up server may be specified by a user or be pre-determined by the migratable framework system. After step 230, the migratable framework then activates TRS1 on S2 in step 240. In one embodiment, all migratable services including instance TRS1 must implement a particular interface. The interface must be registered with the migratable framework and includes migration activate and deactivate methods. In this embodiment, migration is activated when the migratable framework calls the migration activate method of TRS1 currently residing on S2. Then, TRS1 reads and processes the TLOG for S1 in step 250. TRS1 reads S1's TLOG files; instantiates the transactions of the TLOG files, puts them into the transaction map of S2, and schedules resource recovery for S1. As a result, S2 services will read and coordinate the transactions from S1's TLOG. The S2 server becomes the coordinator of previously in doubt transactions and talks to different coordinators and resource managers to resolve transactions. Next, TRS1 performs recovery on behalf of S1 in step 260 while still residing on S2. TRS1 performs recovery on behalf of S1 asynchronously. Meanwhile, the backup server's own transaction manager functions to accept new transactions and perform its own transaction recovery as usual. Thus, there may be more than one instance of TRS activated on a back-up server at any time, the multiple TRS instances originating from different servers. The recovery may include driving prepared transactions to completion and performing resource recovery. Manual migration failover is then complete and operation ends at step 265. Similar manual migration can be performed to migrate

---

```
import weblogic.management.Admin;
import weblogic.management.configuration.MigratableTargetMBean;
import weblogic.management.configuration.ServerMBean;
import weblogic.management.runtime.MigratableServiceCoordinatorRuntimeMBean;
// Obtain the MigratableServiceCoordinatorRuntimeMBean
MigratableServiceCoordinatorRuntimeMBean msc =
Admin.getAdminServer( ).getMigratableServiceCoordinatorRuntime( );
// Obtain the MigratableTargetMBean of the server whose Transaction
Recovery Service needs to be migrated
ServerMBean server1 = (ServerMBean)
Admin.getMBeanHome( ).getConfigurationMBean("server1",
"MigratableTargetConfig");
MigratableTargetMBean mt = server1.getJTAMigratableTarget( );
// Obtain the configuration ServerMBean of the server to which the
Transaction Recovery Service will be migrated
ServerMBean server2 = (ServerMBean)
Admin.getMBeanHome( ).getConfigurationMBean("server2",
"MigratableTargetConfig");
// Perform the migration of Transaction Recovery Service of "server1" to
"server2"
msc.migrateJTA(mt, server2, false /*source up*/, false /*destination up*/);
```

---

the TRS to another available backup server if a backup server fails before completing the transaction recovery actions for the original server.

**[0044]** In one embodiment, failback occurs when a failed primary server restarts and is ready to receive its TRS instance back from a back-up server. The operation of a manual migration failback performed after recovery is completed in accordance with one embodiment of the present invention is shown in diagram 400 of FIG. 4. System operation begins with start step 405. In step 410, an alternate or back-up server S2 completes recovery for a primary server S1. In one embodiment, recovery completion occurs when TRS1 of S1 finishes recovery for S1 while residing on S2. Once TRS1 completes recovery, TRS1 relinquishes control of S1's TLOG files. Next, TRS1 migration back to S1 is initiated in step 420. In one embodiment, an administrator may manually initiate migration of the TRS back to the original server. In another embodiment, migration is initiated when TRS1 contacts the migratable framework and makes a request to migrate TRS1 back to S1. In step 430, the migratable framework completes the migration of TRS1 from S2 back to S1. In one embodiment, the migratable framework first deactivates TRS1 on S2 by calling a deactivation method of TRS1. During the deactivation of TRS1, S2 performs cleanup and removes any remaining transactions of S1 from its internal transaction map. After this deactivation of TRS1, the migratable framework moves TRS1 to S1. Then, the migratable framework activates TRS1 on S1 using a call to an activation method of TRS1. Operation then ends in step 435. When S1 later restarts, S1 will regain ownership of the TLOG corresponding to S1 and will not need to perform further recovery work.

