

Lotus

ClusterProven™

Domino for AS/400

August 2001

ClusterProven Domino for AS/400

Introduction

ClusterProven™ Domino for AS/400 is an iSeries specific enhancement for Lotus Domino that provides OS/400 Version 5 Release 1 with a cluster management (CM) utility to manage iSeries based Domino servers. A Domino server's definition is duplicated on multiple iSeries systems (known as nodes) in an OS/400 cluster, managed by an OS/400 CM utility. The same Domino server is automatically defined on all the nodes in the cluster that can access a switchable independent auxiliary storage pool (IASP) that stores the server's data directory.

This is not the same as Domino clustering support but can be used in conjunction with Domino clustering support. Domino clustering is configured on a database level. Each database can be clustered so that if the primary server is not responding, access to that database is rerouted to a backup server where a copy of that database is automatically maintained. Clustered databases must be replicated to the backup servers as a scheduled event. There is some overhead to do this replication and copies of the database have to be stored on the backup servers.

With OS/400 ClusterProven cluster support the clustering is done on a server's entire data directory and not at a Domino database level. Instead of the databases being replicated, the definition of the Domino server itself is duplicated on other cluster nodes. The cluster supported server's data directory is kept on a switchable disk device. Once configured the same Domino server can be started on any iSeries logical partition (LPAR) or server that can potentially access the server's data directory on that switchable device. An OS/400 cluster management GUI interface available through Management Central can be used to start, stop, and switch between nodes configured for the clustered Domino server. OS/400 APIs and commands can also be used to perform cluster management functions. Cluster middleware products from IBM business partners offer additional resources for managing Domino servers in an iSeries cluster.


Using cluster management to support Domino provides automatic failover to backup iSeries systems or LPARs in the case of a system failure. It can also be used to effectively switch a Domino server from one system to another as a means to provide continuous support for the server on a backup system in the case of system downtime for scheduled maintenance or a system IPL.

Requirements


In order for the same Domino server to be supported on multiple systems it must have its data directory accessible to each system when it is started on that system. The server's data directory has to be stored on a disk storage device that can be switched between all the configured cluster nodes. Only OS/400 V5R1 LPAR systems and V5R1 iSeries systems connected to a switchable disk device through HSL Opticonnect can be configured to access a switchable storage device. V5R1 OS/400 cluster based management utility is required to configure and manage switchable disk storage devices. Domino 5.0.7 is required on every iSeries system in the cluster that will support a replica of a Domino server configured by OS/400 cluster management.

Switchable IASPs and DASD devices


An independent auxiliary storage pool (IASP) is configured on all of the iSeries systems in the cluster that need to access the disk storage represented by that IASP. The IASP represents "independent" disk storage that can be "switched" from system to system. The collection of storage is configured to these systems as a named device. When the device is varied on to a system, the storage for that device is accessible to that system. The device is configured to be switched between specific systems using a device domain. Only one system at a time can vary on the device and get at the disk storage represented by its configured IASP.



A system has its storage configured into ASPs (Auxiliary Storage Pools) and at least one IASP (Independent Auxiliary Storage Pool)




The IASP storage is represented on the system as a device and can be managed "varied off" as a device. The storage is "offline" or not accessible when the device is varied off.



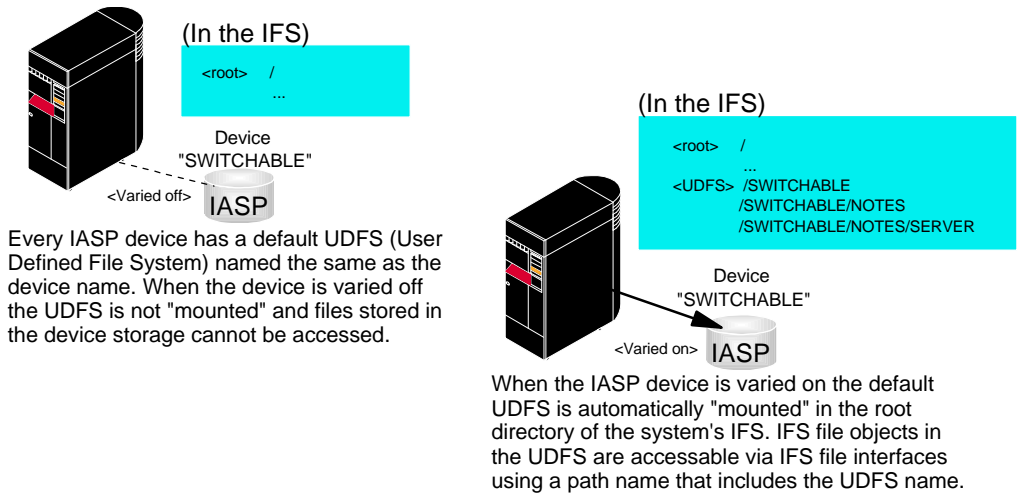
The same IASP can be configured to be switchable between systems. The IASP storage represented by a device is accessible to the system where the device is varied on. A system can define or change objects stored when it has access. These objects can then be accessed by another system when the device is switched and varied on to that system.

A cluster management utility can define the device as part of a device type (*DEV) cluster resource group (CRG) as the means of externally controlling which system currently has access to the storage. Creating a device CRG for the device allows Cluster Management functions to manage the device and switch which cluster node (system) has access to the IASP storage. CM can manage the IASP from any node in the cluster using OS/400 Operations Navigator's Management Central function, system APIs, CL commands, or cluster middleware products. In addition, Cluster Management automatically provides support through device CRGs so that access to the IASP is switched between cluster nodes automatically in the event of a cluster node failure (known as "failover"). Device CRG support is new for OS/400 V5R1.




Using OS/400 Cluster Management, a cluster can be defined for the systems (nodes) that can access the same IASP. A device Cluster Resource Group (CRG) can be configured on the cluster nodes to manage the switchable storage. Cluster Management interfaces can be used to manage CRGs and, through them, control which cluster node can currently access the IASP storage.

Starting with OS/400 V5R1, a default user defined file system (UDFS) is configured for an IASP so that when the IASP is varied on it is automatically "mounted" as a file system to the OS. Once mounted, the UDFS appears as a directory with the name of the UDFS in the root directory of IFS file system on that system. The IFS file objects in this mounted UDFS directory are accessible through the standard IFS interfaces as long as the IASP device is varied on.



The Domino server's data directory is stored on the switchable IASP and so it can be accessed when that device CRG is switched to that cluster node. The data directory is kept in the default UDFS directory which is automatically mounted when the device is varied on to that system. This allows the Domino server to access its data directory on any node in its recovery domain using the same common directory "path" that includes the name of the UDFS in the root of that path. The UDFS is accessible only when the device is varied on and then only when the device CRG that controls the device is active and "primary" on a node within the cluster. When the device CRG is switched to another node (or if the device is just varied off) the UDFS is "unmounted" and the objects in that file system are no longer accessible through IFS.



When a Domino server's data directory is defined to be stored within an IASP's storage (server's directory path includes the default UDFS directory name) then the server's data directory can be accessed from any node where the IASP device is varied on. Control of the device CRG then also controls access to all IASP storage. That means Cluster Management can control which node can access the Domino server's data directory through the device CRG.

There are restrictions on the configuration for switchable IASP devices (and the disk storage they control). OS/400 LPARs can have an IASP defined so that it is accessible from and can be switched between partitions. With a set of partitions in a LPAR environment, the smallest switchable unit is a DASD IOP based set of files.

The other means to create switchable disk storage so that it can be accessed by different systems is with a switchable disk storage tower. These storage tower devices are connected with HSL Opticonnect to pairs of iSeries systems and can be switched between these systems. All the systems that switched IASP devices must at least have OS/400 V5R1 installed. The switchable IASP device must be configured to CM as a device CRG to all the nodes within a cluster where it can potentially be switched. There is no "granular" level of switching the disk storage except at a device CRG level. Switching the CRG between systems means all of the disk storage represented by the IASP is switched. The storage switched could potentially be used by multiple applications on the systems and all the applications would be affected. Between systems connected via HSL cables to a common switchable disk storage tower via HSL cables, the smallest switchable unit is the I/O tower (i.e. model number 5074, 5075 or "half:" of a 5079 tower).


Note When managing the IASP device by using a device CRG, the device does not vary on automatically when the device CRG is started. There is no error message that indicates anything has gone wrong with starting the device CRG. Checking the device CRG (explained below) indicates that it has been started. This creates a

situation where the device CRG is started but its device is not varied on and applications (including Domino servers) cannot access the IFS files stored on the device as the device's default UDFS has not been mounted. Varying the device on for the system where the device CRG is "primary" (explained below) will usually resolve this situation. When a device CRG is switched or fails over to another cluster node, the device associated with the CRG should be configured to automatically be varied on and the default UDFS mounted to the IFS.

Domino for AS/400 Servers Managed as Application Cluster Resource Groups

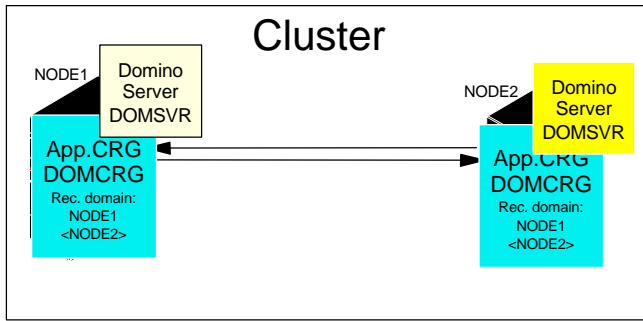
The ability to have a Domino for AS/400 server's data directory accessible from multiple systems means that the server's configuration and every Domino database defined by that server is accessible on these same systems when the IASP device is varied on. Since the device can only be accessed by one system at a time, the Domino server's directory is only accessible to that one system. The same Domino server's definition can be duplicated to LPAR partitions or to systems that can vary on and can be started on the system that currently can access the server's directory. Any number of iSeries systems could potentially be defined to any cluster. The cluster systems (nodes) that can access the same IASP are all candidates for supporting a common Domino server whose definition has been duplicated in this way.

OS/400 Cluster Management can also manage ClusterProven applications as well as device CRGs. The application itself is defined to CM as an application cluster resource group (CRG). Once it is defined to OS/400 Cluster Management as part of an application CRG, the application can be started or stopped using CM interfaces and switched to start on another cluster node (system) where it is also configured. An application CRG can be managed using the same OS/400 Operations Navigator Management Central GUI, CM APIs, CL commands, or cluster middleware products that are used to manage device CRGs.




Like devices, ClusterProven (TM) applications can also have Cluster Resource Groups defined for them so that the applications can be managed as resources within a cluster. Instead of being varied on and off like a device associated with a device CRG, these applications can be started and ended using Cluster Management GUI, APIs or commands using an *application CRG*.

All CRGs when defined to CM have a recovery domain. The recovery domain is a list of systems (nodes) within the cluster where the CRG is defined. The CRG's recovery domain also indicates which node is the current primary node, which nodes in the cluster are backup nodes, and the order that the backup nodes should be tried.



A CRG exists on all the cluster nodes. Every CRG has a recovery domain that is part of its configuration. This shows an example application CRG, DOMCRG, with a recovery domain that has cluster nodes NODE1 and NODE2 in it. The "<NODE2>" notation signifies that at the moment, NODE2 is the "primary" cluster node for the CRG.

A Domino CRG also has a recovery domain. The nodes in the recovery domain represent the systems where that server is configured and can be started. There is an additional requirement for a Domino application CRG as a Domino server requires access to its data directory from the system where it is to be started. The recovery domain of the server's application CRG has to "parallel" the recovery domain of the device CRG that controls access to its data directory on the node in the cluster. (See Creating a Domino application CRG and the documentation on the CFGDOMCRG command below for more information on the recovery domain and creating Domino application CRGs.)



This shows an example application CRG, DOMCRG, and a device CRG, DEVCRG. Both have the same recovery domain and, at the moment, the CRGs are active and NODE2 is the "primary" node for both these CRGs. If the Domino server DOMSVR was to be started on NODE2, the server has access to its data directory on the IASP configured by DEVCRG.

The Basics of How CRGs Work

The nodes of the cluster have to be started with a CM function. When all the nodes defined to the cluster have been started the CRGs defined to that cluster can be started. The CRG is considered "active" when it has been activated on the cluster nodes in its recovery domain. Every CRG has a specific current primary node. Cluster management will always start a CRG on all the active nodes in its recovery domain. The CRG will be active on all these nodes but only the application CRG on the one node where the CRG is in the current primary state is told to actually "start" its application. In the case of a Domino server's application CRG, "starting its application" means doing a STRDOMSVR to start the Domino server on that cluster node.

Application CRGs configure an exit program that is called (submitted within a batch job) by CM when that CRG is created, started, stopped, etc. CM uses the system controlling subsystem (configured with system value QCTLSBSD) to determine where it will submit these exit program jobs. The function of an application CRG exit program submitted by CM depends on the CM request (start, stop, etc.) and what that function means within the exit program for the application being managed. The Domino application CRG exit program QNOTES/QNNHACRG is defined as part of Domino 5.0.7. It is automatically configured when a Domino application CRG is created. This exit program performs the clustering functions required to allow CM to manage a Domino server. For example when a Domino application CRG is started on a cluster node, the

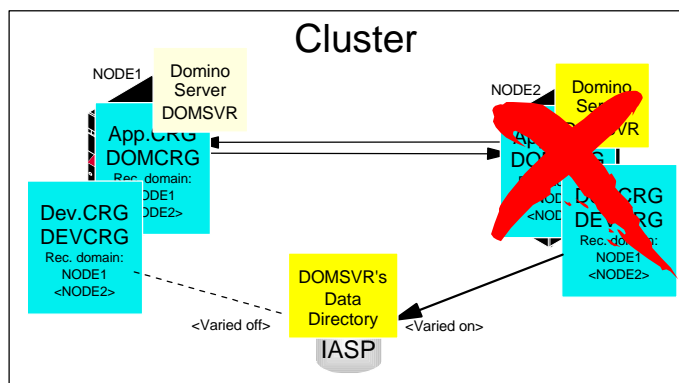
program will attempt to start the Domino server configured by the application CRG. When the application CRG is ended, the exit program is invoked by cluster management and it ends the Domino server.

Every application CRG must be defined with a unique IP address for use by the applications users. This address is associated with the application by CM and is automatically switched (via IP address takeover) between the cluster nodes within an application CRG's recovery domain. External communication with an application CRG is always via this same IP address but the address is switched between nodes in the CRG's recovery domain as each node becomes the current primary active node in the cluster for that CRG. When a Domino server is defined as an application CRG, the server must already be configured to use a specific IP address (see the TCP/IP port options parm (TCPOPT) on the CFGDOMSVR and CHGDOMSVR commands). This IP address cannot already be in use by any other application except for that one Domino server. The Domino server's configured IP address is automatically used for the application CRG's IP address when creating the application CRG for that server. Notes clients that are configured to access the Domino server will use this same IP address. When the Domino server is started on any node in the cluster, the client can access the Domino server using that IP address regardless of which node in the cluster where the Domino server is started.

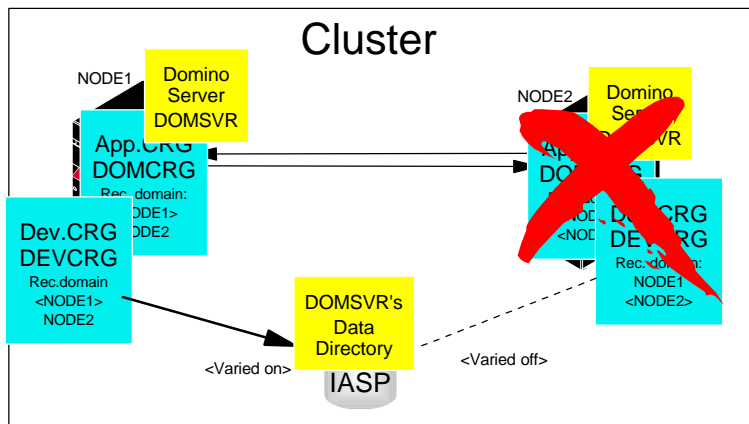
Cluster Management Supported Failover and Switchover

The major function CM provides to an application or device when it is configured as part of a CRG is automatic recovery in the event of a cluster node failure. For any CRG that was active on the failing node, cluster management automatically attempts to restart that CRG on the next available backup node in the cluster by using the CRG's "recovery domain". When a CRG is created it is configured with a list of cluster nodes that is its recovery domain. The list designates which cluster node is started as the "primary" node and the other nodes in the list are the "backup" nodes. The order of the nodes in the recovery domain determines the order of nodes where cluster management will attempt to start the CRG in the event of a failover or switchover of the CRG. When a CRG is initially started by the cluster it first starts it on the primary node as defined in its recovery domain. This is considered the "current" primary node. On a cluster node failure, cluster management makes the next backup node in the list the new "current" primary and the CRG's application is then started on that node. Should this node fail the next node becomes the current primary and so on. When all the backup nodes in the recovery domain have been tried the CM will again start with the original primary node in the recovery domain.

Failover to the next node in a CRG's recovery domain is automatic when a cluster node fails. The current primary node for a CRG can also be "manually" switched. This is known as a "switchover". Essentially the same actions are taken by cluster management in that it uses the CRG's recovery domain and restarts the application CRG on the next available "back up" cluster node as configured in the order the nodes appear in the CRG's recovery domain.

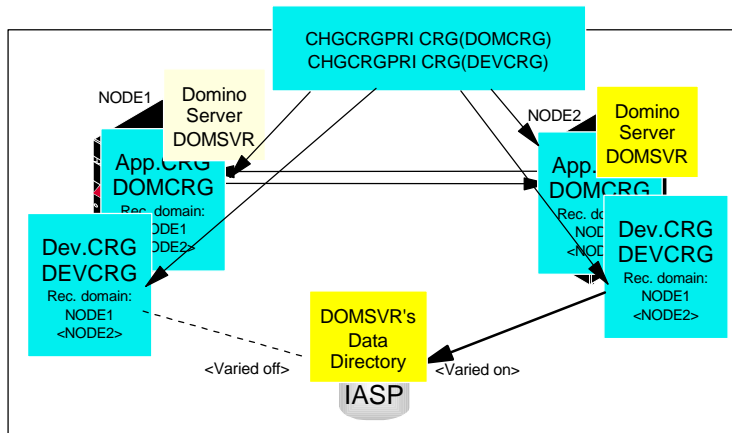


Clustering failover provides a means to automatically recover from the failure of a cluster node by having the cluster resources that are currently active and "primary" on that node switched to a backup cluster node. In this example CRGs are active on cluster node NODE2 that then experiences a failure.

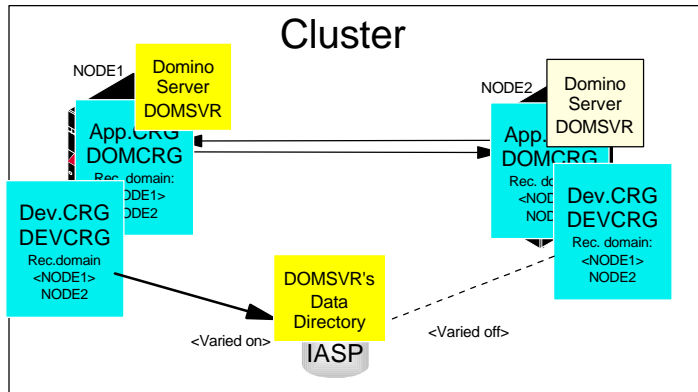


Continuing from the previous example, as a result of the failure of NODE2, Cluster Management function has automatically switched the active CRGs to their backup nodes (NODE1). The CRGs are now primary on that node. In the specific case for the Domino server DOMSVR, it's data directory is now accessible from NODE1 and the server has been started on NODE1.

Using a switchover, a CRG 's application can be started on a backup node in the cluster as a means to provide a controlled way to end the application on one node and (re)start it on a different node. A switchover of an application CRG (a Domino server application CRG for example) would be how cluster management can be used by an operator to start the application on a backup node in order to perform a system IPL or other scheduled maintenance that would not allow the application to remain active on its primary node.



Swiching over a CRG provides a means to control which cluster nodes are the "primary" node for cluster resources, In this example CRGs are active on cluster node NODE2 and a CHGCRGPRI CL command is used to switch the primary nodes for the CRGs.



Continuing from the previous example, as a result CHGCRGPRI commands, Cluster Management function has switched the active CRGs to their backup nodes (NODE1). The CRGs are now primary on that node. In the specific case for the Domino server DOMSVR, its data directory is now accessible from NODE1 and the server has been started on NODE1.

Related References

For more information on iSeries clustering, Cluster Management APIs, and information on switchable IASPs or other related High Availability topics see the following redpaper:

<http://www.redbooks.ibm.com/redpapers/pdfs/redp0111.pdf>

and also the High Availability Web site:

<http://www-1.ibm.com/servers/eserver/iseries/ha/>

For more information on LPAR on iSeries systems see the following Web site:

<http://www-1.ibm.com/servers/eserver/iseries/lpar/index.html>

For more information on Operations Navigator including Management Central see the following Web site:

<http://www-1.ibm.com/servers/eserver/iseries/sftsol/MgmtCentral.htm>

Domino Application Cluster Resource Groups

ClusterProven applications are configured to and managed by OS/400 Cluster Management (CM). CM manages the application as a "cluster resource" just like the switchable device driver. To allow the application to be known to the cluster, an application cluster resource group (CRG) has to be created. An application CRG is a "cluster wide" object and is created on every node in the cluster. The CRG is created as an OS/400 *CRG object on every node in the cluster. There can only be one application CRG with the same name (10 characters) in the cluster. Every Domino server that is to be managed through cluster management has to be defined as an application CRG.

