

IBM Tivoli System Automation for Multiplatforms

Upgrade and migration steps: node by node



This presentation details the procedure for migrating a Tivoli® System Automation for Multiplatforms (TSAMP) environment using a node-by-node approach. This approach allows the underlying managed resources to remain running, as might be required in a production environment.

This approach minimizes the interruption to an actively running production environment; however, there is still a requirement to move active resources from node to node to allow upgrading of all nodes. This move involves downtime that depends on the failover time of the managed application involved. For this reason, you should perform a node-by-node migration during a scheduled maintenance outage or a time period where outages would cause minimal impact.

Refer to Chapter 1 of the Installation and Configuration Guide for IBM Tivoli System Automation for Multiplatforms for complete details about the installation and migration procedures and options.

Table of contents

- Introduction
- Preparation and prerequisite checking
- Steps to perform node-by-node migration
- Verifying the active and installed version
- Completing the migration

This IEA module includes the following topics:

- Introduction
- Preparing and checking prerequisites
- Performing node-by-node migration
- Verifying the active and installed version
- Completing the migration

Introduction

- The migration process involves :
 - Installing a new code level on each node in the domain
 - Activating the new code level after all nodes are at the new installed code level
- There are two different procedures for migrating to a higher code level
 - Migrating an entire domain with it offline
 - Performing a node-by-node migration
- Migrating the nodes of a domain one by one has the advantage that the domain remains online and Tivoli System Automation for Multiplatforms remains available during the migration

I will start by defining what migration means in a TSAMP environment. The software stack is duplicated across every node in a given domain. New code, such as a fix pack, needs to be installed separately on each node in the domain. This is the first part of a migration. To complete the migration, the newly installed code levels for the cluster software (RSCT) and automation software (TSAMP) are activated with the domain in an online state.

The primary purpose of TSAMP is to make an underlying application highly available. You might need to migrate TSAMP while the underlying managed application is kept online and active. If this is the case, a node-by-node migration technique needs to be performed. This technique is also known as a rolling upgrade.

If it is possible to bring down the complete environment across all nodes, namely offline the complete domain, then the entire domain can be migrated at once, which is simpler and quicker.

For the purposes of this module, you are only covering a node-by-node migration procedure.



Preparation and prerequisite checking (1 of 2)

- Download the appropriate fix pack for your version of TSAMP and platform
 - Check the current version
 - [Issamctrl](#)
 - Click the Download link on the TSAMP Support Web site
 - <http://www.ibm.com/software/tivoli/support/sys-auto-multi/>
 - Check the documentation, which includes the Release Notes
 - <http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html>
- FTP (using binary mode) the package so that it is local to each node
- Untar the downloaded package on each node, for example:
 - `tar -xvf 3.1.0-TIV-SAMP-AIX-FP0005.tar`
- Change to the directory that is created from the untar step, for example depending on your platform:
 - `cd SAM3105MPAIX`
 - `cd SAM3105MPLinux`
 - `cd SAM3105MPSunOS`

4

Upgrade and migration steps: node by node

© 2010 IBM Corporation

The majority of migrations involve upgrading to a new fixpack level. In order to download the correct fixpack, you need to know your current TSAMP version, release, and fixpack level. You also need to know your platform.

Select the appropriate download package from the TSAMP Support Web site. If you are unsure which package you need, refer to the latest Release Notes for your TSAMP version by visiting the second URL on this slide.

It is recommended that you have a copy of the downloaded installation package on each node. There have been some installation failures when attempting to install from a common shared disk, typically because of access permission problems. Consider installing from a local package per node if you run into installation problems with error messages that don't make sense for your situation.

The actual TSAMP installer utility is always located in a directory with a name that starts with SAM. This name is created when the tar file is extracted. At the bottom of this slide are three examples, based on the different platforms supported by TSAMP.

Preparation and prerequisite checking (2 of 2)

- Open a root session
- To start the prerequisites check, issue the command:
`./prereqSAM`
- When the prerequisite check is complete, check the following log file for information about missing prerequisites:
`/tmp/prereqSAM.<#>.log`
where <#> is a number that increments each time the prereq check is run; the highest number identifies the most recent log file
- If the prereq check fails on any of the nodes, resolve before upgrading any of the nodes
- Important: For the AIX® platform and TSAMP version 2.3 or lower, you would need to upgrade the RSCT code level (the underlying cluster software) before upgrading the TSAMP code level per the prerequisite requirements documented in the TSAMP Release Notes:
<http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms2.2.html>
- RSCT code level is also typically upgraded when a TSAMP fix pack is installed. This happens on all platforms with fix packs for TSAMP v3.1, but not on AIX for TSAMP versions before 3.1

You must perform preinstallation, installation, and postinstallation activities as the root user.

The installer utility runs a prerequisite checker utility called prereqSAM. Even so, you should prerun prereqSAM on each node first to ensure all nodes meet the prerequisite requirements before starting the upgrade on any of the nodes.

The standard output from prereqSAM indicates if all prerequisites are met. You only need to examine the prereqSAM log if there is a prerequisite failure indicated on standard out.

