IBM WebSphere Transformation Extender

Cluster Support (Sun Solaris)

Version 8.1

October 2006

This edition of this document applies to IBM WebSphere Transformation Extender Version 8.1; and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, email DTX_doc_feedback@us.ibm.com. We look forward to hearing from you.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Chapter 1. About clusters

A cluster is a group of independent network server nodes that operate and appear to clients as if they were a single, highly available system. A cluster typically consists of two nodes. Only one node can be active at a time even though both nodes have the clustering service running.

By enabling one server to take over for another, a clustered configuration is designed to do the following:

- Provide live backup in case a node fails, enhancing stability and minimizing downtime.
- Improve data integrity.

Chapter 2. Cluster support for the Launcher

Sun Cluster support for the Launcher enables the Launcher to provide continuous service. The Sun Cluster 3 framework enables system resources such as devices, file systems, and networks to be shared in cluster failover from one node to another, which in turn reduces or eliminates potential downtime. Provided that a system (**.msl**) is designed properly, transactions and data will not be lost upon the restarting of the Launcher on a new node.

For the Launcher, Sun Cluster 3 delivers high availability through a combination of hardware and software components. A cluster typically consists of a minimum of two nodes, and can have as many as eight. The state of each individual node, hardware, and software component is constantly monitored.

Failover typically completes in a matter of seconds, depending on other applications in the cluster. Once a failure occurs, the primary node switches over to a secondary node.

File triggers

WebSphere Transformation Extender clustering support handles file triggers directly to ensure triggers do not repeat or are not missed when the Launcher is started on the new node. Other source triggers need to be handled by corresponding application cluster or failover support or through a global transaction manager to ensure that those events do not trigger again. For example, with source triggers, when the output is committed and the inputs are removed, there is a window in between where redundancy can occur. This will not happen with file triggers. All in all, the reliability is still better than without cluster support, but he potential of duplicate source triggers is still possible.

Time triggers

Recurring time events that are missed during the failover period are lost; the next occurrence of the trigger is processed without regard to what happened before a failover. (The failover restart has the same functionality as manually stopping and restarting the service). For example, if you have a once-a-week time event that occurs every Tuesday at midnight and you fail over from 11:59 PM Tuesday and start back up at 12:04 AM on Wednesday, that time event will not trigger until the following Tuesday.

Once-only time events that are missed during the failover period will take place after the Launcher service restarts. For example, if a missed time trigger is designed to occur immediately after startup, the trigger will take place after the Launcher service restarts.

Chapter 3. System requirements

IBMscdtx is the WebSphere Transformation Extender package for Sun Cluster support. Before installing, verify the following:

- Sun Cluster 3 software is installed.
- JAVA is included in your $PATH variable.
- The package information utility is used to determine if the packages are already installed.
- You have the appropriate user-access level. You must have root privileges to install the package.

The package is installed to the **/opt/IBMscdtx** directory. Use this package in conjunction with a production Launcher.

About package installation

A package is an application consisting of one or more programs that will perform a particular function. There are two categories of package components: package objects (the application files to be installed) and control files (the files controlling how, where, and if the package is installed). WebSphere Transformation Extender Sun Cluster Support (**IBMscdtx**) is a standard type of package. It can be installed and removed using standard Sun packaging tools (**pkgadd**, **pkgrm**). Use the pkginfo command to view package information.

Chapter 4. User environment

The user environment required to run the Launcher must be defined in the setup script because the environment setup in the user profile file will not be used.

For example, if the **JAVAHOME** directory must be defined, add the following line to the setup script:

```
export JAVAHOME=/usr/java1.2
```

About the Sun Cluster environment

In the Sun Cluster environment, only one node is active at a time. When a failure occurs, Sun's Resource Group Manager (RGM) calls the STOP callback method, which relays a shutdown message to the Launcher.

Callback methods are programs associated with the resource, such as START, STOP, MONITOR_START, MONITOR_STOP, and others.