Creating a Domino Application CRG

An application CRG has to be created for every Domino server that is to be managed by OS/400 CM. Creating a Domino server as an application CRG will automatically create a replica of that server's OS/400 configuration on all the cluster nodes specified in the Domino application CRG when it is created. Once the CRG is created then the server can be managed as a cluster application. Cluster management APIs, CL commands, or the Management Central GUI interface can be used to manage the Domino server and start, stop, or switch the Domino server between nodes in the cluster. The Domino server is started when its application CRG is started on a node that is considered the current primary node for its CRG. Any number of Domino servers can be configured on the cluster nodes and any subset of these servers could be configured as application CRGs and started at various times. While each CRG is managed separately by CM, Domino servers that are configured with an application CRG might have to be managed as a group, if there are many servers that all have their data directories accessible through (and managed by) the same device CRG.

A CL command, CFGDOMCRG, is provided to create the Domino server's application CRG for an existing Domino server. This command has to be invoked on the cluster node where the Domino server specified on the command is already configured. All the nodes in the cluster should be active when this command is invoked. Using the CFGDOMCRG command is the only way supported to create a Domino application CRG that can be used to manage the Domino server as a cluster application. A Domino server can only be configured by one application CRG in the cluster.

A domino server can only be supported as an application CRG in a cluster when its data directory is located on a switchable device. The switchable device allows access to its data directory to be switched to any one cluster configured as part of the device CRG. Device CRGs are managed, including switching nodes, through Cluster Management Operations Navigator or CL commands and APIs. In order to create a Domino application CRG, the device CRG and the switchable device that will allow this access from multiple cluster nodes must exist. The CFGDOMCRG command will check for the device CRG specified in a required command parameter. The command processing program will make sure the Domino server that is to be configured has its data directory configured in a UDFS (user defined file system) that is mapped to this device. In order for the CFGDOMCRG command to make this check, the device CRG must exist and be active and accessible by (switched to) the cluster node where the CFGDOMCRG command is invoked.

The Application CRG's Recovery Domain

Every application CRG has a recovery domain. This is a list of cluster nodes (systems) configured as part of the CRG. The nodes in the list can be any of the nodes defined to the cluster. Not every node in the cluster has to be in a recovery domain for a CRG. When the application CRG is started in the cluster it is started on every active node in its recovery domain. The application itself however is told to start on only one of the nodes in the recovery domain. The recovery domain lists the preferred "primary" node where the application will first be started, and all of the backup nodes. The order of the nodes in the recovery domain control which node becomes the next node where that application would be started in the event of a node failure or a manually controlled "Switchover" to a backup node. Besides the configured preferred primary, there is the notion of a "current" primary that is maintained by CM for the CRG. This is the node that has currently taken over and is where the application is started by CM. Note that in the event a CRG is ended and restarted, it will be restarted on its "current" primary node. The recovery domain of a CRG can be modified after the CRG has been created using CM APIs. Nodes can be added to or removed from the recovery domain and the configured primary and order of backup nodes can also be changed. The CRG cannot be active in the cluster at the time its recovery domain is changed.

Domino application CRGs have a direct dependency on the device CRG that controls the access to its data directory from any node in the cluster. When an application CRG is configured for a Domino server, the CFGDOMCRG command processing program finds the device CRG that controls access to its data directory. The recovery domain for the device is used to create the recovery domain for the Domino application CRG. That's why there is no cluster node list specified on the command. (The name of the device CRG is required as an input parameter on the CFGDOMCRG command.) The Domino application CRG will have the same configured primary and order of backup nodes as does the device CRG that controls access to its data directory. Any number of servers could have access to their data directories controlled by the same device CRG therefore any number of Domino application CRGs can be configured to use the same device CRG.

The recovery domain should always be exactly the same as the recovery domain of the device CRG that controls the access the storage defined by that device. When a cluster node fails the CM function should switch both the device CRG and the Domino server's application CRG over to the same (next) backup node. The same action should be performed when the device CRG and the Domino application CRG is manually switched over to the backup system. When switching over the Domino application CRG has to be switched over first to make the backup node in its recovery domain the new "current" primary node. Then the device CRG it depends on also must be switched over so that it also has the same new current primary node as the server's application CRG. Switching the device CRG means the server can then access its data directory and the Domino server should be able to start up on the backup node. (See the CHGDOMPRI command for more on switching the current primary nodes for these CRGs.)

Once the Domino application CRG is created its recovery domain should never be changed directly using CM APIs or commands. Rather the Domino application CRG should be ended and then the device CRG that controls access to its data directory should have its recovery domain changed first (by any CM interface). Once

the device CRG's new configuration has been tested, the Domino application CRG needs to be changed so that its recovery domain matches the now changed recovery domain of the device CRG. Since CRGs have to be ended to change their recovery domains, the Domino application CRG must be ended as the device CRG has to be ended when its recovery domain is being changed. Once both CRGs again have the same domain they can be restarted and will start with the configured preferred primary node as the node where the Domino server will start. The CHGDOMCRG command processing program will always find the device CRG the Domino server depends on and updates that server's application CRG's recovery domain to match the device CRG's recovery domain. Note that this means the Domino application CRG must be ended when the CHGDOMCRG command is used and will restart on its preferred primary node when the CRG is started. If the recovery domain of the device CRG is ever changed for any reason the CHGDOMCRG command must be used to update the recovery domain of all the Domino application CRG's that were configured to use that device CRG. (See the CHGDOMCRG command for more on changing a Domino application CRG.).

Configuring the Same Domino Server on Other Cluster Nodes

When the CFGDOMCRG command is invoked its command processing program in turn builds and invokes a Cluster Management API called QcstCreateClusterResourceGroup. Calling this API tells the Cluster Management create CRG function to create the new application CRG on all the nodes in the cluster. (CRGs are known to every node in the cluster.) When any application CRG is created, the name of an exit program (QNOTES/QNNHACRG) is passed as a parameter. This program is invoked from cluster management on every node in the application CRG's recovery domain. The Domino application CRG exit program is installed as part of installing Domino for AS/400 5.0.7. That release of Domino must be installed on every node in the CRG's recovery domain or the create of the Domino application CRG will fail on that node.

When cluster management invokes this Domino exit program it automatically indicates that the exit should start the process of creating a replica of the application CRG's server. (The replica of the Domino server will only be known on those cluster nodes that are defined in the application CRG recovery domain.) All the information about the Domino server is not known at the time the CRG is created so the exit program cannot completely create the Domino server's replica on the other nodes. The first time the application CRG is started on a cluster node in its recovery domain, the rest of the information about the Domino server can be found and the Domino server's replica is completely configured and the server is then started on that cluster node.

Domino Application CRG Servers and "*STANDBY" state

Eventually a Domino application CRG will be started on every node in its recovery domain. After it is started the first time a replica of a Domino server is visible on every one of the nodes in the application CRG's recovery domain (even when the application CRG is not active in the cluster). While the server's application CRG is started on the primary node in its recovery domain, that server's configuration is also still visible on all the backup nodes in the CRG's recovery domain. Operations Navigator or the WRKDOMSVR command can be used on the cluster nodes that are only in a "backup" role at the moment. To allow a Domino server to be configured but not able to be used, a server state called "*STANDBY" (stand by) is used. Domino servers are in "*STANDBY" state for the following reasons:

- They are configured as part of an application CRG but the CRG is not currently active for the cluster,
- The CRG that configures them is active but primary (the Domino server is started) on another cluster node,
- The node is in the process of becoming primary as part of starting up the CRG on the node where the server appears in a "*STANDBY" state. The Domino server hasn't yet been started.

Domino servers in a *STANDBY state cannot be started as their CRG is not started on that node. Since it is very likely their data directory is also not active and accessible to that node, there are not many functions that can be performed on a Domino server when it is in *STANDBY state. The server cannot be started and even "stand alone" server functions or BRMS backup of the server cannot work if the server's data directory is not accessible. Once the server's application CRG is started on that cluster node (when the node becomes the current primary for that CRG's recovery domain) the server will be started. Domino servers will go through the *STARTING state to a *STARTED state in a manner similar to using the STRDOMSVR command. Once the node is set to be the current primary for the recovery domain, as long as that server's application CRG is active and the current primary is not changed, the server is like any other Domino server configured on the same system. It can be started and ended. It can have its console displayed and be sent commands like any other

Domino server configured on the same system. The Domino server should not be removed on any cluster node using CFGDOMSVR OPTION(*REMOVE) however as the server is still defined as part of an application CRG. (The CRG has to be deleted first. See the section "Deleting an Application CRG" below.)

When a Domino application CRG is ended, the server will go to a *STANDBY state on all the cluster nodes where it is defined. Starting a Domino server's application CRG may not mean that server will come out of *STANDBY state. When a CRG is started, the cluster management function will only tell the exit program configured to "start the application" on the current primary node. At most one node in the recovery domain of the servers CRG can start be "started" in this way at any given time. Stated another way, the server's application CRG can be primary on only one node in the recovery domain and CRG will only be started on that node. It will remain in *STANDBY on the other nodes.

In addition the exit program that is attempting to start the Domino server will not start it (using STRDOMSVR) until the exit program determines that the device CRG is also in the same primary state on that exit's cluster node. Unless that is true, the server's data directory is not accessible from the exit's node, so there is no reason for the exit program to attempt to start the Domino server. The exit program will keep trying to determine when the device CRG is available and post a message to QSYSOPR message queue when it has tried for about five minutes. Domino servers that are configured as part of an application CRG that cannot start for this reason will remain in the *STANDBY state. The exit program will continue to attempt to start the Domino server on the cluster node until it is ended. Eventually it will post an inquiry message to the QSYSOPR message queue asking if it should keep trying to access the device CRG or if it should cancel. When the device CRG is eventually made primary on that exit's same cluster node, the device will then have been switched to be accessible from the same node as the application CRG. Then the server's data directory also becomes accessible. The Domino application CRG exit program knows that when that happens the server can be started on that node and does the STRDOMSVR. The concept that a Domino server application CRG depends on its device CRG in order to start its Domino server is important in managing both CRGs in the cluster and is discussed at length below in this document.

Restrictions and Limitations of Domino Application CRG configured servers

See the section How to create a Domino application CRG below for more about what is required to be configured and working before a Domino application CRG could be successfully configured. For certain the Domino server has to be created first before the CFGDOMCRG command can be used to configure it as part of an application CRG. In addition to situations that could arise due to the same Domino server being able to start on multiple systems, there are some Domino server functions that are not supported on servers that can be switched to and started on other nodes in a cluster.

OS/400 clustering enforces a restriction that the CRGs can only be switched or failed over to a cluster node that has an OS/400 release level greater than or equal to the system from where the CRG is being switched. Once a CRG has switched to be started on a higher release level of OS/400 it cannot be switched back and started on a system with a lower release level.

OS/400 system specific functions used by a Domino server could behave differently if a server is switched to a system that has a different release level.

Installed PTFs (fixes) for OS/400 or the release and fix levels of products or functions used by a Domino server could be different on the different systems.

OS/400 system specific objects used by Domino (system objects, Notes related objects in system libraries, programs that are invoked as part of a Domino application) might not be duplicated or have the same contents on every system in the cluster where the same Domino server could be started. If a server function or application depended on the existence or the contents of such an object, it could fail or behave unpredictably if that object was not found or was different on other systems where the same server was started. Specifically this could apply to programs called directly from Domino server applications or dependencies on other products that have been installed on some systems and not others. Situations where this occurs might be resolved by identifying what is different on the nodes in the cluster and correcting the differences between the systems.

OS/400 system specific values and attributes that affect how a Domino server works on any given system might not be set the same on all the systems in the cluster where that same Domino server might be started.

Specifically, system languages might not be installed on all the systems to which a Domino server could be switched.

The same release of Domino should be installed on all the systems where a server could be started. Domino application CRG support does not enforce this. Failing over to or switching a Domino server between systems running different releases of Domino is a "user beware" situation. The first release of Domino that could support a server as an application CRG is 5.0.7 so that version of Domino has to be installed on all the nodes in the cluster that are part of the Domino server application CRG's recovery domain or it cannot be created for the recovery domain. When the next release level of Domino is installed it is recommended that all the Domino application CRGs be ended (the Domino servers will end) and the new release of Domino be installed on all the cluster nodes where any Domino server could be started.

It is recommended the fix level of Domino code should be kept the same on all the systems where a server could be started. This is not enforced by Domino server CRG support and is a "user beware" situation. While this could possibly be used to test a Domino fix's effect on a server function, depending on the nature of the fix it could change the content of the information in a server's data directory. The server might not function correctly when it is then switched to another system that does not have the same Domino fix level installed.

Domino for AS/400 server specific functions to support directory synchronization (CFGDOMSVR parameter DIRSRV set to *SYSDIR or *ALL) will not work correctly when a server is switched to and started on other systems. This function should not be specified when creating the server to be configured as part of an application CRG. If the server specifies this attribute it could be changed to remove this from its configuration before it is configured as part of an application CRG.

Domino for AS/400 server specific functions to support routing e-mail through the OS/400 Mail Server Framework (CFGDOMSVR parameter SMTP set to *MSF) will not work correctly when a server is switched to and started on other systems. This function should not be specified when creating the server to be configured as part of an application CRG. If the server already specifies this attribute it could be changed to remove this from its configuration before it is configured as part of an application CRG.

Effects on the Installation of Another (Next) Release Level (QMR) of Domino

As mentioned above, the Domino CRG support does not enforce that the same release level of Domino must be installed on all the cluster nodes where that server could be started. This affects the procedure for installing a new release level (such as a QMR) of Domino in an environment that is using Domino servers that are configured by application CRGs. In general the only supported situation is that all the systems have the same release of Domino installed. The correct procedure would be to end all the Domino application CRGs on all the cluster nodes and then install the QMR on all these nodes. When the new level of Domino has been successfully installed on all the cluster nodes where a CRG configured server could start, the Domino application CRG can be restarted. The first release of Domino that could support a server as an application CRG is 5.0.7 so initially that version of Domino has to be installed on all the nodes in the cluster that are part of the Domino server application CRG's recovery domain or it cannot be created for the recovery domain.

The Domino servers configured with an application CRG will likely need to have Domino templates in their data directories updated when the new (next) release level of Domino after 5.0.7 is installed. All the Domino server application CRGs should be ended on the node where the new release is being installed. (All the Domino servers on that node should be ended.) The device CRG(s) that control the access to these server's data directories should remain active and all the server's data directories should be accessible to a cluster node while the next release of Domino is installed on that node. The Domino for AS/400 installer program will find and update the templates in these data directories on the IASP device as it would any other Domino server's when a new release level of Domino is installed on that cluster node. Once the new release of Domino is installed all the templates for the data directory(s) are updated. This only has to be done once. There is only one copy of the data directories on the IASP device. It is the access to these now updated directories that is shared. All the directories should have been updated to the next release by the install of the release. The same release of Domino should be installed on all the other nodes in the cluster before restarting the Domino servers application CRGs. The IASP device does not have to be switched over or started on these other nodes in order to do that install as the data directories have been updated once already.

Domino Application CRG Commands

There are specific configuration requirements to correctly configure a Domino server to cluster management as an application CRG. There are three commands provided to allow Domino application CRGs to be configured, changed, and deleted. Note that these commands are "server oriented" meaning that the commands assume that users of Domino application CRGs might be more familiar with managing Domino servers by their names rather than the application CRG that is configured for them or the device CRG that controls access to their data directory.

The only supported way to create a Domino server application CRG is to use the CFGDOMCRG command. There is error checking done, including checking the configuration of the server and validating that it has a data directory managed by a device CRG. This is done as part of the command processing program to assure the application CRG being created has a chance of functioning. It is recommended that the CHGDOMCRG and DLTDOMCRG commands be used when a Domino application CRG is changed, or removed; however, the clustering Management Central GUI, clustering APIs, or commands could be used to change or remove a Domino application CRG with no error provided that the Domino server application CRG exit program name and exit program data are never changed via these interfaces. There are also three commands provided that can be used to stop, start, and change the primary node (switchover) a Domino application CRG. Using these three commands to manage a Domino application CRG is optional. Domino application CRGs and the device CRGs they are associated with can be started, stopped, and switched with the cluster management GUI (using AS/400 Operations Navigator and Management Central), the system APIs that are supported by OS/400 cluster management, or by the QUSRTOOL shipped cluster management commands.

These commands must be created after installing Domino for AS/400 5.0.7.

How to create commands

To create the six commands enter the following on an AS/400 command line:

```
CALL QNOTES/QNNHACCMD PARM('<user library name>')
```

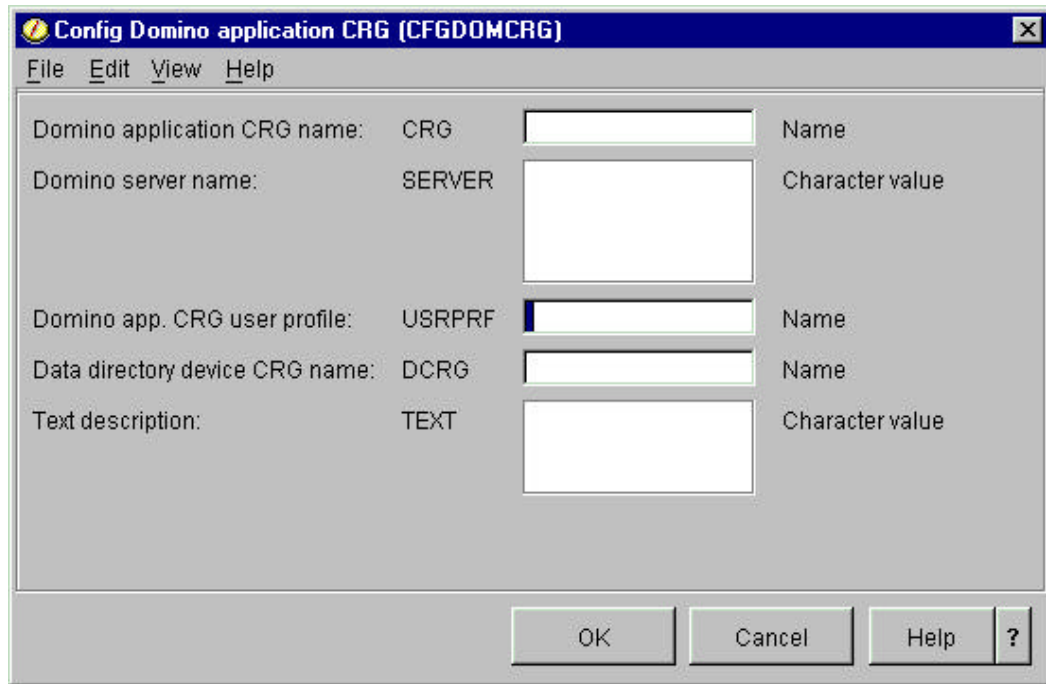
This will call a program that is installed as part of Domino for AS/400 5.0.7 that creates the six commands that are associated with configuring and managing Domino application CRGs. The commands are created in the library <user library name> named in the parameter passed into the program. The program creates these commands with QSYS as the owner and *PUBLIC *EXCLUDE authorities. Note this program is shipped with QSYS as its owner and *PUBLIC *EXCLUDE authority so the caller of the program will need the appropriate authority in order to run the program to create the commands.

This program can be called to create the six commands into as many libraries (including QSYS) as needed. The commands created this way are not part of the Domino product (LNT5769) and are not saved when the Domino product is saved.

The following sections explain what each command does.

CFGDOMCRG - Configure a Domino server application CRG

Configures a Domino application CRG given a Domino server name.



This command will attempt to create a OS/400 cluster management application CRG object (type *CRG). The command builds the parameters required and then invokes the QcstCreateClusterResourceGroup system API. Any errors that result from invoking this API are resigalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

The cluster must be configured before this command will work. This command will only work on a system that is defined as a cluster node. When the command completes successfully the application CRG object named by the CRG parm will exist on every node in the cluster. There can only be one CRG (of any type) with this name in the cluster.

Command parameters:

CRG - The name (up to 10 characters) of the application CRG that is to be configured in the node's cluster. If this CRG already exists the command will fail.

SERVER - The name of the Domino server that is going to be configured and managed as part of the application CRG being configured. This Domino server must exist on the system where this command is invoked. Only "plain" Domino type servers are allowed to be specified. The Domino server must be ended (ENDDOMSVR or use WRKDOMSVR option 6=End server) before changing it to be managed as part of an application CRG. If the server specified is already part of another application CRG the command will fail.

As part of creating a Domino application CRG for this Domino server, CFGDOMCRG will attempt to create a replica of this Domino server on every cluster node named in the recovery domain of the application CRG being created. If the Domino server is already configured on any of these nodes in the cluster the command will fail.

When creating an application CRG a unique, unused, TCP/IP address must be configured as part of the application CRG. The IP address used by CFGDOMCRG will be taken from this Domino server's configuration. If the server is not configured to use a unique IP address, the command will fail to create the application CRG. An existing Domino server can be changed and configured to use a specific IP address (see the TCP/IP port options parm (TCPOPT) on the CHGDOMSVR command).

USRPRF - The user profile that is configured on every node in the cluster that is to be used to manage the Domino application CRG. If this user profile is not configured on the cluster nodes, the command will fail.

This user profile is configured as part of the application CRG and is used by cluster management to submit the application CRG exit program that is also configured as part of the application CRG. The application CRG's exit program is configured as part of the application CRG. The same exit program is used for all Domino

application CRGs. It is always QNOTES/QNNHACRG and cannot be changed. When cluster management wants the application CRG to perform a clustering function (start, stop, or switchover for examples) it submits a batch job under the user profile configured with the USRPRF parameter. This job that runs this exit program for the application CRG is submitted on every node in the application CRG's recovery domain.