**[0045]** Manual migration failback may also be performed before recovery is complete. Operation of manual migration failback performed before recovery is completed in accordance with one embodiment of the present invention is shown in diagram 500 of FIG. 5. Operation begins at start step 505. In step 510, S1 is restarted. Up until just before S1 restart, S2 is still performing recovery work for S1. During S1 startup, S1 notifies S2 that S1 is now operational. In one embodiment, the notification is in the form of an administrative MBean event sent from S1 to S2. Next, TRS1 migration back to S1 is initiated in step 520. In one embodiment, TRS1 residing on S2 sends a request to the migratable framework to migrate TRS1 back to S1. Then, TRS1 migrates from S2 to S1 in step 530. In one embodiment, an administrator may manually migrate TRS1 back to S1 from S2. This may be performed when the back-up server fails to implicitly migrate TRS1 back to the original server S1. During this manual migration, the migratable service framework deactivates the TRS1 on S2. The deactivation of TRS1 suspends recovery for S1 and allows S2 to perform cleanup and remove any remaining transactions of S1 from its internal transaction map. In another embodiment, the migratable framework first deactivates TRS1 on S2 by calling a deactivation method of TRS1. The deactivation of TRS1 on S2 suspends recovery processing for S1. Thus, S2 may checkpoint the TLOG for S1, purge transactions in its transaction map originating from S1's TLOG, and stop resource recovery performed for S1. During the deactivation of TRS1, S2 performs cleanup and removes any remaining transactions of S1 from its internal transaction map. S2 then relinquishes control of S1's TLOG. After this deactivation of TRS1, the migratable framework moves TRS1 to S1. Then, the migratable framework activates TRS1 on S1. TRS1 is activated by issuing a call to an activation

method of TRS1. Operation then ends in step 545. Though falling under the category of manual migration, no administrator intervention is required for manual migration failback before recovery is done. Once S1 regains ownership of TRS1, it restarts and completes the remaining transaction recovery work.

**[0046]** Automatic migration occurs without any administrative intervention required. Automatic failover and failback migration occur without any input from a user. In one embodiment, migration occurs seamlessly and without notification to the user. Operation of automatic migration failover in accordance with one embodiment of the present invention is shown in diagram 600 of FIG. 6. Operation begins at start step 605. Server failure of S1 occurs at step 610. Next, TRS1 is migrated to S2 in step 620. In one embodiment, TRS1 migration to S2 is triggered when the migratable framework detects the failure of S1. The migratable framework then migrates TRS1 from S1 to S2. In one embodiment, a user may specify a preferred order of back-up servers. A preferred server list as indicated by a user may be stored in the migratable target MBean. The migratable framework will then attempt to migrate TRS1 to the preferred back-up servers in the order specified by the user. Then, the migratable framework activates TRS1 on S2 in step 630. In one embodiment, TRS1 is activated when the migratable framework calls a migration activation method of TRS1. Next, S1's TLOG is read and processed in step 640. In one embodiment, during activation on S2, TRS1 reads S1's TLOG files regarding S1 transactions and configures S2 accordingly. In one embodiment, TRS1 instantiates the S1 TLOG files, places the files in S2's transaction map, and schedules resource recovery for S1. Thus, S2 is configured to be the coordinator of the transactions read from S1's TLOG. Next, TRS1 performs recovery on behalf of S1 in step 650. In one embodiment, recovery includes driving prepared transactions to completion and performing resource recovery. Automatic migration failover operation then ends in step 655. S2's own transaction manager and TRS2 function as usual during automatic migration failover. Similar manual migration can also be performed to migrate TRS1 to another available backup server if a backup server S2 fails before completing the transaction recovery actions for the original server S1.