The shared drive in which maps, data, and systems are stored can be accessed from all nodes at the same time.

Chapter 5. Sun Cluster 3 package (IBMscdtx) installation

Prior to installation, verify that all system requirements have been met.

You must install WebSphere Transformation Extender on each cluster node. Use the standard Sun maintenance command for software installation:

```
pkgadd
```

To install the IBMscdtx package:

1. From the command prompt, enter the following command: # pkgadd -d *<directory containing the package>* IBMscdtx
2. When you are prompted to continue with the installation of **IBMscdtx**, enter **y**.

   The following message appears:

   ```
   Installing WebSphere Transformation Extender
   resource type for Sun Cluster as IBMscdtx.
   ```

   As the pre-install scripts begin to execute, more messages are displayed.

Multiple installations

If you install more than one instance of WebSphere Transformation Extender, total failover support is available for all Launchers that are online. You can also individually take Launchers online or offline.

For multiple installations, you must install one package only and then complete multiple registrations. An individual Launcher has a resource type with its own group and resource. This mechanism allows control for each instance of the Launcher to be separate, yet for a cluster as a whole.

For multiple installations, there will be multiple entries for each resource type, resource, and group for each instance. Each one is uniquely named based on input from the setup program.

Setup example

The following is a setup scenario for two instances of WebSphere Transformation Extender. To register multiple instances, the cluster administrator must know the following information:

- DTX_HOME_DIR paths
- The logical host name to use for each instance
- Unique instance IDs

   Note: Unique instance IDs are user-defined, arbitrary values.

Install_dir is the path in which WebSphere Transformation Extender is installed. The information in the following table pertains to this example.

	Instance 1	Instance 2
Installation directory	*install_dir*/system1	*install_dir*/system2

	Instance 1	Instance 2
DTX_HOME_DIR=	*install_dir*/system1	*install_dir*/system2
Logical host name	clusterhost1 (joined to instance 1)	clusterhost2 (joined to instance 2)
Unique instance ID	EDISYS	CM

Based upon the information in the preceding table, use the following commands to register each instance, substituting your actual installation path for the *install_dir* value below.

Instance 1

```
register_scdtx -h clusterhost1 -m /install_dir/system1 -i
EDISYS
```

Instance 2

```
register_scdtx -h clusterhost2 -m /install_dir/system2 -i CM
```

As a result, the following groups and resources are created:

	Instance 1	Instance 2
Group name	TransformationExtender-rgEDISYS	TransformationExtender-rgCM
Resources	clusterhost1 (type=LogicalHostName)	clusterhost2 (type=LogicalHostName)
	Launcher-rsEDISYS (type=IBM.scdtx)	Launcher-rsCM (type=IBM.scdtx)

Use the following commands to unregister each instance of WebSphere Transformation Extender:

* **Instance 1**

  ```
  unregister_scdtx -h clusterhost1 -i EDISYS
  ```
* **Instance 2**

  ```
  unregister_scdtx -h clusterhost2 -i CM
  ```

Chapter 6. Removing cluster support

The procedures in this section undo the configuration steps as described in
"Installation commands" . To remove theWebSphere Transformation Extender Sun
Cluster Support (**IBMscdtx**) package, use the following command:

```
pkgrm IBMscdtx
```

The removal procedure is described in this section. Before removing the package
from your system, verify that the **IBM.scdtx** resource type is not registered as a
Sun Cluster resource type.

To verify that **IBM.scdtx** is not registered as a Sun Cluster resource type, use the
`scrgadm` Cluster administrator command.

To verify the resource type:
* At a command prompt, enter the following:

  ```
  # scrgadm -p -t IBM.scdtx
  ```
  The following text is output:
  ```
  Res Type name: IBM.scdtx
  ```
  ```
  Res Type description: WebSphere Transformation Extender for Sun Cluster
  ```

To remove the IBMscdtx package from the system:
1. At the command prompt, enter the following command: # pkgrm IBMscdtx.