The user profile specified must have job control authority (*JOBCTL) on all these systems and must be able to switch to QNOTES user profile as there are objects owned by QNOTES that are required to be created when the Domino server and its application CRG are created on the other nodes in the cluster. In addition, the user profile must be able to access the UDFS directory.

Note Even though it would already be defined on every cluster node, the QNOTES user profile should NOT be used here. It is not created with *JOBCTL authority with Domino for AS/400. For security reasons QNOTES should never be modified to be given *JOBCTL authority.

The user profile specified must be able to read and write to the default UDFS associated with device CRG specified by the DCRG command (user profile must have read and write authority to the UDFS directory) or the command will fail.

DCRG - The name of the device CRG that contains the Domino servers data directory. A Domino server can only be supported on multiple cluster nodes if its data directory is also accessible from these same cluster nodes. When creating a Domino application CRG the device CRG that is going to be used to allow the storage for a servers data directory to be switched between cluster nodes must already exist. It is recommended that the ability to switch the device CRG between all the cluster nodes in its recovery domain be tested first before using it to create a Domino CRG.

With CFGDOMCRG, the device CRG must be active ("primary") on the cluster node where the server specified by the SERVER parm above is configured and the CFGDOMCRG command is invoked. The command processing program will look for the data directory configured for the Domino server specified by the SERVER parm. If the server's data directory is not in the UDFS configured as the default UDFS for this device CRG, the command will fail.

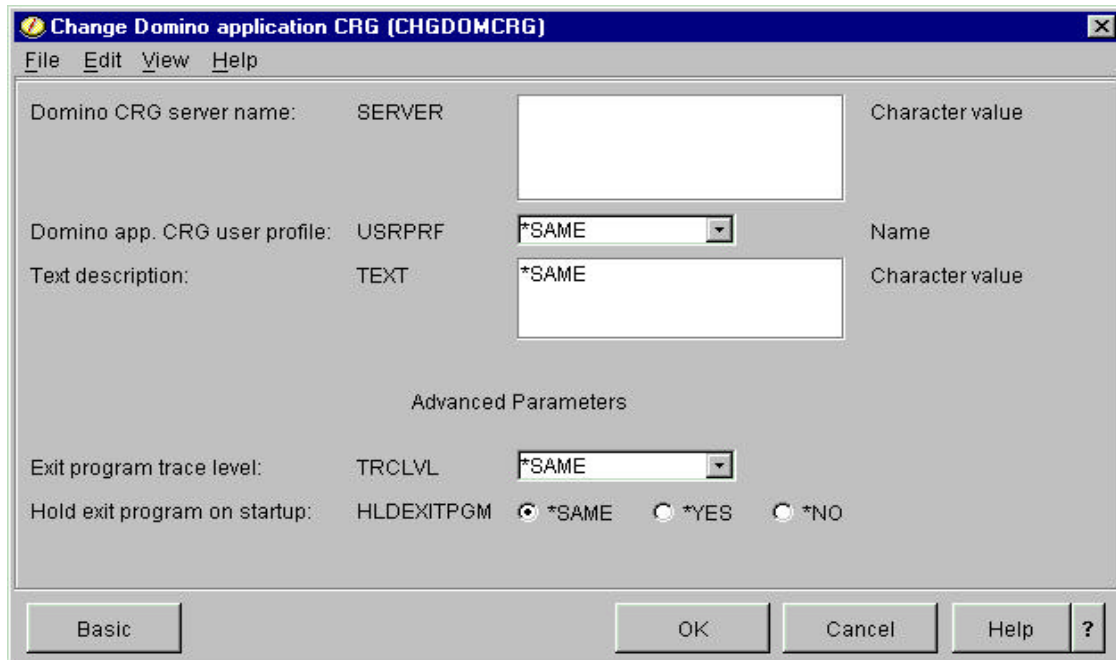
Any number of Domino servers can use the same device CRG to provide access to the server's data directory to other cluster nodes. When configured this way, all the servers have to be "switched over" (or their application CRGs have to be ended) when the shared device CRG is switched over to another node. Note that a device CRG cannot be switched over (it will fail to switch with an error) if any Domino servers that have their data directories on that switchable device are active.

Note The recovery domain used to create the new Domino application CRG is not configurable. Instead it is taken directly from the recovery domain of the device CRG that manages which node can access its data directory. The recovery domain of the application CRG should not be changed unless the recovery domain of this device CRG is also changed. If the device CRG recovery domain is ever changed, the CHGDOMCRG command can be used to automatically reset the recovery domain of a Domino application CRG so that it will again match up with the device CRG's recovery domain.

TEXT - Up to 50 characters of text are used when creating the Domino application CRG. This text is optional but can be useful to help identify the Domino server associated with this application CRG. The cluster management Management Central GUI interface will allow application CRGs to be managed but does not show which are Domino application CRGs and which Domino application CRGs are associated with which Domino servers in the cluster.

CHGDOMCRG - Change Domino server application CRG

Changes a Domino application CRG given a Domino server name.



This command will attempt to change an OS/400 cluster management Domino application CRG object (type *CRG). In this case the "change" is only changing a few select application CRG attributes. (Use the CHGDOMPRI to change the current primary node for the application CRG.) The command builds the parameters required and then invokes the QcstChangeClusterResourceGroup system API. Any errors that result from invoking this API are resignalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

Command parameters:

SERVER - The name of the Domino server that is to have its application CRG changed. This Domino server must exist and be configured with an application CRG on the system where this command is invoked. Since a Domino server is configured on all the nodes in its application CRG's recovery domain, it's CRG can be changed with this command from any of these nodes.

The Domino server (and its application CRG) do NOT have to be ended when they are changed however some changes will not take effect until the changed application CRG is ended and restarted. If the server specified does not exist or is not configured with an application CRG the command will fail.

The application CRG is a cluster level object and is changed on every node in the cluster. The Domino server associated with the application CRG is not changed.

When changing an application CRG the associated Domino server's IP address used by CHGDOMCRG will be taken from this Domino server's configuration. If the Domino server's IP address has been changed since the application CRG was configured the application CRG is automatically changed to use the new Domino server IP address. If the server is not configured to use a unique IP address, the command will fail to change the application CRG. An existing Domino server can be changed and configured to use a specific IP address (see the TCP/IP port options parm (TCPOPT) on the CHGDOMSVR command).

Note Whenever the Domino application CRG is changed with this command, the recovery domain for the CRG is always automatically reset. It is taken directly from the recovery domain of the device CRG that manages which node can access its data directory. This is done so that the application CRG's recovery domain will always match up with the device CRG's recovery domain. The recovery domain of the application CRG should not be changed unless the recovery domain of this device CRG is also changed.

USRPRF - The user profile that is configured on every node in the cluster that is to be used to manage the Domino application CRG. If this user profile is not configured on the cluster nodes, the command will fail. This is provided as a means to change the user profile a Domino application CRG uses without having to delete and recreate the application CRG.

TEXT - Up to 50 characters of text that are used when changing the Domino application CRG.

Additional/Advanced parameters:

TRCLVL - Allows the trace level of the application CRG to be set for trouble shooting or testing purposes. Whenever the application CRG exit program is invoked it does a certain level of logging of information in its job log. This parm sets the level of information that is logged. If not set with a CHGDOMSVR command, the default level when a Domino application CRG is configured is *ERROR. Parm values are:

- *SAME - (default) keep the same value for this CRG attribute
- *NONE - don't do any logging in the exit program job log
- *ERROR - log errors that might happen when the exit program fails to perform a function
- *INFO - log errors and other information that might help someone figure out what the exit program is doing as it performs a function
- *ALL - log everything possible

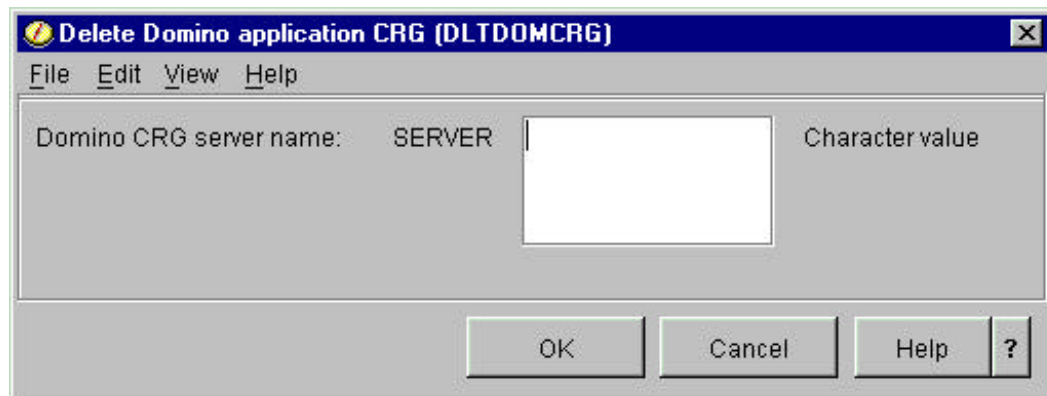
HLDEXITPGM - Allows the Domino exit program QNOTES/QNNHACRG to be held (halted) whenever it is submitted. This is to be used for allowing a Domino application exit program job to be serviced. If not set with a CHGDOMSVR command, the default when a Domino application CRG is configured is *NO . Parm values are:

- *SAME - (default) keep the same value for this CRG attribute
- *YES - Halt the exit program job on every cluster node in the Domino application CRG's recovery domain every time that job is submitted on those nodes to run the application CRG's exit program to perform a clustering function for the application.
- *NO - Do not halt the exit program job every time it is submitted.

Caution - When this option is turned on the next time the application CRG is started, the exit program's job is submitted and halted on every node in the application CRG's recovery domain. The application CRG will not completely start, end, switchover, etc. on any of the cluster nodes until every halted exit job is released for every cluster function on every node in the recovery domain.

DLTDOMCRG- Delete Domino server application CRG

Deletes a Domino application CRG given a Domino server name.



Note This does not remove the Domino server configured as part of the application CRG. To completely remove the server, the CFGDOMSVR command with the OPTION(*REMOVE) must be used. The command builds the parameters required and then invokes the QcstDeleteClusterResourceGroup system API. Any errors that result from invoking this API are resignalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

This command will attempt to delete an OS/400 cluster management Domino application CRG object (type *CRG). The command builds the parameters required and then invokes the QcstDeleteClusterResourceGroup system API. Any errors that result from invoking this API are resignalled to the command's caller.

Command parameters:

SERVER - The name of the Domino server that is to have its application CRG deleted. This Domino server must exist and be configured with an application CRG on the system where this command is invoked. Since a Domino server is configured on all the nodes in its application CRG's recovery domain, its CRG and all the replica server configurations can be deleted with this command from any of these nodes.

The application CRG is a cluster level object and is deleted from every active node in the cluster. The Domino server's "replicas" and the application CRG are removed from all the nodes in the cluster. (It is highly recommended all these nodes are active in the cluster when the CRG is removed.) The original Domino server associated with the application CRG is NOT removed from the node where it was originally configured (where the CFGDOMCRG command was run). The Domino server will remain configured on the original cluster node where the application CRG was created with the CFGDOMCRG command. CFGDOMSVR OPTION (*REMOVE) has to then be used to completely remove the server.

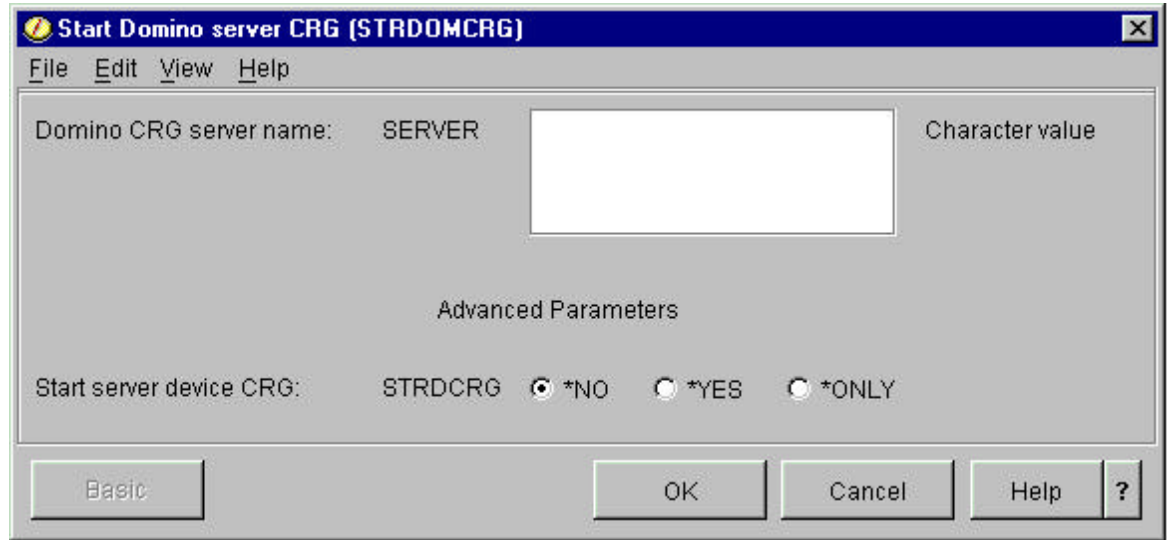
Note A Domino server should not be removed with the OPTION(*REMOVE) if configured with an application CRG. The server's Domino application CRG must be deleted first with the DLTDOMCRG command.

Deleting a Domino server's application CRG has no effect on the server's data directory. The server's data directory would still be configured on a switchable DASD device configured and managed as part of a device CRG.

Once a Domino server's application CRG has been deleted it can be reconfigured as part of a different, new, application CRG using the CFGDOMCRG command.

STRDOMCRG - Start Domino server application CRG

Starts a Domino application CRG given a Domino server name.



This will start the Domino application CRG on all the nodes in its recovery domain. It will only attempt to start the Domino server on the cluster node in the current primary node in the application CRG's recovery domain. If the device CRG associated with the application CRG is also active as a current primary on the same node, the application CRG starts the Domino server (STRDOMSVR) on that same node. The application CRG will NOT start the Domino server if the device CRG that controls access to its data directory is not also started and the "primary" node, for that device CRG is the same as that for the Domino server's device CRG.

The command builds the parameters required and then invokes the QcstStartClusterResourceGroup system API for every CRG started. Any errors that result from invoking this API are resignalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

Command parameters:

SERVER - The name of the Domino server that is to have its application CRG started. This Domino server must exist and be configured with an application CRG on the system where this command is invoked. Since a Domino server is configured on all the nodes in its application CRG's recovery domain, its CRG (and its associated device CRG) can be started with this command from any of these nodes.

If the application CRG is already started, the command will fail but does not effect the active CRG or the Domino server. The command can be used to start only the device CRG associated with the Domino server even though the application CRG might already be active.

Additional/Advanced parameters:

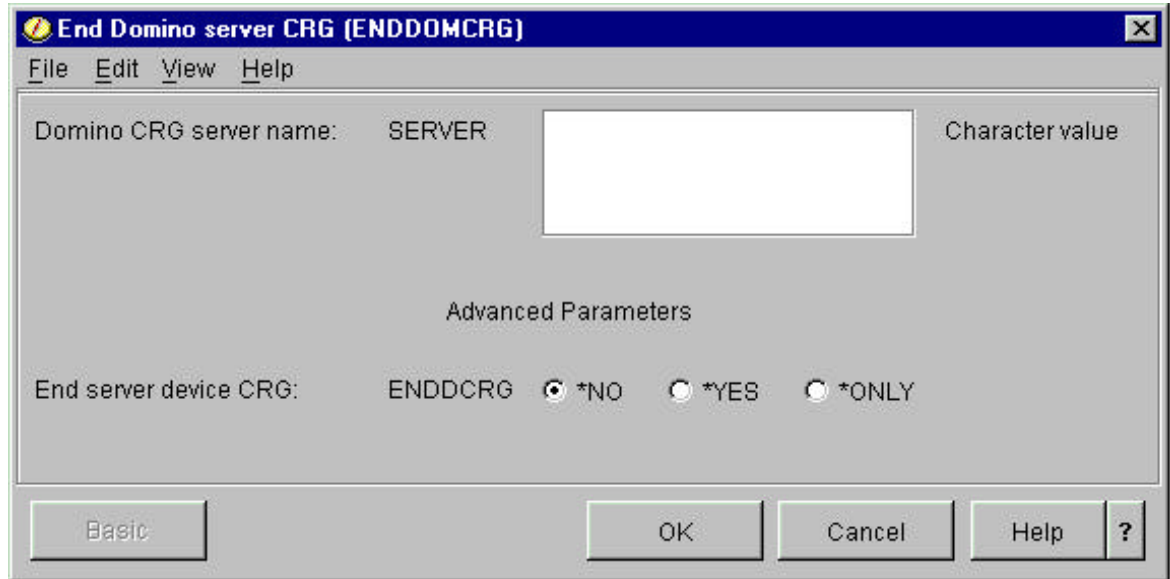
STRDCRG - Starts the device CRG that controls access to the server's data directory. The name of this device CRG is kept as part of the internal configuration of a Domino application CRG. The device CRG for a Domino application CRG server can be started from any node in the cluster where a replica of the server is configured. If not set with a STRDOMSVR command STRDCRG parm the default when a Domino application CRG is started is *YES . Parm values are:

- *YES - (default) start the device CRG that was configured as part of this Domino server's application CRG after starting the Domino server's application CRG.
- *NO - don't start the device CRG.
- *ONLY - only start the device CRG that was configured as part of this Domino server's application CRG. Do NOT also start the Domino server's application CRG.

Note This parameter will start the device CRG if it is not already started. It will not set the current primary node for the device CRG (where the Domino server's data directory could be accessed) to be the same current primary for the Domino server's application CRG.

ENDDOMCRG - End Domino server application CRG

Ends a Domino application CRG given a Domino server name.



This will end the application CRG on the cluster node where it is active. Ending a server's Domino application CRG always ends the Domino server (ENDDOMSVR) if it is running.

The command builds the parameters required and then invokes the QcstEndClusterResourceGroup system API for every CRG ended. Any errors that result from invoking this API are resigalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

Command parameters:

SERVER - The name of the Domino server that is to have its application CRG ended. This Domino server must exist and be configured with an application CRG on the system where this command is invoked. Since a Domino server is configured on all the nodes in its application CRG's recovery domain, its CRG (and its associated device CRG) can be ended with this command from any of these nodes.

If the application CRG is already ended, the command will fail but does not effect the CRG or the Domino server. The command can be used to end only the device CRG associated with the Domino server even though the application CRG might already be ended.

Additional/Advanced parameters:

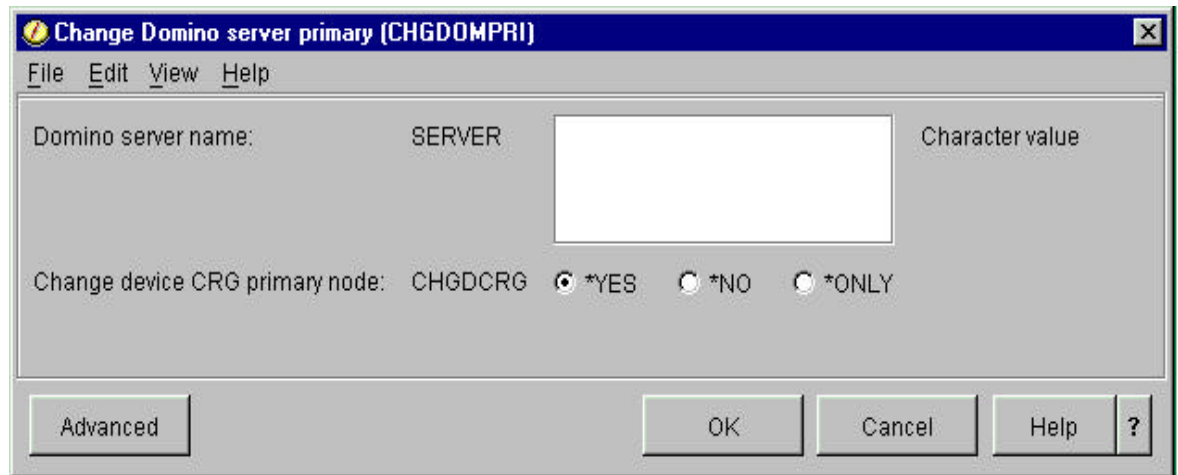
ENDDCRG - Ends the device CRG that controls access to the server's data directory. (The order of ending is always that the application CRG has to be ended first, then the device CRG.) The name of this device CRG is kept as part of the internal configuration of a Domino application CRG. The device CRG for a Domino application CRG server can be ended from any node in the cluster where a replica of the server is configured. If not set with a ENDDOMSVR command ENDDCRG parm the default when a Domino application CRG is ended is *NO . Parm values are:

- *NO - (default) don't end the device CRG that was configured as part of this Domino server's application CRG.
- *YES - also end the device CRG after ending the Domino server application CRG.
- *ONLY - only end the device CRG that was configured as part of this Domino server's application CRG. Do NOT also end the Domino server's application CRG. This parm option is provided as a way to end this device CRG only when the configured Domino server's application CRG is already ended.

Note This parameter will end the device CRG if it is not already ended. This parameter should be used with caution. A device CRG could configure more than one Domino server's data directory or the storage controlled by that device could be used by any number of applications on the node where it is being ended.