It is important to note that TSAMP installs and upgrades RSCT (the underlying cluster component) on all versions of Linux® for all versions of TSAMP. However on AIX, RSCT packages were not included in the TSAMP installation or upgrade process until version 3.1. Prior to version 3.1 and on AIX, the process of installing or upgrading RSCT to the prerequisite level was a manual step, performed independently of the TSAMP migration process.

Steps to perform node-by-node migration (1 of 3)

Start with your standby node, where your highly available application is not currently active, for example, node2, and perform the following steps as the root user

- 1 Exclude the node from automation to ensure that resources that must be kept available are forced to move to another node in the peer domain

```
samctrl -u a node2
```

Note: The move operations might take a moment

- 2 Do not update the installed code level while the node is online in a domain. Offline the node from the cluster perspective:

```
stoprnode node2  
lsrnode
```

- 3 To upgrade the node, run the script ./installSAM from the installation directory for the fix pack that was extracted, for example :

```
<your_media_location>/SAM3105MPAIX/installSAM
```

To ensure a minimum interruption to your managed applications, start the upgrade process on the node that is currently considered your standby server. This server is the one where you do not have online resources currently running or where offlining resources have the least impact.

You might need to move the online resources to another node if they are defined to run on different servers. Otherwise, you might need to offline resources that are fixed to run on the one server. The point is that all resources on the node that you are about to upgrade need to be in an offline state.

One of the ways of doing this is to exclude the node from automation using the samctrl command shown on this slide in step 1. Alternatively, you can issue a resource group move request using the 'rgreq -o move ...' command. For fixed resources, you need to change their resource group to offline using the 'chrg -o offline ...' command.

Confirm all resources are offline on the node you are about to stop by checking the output of the 'lssam' command. Only then would you stop the node from participating in the domain by using the 'stoprnode ...' command in step 2. Confirm the node is offline in the domain by using the 'lsrnode' command, issued on the opposite node. If you run 'lsrnode' on the node that you have stopped, you will see an error message returned.

With the node offline, you can start the TSAMP installation process by starting the installSAM utility, shown as step 3 on this slide.

Steps to perform node-by-node migration (2 of 3)

4 When the installer is complete, check the following log file if it indicates an error occurred:

```
/tmp/installSAM.<#>.log
```

where <#> is a number that increments each time installSAM is run; the highest number identifies the most recent log file

5 Start the node so that it can rejoin the domain:

```
startprnode node2 [issued from node1]
```

or

```
startprdomain [issued from node2]
```

6 If you had previously excluded the node from automation, include the node in automation again:

```
samctrl -u d node2
```

At this point, TSAMP (and RSCT) code level on node2 is upgraded, though it is still using the older code level for both TSAMP and RSCT. The new code levels have not been activated and cannot be activated until all nodes are at the same new levels

Step 4 is optional. You only need to examine the installSAM log if there were errors displayed during the installation process. If unsure, check it. You should save the installSAM logs if your system is configured to delete all contents of /tmp on reboot.

Assuming a successful upgrade installation, the node can be restarted as shown in step 5. You can either issue the 'startprnode ...' command from a node that is currently online, or you can issue the 'startprdomain ...' command on the currently offline node. It rejoins the existing online domain. Either method yields the same results. Neither method affects the online status of the overall domain, and therefore the status of your managed applications is untouched.

You only perform step 6 if you excluded the node at the beginning of this migration process (step 1 in the previous slide).

This node (node2 as shown in the contents of this slide and the previous slide) now has the new code level for TSAMP (and potentially RSCT). The code is not active yet. The node is still running at the older code level.

Verifying the active and installed version

- When upgrading the TSAMP code level, the previous and new code levels coexist until the migration is complete
- The following command shows you the active version number AVN (3.1.0.0 in the sample below) and the installed version number IVN (3.1.0.5 in the sample below) of the product:
`lssrc -ls IBM.RecoveryRM`
- When IVN and AVN are the same, migration is complete. Here is an example of the output:

```
Subsystem : IBM.RecoveryRM
PID : 27973
Cluster Name : ws
Node Number : 1
Daemon start time : Wed Dec 9 08:09:10 2009
Daemon State:
  My Node Name : node2
  Master Node Name : lnxcm3x (node number = 1)
  Our IVN : 3.1.0.5
  Our AVN : 3.1.0.0
  Our CVN : 11082527751 {0x140861007}
  Total Node Count : 1
  Joined Member Count : 1
  Config Quorum Count : 1
  Startup Quorum Count : 1
  Operational Quorum State: HAS_QUORUM
  In Config Quorum : TRUE
  In Config State : TRUE
  Replace Config State : FALSE
```

To verify the installed and active versions of TSAMP, you can use the 'lssrc -ls IBM.RecoveryRM' command to check the Installed Version Number (IVN) and compare it to the Active Version Number (AVN). You should see the IVN set to the version you just installed and the AVN set to the previous version.

When the IVN and AVN differ, the migration is not considered complete. However, there is nothing more that can be done on this node until all the other nodes contain the same installed version, the same IVN.

Time to move onto repeating these upgrade steps on the other node or nodes.