   Note: This command does not remove the IBMscdtx package if the IBM.scdtx
   type is registered on the cluster.

   The following message appears:
   ```
   The following package is currently installed:
   IBMscdtx IBM Transformation Extender Cluster Support
   ```
2. When prompted to remove the package, enter y.
3. When prompted to continue with the removal of the package, enter y.

   A series of processing messages are displayed, followed by a final confirmation
   message:
   ```
   Removal of <IBMscdtx> was successful.
   ```

 11

Chapter 7. Configuring the Launcher for Sun Cluster 3 support

After WebSphere Transformation Extender is installed on each cluster node or on a common node, you can begin configuration. Configuration for Sun Cluster 3 involves the following high-level steps:

- Installing Launcher maps and other data on a common node drive, which is the common drive that the one active node can see at any time. WebSphere Transformation Extender binaries and systems can be either local or on the node-common drive.
- Enabling the automatic startup option in the Launcher Administration interface. Upon failover, this option allows the Launcher to restart on a new node.
- Enabling the journal function to provide transactional file support
- Pre-registry testing
- Registering the Launcher with Sun Cluster 3 (using the **register_scdtx** script).

Common node drive

In a cluster environment, you should place maps and data that the Launcher uses on one of the shared cluster drives so that the data remains accessible to the active node.

Launcher administration

Before starting the Launcher service, you must ensure that the Launcher is set to start automatically. You can complete this task from the Launcher Administration interface. The procedure is listed below.

To enable the automatic startup option

1. Open Launcher Administration.

   The **Launcher Administration** - **General** tab is displayed.
2. Enable the **Automatic startup** check box.
3. Click **OK** to save changes and close the application.

File triggers and recovery

For catastrophic failover recovery and to prevent lost or duplicated transactions, specify a directory (available to active nodes of a cluster) in which the journaling of file-triggered inputs will occur. This file is specified in the **dtx.ini** file under the **M4File** heading and must be duplicated for all nodes upon which the Launcher is installed.

Dtx.ini is located in the product installation directory.

Setting up the journaling directory

You can specify the journaling directory in the **dtx.ini** file under the **M4File** heading. This directory belongs to the global file space (a shared cluster drive). For example, if the global directory is a mount point for a shared drive, the entry will look similar to the following:

```
[M4File]
Journaling=/global/install_dir/journaldir
```

where *install_dir* is the directory in which WebSphere Transformation Extender is installed.

Be sure to confirm that the directory exists.

About journal files

The journal files are created by means of the **M4File** during map execution. The lifespan of a journal file is from map initialization through map termination.

The journal file contains information about the state of the input files triggered, the map, and any output files. During the processing of maps that are triggered by files, a journal file is created to track the state of the mapping process. Journal files are created only if the "delete on success" option is set for any input file trigger and there are output files. The "delete on success" setting refers to the **OnSuccess** map setting. All journal files are created in the user-specified directory, and have the following name syntax:

```
<instance>.log
```

The journal file is removed when the mapping process has completed.

For information about the **OnSuccess** setting, see the Map Designer documentation.

Failover recovery

Upon Launcher startup, if journal files are present in the user-defined directory (as specified in **dtx.ini**), each file is opened and the input and output status is checked before proceeding to complete the action or rollback.

During the recovery process, a unique log file is created, logging the steps for each journal file. The journal file is not readable by a user; therefore, a log file is generated. Only one log file is created per startup in the journaling directory with a **.log** filename extension.

```
<time_t-value>.log
```

Time_t-value is the number of seconds elapsed since midnight January 1, 1970.

The log file contains information about outputs and inputs and any actions taken to correct the previous failure. It also includes the number of journals processed.

Any remaining journal files that exist upon the startup of the Launcher are processed according to the state of the mapping described in the journal file.