**[0047]** Automatic migration failback is similar to automatic migration failover in that no administrative intervention is required. Operation of automatic migration failback after recovery is complete in accordance with one embodiment of the present invention is shown in diagram 700 of FIG. 7. Operation begins at start step 705. Next, S1 recovery is completed in step 710. In one embodiment of the present invention, recovery is completed when TRS1 finishes recovery for S1 while located on S2. The TRS1 checkpoints S1's TLOG files and relinquishes control of S1's TLOG files. Then, TRS1 migration back to S1 is initiated in step 720. In one embodiment of the present invention, the migration is initiated when TRS1 requests the migratable framework to migrate TRS1 back to S1. Next, the migratable framework completes migration of the TRS to S1 in step 730. The migratable framework first deactivates TRS1 on S2 by calling a deactivation method of TRS1. TRS1 deactivation results in S2 relinquishing control of S1's TLOG. During the deactivation of TRS1, S2 performs cleanup and removes any remaining transactions of S1 from its internal transaction map. After this deactivation of TRS1, the migratable framework moves TRS1 to S1. Then, the migratable framework activates TRS1

on S1. TRS1 is activated by issuing a call to an activation method of TRS1. Once migration is complete, operation ends in step 735. When S1 is restarted, S1 regains ownership of it=s TLOG as a result of TRS1 residing on S1. S1 does not need to perform additional recovery work.

**[0048]** Automatic migration failback may also occur before recovery of the failed server is complete. Operation of automatic migration failback before recovery is complete in accordance with one embodiment of the present invention is shown in diagram 800 of FIG. 8. Operation begins with start step 805. Next, S1 is restarted in step 810. At the time of S1 restart, TRS1 residing on S2 has not completed performing recovery on behalf of S1. TRS1 migration is then initiated in step 820. In one embodiment, the migratable framework initiates migration upon detecting that S1 has performed startup. The migratable framework may detect the failure of the server itself or be notified of the server startup by an outside source. In one embodiment, S1 informs S2 that S1 has restarted. After migration has been initiated in step 820, the TRS1 migrates to S1 in step 830. In one embodiment, the migratable framework first deactivates TRS1 on S2 by calling a deactivation method of TRS1. The deactivation of TRS1 on S2 suspends recovery processing for S1 by TRS1 on S2. The deactivation includes checkpointing the TLOG for S1, purging transactions in its transaction map originating from S1's TLOG, and stopping resource recovery performed for S1. During the deactivation of TRS1, S2 performs cleanup and removes any remaining transactions of S1 from its internal transaction map. S2 then relinquishes control of S1's TLOG files as TRS1 migrates back to S1. After this deactivation of TRS1, the migratable framework moves TRS1 to S1. Then, the migratable framework activates TRS1 on S1. TRS1 is activated by issuing a call to an activation method of TRS1. Once migration is complete, operation ends in step 835. Once S1 regains ownership of TRS1 and restarts, S1 performs the remaining recovery work.

**[0049]** A highly available transaction recovery service migration system in accordance with one embodiment of the present invention implements a server=s Transaction Recovery Service as a migratable service. In one embodiment of the present invention, the TRS is a server instance or software module implemented in JAVA. Each server in a cluster has a corresponding TRS, which maintains ownership of the server=s TLOG. When a primary server fails, the failed server=s TRS migrates to an available back-up server that resides in the same cluster as the failed server. The primary server and back-up server share access to the same memory disk. While residing on the back-up server, the migrated TRS obtains access to the TLOG of the failed server, reads the transaction log, and performs transaction recovery on behalf of the failed server. The migration may occur manually or automatically on a migratable services network. Automatic migration requires the TRS be deployed on the migratable service framework. The TRS of the failed server migrates back to the primary server in a fail back operation once the failed primary server is restarted. Failback operation may occur whether recovery is completed or not. This expedites recovery and improves the availability of the failed server thereby preserving the efficiency of the network and other servers.