CHGDOMPRI - Change Domino server application CRG Primary node

Changes a Domino application CRG primary node (does a "switchover" to the next node in the CRG recovery domain) given a Domino server name.



The CRG(s) that are changed with this command must be active (started) first. When an application CRG is switched over to a new primary, it is automatically ended on the current primary node and then started on the new primary node. In the case of a Domino application CRG that means that the Domino server will first be ended and then attempted to be started on the new node.

Note The device CRG must also be active and the same node in its recovery domain also the current primary as the Domino server's application CRG.

The new current primary node in the cluster is determined by the recovery domain(s) of the CRGs that are changed. The "next" active node in the list of nodes in the CRG's recovery domain should become the new primary. The command can be used to change only the device CRG associated with the Domino server even though the application CRG for that server might be ended.

The command builds the parameters required and then invokes the QcstInitiateSwitchOver system API for each CRG that is changed. Any errors that result from invoking this API are resignalled to the command's caller. The description of this API includes documentation on the errors that can occur. This information will be helpful in diagnosing these errors.

Command parameters:

SERVER - The name of the Domino server that is to have its application CRG's current primary node changed. This Domino server must exist and be configured with an application CRG on the system where this command is invoked. Since a Domino server is configured on all the nodes in its application CRG's recovery domain, its current primary node (and/or the current primary of its associated device CRG) can be changed with this command from any of these nodes.

CHGDSCRG - Changes the current primary node to the next node in the recovery domain of the device CRG that controls access to the server's data directory. (The order of changing is always that the application CRG has to be changed first, then the device CRG.) The name of this device CRG is kept as part of the internal configuration of a Domino application CRG. The device CRG for a Domino application CRG server can have its current primary node changed from any node in the cluster where a replica of the server is configured. If not set with a CHGDOMPRI command CHGDSCRG parm the default when a Domino application CRG's primary node is changed is *Yes . Parm values are:

- *YES - (default) also change the current primary node for the device CRG after ending the Domino server application CRG.
- *NO - don't change the current primary node for the device CRG.
- *ONLY - only change the current primary node for the device CRG that was configured as part of this Domino server's application CRG. Do NOT also change the primary node for the Domino server's application CRG.

This parm option is provided as a way to change the current primary node of this device CRG only when the configured Domino server's application CRG is ended or is already the current primary on a desired cluster node and the device CRG is not. For example this function can be used to change the primary for the device CRG so that it will match the current primary for the Domino server application CRG. Recall that the application CRG could be currently active and waiting for the device CRG to be available on the node where the server's application CRG has already been started.

Note This parameter will implicitly end the device CRG on one node and start it on another. A device CRG could configure more than just one Domino server's data directory. In addition the storage controlled by that device could be used by any number of other applications on the current primary node.

The recovery domains of the Domino application CRG and the device CRG configured with it should be the same when the Domino CRG was originally configured. However, there is no guarantee they already have the same current primary when this command is used. There is no guarantee that the current primary node for the Domino server's application CRG is the same as the current primary for the device CRG it needs. This command will only set the new primary nodes for both CRGs to the next node in their respective recovery domains.

Preparing to Create a Domino Application CRG

Before attempting to create an application CRG there are some prerequisite steps that should be done. The most important of these is to make sure all the systems that are going to be part of a recovery domain for a Domino application CRG have Domino for AS/400 5.0.7 (or later) installed. The Domino application CRG will fail to create if there is even one node in the device domain that does not have this release of Domino installed.

There is a considerable amount of preparation to create an application CRG. Some of it might seem "trivial" in a cluster of two nodes and one device CRG. The idea behind this preparation is to be able to create a working Domino server application CRG the first time. Some of these steps don't have to be repeated every time a new Domino application CRG is created. Where these steps should be repeated is if a new node is added to the cluster and the intention is to use this node for Domino application CRGs or to add the node to an existing CRG. It should not be assumed that the new cluster node is configured correctly to be used with a Domino application CRG.

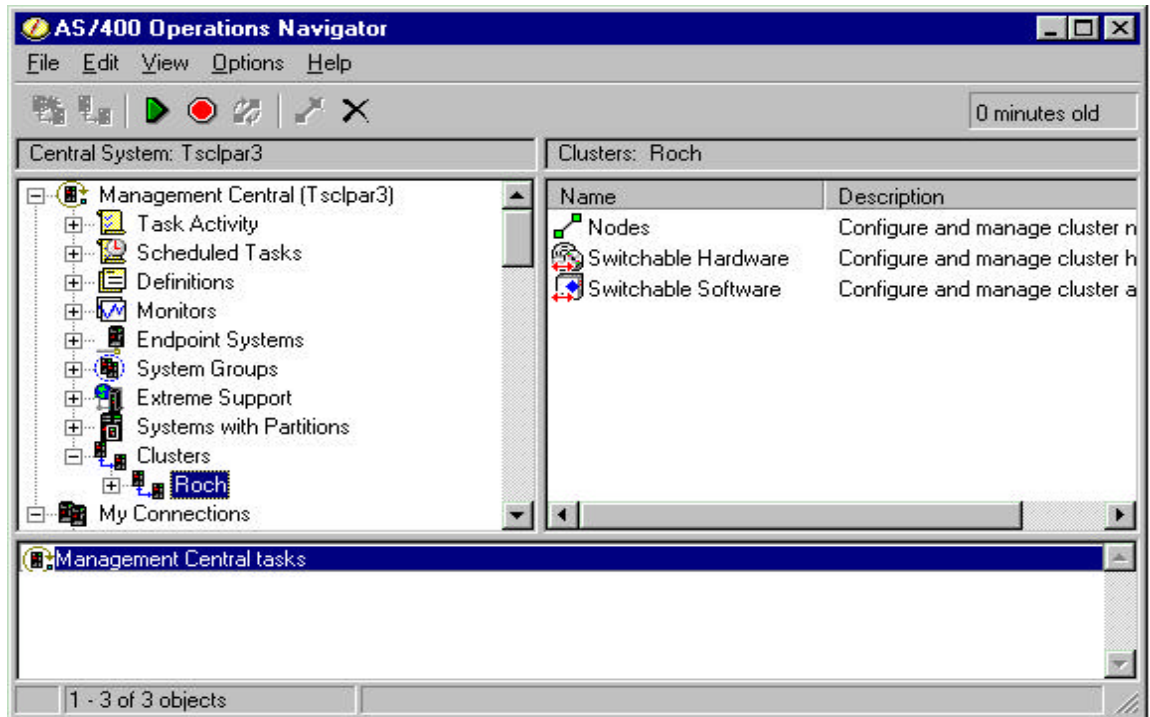
Checking the Cluster

The next major requirement is to have a working cluster. Recall that a working cluster can have as few as two nodes and that these can be LPAR systems in a single iSeries system. Note that High Availability business partners can provide clustering related applications that allow clusters, nodes, and CRGs to be managed. These are not discussed here and the documentation for these applications should be consulted to determine their requirements and functions that might be useful to help verify the cluster configuration before attempting to configure a Domino application CRG.

Checking cluster with Management Central

If the intention is to use Operations Navigator and Management Central to manage Domino application CRGs and their associated device CRGs, this should be tested to make sure the cluster being used is accessible from the Management Central system. If a cluster has been configured it should show up under the Management Central system.

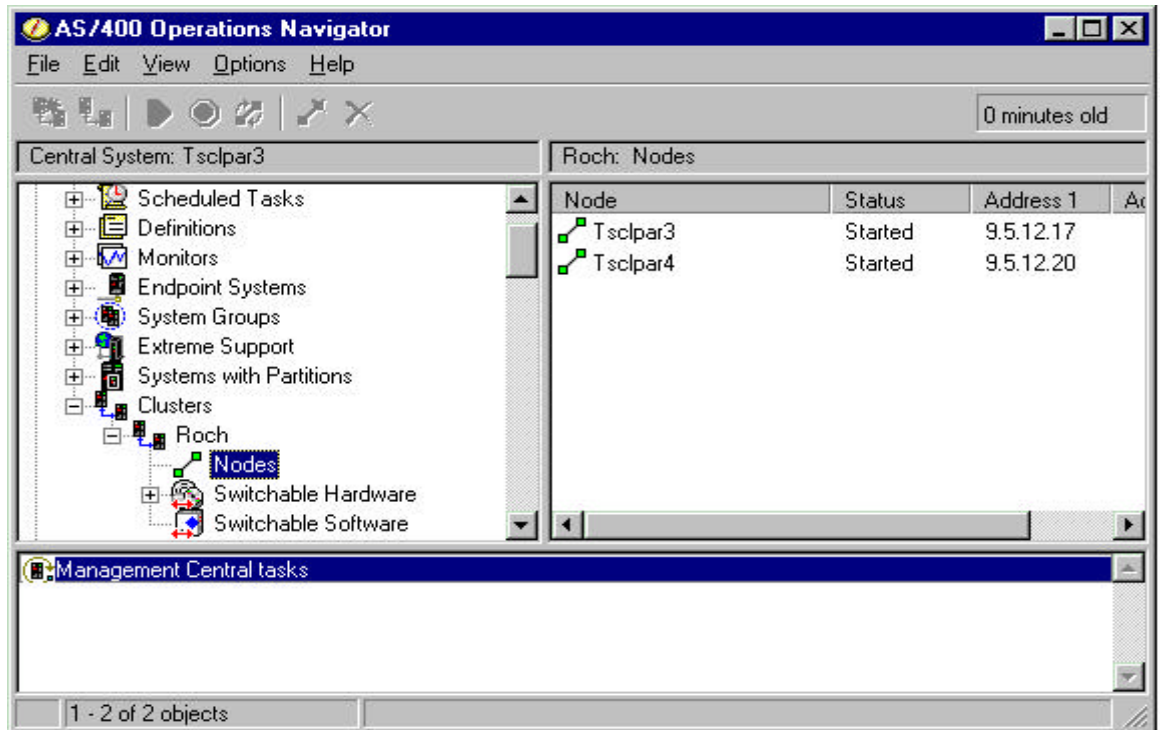
If the cluster does not exist, it can be created from either the cluster management utility shown here under Operations Navigator Management Central or use one of the available cluster management applications provided by a High Availability business partner. This document will not show how to create or add nodes to a cluster using these interfaces. It will not show the creation or management of a Domino application CRG using cluster management interfaces provided by these business partners.



This shows Management Central with a cluster called "ROCH" selected.

Note Operations Navigator does not display the real OS/400 names of the objects but rather shows the names in lower case with the first letter capitalized.

All cluster nodes defined for that cluster and their current cluster activity state can also be seen by selecting the Cluster icon for "Nodes."



This shows the same "ROCH" cluster after selecting the "Nodes" icon. There are two nodes shown, TSCLPAR3 and TSCLPAR4. Both are in the "started" state which means they are active in the cluster.

Checking cluster with Cluster Management commands

If using OS/400 clustering commands any node in the cluster should be able to display its cluster's name with the PRTCLUNAM command. If a cluster is configured, this will generate a spoolfile with the node's cluster name.

```
Cluster Name: ROCH
Requesting Node Id: TSCLPAR3
Current Cluster Version: 1
Current Cluster Version Mod Level: 0
Potential Node Version: 2
Potential Version Mod Level: 0
```

This is the output from the PRTCLUNAM command. The cluster name is ROCH in this example.

For more information about what nodes are known to the cluster, use the cluster command PRTCLUINF. Given a cluster's name, this command will generate a spoolfile listing all the cluster nodes and their current state.

```
User Space Name: QCSTUSRSPC
User Space Library Name: QUSRSYS
Cluster Name: ROCH
Format Name: LCTI01D0
Node Id to retrieve: *ALL
-
Cluster Information:
Cluster Resource Services Status: Not Active
Current Cluster Version: 1
Current Cluster Modification Level: 0
Number of Nodes: 2
-----
1. Node Id: TSCLPAR3
   Node Status: Inactive
   Potential Node Version: 2
   Potential Node Modification Level: 0
   Device Domain Name: ROCH
   Number of Interfaces: 1
   IP address: 9.5.12.17
-----
2. Node Id: TSCLPAR4
   Node Status: Inactive
   Potential Node Version: 2
   Potential Node Modification Level: 0
   Device Domain Name: ROCH
   Number of Interfaces: 1
   IP address: 9.5.12.20
-----
```

This is the output from the PRTCLUINF command on cluster ROCH. The two nodes in this cluster are TSCLPAR3 and TSCLPAR4. Clustering on these nodes happened to be inactive at the time the command was run.

These interfaces can be used to display this information from any node in the cluster but if the node is not currently active, the information reflects what was known about the cluster when the node was last active. The cluster nodes should be started to make sure the information about the cluster returned by these commands is current.

These interfaces can be used to make sure the nodes that are going to be part of an application CRG's recovery domain are configured, active, and manageable by the cluster.

Checking the Device CRG

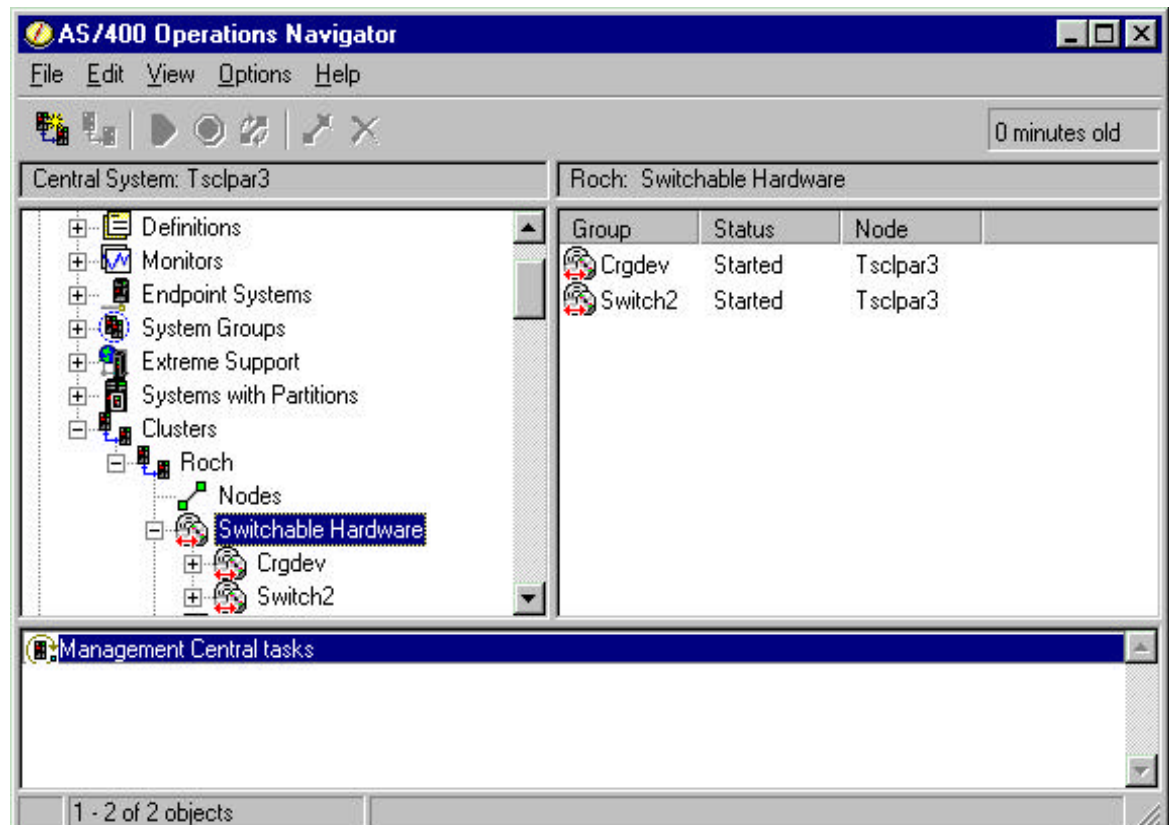
The other major prerequisite for creating a Domino server application CRG is that the server's data directory has to be stored on switchable disk storage configured on the cluster nodes as a device. In order for the



Domino server to be managed (and switchable) as part of the cluster, this disk device also has to be managed as part of a device CRG that is known to the cluster.

In any given cluster there could be many device CRGs configured and available to be used, each representing some set of switchable disk storage. Choosing the right one requires some planning and the choice should consider what, if any, other applications, users, or Domino application CRG servers might also store data in IFS files on the same switchable device. Every time this device is switched between cluster nodes, all the users of IFS files in that device are affected. The other consideration regarding what device CRG that a Domino server should use to control access to its data directory is the recovery domain of the device CRG. This list of cluster nodes will automatically become the recovery domain of the Domino application CRG when it is created. The choice of device CRG might mean that the recovery domain needs to be adjusted. It might mean that IFS file objects controlled by the device and used by a specific application need to be moved into other storage or managed by other device CRGs. Deciding on which device CRG to use to manage access to a data directory for a new Domino application CRG server is a simple task in a cluster of two nodes and one device CRG. However it can be quite complicated and require a great deal of planning in a large cluster with multiple device CRGs that each control access to many IFS files used by multiple applications, including existing Domino application CRG configured servers.

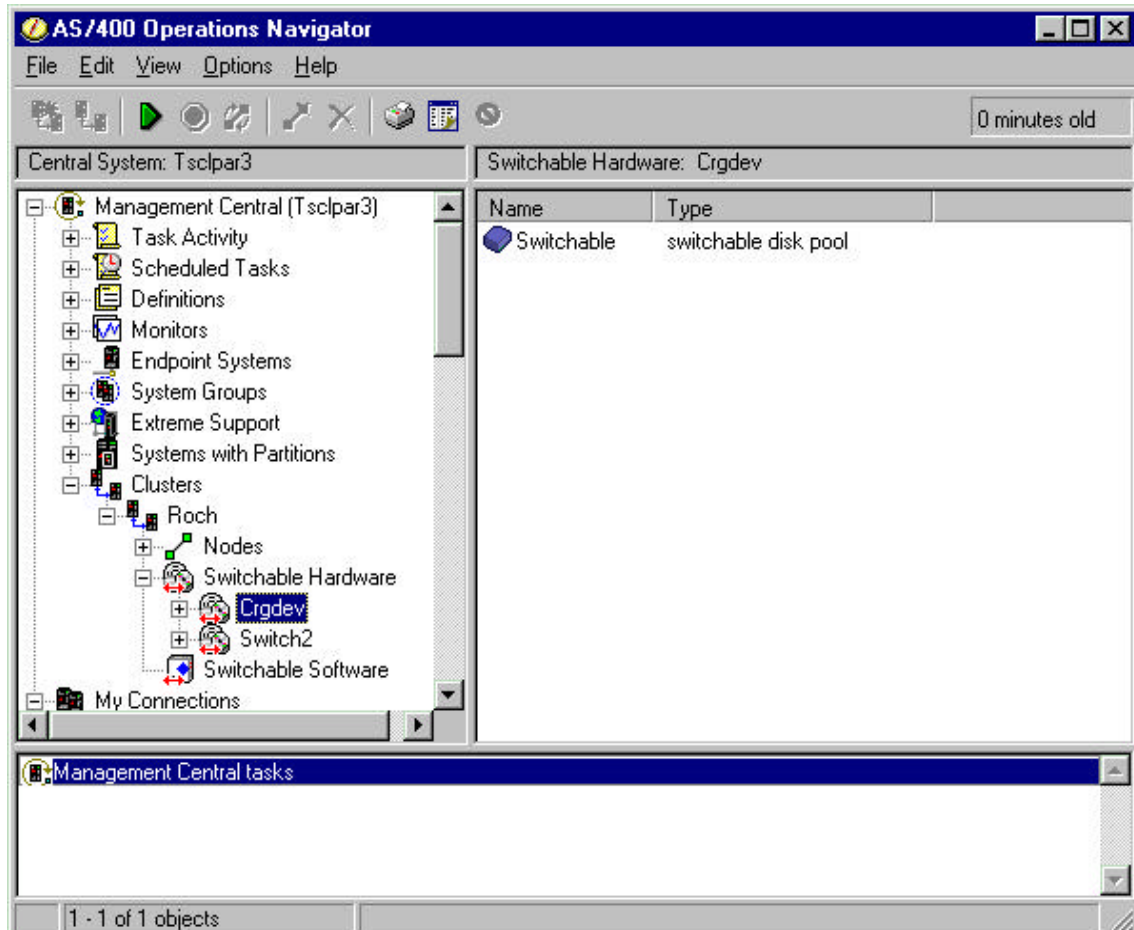
Checking device CRGs with Management Central

Management Central can be used to view the configured device CRGs for a cluster. This can be used to determine which existing device CRG should contain (and control the access to) a Domino server's data directory.



This shows the same "ROCH" cluster after selecting "Switchable Hardware" icon . There are two device CRGs (Cluster Resource Group is shortened to "Group" here). Small "switched disk" icons  represent switchable disk storage. CRGs CRGDEV and SWITCH2 are shown. Both are in the "started" state which means they are active in the cluster. The Node listed is the node that is the current primary for the device CRG (and the cluster node where the disk storage configured by that device CRG can be accessed). Both CRGs have TSCLPAR3 as their current primary node.

Once a decision as to which device CRG is a good candidate for a managing a Domino application CRG data directory, it can be selected to find the device name associated with the device CRG.



This shows the same device CRGs but CRGDEV has been selected. Operations Navigator shows the configured device name for device CRG CRGDEV. In this case it is called "SWITCHABLE". That means there is a device called "SWITCHABLE" defined on all the cluster nodes in the recovery domain of device CRG CRGDEV. There is also a default UDFS called SWITCHABLE that is always automatically mounted in the IFS file system of the cluster node where this device is varied on. If planning to use CRGDEV for a Domino application CRG, this name is important information as the data directory for that Domino server has to be located within this directory. (The server's directory name must have this UDFS name, Switchable configured in the file path name for the server.)