Steps to perform node-by-node migration (3 of 3)

- Switch to a root session on the other node, for example, node1, to upgrade it to the same new code level as node2
- The node must be offline.
 - Move the active/online resources to the other node, for example:


```
rgreq -o move <your_floating_resource_groups>
```

Note: Do not use rgreq in DB2® V9.5/9.7 HADR environments. Instead, use the native 'db2 takeover hadr ...' (executed on the standby node, node2) command to move the active HADR database to the standby node
 - Offline the fixed resources that only run on this node, for example:


```
chrg -o offline <your_fixed_resource_groups>
```

Note: DB2 v9.5 and 9.7 HADR environments should use the native 'db2stop force' command to stop the local DB2 instance instead of offlining the resource group for the DB2 instance
- To upgrade node1, repeat the same steps 1 - 6 (slides 6 and 7) that were performed on node2
- Also, perform the same IVN/AVN verification as described on the previous slide (slide 8)

The upgrade process is repeated for each node. The requirements are the same: stop all online resources or move them to other nodes so that the 'lssam' command shows all resources on this node in an offline state. In this example, you are now upgrading node1.

Consider the case where DB2 is the underlying managed application, in particular DB2 9.5. For the HADR database resources, you move these resources to an alternative DB2 server by performing controlled takeovers. Instead of using the TSAMP command 'rgreq -o move ...', use the native 'db2 takeover ...' command for each database. In the case of the locally running DB2 instance, you stop the instance after the HADR takeovers have completed. Use the native 'db2stop force' command instead of offlining the instance's resource group.

When repeating steps 1 to 6, make sure you substitute the node name for the one you are currently working on. In this example, it would now be node1 where you see node2. In one case, you would be using node2 where the previous instructions described an action on node1.

Completing the migration (1 of 2)

To check and complete the migration, perform the following steps

- 1 Make sure that the domain is started and that all nodes in the domain are online
- 2 Issue the **lsrpdomain** command to display the version of RSCT that is active in the peer domain, and the mixed version status:

```
Name           OpState  RSCTActiveVersionMixedVersions  TSPort  GSPort
SA_Domain      Online   2.4.9.2                          Yes      12347   12348
```

- 3 Issue the **lsrpnode** command to display the version of RSCT that is installed on each node. All nodes must be online:

```
Name  OpState  RSCTVersion
node1 Online   2.4.11.5
node2 Online   2.4.11.5
```

Note: It can take a couple minutes before the RSCTVersion is updated to the newly installed version and before the domain reports MixedVersions is Yes. Wait until this occurs before proceeding.

At this point, each node has been upgraded to the same code level, the same IVN. Now you can complete the migration by following the next five steps.

Step 1: Ensure all nodes in the domain are online.

Step 2: List the domain information and note that the MixedVersions show Yes.

Step 3: List the node information and check that all show the same RSCT level and the new RSCT code level.

The main point of steps 2 and 3 is to compare the RSCT level shown in the node listing with the "RSCTActiveVersion" from the domain information. They should differ. If not, wait a couple minutes and re-issue each of the **lsrpdomain** and **lsrpnode** commands until you see the difference. Do not proceed until you do.

Completing the migration (2 of 2)

4 If the RSCT peer domain is running in mixed version mode (MixedVersions = Yes) and all nodes show the new RSCT level (all nodes have been upgraded to the new release of TSAMP), you must update the active RSCT version by running the RSCT CompleteMigration action on one of the nodes. To update the RSCTActiveVersion, make sure that all nodes are online, and issue the following command on one of the nodes :

```
runact -c IBM.PeerDomain CompleteMigration Options=0
```

To verify that the active RSCT version has been updated, issue the lsrpdomain command again:

```
Name      OpState  RSCTActiveVersion  MixedVersions  TSPort  GSPort
SA_Domain  Online   2.4.11.5           No             12347   12348
```

5 To finish the migration, activate the newly installed TSAMP code level by issuing the following command on one of the nodes :

```
samctrl -m
```

The code version of the Active Version (AVN) and the Installed Version (IVN) of TSAMP should now be the same for all nodes :

```
lssrc -ls IBM.RecoveryRM
```

Until this is true, the new code features have not been activated and cannot be used

The last two steps complete the migration by activating the new code levels for both RSCT and TSAMP components.

Step 4: Use the runact command to activate the new RSCT level. List the domain information using 'lsrpdomain' to confirm that the newer code level is now the "RSCTActiveVersion" and that "MixedVersions" changed to No.

Step 5: Use the 'samctrl -m' (lower case "m") command to activate the newer code level for TSAMP. As described on slide 8, the IVN (installed version) and AVN (active version) can be compared using the 'lssrc -ls IBM.RecoveryRM' command. The activation of the new code level of TSAMP is complete when the AVN matches the IVN.

When both RSCT and TSAMP activation steps have been performed, the migration is considered complete and successful.



Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_node-by-node-upgrade.ppt

This module is also available in PDF format at: [../node-by-node-upgrade.pdf](http://node-by-node-upgrade.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, AIX, DB2, and Tivoli are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2010. All rights reserved.