**[0050]** With systems of the prior art, transaction recovery cannot take place before a failed server restarts. This limits the availability of transaction recovery of the failed server and thus the availability of other XA resources (e.g. JMS back-

ends). In the present invention, we achieve highly available transaction recovery of a server within a cluster by migrating the Transaction Recovery Service to another available server in the same cluster. This allows the backup server to read the transaction log and perform recovery on the behalf of the failed server.

**[0051]** Transaction Recovery Service depends on the migratable service framework for manual and automatic migration support. JMS backends, which are XA resources, in turn depend on Transaction Recovery Service migration to recover their resources when they are migrated.

**[0052]** Administrators can manually migrate the Transaction Recovery Service of a failed server (i.e. the original server) to another available server in the same cluster. Before the original server restarts, the administrator may also need to manually migrate the Transaction Recovery Service back to the original server. Manual migration is recommended if the server has failed and is not expected to be restarted soon. In the absence of migration, a server will perform its own transaction recovery when it is restarted after a failure.

**[0053]** There are two aspects of manual migration: fail-over and fail-back. When a clustered server (Transaction Coordinator) fails, the Transaction Recovery Service associated with the failed server (i.e. the original server) can be manually migrated to another available server (the backup server) in the same cluster, via either the Administration Console or the JMX API. During the migration, the migratable service framework activates the Transaction Recovery Service on the backup server. During activation, the Transaction Recovery Service reads the transaction log of the failed server and initializes the transaction recovery asynchronously. Meanwhile, the backup server=s own transaction manager functions (accepting new transactions and performing its own transaction recovery) as usual. Note that there may be more than one instances of Transaction Recovery Service (that originates from different servers) activated on a backup server at the same time.

**[0054]** Similarly, manual migration can also be performed to migrate the Transaction Recovery Service to another available backup server if a backup server fails before completing the transaction recovery actions for the original server.

**[0055]** Fail-back happens when migrating the Transaction Recovery Service from a backup server to the original failed server. Note that fail-back is implicit and does not require administrator intervention unless a migration error occurs, as described below.

**[0056]** When a backup server finishes transaction recovery before the original server restarts, it gives up ownership of the Transaction Recovery Service and migrates it back implicitly to the original server. No administrative action is needed in this case. Subsequently, when the original server restarts, it regains ownership of its Transaction Recovery Service. However, it does not need to do further recovery work.

**[0057]** When the resources to be recovered are not available, transaction recovery for the original server may not be finished when the original server is to be restarted. In this case, the backup server, on detecting that the original server is coming up, suspends the ongoing transaction recovery, performs some internal cleanup, and implicitly migrates the Transaction Recovery Service back to the original server. No administrative action is needed in this case. The original server, once it regains the ownership of its Transaction Recovery Service, restarts successfully and finishes the remaining transaction recovery work.

**[0058]** If the backup server fails to implicitly migrate back to the original server, the original server will fail to boot. In this case, the administrator needs to manually migrate the Transaction Recovery Service from the backup server to the original server before rebooting the original server. During the manual migration, the migratable service framework deactivates the Transaction Recovery Service on the backup server. The deactivation suspends the transaction recovery, performs some internal cleanup and gives up ownership of the Transaction Recovery Service. Subsequently, when the original is restarted, it regains ownership of its Transaction Recovery Service and finishes the remaining transaction recovery work.

**[0059]** Under automatic migration, no administrator intervention is needed, and the migration of Transaction Recovery Service happens transparently under the control of migratable service framework.

**[0060]** There are two aspects of automatic migration: fail-over and fail-back. When the migratable service framework detects that a clustered server (Transaction Coordinator) has failed, it automatically migrates the Transaction Recovery Service associated with the failed server to the next available server (the backup server) in the preferred server list of the migratable target MBean. During the migration, the migratable service framework activates the Transaction Recovery Service on the backup server. During activation, the Transaction Recovery Service reads the transaction log of the failed

Transaction Recovery Service back to the original server. The original server, once it regains the ownership of the Transaction Recovery Service, restarts successfully and finishes the remaining transaction recovery work.