Sample log file

The log file is created in the journaling directory and describes the processing that occurs to prevent lost or duplicated transactions. After a node fails, the journaling results will be written to a log file containing information similar to the following:

```
Initializing M4File on Mon Dec 10 11:24:39 2003
Searching for previous journal files
Processing c:\your_journal\M4F150.txt:
 output was written: c:\out\wally.10104
 input not found: c:\stage\wal850.10104
 removed journal file
Processing c:\your_journal\M4F172.txt:
 output was written: c:\out\wally.10120
 input found with same size and time stamp -- removing: c:\stage\wal850.10120.
 removed journal file
Processing c:\your_journal\M4F211.txt:
 output c:\out\wally.10150 not written.  removing tmp file: c:\out\Merd3.tmp
 removed journal file
Processing c:\your_journal\M4F213.txt:
 0 byte journal file...
 removed journal file
Processed 4 journal files
```

Testing

Before registering the Launcher to the cluster, perform a single node test on the Launcher. If the test is successful (which means that it appears as a single-node box without clustering), you can register the Launcher to be under the cluster's control.

Registry

You must register the Launcher with Sun Cluster 3. The registration process should be completed one time only.

Before WebSphere Transformation Extender can be registered to the cluster, it needs to be installed on the cluster. There are two different scenarios for WebSphere Transformation Extender installation:

- One instance of WebSphere Transformation Extender is installed on the common node.
- Multiple instances of the same WebSphere Transformation Extender application are installed on each cluster node into the local file space.

Only resource groups can be registered for each unique host name, for which the **_CM** identifier is used. One instance of WebSphere Transformation Extender can be installed once on common cluster space or can be installed on the private space of each node. For the latter case, the installation directory must be the same (DTX_HOME_DIR).

If a file named **user_startup.sh** is present in the installation directory and is executable, the pending cluster node runs the file as the user startup script. This script must complete with a successful return value (0) for the Launcher to be started.

For multiple installations, DTX_HOME_DIR must not be the same.

Registering WebSphere Transformation Extender to the Sun Cluster environment

The procedures in this chapter are based upon the following configuration criteria:

- The cluster has two nodes: node1 and node2.
- Two instances of WebSphere Transformation Extender are installed into the local file space (**DTX_HOME_DIR**).
- The WebSphere Transformation Extender **IBMscdtx** package is installed into the default directory: **/opt/**.
- The requested logical host name is **scIBM**.

To register the first instance of WebSphere Transformation Extender to the Sun Cluster environment:

At the command prompt of the installation directory, enter the following command:

```
/opt/IBMscdtx/util/register_scdtx -h scIBM -m
/install_dir -i _CM
```

This registers WebSphere Transformation Extender resources to the Sun Cluster. The following output is displayed:

```
Registering resource type <IBM.scdtx>...done.
Creating failover resource group <TransformationExtender-rg_CM>...done.
Creating logical host resource <scIBM>...done.
Creating resource <Launcher-rs_CM> for the resource type <IBM.scdtx>...done
```

After the registration process is complete, use the Sun Cluster Manager Monitor to manage resources. The Sun Cluster Manager utility is a web-based tool that allows you to control cluster resources and perform other administrative jobs.

Sun Cluster Manager Monitor is included in the Sun Solaris SUNWscvw package.

The following table is the resources layout as shown in the Sun Cluster Manager Monitor utility:

Resource Group / Resource	Type	Node1	Node2
TransformationExtender-rg_CM		Offline	Offline
scdtx	LogicalHostname	Offline	Offline
Launcher-rs_CM	IBM.scdtx	Offline	Offline

After the registration process is complete, new resources are available on the Sun Cluster. The resources are in an off-line status by default. An authorized user can use the Sun Cluster Manager to activate resources (that is, start Launcher services) or use the command line administrator commands.

For more information about the Sun Cluster Manager utility, refer to your Sun Cluster Manager documentation.