Checking device CRGs with Cluster Management commands

The PRTCRCG Cluster Management command can be used to get information about configured CRGs. Given a cluster name, this command will create a spool file that lists all the CRGs configured for a cluster along with the CRG types.

```

User Space Name: QCSTUSRSPC
User Space Library Name: QUSRSYS
Format Name: CRGLD100
Cluster Name: ROCH
Cluster Resource Group Information
-----
Cluster Resource Services Status: Not Active (Informatic
API may not be consistent with all active nodes in the
Resolution:
  Call this API on an active node in the cluster, if t
  node running the API is not active
  Start Cluster Resource Services on the node running
  API if it is not active
  Call the End Cluster Node and Start Cluster Node API
  the node running the API, if the node is active
Number of CRGs: 3
-----
1. Crg Id:      CRGDEV
   Crg type:    3 (QcstCrgDevResiliency)
   CRG Status: 20 (inactive)
   Primary Node ID: TSCLPAR3
-----
2. Crg Id:      IASPTESTER
   Crg type:    3 (QcstCrgApplResiliency)
   CRG Status: 20 (inactive)
   Primary Node ID: TSCLPAR4
-----
3. Crg Id:      SWITCH2
   Crg type:    0 (QcstCrgDevResiliency)
   CRG Status: 20 (inactive)
   Primary Node ID: TSCLPAR3
-----

```

This shows the output of PRTCRG on the cluster ROCH. We are interested in knowing more information about the device CRG called CRGDEV.

The PRTCRGINF command can be used to see more detailed information about a CRG. Given the cluster and CRG names, this command will list the configured cluster information for the CRG, including its recovery domain.

```

User Space Name:          QCSTUSRSPC
User Space Library Name:  QUSRSYS
Cluster Name:            ROCH
Format Name:             IRG10100
Cluster Resource Group Listed: CRGDEV
-
Cluster Resource Group Information
-----
Retrieved Information Consistency: 1 | Information retriev
Resolution:
  Call this API on an active node in the cluster,
  Start Cluster Resource Services on the node runnd
  Call the End Cluster Node and Start Cluster Node
Cluster Resource Group Type:      3 (Device Resiliency
Cluster Resource Group Status:    20 (Inactive)
Exit Program Name:                *NONE
Exit Program Library:
Exit Program Format Name:
User Profile:
Exit Program Data:
Take Over IP Address:
Job Name:
Allow Application Restart:        0 (Do not attempt to
Number Of Restarts:               0
Configure Takeover IP Address:    0xDD (crs)
Previous CRG Status:              570 (Switch Over Pen
Reporting Node ID:                TSCLPAR3
Text description:
DI User Queue Name:               *NONE
DI User Queue Library:
Number of Nodes:                  2
-----
1. Node Id:      TSCLPAR3
   Current Node Role: 0 |Primary node|
   Membership Status: 0 |Active|
   Preferred Role:   1 |Backup node|
-----
2. Node Id:      TSCLPAR4
   Current Node Role: 1 |Backup node|
   Membership Status: 1 |Inactive|
   Preferred Role:   0 |Primary node|
-----

```

This shows the PRTCRGINF for the device CRGDEV. It has two nodes, TSCLPAR3 and TSCLPAR4 in its recovery domain. The preferred role for TSCLPAR3 is as the backup. For TSCLPAR4 the role is the preferred primary. Although this shows that these nodes are both inactive, the current primary for CRGDEV is TSCLPAR3 and will be the primary node where the device managed by CRGDEV will be available when next started.

A device CRG has to be used for the Domino server's data directory when it is managed as part of an application CRG. If no device CRGs are available or suitable to be used for the server's data directory, a new device CRG will have to be configured before a Domino server's application CRG can be created. If a new device CRG is configured it should be tested before attempting any further steps in the creation of the application CRG that will depend on it. This is especially true if this device CRG has never been used from all the nodes in its recovery domain.

One reason to test the device CRG is to verify that it was created to automatically vary on the switched storage device when the device CRG becomes active and primary on a cluster node. This device CRG attribute is not visible using the Management Central GUI. The PRTCRGINF command with the OPTION(*CFGOBJ) specified, given the cluster and CRG names, will produce a spoolfile that shows this attribute (see the line "Configuration Object Online" in the spool file output below).

```

User Space Name:          QCSTUSRSPC
User Space Library Name:  QUSRSYS2
Cluster Name:            ROCH
Format Name:             LRGI0200
Cluster Resource Group Listed: CRGDEV
-
Cluster Resource Group Information
-----
Retrieved Information Consistency: 1 (Information retrieved)
Resolution:
    Call this API on an active node in the cluster,
    Start Cluster Resource Services on the node running
    Call the End Cluster Node and Start Cluster Node
Cluster Resource Group Type:      3 (Device Resiliency)
Cluster Resource Group Status:    20 (Inactive)
Exit Program Name:                *NONE
Exit Program Library:
Exit Program Format Name:
User Profile:
Exit Program Data:
Take Over IP Address:
Job Name:
Allow Application Restart:        0 (Do not attempt to res
Number Of Restarts:              0
Configure Takeover IP Address:    0x00 (CIS)
Previous CRG Status:             570 (Switch Over Pending)
Reporting Node ID:              TSCLPAR3
Text description:
DI User Queue Name:              *NONE
DI User Queue Library:
Number of Config object entries:  1
-----
1. Object Name: SWITCHABLE
   Configuration Object Type: 1 (Device description)
   Device Type: 1 (Auxiliary storage pool)
   Configuration Object Online: 1 (Is varied on)
-----

```

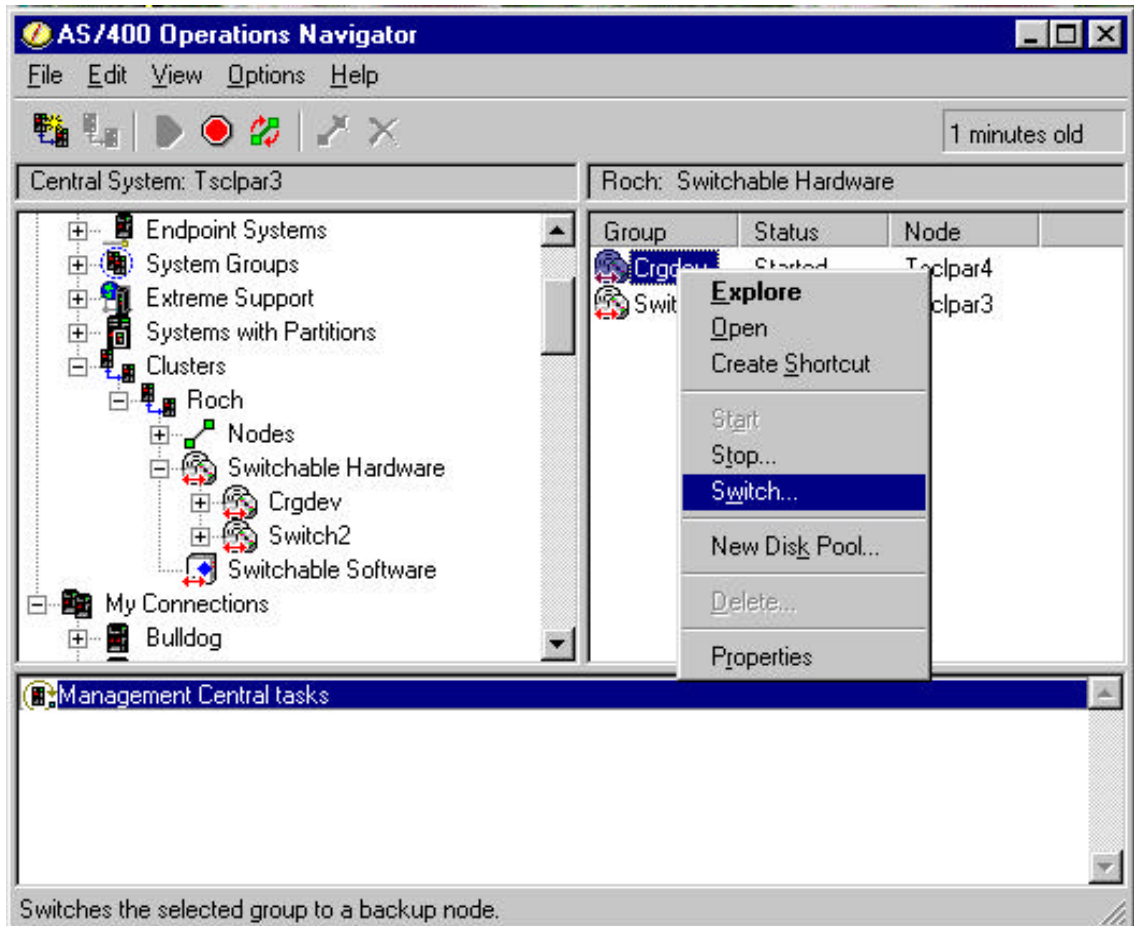
This also shows how PRTCRGINF with the OPTION(*CFGOBJ) specified can be used to display the device object configured for any given device CRG. In this example the device CRG CRGDEV manages the device SWITCHABLE. The last line indicates the device is currently varied on. SWITCHABLE is also the default UDFS directory name that will need to be used in the Domino server's application CRG data directory path name.


Either use the Management Central GUI, clustering APIs, or commands to start the device CRG and then switch to each cluster node specified in its device domain. (If this is the first time these interfaces have been used this will be good practice on how to use them.) As the device CRG is switched between nodes, the device controlled by the device CRG should be varied on and the default UDFS directory should be mounted on each node as it becomes the "current" primary. If this is a new device CRG and switchable disk storage configuration the idea is to make sure that IFS files in the UDFS are readable and writable when that device is


varied on for each cluster node in its recovery domain. This should be attempted at least once before the device CRG and its switchable disk storage has to actually be used by a Domino server to support access to its data directory. Note the device CRG and all the cluster nodes in its recovery domain might have to be started in order to test the CRG.


Switching device CRGs with Management Central

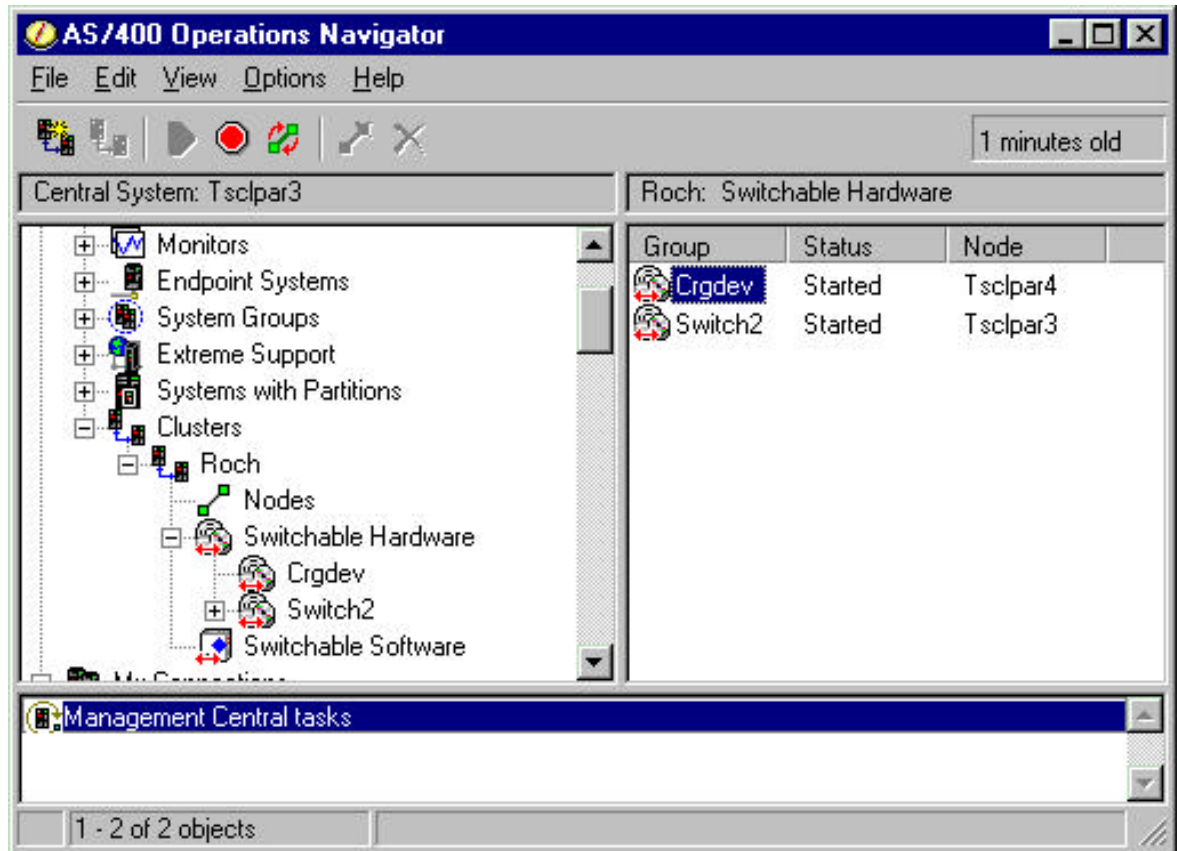
Management Central can be used to test the device CRG by switching it to other nodes in its recovery domain.



This shows the same device CRGs as before. To switch a device CRG to the next node in its recovery domain, right click on the CRG and select "Switch". The "switch" icon on the menu bar  can also be used to switch CRGs. This interface does not show if there are applications that are currently using IFS files that are defined within the switchable storage managed by the CRG. The device CRG will not be able to be switched if any IFS files that are defined in its storage are in use.

Note The "stop sign" icon  can be used to end the device CRG or right click on the CRG and select "Stop". The device associated with a device CRG is not automatically varied off when that device CRG is stopped.

If a CRG is already stopped, it can be started again with the start icon  on the menu bar or by right clicking on the CRG and selecting "Start".



This shows the same device CRGs as before, CRGDEV has been switched and TSCLPAR4 is now its current primary cluster node.

Switching device CRGs with Cluster Management commands

An alternate Cluster Management API and CL command are provided that allow the current primary for a given CRG to be changed. The CHGCRGPRI command, when given the name of the cluster and device CRG, can be used to change the current primary node for the device CRG.

The PRTCRGINF command will produce a spoolfile that will show the current state of the device CRG, given the Cluster and CRG names. The spoolfile output will show the CRG's new current primary node after each switch has been completed using the CHGCRGPRI command.

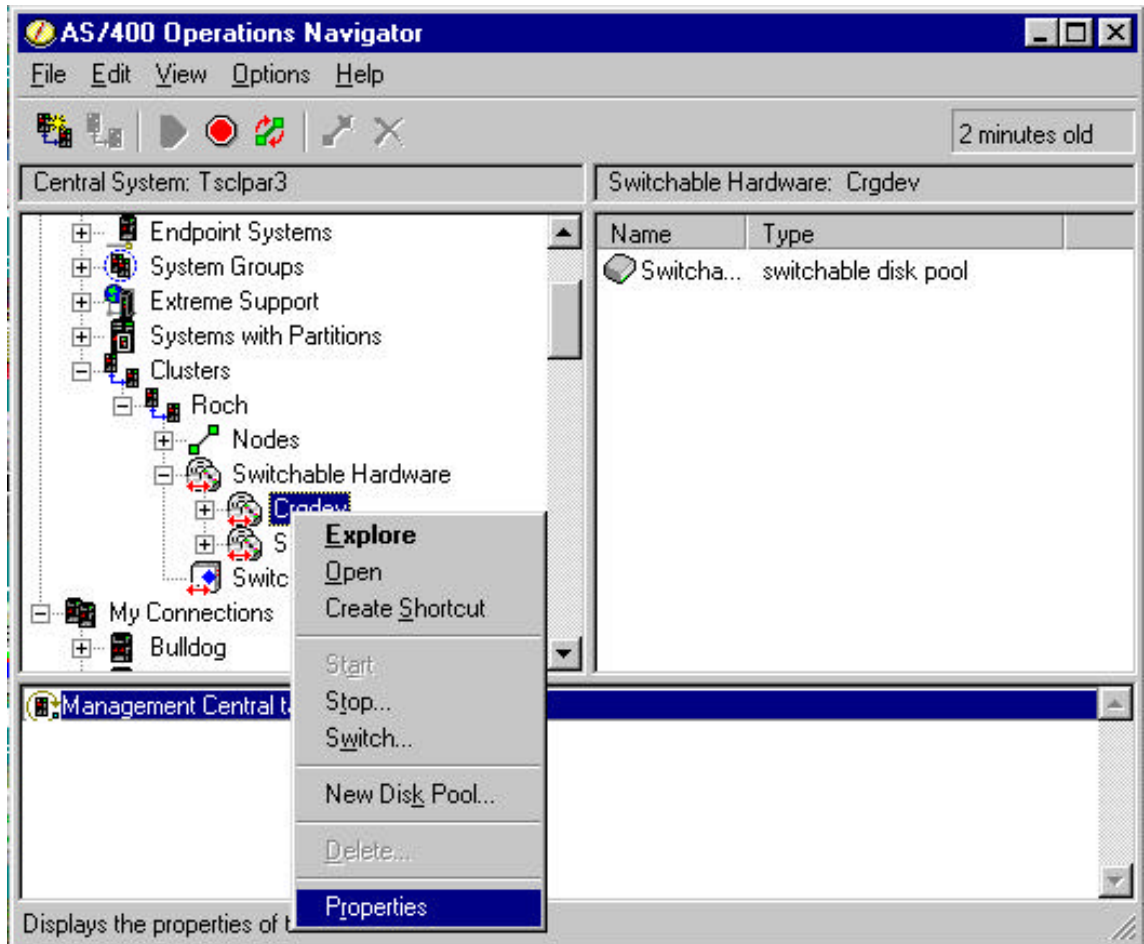
Application CRGs (under the Switchable Software icon for Management Central) can be managed in a similar manner to device CRGs. Managing a Domino server's application CRG is discussed below.

The Device CRG's Recovery Domain

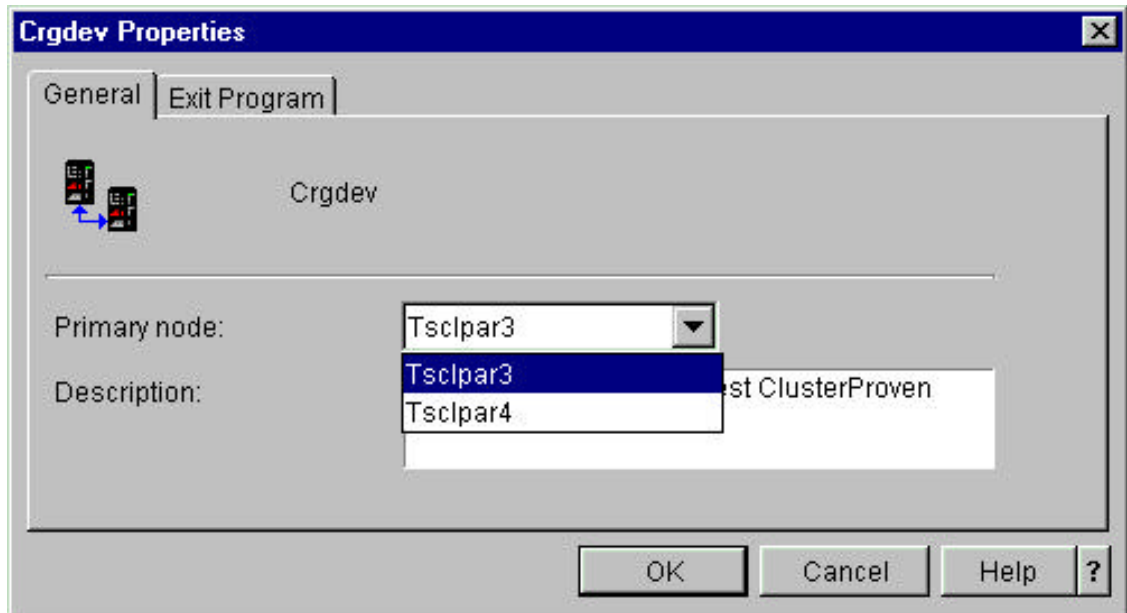
Once a device CRG has been selected for a Domino server's data directory, an important piece of information about that device CRG is its recovery domain. When a Domino application CRG is configured using the CFGDOMCRG command, the name of the device CRG that will manage access to its data directory must be specified (see the DCRG parm on the description of the CFGDOMCRG command). The recovery domain of the device CRG will be used as the recovery domain of the Domino server's new application CRG.

Determining the CRG's recovery domain with Management Central

To determine the recovery domain of a device CRG using Management Central, right click on the device CRG and select "Properties."



The properties of the device CRG will be displayed in a new window.



This shows the "Properties" window for CRGDEV. Selecting the down arrow on the "Primary node:" field will show the nodes in its recovery domain. In this example CRGDEV has two cluster nodes, TSCLPAR3 and TSCLPAR4, in its recovery domain. TSCLPAR3 is the configured primary node and TSCLPAR4 is a back up node.

This field shows the configured preferred primary node of the device CRG as the first node listed. This is not the current primary node for the device CRG which could be any node in its recovery domain. If a Domino application CRG is created with CRGDEV to manage access to its data directory, the application CRG would have the same recovery domain as CRGDEV.

Determining the CRG's recovery domain with Cluster Management commands

Alternatively the Cluster Management API or the PRTCGRINF command can be used to see the recovery domain for a CRG. Given a cluster and CRG name a spoolfile will be produced that will show the recovery domain for the CRG.