**[0065]** Embodiments may provide Administrative means to configure, migrate and monitor Transaction Recovery Services.

**[0066]** Embodiments can provide configuring Transaction Recovery Service. Each server in a cluster is associated with a Transaction Recovery Service, and each Transaction Recovery Service is associated with a Migratable Target. If the Transaction Recovery Service is not configured for a particular server, no migration will be enabled. In this case, if the server fails, transaction recovery will only be performed after it restarts.

**[0067]** Embodiments can provide configuring Transaction Recovery Service via the Console. A new JTA Recovery tab for the server will be provided, from which administrators can specify various attributes of the Migratable Target and perform manual migration of the Transaction Recovery Service associated with the server.

**[0068]** Embodiments can provide configuring Transaction Recovery Service via config.xml. Administrators can configure a JTAMigratableTarget element for a clustered server. The following is an example of the JTAMigratableTarget configuration:

---

```
<Server Name="server1" Cluster="mycluster" ListenAddress="campton-1" ListenPort="7001">
  <JTAMigratableTarget Name="server1" ConstrainedCandidateServers="server1,server2"
/></Server>
```

---

server and initializes the transaction recovery asynchronously. Meanwhile, the backup server's own transaction manager (including its own transaction recovery) functions as usual.

**[0061]** Similar automatic migration sequences can also happen to migrate the Transaction Recovery Service to another backup server if a backup server fails before completing the transaction recovery actions.

**[0062]** Failback happens when migrating the Transaction Recovery Service from a backup server to the original failed server. There are also two cases: Fail back without automatic migration and Failback with automatic migration.

**[0063]** Fail-back without automatic migration is similar to fail-back when recovery is done for manual migration as described above. When a backup server finishes transaction recovery before the failed server restarts, it gives up ownership of the Transaction Recovery Service and migrates it back implicitly to the failed server. Subsequently, when the failed server restarts, it regains ownership of its Transaction Recovery Service. However, it does not need to do further recovery work.

**[0064]** Fail-back with automatic migration happens if the backup server has not finished transaction recovery when the original server restarts. In this case, the backup server, on detecting that the original server is coming up, suspends the ongoing transaction recovery, performs some internal cleanup, and gives up the ownership of the Transaction Recovery Service of the original server. The migratable service framework will then transparently migrate the Transac-

**[0069]** Embodiments can provide Monitoring JTA recovery service. Each server maintains runtime information of all Transaction Recovery Service instances that it hosts. The runtime information is available from a new JTA runtime MBean: JTARecoveryRuntimeMBean, which can be obtained from the existing JTA runtime MBean. Two new methods are added to the existing JTARuntimeMBean MBean to provide access to the JTARecoveryRuntimeMBean.

**[0070]** The first, JTARecoveryRuntimeMBean[ ] getRecoveryRuntimeMBeans() method returns an array of JTARecoveryRuntimeMBean MBeans that corresponds to the Transaction Recovery Service instances that are deployed on the current server.

**[0071]** The second, JTARecoveryRuntimeMBean getRecoveryRuntimeMBean(String serverName) method returns the JTARecoveryRuntimeMBean MBean that is associated with the specified server. If the corresponding JTARecoveryRuntimeMBean MBean is not deployed on this server, null is returned.

**[0072]** The JTARecoveryRuntimeMBean MBean has the following methods:

A boolean is Active() method returns whether the Transaction Recovery Service is currently activated on the server. An int getInitialRecoveredTransactionTotalCount() method returns the total number of transactions that are read from the transaction log by the Transaction Recovery Service. The administrator may use this information to increase the value of the MaxTransactions attribute of the JTAMBean MBean as

appropriate. An `int getRecoveredTransactionCompletionPercent()` method returns the percentage of the recovered transactions that are completed by the Transaction Recovery Service. Note that the name of the `JTAReccoveryRuntimeMBean` MBean is the name of the original server of the Transaction Recovery Service.