Chapter 8. Working with Sun Cluster 3

This section discusses startup scripts, removing and changing the state of resources, and system message logging.

Startup scripts

You can create your own pre-named start up scripts that will execute during start method, but before Launcher services. To do so, use the following script, which must be located in DTX_HOME_DIR:

```
user_startup.sh
```

Because there is no user interface, the script must return non-error status (0) for the start method to continue. If this script returns error status, the start method aborts and the Launcher service is not started. This script is not mandatory; therefore, if it does not exist, no errors will occur.

The timeout for this script is 50 percent of the START_TIMEOUT value. The default value for START_TIMEOUT is 200 seconds. A cluster administrator can change the timeout value using Sun Cluster 3 tools. If the timeout period expires, the **user_startup.sh** script is aborted and the startup method fails. If the timeout value expires, the following message is saved to the cluster error log file:

```
Nov 20 11:40:38 sun9 [IBM.scdtx,TransformationExtender-rg,Launcher-rs]:
[ID 702911 daemon.error]
USER STARTUP SCRIPT - timeouted (TIMEOUT=100)
```

If the **user_startup.sh** script returns an error status, the following message appears:

```
Nov 20 11:13:36 sun9 [IBM.scdtx,TransformationExtender-rg,Launcher-rs]:
[ID 702911 daemon.error]
USER STARTUP SCRIPT ERROR=1
```

The Sun Cluster Manager will report the following:

```
Resource / Resource Group Status
Running /usr/cluster/bin/scswitch -Z -g TransformationExtender-rg
Failure
scswitch: On node cluster-node1, resource group TransformationExtender-rg
 is online but resource Launcher-rs failed to start
scswitch:  Unable to bring resource group "TransformationExtender-rg" online.
```

Command:

```
scswitch -Z -g TransformationExtender-rg
```

Returns:

```
scswitch: On node cluster-node1, resource group TransformationExtender-rg
 is online but resource Launcher-rs failed to start
scswitch:  Unable to bring resource group "TransformationExtender-rg" online.
```

User script failure

If the **user_startup.sh** script fails, the standard error log (**/var/adm/messages**) contains information about the **user_startup.sh** error exit value. A message appears that is similar to the following example:

If **user_startup.sh** returns 44:

```
Jan 18 17:09:02 sun9 [IBM.scdtx,TransformationExtender-rgVP,Launcher-rsVP]:
[ID 702911 daemon.error]
USER STARTUP SCRIPT ERROR=44
```

The **Launcher** error exit value is saved in the system log file (which reports an exit code of 1).

The user script return value is presented in the logged error message. The value of "99" is reserved for a startup timeout condition and should be avoided by the user script.

Removing resources from the cluster

To remove resources from Sun Cluster:

From a command prompt, run the following command once on any available node:

```
/opt/IBMscdtx/utils/unregister_scdtx -h scIBM -i _CM
```

The following output text is displayed:

```
Disabling the resource <Launcher-rs_CM> ...done.
Removing the resource <Launcher-rs_CM> ...done.
Disabling the network resource <scIBM> ...done.
Removing the resource <scIBM> ...done.
Removing the resource group MTransformationExtender-rg_CM> ...done.
Removing the resource type <IBM.scdtx> ...done.
```

This removes previously registered resource group from all cluster nodes.

Changing the state of resources

After the resources were registered using utility scripts, the cluster administrator or other authorized user can perform various operations with the WebSphere Transformation Extender resource group (TransformationExtender-rg[ES_ID]) and with WebSphere Transformation Extender resources (Launcher-rs[ES_ID]). Various operations include setting online or offline, setting to manage or unmanage states, selecting primary nodes, or deleting all resources.

An WebSphere Transformation Extender resource is a process or service that is being managed by the cluster RGM facility. The resources for Sun Cluster support for the Launcher are Launcher and Logical Host Name.