Note In this simple example there are only two nodes in the cluster and both nodes appear on the recovery domain of the device CRG. It is possible that a given cluster could consist of many nodes and any subset of these nodes could appear in the recovery domains of the CRGs configured for that cluster.

Unused IP Address Defined as Interface for Every Cluster Node

One powerful function provided to an application CRG by cluster management is the support of a takeover IP address. When the application fails over or is switched over to another node in its device domain, the clustering support does the automatic IP address take over for the IP address configured as part of the application CRG. For Domino servers and its Notes clients this means that server's clients will automatically reconnect to the Domino server when it is started on the new current primary node in the event the application CRG is switched over or failed over. Domino servers that are going to be configured with an application CRG need to have their own IP address (see the TCPOPT of CFGDOMSVR command for more information). This IP address is automatically retrieved from the Domino server's configuration and used to as a parameter when creating its application CRG.

The IP interface might already be active since it may have been configured to be used by the Domino server when the server was last started. If active, the IP interface should be ended before the CFGDOMCRG is done on that Domino server. Use Operations Navigator or ENDTCPIFC command to end the interface.

Note The Domino server should also be ended first when the IP interface is ended.

Normally when an IP address is defined in a TCP/IP interface for the Domino server it is configured to start whenever the system was started. When an IP interface is used for an application CRG, even if it was already being used by a Domino server, it must be set to NOT start when the system is started. If the interface has this attribute then use Operations Navigator or the CHGTCPICF command with AUTOSTART to change the interface to not start (set to *NO) when system is started.

The server's IP address is something that will be used on multiple cluster nodes as the Domino application CRG fails over or is switched over to these nodes. This same IP interface must be configured on all the nodes in the cluster wherever the Domino server application CRG could be started, namely the nodes in a CRG's recovery domain. This is also every node in the recovery domain of the device CRG that will contain the Domino server's data directory. (The recovery domain of the device CRG is used to create the Domino server's recovery domain when it is created with CFGDOMCRG command.)

Note The cluster nodes that are being supported could possibly be different subnets. This kind of configuration requires functions such as virtual IP and "RIP2" capabilities on the subnets routers. This level of IP function and configuration is not covered in this document.

Application CRG User Profile Configured on Every Cluster Node in the Recovery Domain

Part of the configuration required when the Domino application CRG is created requires a user profile (see USRPRF parm description on the CFGDOMCRG command description above). This user profile is used by Cluster Management to submit jobs on each cluster node. These jobs invoke the CRG exit program and pass it the exit program data configured as part of the application CRG. This profile has to exist on every cluster node that will support the Domino server's application CRG.

Note QNOTES cannot be used for this role in the application CRG's function.

A user profile does not necessarily have to be created just to create and manage the Domino server application CRG. An existing user profile can be used or there could possibly be a user profile already defined on every cluster node specifically to be used for CRGs. However the profile must be defined on every cluster node to have *JOBCTL authority and the ability to switch to QNOTES on each node in the cluster to perform the

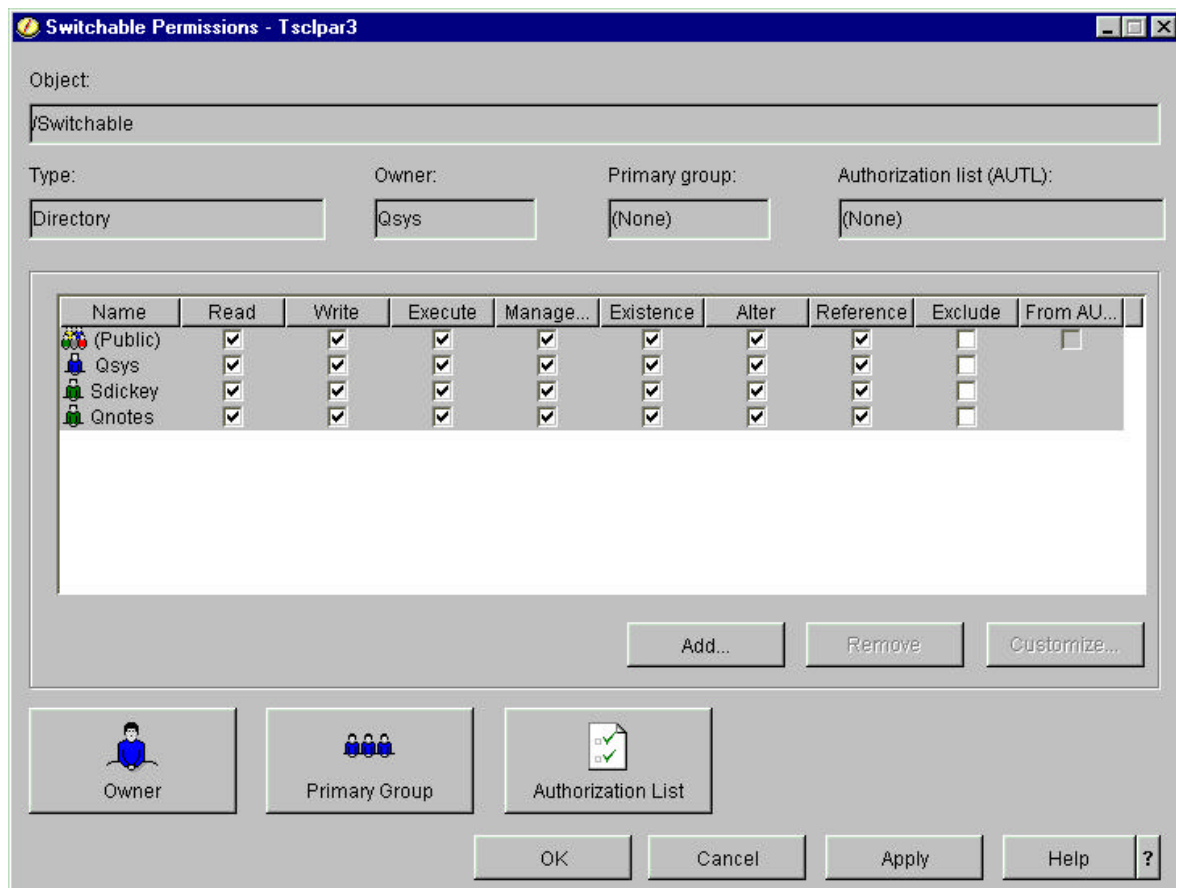
functions of the Domino CRG exit program. Since the exit program will have to create and change some internal Domino configuration objects this profile will also need *ALLOBJ authority on each node in the cluster. Operations Navigator can be used to correctly configure and set the attributes for the user profile on all the systems in the cluster that will be in the Domino server application CRG's recovery domain. Note that the same user profile could be used for multiple CRGs as long as it is correctly configured on all the cluster nodes where it is to be used. This profile also needs to be able to access objects owned by QNOTES. Lastly, this user profile must have at least *USE authority to the STRDOMSVR and ENDDOMSVR commands on all the cluster nodes where it is configured.

Specific Permissions Required to the UDFS of the Switchable Disk Device

The user profile used for the application CRG needs to have the authority to read and write files and to create new directories within the default UDFS for the switchable disk device. (Recall this default UDFS is named the same as the device and is always mounted to the IFS file system on the system where it is varied on.) The UDFS directory authorization should be examined with the Operations Navigator GUI or using the WRKAUT command to view the current directory authorities on the cluster node where the device is currently varied on.

The user profile that is to be configured as part of the Domino server's application CRG should be listed as one of the user profiles that are authorized to read, write, and add directories. If required, the Operations Navigator "Permissions" interface or CHGAUT command can be used to add the appropriate user profile to this list and/or have this profile's authorities adjusted so that it has the appropriate authority to the UDFS directory. Note that this must be done from the system that is the current primary for the device CRG. Once configured for the UDFS on the current primary system, the same permissions will apply to the UDFS directory on all the systems where the device can be varied on.

While this is being done, the QNOTES user profile must also be granted the same authorities to the same UDFS directory. The subsystem jobs that run under the QNOTES user profile will need access to the server's data directory.



This shows the Operations Navigator interface for the permissions to the UDFS Switchable directory. QNOTES is configured to have appropriate access to the UDFS directory "Switchable". (In this case this

authority wasn't technically required as (Public) already had all the authorities required to Switchable on this system.) Operations Navigator can be used to add new profiles as required or adjust the profiles' authorities.

```

                                Work with Authority
Object . . . . . : /SWITCHABLE
Owner . . . . . : QSYS
Primary group . . . . . : *NONE
Authorization list . . . . . : *NONE

Type options, press Enter.
  1=Add user   2=Change user authority   4=Remove user

      Data
Opt  User      Authority  --Object Authorities--
                                Exist  Mgt  Alter  Ref
*PUBLIC  *RWX      X      X      X      X
QSYS     *RWX      X      X      X      X
SDICKEY  *RWX      X      X      X      X
QNOTES   *RWX      X      X      X      X

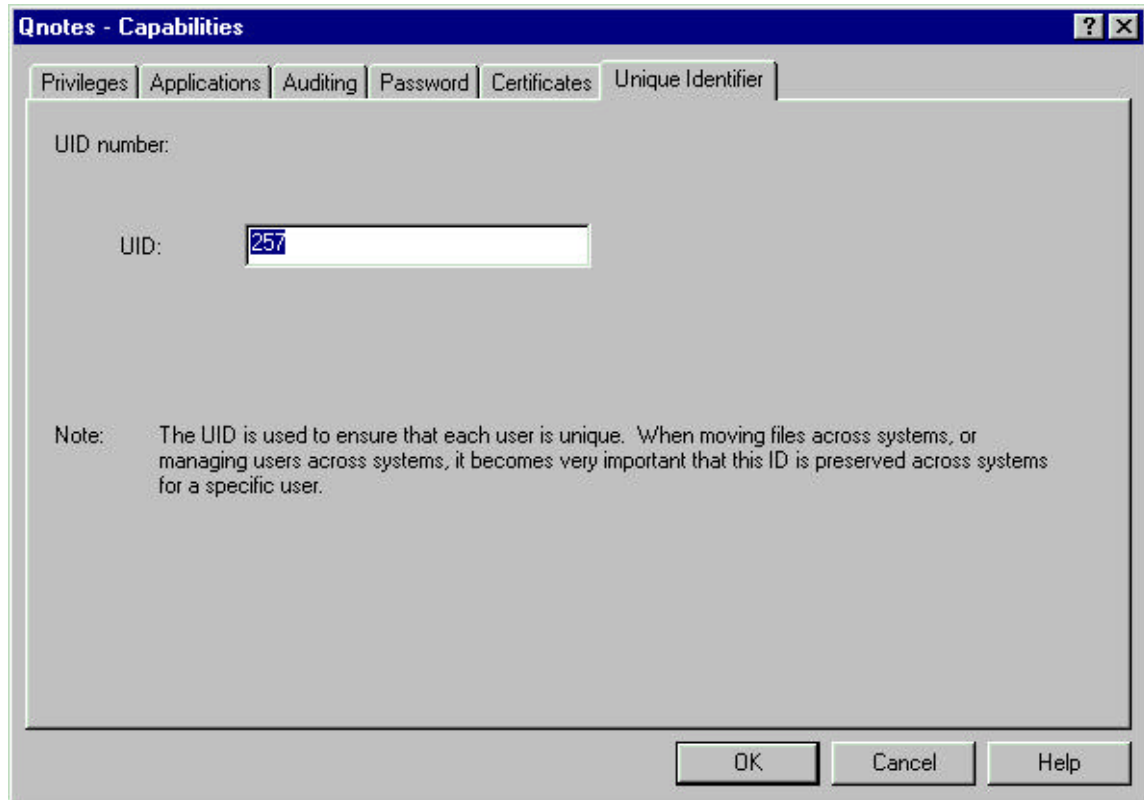
Parameters or command
====>
F3=Exit   F4=Prompt   F5=Refresh   F9=Retrieve
F11=Display detail data authorities   F12=Cancel   F24=More keys
(C) COPYRIGHT IBM CORP. 1980, 2000.

Bottom
```

This shows the screen from the WRKAUT interface for the same UDFS Switchable directory. This interface can be used to add or adjust the profiles' authorities.

QNOTES UID "Normalized" on all the Cluster Nodes

QNOTES will own a large number of IFS files on the switchable disk device as the entire data directory of the Domino server will be kept there. The performance of the function that switches the device (and the access to the IFS file objects owned and accessible by the same user profile on another system) will be much better if the QNOTES user profile user IDs (UIDs) are "normalized" so that they all have the same UID value on every node in the cluster. Operations Navigator or the DSPUSRPRF USRPRF(QNOTES) TYPE(*ALL) command can be used on each node to determine the current UID value for QNOTES.



This shows the QNOTES Capabilities panel under Operations Navigator under the QNOTES Properties. The UID for QNOTES on this system happens to be 257.

Since QNOTES owns IFS directories on every system where Domino for AS/400 is installed, its UID cannot be changed with the CHGUSRPRF command. Part of clustering support provides a program that will set the UID of QNOTES to a given input value. The QNOTES user profile should be set with this provided interface to have the same UID value on all the cluster nodes. Note that the QNOTES user profile cannot be in use when its UID is changed (all Domino servers on that system will have to be ended). Also changing the UID means that every object owned but the QNOTES user profile will be accessed as part of the change. Typically QNOTES owns quite a few objects especially if there are a number of Domino servers configured. Depending on the number of objects QNOTES owns, changing the UID of the QNOTES user profile could take quite a long time.

To set the UID for QNOTES user profile call the program QNNHACHGID with 1 parameter, the 4 byte hexadecimal value of the new UID. For example to change the UID to 1234 (hex 4D2), enter the following command:

```
CALL QNOTES/QNNHACHGID PARM(X'00004D2')
```

The caller of the program must have *ALLOBJ and *SECADM special authorities. The QNOTES USRPRF can not be active or in use when the program is called (that means all Domino servers on that system have to be ended).

If coded in as a CL command, the command source to call the program would look as follows:

```
CMD PARM KWD(UID) TYPE(*INT4) PROMPT(UID)
```

The UID of QNOTES will have to be set to the same value on every system in the device CRG's recovery domain using this method. Once set, the UID will not change unless the QNOTES user profile is deleted and recreated on some cluster node. If the QNOTES user profile is ever deleted and recreated the UID for the new QNOTES user profile should again be set to the same value as it is on the rest of the cluster nodes.

Other Configuration to Check Before Creating a Domino Application CRG

The cluster management function needs to submit batch programs on cluster nodes to allow CRG's exit program to perform the actions that are part of an application CRG's function. OS/400 Cluster Management submits these jobs to the QBATCH JOBQ on each system. The jobs will run under the subsystem that has this JOBQ allocated. Care must be taken to make sure the QBASE subsystem does not have the QBATCH JOBQ allocated. OS/400 is shipped with the "QBASE" specified for the QCTLSBSD system value. If QCTLSBSD is set to "QBASE" it has to be changed or the Domino application CRG will not function correctly. DSPSYSVAL SYSVAL(QCTLSBSD) command can be used to display its current value. If QCTLSBSD is set to "QBASE" one possibility is to change this value to "QCTL" as that will work correctly. The QCTLSBSD system value should be checked on every cluster node and changed if required.

Note that changing the QCTLSBSD value only has an effect the next time the system is IPL'ed. An alternative adjustment is to change the job queue entry for QBASE using the following command:

```
CHGJOBQE SBSDB(QBASE) JOBQ(QBATCH) MAXACT(*NOMAX)
```

As shipped the JOBQ QBATCH would only allow one job to be active at a time and clustering cannot support device CRGs and the Domino server application CRG if this is not changed to *NOMAX.

Creating a Domino Application CRG

Doing the prerequisite checks described above should make the creation of a Domino server application CRG less problematic.

At this point we should have made sure of the following:

- The device CRG can be switched to every node in its recovery domain and accessing it on that domain, we know the cluster nodes in the device CRG's recovery domain are working. We also know the Domino server's data directory should be accessible when its CRG is switched between these nodes since they will be the nodes in the new application CRG's recovery domain.
- IP interface configuration has been done on all the device CRG's cluster nodes (and that IP address is configured for the Domino server that is to be used to create the application CRG).
- User profile that will be configured with the application CRG is correctly configured on each of the device CRG's cluster nodes.
- Access to the UDFS directory has been given to the user profile that will be configured to the application CRG and to QNOTES.
- QNOTES UID has been "normalized" to the same integer value on each of the device CRG's cluster nodes.
- QCTLSBSD is set to an appropriate value for each of the device CRG's cluster nodes.

Having done all this, the series of steps to create a Domino server application CRG are as follows:

1. A new Domino server is created with its data directory created within the default UDFS of a switchable disk device managed by a device CRG. If an existing server is to be used, its data directory has to be copied to the UDFS and the server directory name "path" adjusted to reference its data directory on the UDFS.
2. The server should be started to make sure the data directory is accessible using the configured path. Server applications should be tested if this is a new server.
3. The server has to be ended and its IP interface ended.

4. Create the server's application CRG with the command CFGDOMCRG.
5. Start the server and switch it over to every node in its application CRG's recovery domain. (Also must switch over its device CRG while this is done.) This will configure a "copy" of that Domino server on all the cluster nodes listed in the recovery domain.

Note There are some functional restrictions for Domino servers that are configured by an application CRG. Also if the server has not been backed up and its data directory saved, this should be done before configuring the server as part of an application CRG. "Stand alone" operations (such as a scheduled back up of the server's data directory with BRMS) should not be attempted while the server's application CRG is being created.

There is no Operations Navigator Management Central means to create the Domino server application CRG. The CFGDOMCRG command provided must be used. There are some error conditions tested by the CPP of the CFGDOMCRG command (see below) before a Domino server will be configured as part of an application CRG.

Creating an application CRG for a Domino server means that the Domino server's OS/400 definition will automatically be copied to every cluster node in the application CRG's recovery domain. The new application CRG's recovery domain is copied from the recovery domain of the device CRG specified to manage the server's data directory access on these nodes. The recovery domain of the device CRG can be determined using the methods described previously.

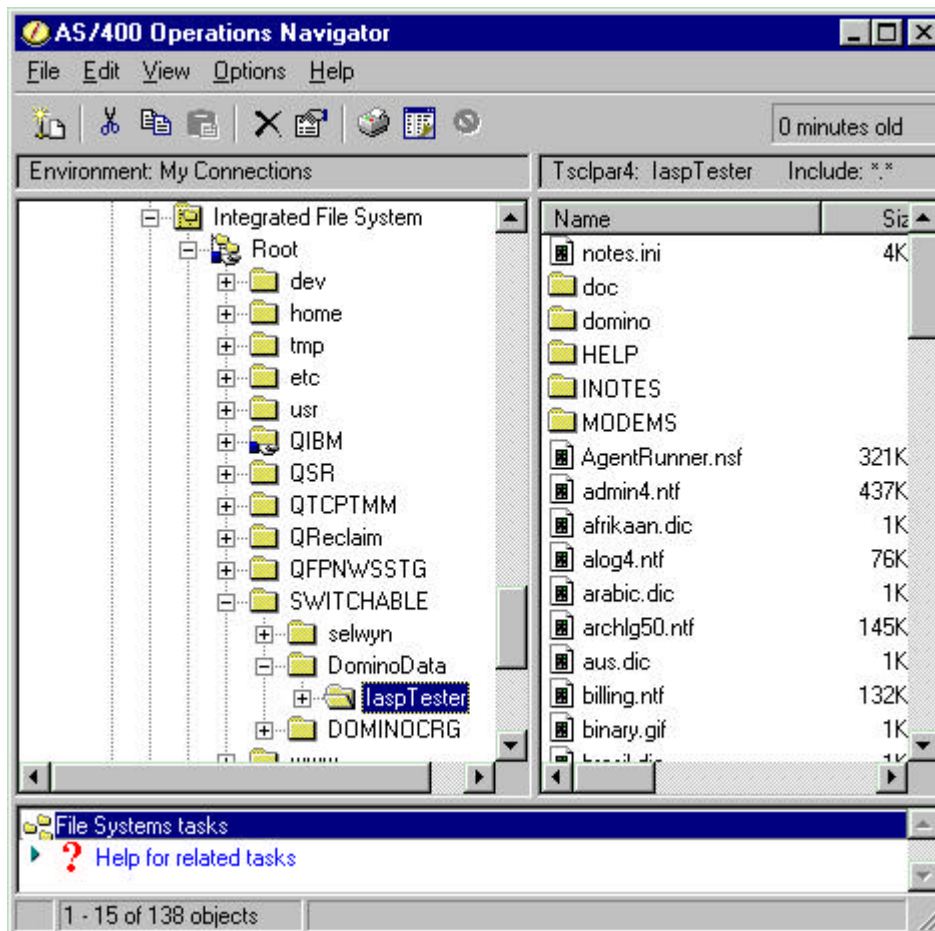
Before using the CFGDOMCRG command to create each new Domino server application CRG, one last series of checks should be made to determine if:

- The device CRG specified must be started in the cluster and the current primary node set to the cluster node where the Domino server is configured. This will also be the system where the CFGDOMCRG command must be issued. The instructions above show how to use Management Central or the Cluster Management commands to find, start, and switch device CRGs if this device CRG is not active against the Domino server's cluster node.
- All the cluster nodes in this device CRG's application domain are active. Use Management Central or the STRCLUNOD clustering commands to start any node not active.
- There is not an existing Domino server with this same name on any of the other cluster nodes. (The Domino server is about to be configured on all the systems in the device CRG's recovery domain as part of creating its application CRG.)