**[0073]** In addition to an embodiment consisting of specifically designed integrated circuits or other electronics, the present invention may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present disclosure, as will be apparent to those skilled in the computer art.

**[0074]** Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application specific integrated circuits or by interconnecting an appropriate network of conventional component circuits, as will be readily apparent to those skilled in the art.

**[0075]** The present invention includes a computer program product which is a storage medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data.

**[0076]** Stored on any one of the computer readable medium (media), the present invention includes software for controlling both the hardware of the general purpose/specialized computer or microprocessor, and for enabling the computer or microprocessor to interact with a human user or other mechanism utilizing the results of the present invention. Such software may include, but is not limited to, device drivers, operating systems, and user applications. Ultimately, such computer readable media further includes software for implementing Node Managers.

**[0077]** Included in the programming (software) of the general/specialized computer or microprocessor are software modules for implementing the teachings of the present invention, including, but not limited to, separating planes of a source image, averaging at least one of foreground and background colors, replacing colors, and compensating for error introduced by color replacement in one plane by feeding error into a second plane, storage, communication of results, and reconstructing an image according to the processes of the present invention.

**[0078]** Other features, aspects and objects of the invention can be obtained from a review of the figures and the claims. It is to be understood that other embodiments of the invention can be developed and fall within the spirit and scope of the invention and claims.

**[0079]** The foregoing description of preferred embodiments of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Obviously, many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the

principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the following claims and their equivalence.

1. A transaction recovery apparatus, comprising:
  - a shared storage to store a first transaction log corresponding to a first server and a second transaction log corresponding to a second server,
  - an interface to the first server, the first server having a first transaction recovery service to failover migrate from the first server to the second server when the first server fails; and
  - an interface to the second server; the second server having a second transaction recovery service to coordinate transactions from the first transaction log upon the failover migration of the first transaction recovery service.
2. The transaction recovery apparatus of claim 1 wherein the second server instance is configured to receive more than one transaction recovery service.
3. The transaction recovery apparatus of claim 1 further including an input apparatus, the input apparatus enabling receiving a migrate command to initiate the failover migration.
4. The transaction recovery apparatus of claim 1 wherein the first server instance, the second server instance, and the shared storage reside on a single cluster.
5. The transaction recovery apparatus of claim 1 wherein the first transaction recovery service is configured to be moved from the first server instance to the second server instance and activated on the second server instance.
6. The transaction recovery apparatus of claim 1 wherein the first transaction recovery service is configured to failback migrate from the second server instance to the first server instance upon detecting the startup of the server instance startup.
7. The transaction recovery apparatus of claim 7 wherein the first transaction recovery service is configured to be deactivated, moved from the second server instance to the first server instance, and activated on the first server instance.
8. The transaction recovery apparatus as claimed in claim 3 wherein the input apparatus includes an input configured to be used with a Graphical User Interface (GUI), the GUI configured to accept manual input of a migration command.
9. The transaction recovery apparatus as claimed in claim 3 wherein the input apparatus includes a keyboard, the keyboard configured to accept manual input of a migration command entered on a command line.
10. The transaction recovery apparatus as claimed in claim 9 wherein the input apparatus includes a keyboard, the keyboard configured to accept manual input of a migration command by using a JMX MBean.
11. The transaction recovery apparatus as claimed in claim 1 further comprising a migratable framework, the migratable framework configured to automatically failover migrate the first transaction recovery service from the first server to the second server.
12. The transaction recovery apparatus as claimed in claim 11 wherein the migratable framework is configured to automatically failback migrate the first transaction recovery ser-



vice from the second server to the first server when the second server has completed coordinating transactions from the first transaction log.

13. The transaction recovery apparatus as claimed in claim 11 wherein the migratable framework is configured to auto-

matically failback migrate the first transaction recovery service from the second server to the first server when the first server has restarted.

\* \* \* \* \*