The WebSphere Transformation Extender resource group is the name of the resource group in which the resource has been configured. Each instance of WebSphere Transformation Extender creates one group. When failover occurs, all resources from the same group are moved together to a secondary node.

There are two ways to manipulate resources:
- Use the web-based Sun Cluster Manager utility.
- Use Sun Cluster administrative commands.

The following table contains important actions and appropriate commands:

What do you want to do?
Command(s)

Put resource group on-line.
scswitch -Z -g TransformationExtender-rg

Put resource group off-line.
scswitch -F -g TransformationExtender-rg

Put Launcher and logical host resources to ″disable state″ and resource group to ″unmanaged state″.
scswitch -n -j Launcher-rs

scswitch -n -j scIBM

scswitch -F -g TransformationExtender-rg

scswitch -u -g TransformationExtender-rg

Enable all resources and put resource group to ″managed state″.
scswitch -e -j scIBM

scswitch -e -j Launcher-rs

scswitch -o -g TransformationExtender-rg

Change primary node of the running resource to secondary node (clusternode2).
scswitch -z -g TransformationExtender-rg -h clusternode2

System message logging

The Sun Cluster resource group manager (RGM) uses standard UNIX logging functionality (**syslogd**). Callback methods (such as start/stop/monitor/probe/validate/update) also use standard UNIX system logging.

The number of the **syslog** facility is defined at the cluster level, with a default value of **DAEMON**.

To get the value of the syslog facility on the Sun Cluster:

At a command prompt, enter:
```
SYSLOG_FACILITY=`scha_cluster_get -O SYSLOG_FACILITY`
```

Each message has the following tag as an identifier:
```
[ResourceTypeName,ResourceGroupName[ES_ID], ResourceName[ES_ID]
```

Chapter 9. Troubleshooting

You can use the examples of utility script and registry error messages for
troubleshooting purposes.

Utility script error messages

To view a detailed output message, utility scripts can be run using the -v option.
The text of error or warning message explains the reason for error. Common error
messages related to the **register_scdtx** script are listed in this section.

Register_scdtx called without parameters

When the **register_scdtx** utility is called without parameters, the following usage
message is reported:

```
# /opt/IBMscdtx/util/register_scdtx
The logical hostname must be specified.
Usage: register_scdtx -h <logical host name>
        -m <Transformation Extender home dir>
        [ -i <Launcher identificator>]
        [ -d <debug level 0-nodebug 1-debug 2-full debug>]
        [ -n <nafogroup/adapter list>]
        [-v]
```

Invalid value for the DTX_HOME_DIR variable

The specified installation directory is either invalid or does not exist.

```
# /opt/IBMscdtx/util/register_scdtx -m BAD/IBM/HOMEDIR -h scIBM
DTX_HOME_DIR=BAD/IBM/HOMEDIR is not a directory.
```

Java class file not found

The Java Launcher jar file (**launcher.jar**) cannot be found in the DTX_HOME_DIR.
This might be the result of a previously broken or incomplete installation or
JAVAHOME has not been defined.

```
# /opt/IBMscdtx/util/register_scdtx -m /global/home/user_name/ascl60 -h scIBM
Exception in thread "main" java.lang.NoClassDefFoundError:
com/ascential/eventserver/core/JNILauncherAdminInfo
Error getting Java services port number - error=0
```

Logical host name cannot be created

The given logical host name (using the -h option) cannot be resolved to the IP
address because it is either in use or it was not registered to the DNS table or local
host table.

```
# /opt/IBMscdtx/util/register_scdtx -m /global/home/
user_name/ascl60 -h scmctr11
Creating a failover instance...
No nafo groups or network adapters specified ...
will try to auto-discover the network adapters and
configure them into nafo groups.
Registering resource type <IBM.scdtx>...done.
Creating failover resource group <TransformationExtender-rg>...done.
```

```
Creating logical host resource <scmctr11>...scmctr11 cannot be mapped
to an IP address.
FAILED: scrgadm -a -L -g TransformationExtender-rg -l scmctr11
```

Unregister_scdtx script error messages

The common errors related to **unregister_scdtx** scripts are discussed.