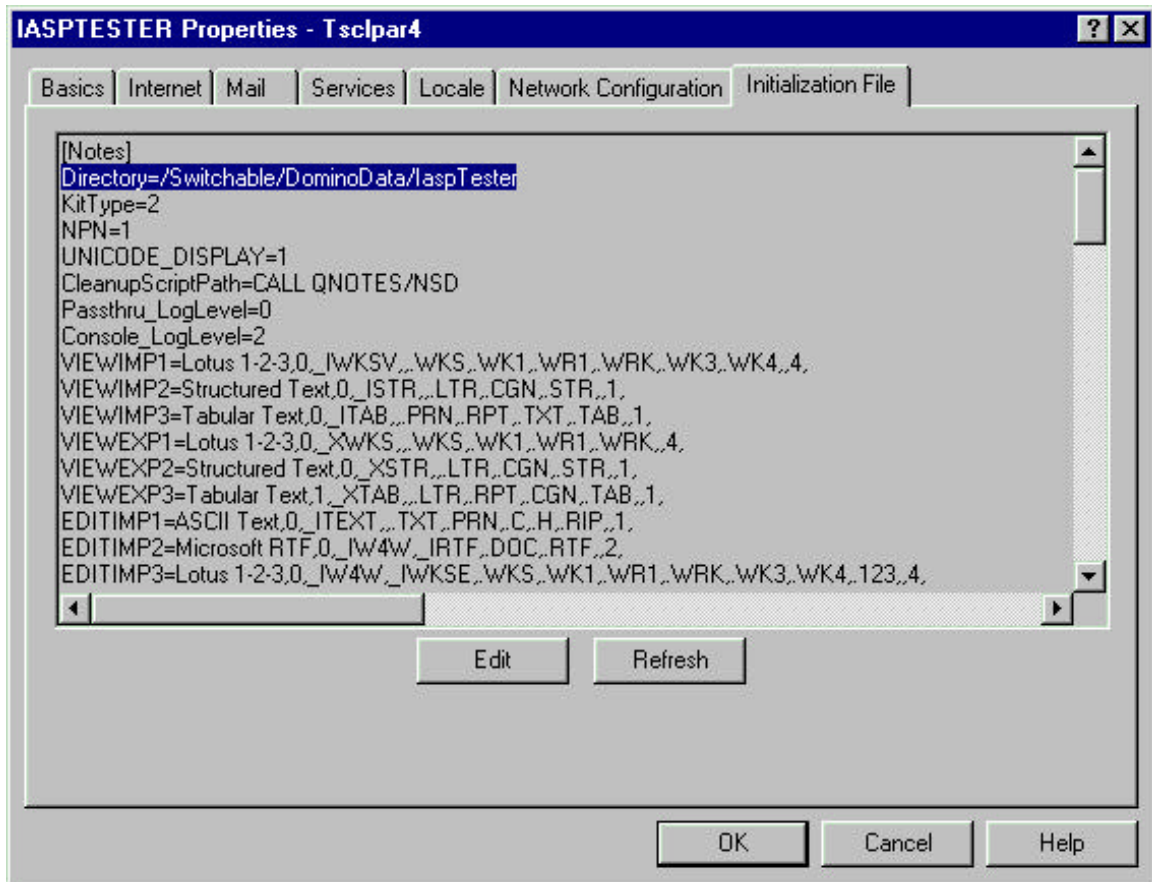
Once these conditions are verified, the next step to creating a Domino server application CRG is to get its data directory configured in the default UDFS defined by the device CRG. Since we've verified that the device CRG is started and the current primary is set to Domino server's node above, once the device that this CRG represents is varied on, the UDFS should be mounted and accessible in the IFS file system.

Getting a Domino server's data directory in the UDFS directory is simple when creating a new Domino server. When the CFGDOMSVR command or Operations Navigator interface is used to create the Domino server, include the UDFS in the path name specified for its new data directory (see the DTADIR parameter on the CFGDOMCRG command for more information). CFGDOMSVR creates the new server's data directory using this path name.

The device and default UDFS always have the same name. In this example they're called "SWITCHABLE". This is also the directory name used in the path name for the data directory when the Domino server was created. CFGDOMSVR creates the Domino server's data directory in the storage managed by the device CRG.



This shows an Operations Navigator view of the UDFS "Switchable". The device associated with the device CRG, called CRGDEV, is active and primary on this system. When the device Switchable is varied on, the default UDFS is mounted within the Integrated File System under the Root directory for that system. In this case, while "SWITCHABLE" was mounted, the Domino server was created with the new server's data directory path specified (DTADIR parm on the CFGDOMSVR command) as "/Switchable/DominoData/IaspTester". The contents of the data directory for this server can be seen in the SWITCHABLE UDFS directory in the right side of the panel.



This shows an Operations Navigator view of the notes.ini file under the Properties as configured for the Domino server IaspTester when it was configured. The highlighted directory path in the .ini includes the UDFS directory "Switchable".

```

Work with Domino Servers
System: TSCLPAR3

Type options, press Enter.
1=Start server      2=Change server    5=Display console  6=End server
7=Submit command   8=Work console     9=Work server jobs
11=Change current  12=Work object links 13=Edit NOTES.INI

Opt  Domino
     Server
     *HTTPSETUP
     IASPTESTER

Subsystem
QDOMINOHT
IASPTESTER

Path
/QIBM/UserData/HTTPSetup/Notes
/Switchable/DominoData/IaspTes

Parameters or command
===>
F3=Exit      F4=Prompt  F5=Refresh  F9=Retrieve  F11=Display status
F12=Cancel   F17=Top    F18=Bottom  F22=Display entire field  F24=More keys

```

The WRKDOMSVR command F11 key will show the data directory path for a server as configured to that system. In this case the server IASPTESTER has a path that includes the UDFS directory Switchable.

Moving the Data Directory for an Existing Domino Server

It may not be possible to create a Domino server so that its data directory is in the default UDFS file system. An existing server will need to move its data directory so that the directory access is managed with a device CRG. To do this, the server's entire data directory has to be copied on to the device CRG and the new data directory path name has to be configured for the Domino server. This section describes how to do this correctly.

Caution Before proceeding with the following instructions, the existing server's data directory and system configuration should be backed up. The Domino server should not be active and no "stand alone" Domino functions should be run against the server while this procedure is being followed.

The major steps involved are as follows:

1. Make sure the device CRG is started and that it is switched so that its current primary is the node where the data directory is to be copied.
2. Copy the existing Domino server's data directory to the device CRG's IASP managed device. This has to be within the default UDFS directory for that device. In this example the UDFS name is "Switchable". The Domino server's data directory name is /DominoData/IaspTester .

The copy command for this example is:

```
CPY OBJ('/DominoData/IaspTester/*') SYMLNK(*YES)
TODIR('/Switchable/DominoData/IaspTester/') SUBTREE(*ALL) OWNER(*KEEP)
```

Note SYMLNK(*YES), SUBTREE(*ALL), and OWNER(*KEEP) must be specified on the copy as all the symbolic links and sub directories are to be copied. The owner of all the data directory objects must remain as QNOTES in the copy.

3. Change the server's internal OS/400 configuration path names to reference the copied data directory on the device.

To change the data directory of the server, IaspTester, to /Switchable/DominoData/IaspTester, call the program QNOTES/QNNHACHIR as follows:

```
CALL QNOTES/QNNHACHDIR PARM(IaspTester '/Switchable/DominoData/IaspTester')
```

*ALLOBJ authority is needed by the caller of the program. Calling this program will result in changing where the Domino administration commands find the servers data directory. It will not change where the Domino server itself finds the data directory. This will require modifying the notes.ini, located in the new data directory, and changing every reference to the old data directory path name to the new data directory path name in the notes.ini.

To verify that the program call changed the data directory, use the WRKDOMSVR command. Use function key PF11 to view the data directory of the servers. Verify that the directory displayed is the new directory name entered on the program call to QNOTES/QNNHACHDIR.

4. Change the server's notes.ini references to any data directory object to use the new configuration path name. Change EVERY reference to any object with a path name "/DominoData/IaspTester" to have the path name "/Switchable/DominoData/IaspTester" in this example. There is a "Directory=" line and also at least two other .ini references to .id files that have complete path names that have to be updated: look for "KeyFilename=" and "CertifierIDfile=" as they have to be changed. There could be more directory path name references depending on what is configured in the .ini for the server.
5. Start up the Domino server to verify that it can start and use the copied data directory and the path names to that directory have been correctly revised.
6. <Optionally> Remove the original copy of the server's data directory.

Using the CFGDOMSVR Command

When the Domino server's data directory is located in the UDFS directory and the device CRG has its current primary set to the node where the Domino server has been created, the server's application CRG can be created. The CFGDOMCRG command has to be used on the system where the Domino server has been created.

The command is used as follows:

```
CFGDOMCRG CRG(<name of CRG to create>)  
SERVER(<your server name>) USRPRF(<User profile to submit CRG exit job>)  
DCRG(<device CRG that has SERVER's data directory configured>)  
TEXT(<(Optional) 50 character description of application CRG.>)
```

(Also see description of CFGDOMCRG command and its parameters above.)

The CRG name specified cannot be the name of an existing CRG for the cluster. The CRG can have the same name as the Domino server it configures but can only be 10 characters long and will be uppercased when the CRG is created.

The SERVER has to be a Domino server and configured on the system where the command is run.

The USRPRF has to exist on the system where the command is run as well as every system in the device CRG's device domain.

The device CRG specified by DCRG has to be active and primary on the system where the command is run. This device CRG will be checked to make sure the data directory of the SERVER is located there. If it is, the device CRG's recovery domain is copied and used to set the recovery domain for the application CRG about to be created.

The CFGDOMCRG command will call the command's CPP to internally build the parameters required to create the application CRG. The command program calls the OS/400 Cluster Management QcstCreateClusterResourceGroup system API to create the CRG. The command processing program then waits on an internal message queue where this Cluster Management API will return a message that the application CRG was either created or failed to create. The CPP will not return to the caller of the CFGDOMCRG command until a message from Cluster Management is posted to the internal message queue. (For more information on the QcstCreateClusterResourceGroup system API see the reference for OS/400 Cluster Management).

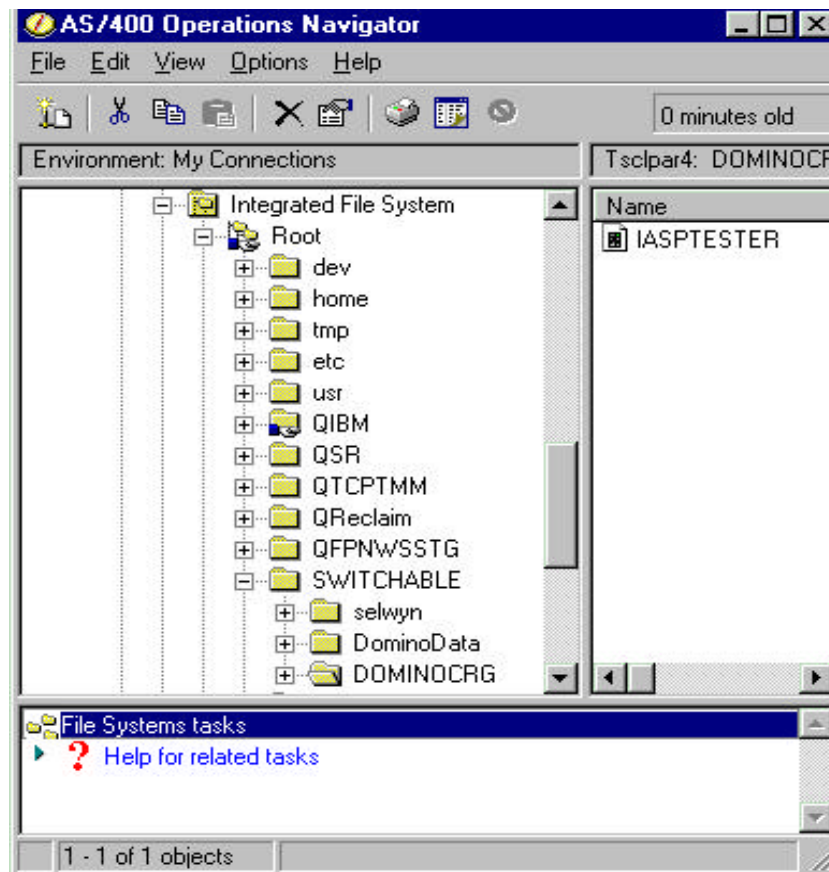
What Happens While the Application CRG is Being Created

The CFGDOMCRG command passes the CRG user profile name and the Domino application CRG exit program name, QNOTES/QNNHACRG as parameter values when it invokes the QcstCreateClusterResourceGroup system API. It also sets some CRG exit program data that includes the SERVER name and the DCRG name. Calling the QcstCreateClusterResourceGroup system API activates the OS/400 Cluster Management function that starts to create the application CRG on all the cluster nodes. Providing that the application CRG can be created, Cluster Management will submit a batch job under the CRG user profile that will call the Domino CRG exit program on every node in the recovery domain of the application CRG being created. This Domino exit program will be passed the exit program data that contains the SERVER and the device CRG names. This batch job is only submitted by cluster management on nodes that are part of the new application CRG's recovery domain. This is true even though the application CRG object is actually created on all the nodes in the cluster, not just those nodes in the recovery domain. These batch jobs run asynchronously. How fast they are submitted and executed on the individual cluster nodes depends entirely on the current conditions of each cluster node.

The Domino exit program will make sure that there is no server currently configured on its node with the SERVER name and if there is none, creates some internal configuration that will "reserve" the SERVER's name on each node. The exit does NOT create the Domino server on its cluster node at this time. There is configuration information about the server (data directory name, subsystem name, etc.) that is not known to every node in that cluster.

If the creation of the new application CRG fails on any cluster node in its recovery domain for any reason (i.e. Cluster Management has an error like the CRG user profile was not found or the Domino server exit fails if there is already a server with the same name configured on that node) the creation of the CRG fails entirely on all the nodes in the cluster. Any configuration information about the Domino server that might have been made by the Domino exit program are removed from all the cluster nodes in the recovery domain.

The one node in the cluster that does know all the information required to create a Domino server is the node where the CFGDOMCRG command was run. The Domino server had to exist on that node when the CFGDOMCRG command was invoked. When the Domino CRG exit job runs on that specific cluster node, it gets the information about the Domino server configured and creates an object in a directory called DOMINOCRГ that is under the default UDFS directory.



This shows an Operations Navigator view of the DOMINOCRГ directory under the Switchable UDFS. The Domino application CRG exit program has created a file called IASPTESTER (named after the application CRG being created). This file contains the information about how to configure a copy of the Domino server that will eventually be configured on all the nodes in the application CRG's recovery domain.

Create the application CRG

The following examples show information about a Domino server application CRG that was created with the following command:

CFGDOMCRG

In this use, CFGDOMCRG specified that the application CRG "IASPTESTER" was to be created for Domino SERVER "Iasptester". (They don't need to be the same name.) The user profile this application CRG exits will run on is "HAROLD". The device CRG (DCRG) that controls the access to server "Iasptester's" data directory is called "CRGDEV".

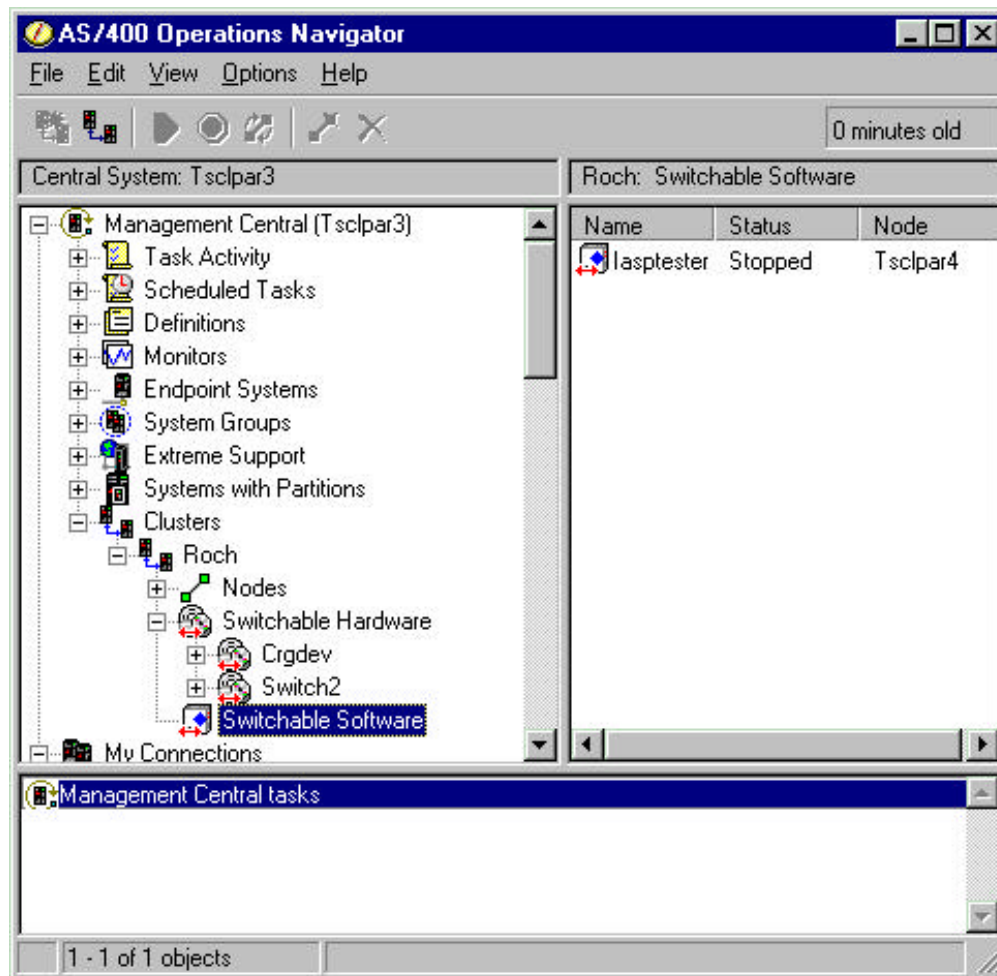
Label	Type	Value	Description
Domino application CRG name:	CRG	IASPTESTER	Name
Domino server name:	SERVER	Iasptester	Character value
Domino app. CRG user profile:	USRPRF	HAROLD	Name
Data directory device CRG name:	DCRG	CRGDEV	Name
Text description:	TEXT	*NONE	Character value

Verifying the application CRG was created correctly

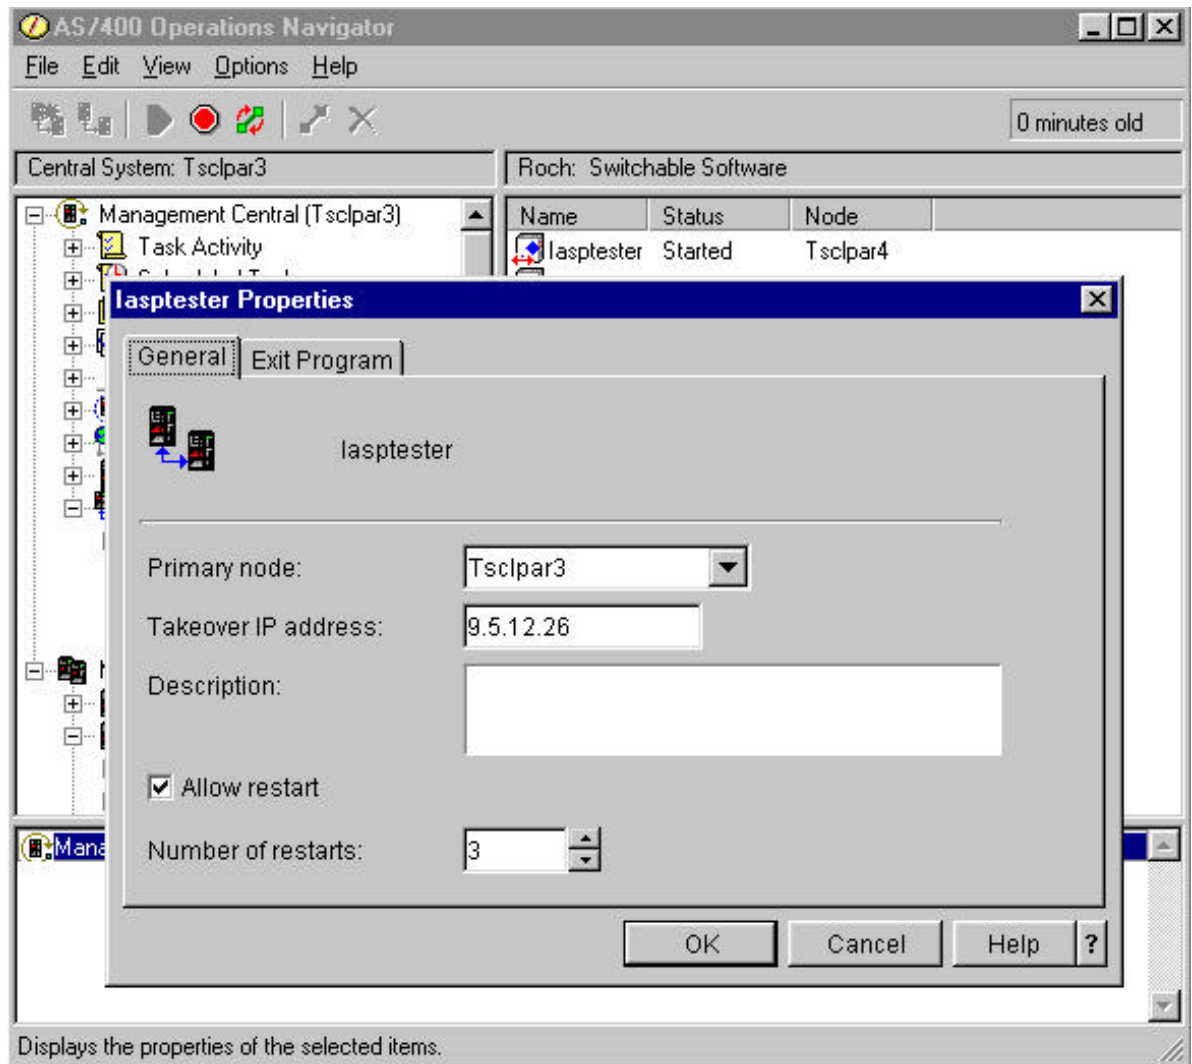
When the CFGDOMCRG command returns it will indicate if the Domino server's application CRG was created correctly. Note that the actual creation of the CRG and the Domino CRG exit program happens asynchronously on the nodes in the cluster. The CPP will not return to the caller of the CFGDOMCRG command until some indication has been returned from Cluster Management that the CRG failed to create on at least one node or successfully created on all the nodes. The CPP for the CFGDOMCRG command waits for a message on an internal message queue signalled to that queue from the OS/400 Cluster Management function. The message will indicate if the CRG was successfully created on all the cluster nodes. When the CPP gets that message it will return to the caller. A message, CPCBB01 (Cluster Resource Services API QcstCreateClusterResourceGroup completed in the job log will indicate that the QcstCreateClusterResourceGroup API completed. A return code in this message indicates if the application CRG was successfully created.