Resource group does not exist

The resource group with unique ID = _CM does not exist because it was not
created or was previously deleted.

```
# ./unregister_scdtx -h scIBM -i _CM
Resource group name <TransformationExtender-rg_CM> does not exist
Removing the resource type <IBM.scdtx> ...done.
```

Logical host name resource does not exist

The given logical host name resource does not exist.

```
./unregister_scdtx -h scIBM
Disabling the network resource <scIBM> ...scswitch:
scIBM: Invalid resource
FAILED: scswitch -n -j scIBM
```

Run without parameters

In the following example, provided that **<TransformationExtender_rg>** exists as a
resource group, only the resource group will be deleted because other resources do
not exist.

```
# ./unregister_scdtx
Resource name <Launcher-rs> does not exist
Removing the resource group <TransformationExtender-rg> ...done.
Removing the resource type <IBM.scdtx> ...Resource type IBM.scdtx is not registered
FAILED: scrgadm -r -t IBM.scdtx
```

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
577 Airport Blvd., Suite 800
Burlingame, CA 94010
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Warning: Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks and service marks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries, or both:

AIX
AIX 5L
AS/400
Ascential
Ascential DataStage
Ascential Enterprise Integration Suite
Ascential QualityStage
Ascential RTI
Ascential Software
Ascential
CICS
DataStage
DB2
DB2 Universal Database
developerWorks
Footprint
Hiperspace
IBM
the IBM logo
ibm.com
IMS
Informix
Lotus
Lotus Notes
MQSeries
MVS
OS/390
OS/400
Passport Advantage
Redbooks
RISC System/6000
Roma
S/390
System z
Trading Partner
Tivoli

WebSphere
z/Architecture
z/OS
zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

This product includes software developed by the Eclipse Project (http://www.eclipse.org/).

IBM WebSphere Transformation Extender, Version 8.1

Index

Special characters

.log   14

C

cluster support
   removing   11
clusters   1
commands
   for installation   9
   for resource type verification   11
commands, Sun Cluster administrative   18

D

DTX_HOME_DIR   9, 17, 21
dtx.ini   13
   M4File heading   13

E

error log files   17
error messages
   warnings   21
examples
   appending a setup script   7
   resource group   22
   setup for multiple installations   9
exit codes   18

F

filename extensions
   .log   14

I

installations, multiple   9
   setting up example   9

J

jar files
   launcher.jar   21
Java class file   21
journaling directory   14

L

Launcher
   Administration interface   13
   multiple installations of   9
   removing cluster support   11
launcher.jar   21
log files   14
   created during recovery process   14
   example   15

log files *(continued)*
   exit codes in   18
   time_t-value.log   14
logging   19
logical hostname resource   22

M

M4File   14
   heading of   13
map settings
   OnSuccess   14

N

node-common drives   13

O

OnSuccess map setting   14

P

packages   5
   viewing information about   5
procedures
   enabling the automatic startup option   13
   installing IBMscdtx   9
   verifying resource type   11

R

recovery process   14
resource groups   22

S

scripts
   error messages generated by   21
   for startup   17
   register_scdtx   21
   return values   18
   unregister_scdtx   22
   user_shartup.sh
      failures   17
   user_startup.sh   17
shared drives See node-common drives   13
START_TIMEOUT value   17
startup option, automatic   13
Sun Cluster 3
   administrative commands   18
   support for the Launcher   3
Sun Cluster Manager utility   16, 18
SUNWscvw package   16
syslogd   19
system message logging   19

T

testing
   single node   15
time_t-value.log file   14
triggers
   file   13
troubleshooting
   resource groups   22

U

unique ID   22
UNIX
   system logging   19
unregister_scdtx script   22
user_startup.sh script   17

IBM ®

Printed in USA