Determining information about an application CRG using Management Central

To see information about a Domino server's application CRG the Operations Navigator Management Central interface can be used. This interface can be used to manage the application CRG (see documentation below).



This shows the same "Roch" cluster example as before after selecting "Switchable Software" icon. Application CRGs defined for the cluster, including those associated with Domino servers, will appear in the list on the right hand side of the panel. The application CRG in this example is called Iasptester. (It just happens to be the same name as the Domino server it configures.) It is currently stopped. The current primary node is Tsc1par4. The "stopped" state means the CRG is not active in the cluster. Successfully creating a CRG does not automatically start the CRG. It has to be started as a separate step. Application CRGs can be managed using this interface in a similar manner to how device CRGs are managed. More information about an application CRG can be displayed by right clicking on the name and then selecting "Properties".



This shows the "Properties" window for the application CRG IASPTESTER. Note that the takeover IP address taken from the Domino server's configuration is shown here. The "Exit Program" tab will show information about the configured exit program and exit program data. The exit program data will have the Domino server's name in the front of the data field and the device CRG that controls access to that server's data directory at the end of the data field.

Caution Do not change any of the exit program information using this interface. It would likely disable the application CRG so that it can no longer be used.

Determining information about an application CRG using OS/400 Cluster Management CL commands

The PRTCRCG command can be used to find out information on all CRGs known to the cluster given the cluster's name. All the configured CRGs, including application CRGs, will appear in a spoolfile generated by this command. PRTCRCGINF can be used to display specific information about the application CRG. This command will generate a spoolfile with the specific CRG's cluster configuration information.

```

User Space Name: QCSTUSRSPC
User Space Library Name: QUSRSYS
Format Name: CRGLD100
Cluster Name: ROCH
Cluster Resource Group Information
-----
Cluster Resource Services Status: Not Active (Informatic
API may not be consistent with all active nodes in the
Resolution:
  Call this API on an active node in the cluster, if t
  node running the API is not active
  Start Cluster Resource Services on the node running
  API if it is not active
  Call the End Cluster Node and Start Cluster Node API
  the node running the API, if the node is active
Number of CRGs: 3
-----
1. Crg Id:          CRGDEV
   Crg type:        3 (QcstCrgDevResiliency)
   CRG Status:      20 (inactive)
   Primary Node ID: TSCLPAR3
-----
2. Crg Id:          IASPTESTER
   Crg type:        3 (QcstCrgApplResiliency)
   CRG Status:      20 (inactive)
   Primary Node ID: TSCLPAR4
-----
3. Crg Id:          SWITCH2
   Crg type:        0 (QcstCrgDevResiliency)
   CRG Status:      20 (inactive)
   Primary Node ID: TSCLPAR3
-----

```

This is a spoolfile generated by the PRTCRG command for this example. The Domino application CRG IASPTESTER appears as the second CRG in this listing.

```

User Space Name:          QCSTUSRSPC
User Space Library Name:  QUSRSYS
Cluster Name:            ROCH
Format Name:             LRGID100
Cluster Resource Group Listed: IASPTESTER
-
Cluster Resource Group Information
-----
Retrieved Information Consistency: 1 (Information retrieved
Resolution:
  Call this API on an active node in the cluster,
  Start Cluster Resource Services on the node running
  Call the End Cluster Node and Start Cluster Node
Cluster Resource Group Type: 2 (Application Resiliency)
Cluster Resource Group Status: 20 (Inactive)
Exit Program Name:         QNNFACRG
Exit Program Library:      QNOTES
Exit Program Format Name:   EXTPO100
User Profile:              HARCLD
Exit Program Data:         IASPTESTER
Take Over IP Address:      9.5.12.26
Job Name:                  IASPTESTER
Allow Application Restart: 1 (Attempt to restart th
Number Of Restarts:        3
Configure Takeover IP Address: 0x01 (user)
Previous CRG Status:       570 (Switch Over Pending)
Reporting Node ID:        TSCLPAR3
Text description:
DI User Queue Name:       *NONE
DI User Queue Library:
Number of Nodes:          2
-----
1. Node Id:          TSCLPAR4
   Current Node Role: 0 (Primary node)
   Membership Status: 1 (Inactive)
   Preferred Role:    1 (Backup node)
-----
2. Node Id:          TSCLPAR3
   Current Node Role: 1 (Backup node)
   Membership Status: 0 (Active)
   Preferred Role:    0 (Primary node)

```

This is information about the Domino server application CRG IASPTESTER that was printed with the PRTCRGINF command. It shows the exit program name that was set in the CFGDOMCRG CPP (QNOTES/QNNHACRG), the user profile associated with the CRG (HAROLD), the takeover IP address that was taken directly from the Domino server configuration, and the recovery domain that was set directly from the recovery domain of the device CRG called CRGDEV.

Note The Domino server's name will always appear in the Exit Program Data field.

In this case the Domino server's name is "IaspTester". Not shown on this cut out segment of the spoolfile is that the device CRG's name that controls access to that Domino server's data directory will also appear at the end of the Exit Program Data field. In this example the "Text description:" field for the CRG has not been set but that description could be set using the TEXT parm on the CFGDOMCRG or CHGDOMCRG commands.

Once an application CRG has been configured for a Domino server, the server should not be removed using the Operations Navigator interface or the CFGDOMSVR OPTION(*REMOVE) CL command on any of the cluster nodes where that server is configured. The Domino server's application CRG has to be removed first. (See the section on Deleting a Domino Server's Application CRG.)

The Configured Primary of the New Application CRG vs. the Current Primary of the Device CRG

Even though the Domino Server's new application CRG was created with the exact same recovery domain as that of the device CRG that manages its data directory, there can still be a mismatch as to where the Domino application CRG will initially be started. The reason is that the current primary node of the device CRG is not necessarily the preferred (configured) primary. For any CRG there is the configured, preferred primary node. The Domino server's application CRG copies the recovery domain of the device CRG so it will start the very first time on the device CRG's preferred primary node. That node will then be the first current primary for the Domino server's application CRG. The device CRG had to be active on the node where the CFGDOMCRG command was run but it might not be currently active on its configured preferred primary node.

This concept that an application CRG can have a current primary node (and be trying to start a Domino server on that node) that is not the same current primary node of the device CRG that manages access to its data directory, is an important one to understand as it will eventually be encountered when managing Domino application CRGs within a cluster. The current primary node of a device CRG is controlled independently of the current primary of application CRG that might depend on that device CRG. The device CRG could have a different current primary node. The Domino application server CRG exit will still attempt to start on its new current primary but will not be able to STRDOMSVR since it cannot access the server's data directory from that node. Once the device CRG primary is also active and primary on that same node then the Domino exit program can attempt the STRDOMSVR command to start the Domino server. The Domino application CRG exit program detects and handles this situation.

To better explain this concept of "independent" primary nodes consider that this situation happens for a moment every time a Domino application CRG is switched to a new current primary. An active Domino server that is using IFS files defined in an IASP prevents that IASP device from being varied off on a system. That means that if a Domino server is started and using its data directory, the device CRG that controls access to that data directory cannot be switched to a new current primary node. The Domino server application CRG (and any other application that references any IFS files) has to be switched first so that server will stop using IFS files on the switchable device. Otherwise the device cannot be varied off and the device CRG cannot be switched.

To get out of this dilemma the Domino server's application CRG is switched first. This means the server is ended (ENDDOMSVR) on its current primary and then will be restarted on its new primary node. As this is happening with the Domino application CRG, there will almost always be some further delay in switching the device CRG the server depends on to the same new primary node. While the device is in the process of switching over however the Domino application CRG has already been started in that node and the Domino CRG exit program has to handle this situation. The Domino CRG exit cannot do the STRDOMSVR (it would fail to start without access to its data directory). This situation eventually should resolve itself - the device CRG is eventually switched over to the same cluster node and the device varied on to that node. Once that happens the Domino server can be started on that node. Sometime this resolution might be delayed for a few seconds but could be much longer depending on how long it takes to first vary off the device, then switch it

over and then vary on the device associated with the device CRG on the new node. Only when that has happened will the server's data directory be available and only then can the Domino CRG exit STRDOMSVR to get the Domino server started. (More on this situation is documented in the "Managing a Domino Application CRG" section below.)

A similar situation might occur if a Domino server's application CRG is started and the device CRG that it depends on is not started at all. Note that ending the device CRG does not automatically vary off the switched storage device it controls. It is possible the device CRG is not active at all in the cluster but the Domino server application CRG is active and can access its data directory on its current primary. A switch over of the server's CRG in this case would result in the Domino server having a new current primary node but it could not access its data directory on that new node. The device CRG was not active at all. CRGs cannot be switched over (and be assigned a new primary) if not active. The device CRG has to be started eventually and then it would be possible to change its current primary node to have the same current primary as the Domino application CRG.

Managing the relative current primary nodes of Domino application CRGs and the device CRGs on which they depend is a skill that operators that support a cluster with Domino application CRGs must develop. The Domino application CRG exit program QNOTES/QNNHACRG is built to attempt to help handle the situations where a Domino application CRG is active and primary on a cluster node where the device CRG on which it depends is not. How this can be done is explained below in the "Managing a Domino Application CRG" section below.

How does the application CRG's Domino server get configured on the other nodes in its recovery domain?

The same Domino server has to be automatically configured on every node in the application CRG's recovery domain. There is not enough information about the Domino server that is known to create the servers at the time the CRG is created on the other cluster nodes. In addition, since the create of the CRG is asynchronous, it is possible for the create of the CRG to succeed on some nodes and not on other nodes. If the create fails on any node then anything that might have been added to the other nodes in the recovery domain due to the create CRG process would have to be removed. Since the Domino server isn't fully configured until the CRG is started, the server's configuration doesn't have to be removed should the CRG fail to create.

The Domino server is completely configured on each cluster node in the recovery domain as the application CRG is started for the first time on that cluster node. There is some internal information that is kept when the application CRG is first created on that node, but when the application CRG is started as a primary node, the device CRG also becomes primary on the same node. This allows the Domino exit program access to information about the Domino server associated with the new server that is stored in the directory in the default UDFS for the device to become available to the exit program. With this additional information accessible, the exit program can completely configure the Domino server on that node prior to the exit starting up the server (STRDOMSVR) for the first time. Until the exit program successfully starts the application CRG for the first time the Domino server will not be visible from the Operations Navigator interface that shows the configured Domino server on that node or using the WRKDOMSVR interface.

When a domino server is configured on the other cluster nodes, it is given a subsystem name to be used for the server's jobs when the server is started on the node. The same subsystem name configured for the server on the original node (when it was configured with CFGDOMSVR) will be used on the other cluster nodes. There is a possibility however there is already a subsystem with that same name already defined on the cluster node at the time the server is being created (the first time it is started up on that node). If that happens, a default subsystem name will be used (i.e. a subsystem name like "DOMINOnn" where nn is a number between 01 and 99) to configure the Domino server's subsystem on that node. The copies of the Domino server might not all end up with the same subsystem name on every node in the recovery domain where the server gets configured.

Trouble Shooting Domino Server Application CRG Creation

Some of the recommended checking that should be done up front will reduce the chance of something going wrong when a Domino server's application CRG is created. The most crucial pre-checks are to make sure the device CRG has the Domino server's data directory on it, that the device CRG is active, that the cluster nodes

in the device CRG's recovery domain are all active in the cluster, and the device CRG's current primary node is set to the cluster node where the CFGDOMCRG command was run.

Errors that are likely to occur when a Domino server's application CRG is created are divisible into two major categories. Diagnosing and recovering from the two categories of error is different. The first step is to attempt to figure out which category of error has occurred. There should be an error message that followed the CFGDOMCRG command.

If this error message is an LNT1210 escape message, then the error has something to do with the checks made by the CFGDOMCRG CPP program. This message should have a return code that should help determine why the CPP would not create the CRG.

The program has to access information about the Domino server that was passed in with the SERVER parameter so the first thing it does is to find the server's configuration on that system. If the Domino server specified by the SERVER parm is not found on that system the LNT1210 error code will indicate the server wasn't found.

It makes sure SERVER is a Domino server (not HTTP, Quickplace, Sametime, or other server type)

It makes sure SERVER is not already defined as part of an application CRG. The application CRG could be deleted and then a new CRG created.

It attempts to find the configured IP address of the server so that the address can be used to create the application CRG. It is possible that the server wasn't configured with a specific IP address and the server must be changed (can use CHGDOMSVR command) before it can be managed as part of an application CRG.

It makes sure the server is a "normal" Domino server and it attempts to locate the server's data directory in the UDFS directory of the device that is managed by the device DCRG. There could be multiple reasons why this might not work. The UDFS directory should be examined to make sure the data directory for the server is there. The path configured for that Domino server should be the path on the UDFS where this directory is located (use WRKDOMSVR function key PF11 to see the configured directory name for that server.)

The device CRG has to be started and should be active and primary on the node. Use the Management Central GUI or the cluster management STRCRG and CHGCRGPRI commands to start and switch the current primary node for the device CRG.

If the error message is an LNT0998 escape message then the error falls into a different category of configuration failure. All the above information was checked by the CPP and was found to be correct. It is likely the problem instead has something to do with the execution of the QcstCreateClusterResourceGroup Cluster Management API. The CPP builds the parameters for this command and calls this API in an attempt to create the application CRG. There should be an error message (or multiple error messages) in the job log ahead of the LNT0998. This message will help diagnose the problem and suggest a recovery action but it might not be obvious what is wrong. Creating a CRG in a cluster is an asynchronous process that takes place on all the active cluster nodes. The Cluster Management documentation regarding creating application CRGs with the QcstCreateClusterResourceGroup Cluster Management API might be helpful in determining why the CRG failed to create. If the creation fails on any of the nodes in the cluster the CRG will not be created. There could be problems with the ability to create the CRG. For example there might already be a CRG with the same name defined in the cluster or the user profile specified by the USRPRF parameter was not found on at least one cluster node. Another possibility is that the Domino exit program found a Domino server with the same name already configured on at least one cluster node.

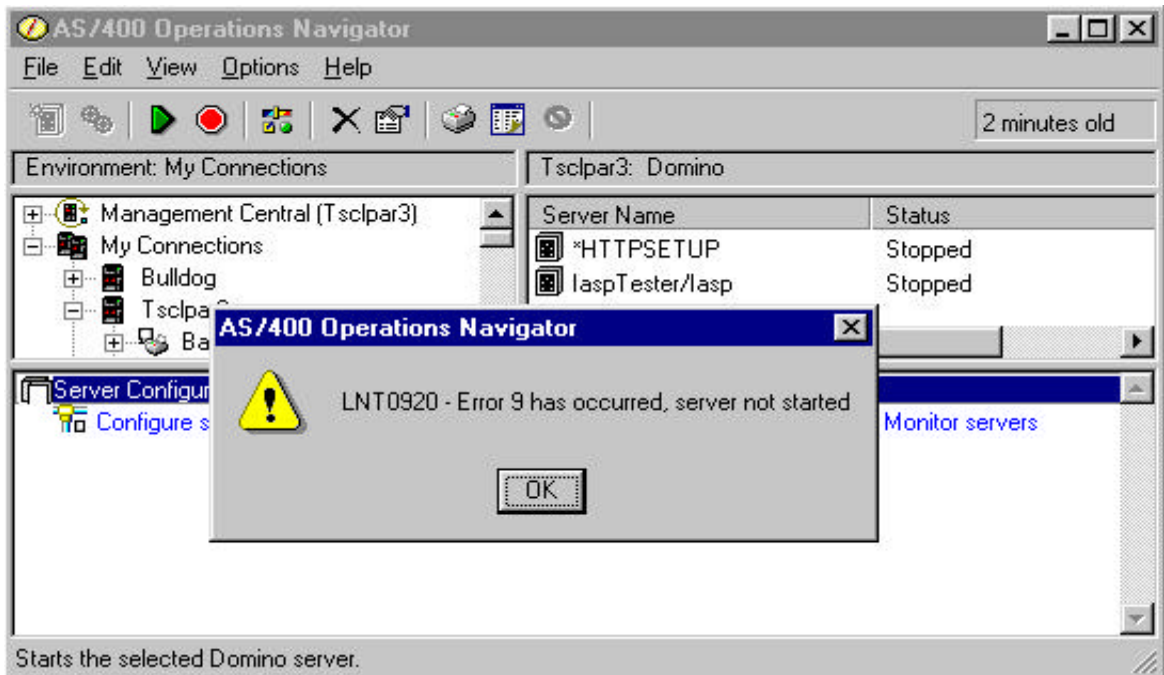
Managing a Domino Application CRG

Domino application CRGs and the device CRGs they are associated with can be started, stopped, and switched with the cluster management GUI (using AS/400 Operations Navigator Management Central), the system APIs that are supported by OS/400 cluster management, or by the QUSRTOOL shipped OS/400 Cluster Management CL commands. Specific commands STRDOMCRG, ENDDOMCRG, and CHGDOMPRI are also provided as part of Domino for AS/400 to manage Domino application CRGs using a Domino server name. The Domino application CRG can be started, ended, or have its primary node switched from any node in its recovery domain. These commands can be used on any cluster node that is in the application CRG's

recovery domain. The node on which the command is used does not have to be the server application CRG's current primary node. The command interface is designed so that even if the Domino server's CRG name and the device CRG name it depends on are not immediately known, its application CRG and/or the device CRG used to manage access to its data directory can be started, ended, or switched from any of the nodes in the CRG's recovery domain. In this example the STRDOMCRG command could be used to start the IASPTESTER application CRG from either node in the cluster ROCH.

Note This section contains some documentation on managing device CRGs. When using any interface to manage device CRGs the operator has to consider that other applications or users might also currently be accessing or depend on access to files stored on the switchable storage managed by a device CRG. If the files are currently in use, the device CRG cannot be switched. If the device CRG is switched to another cluster node then it is possible that an application on some specific cluster node could later fail if it depended on access to the IFS files and that access is no longer available.

A Domino server that is configured as part of an application CRG cannot be started on a cluster node unless the application CRG itself has been successfully started. Servers that are configured as part of an application CRG are in a "Stopped" status when viewed with Operations Navigator.



This shows the error when a server in the *STANDBY state is started from Operations Navigator. "Error 9" for message LNT0920 is a failure to start due to server being in a *STANDBY state.

```

Work with Domino Servers
System: TSCLPAR3

Type options, press Enter.
1=Start server      2=Change server      5=Display console    6=End server
7=Submit command   8=Work console      9=Work server jobs
11=Change current  12=Work object links 13=Edit NOTES.INI

Opt  Domino          Subsystem           Domino
     Server          QDOMINOHT          Status
1    IASPTESTER      IASPTESTER        *STANDBY

Parameters or command
====>
F3=Exit      F4=Prompt  F5=Refresh  F9=Retrieve  F11=Display path
F12=Cancel   F17=Top    F18=Bottom  F22=Display entire field  F24=More keys
Error 9 has occurred, server not started
    
```

This shows the same condition but WRKDOMSVR option 1=Start server was used to start a server in a *STANDBY state. The same error, LNT0920, is reported with "Error 9".

When starting, the Domino application CRG exit program will always determine if the device CRG that its server depends on to access its data directory is primary on the same node as the Domino server's application CRG. The exit also makes sure that the device associated with that device CRG is varied on. The Domino application CRG exit will always automatically start the Domino server once it has verified these conditions have been met. Once a Domino server application CRG is started, the server can be started and ended using the STRDOMSVR and ENDDOMSVR commands or the Operations Navigator interface in the same way as any other Domino server on that system. The server will transition between *STARTED and *ENDED states as would any other Domino server when it is started and ended while its application CRG is active and the current primary node is the system where the server is being managed.

The difference between using STRDOMCRG (or Management Central or STRCRG) and STRDOMSVR in this regard is that STRDOMCRG starts the CRG associated with a Domino server first and then attempts to start the Domino server. The server will attempt to start only on the application CRG's current primary cluster node. STRDOMSVR (or Operations Navigator "Domino" server interface) will just try to start the Domino server on the node where the command (or Operations Navigator "Domino" server interface) is issued. In the case of STRDOMSVR, it will only be able to start a server if the server is not in a *STANDBY state on the node where the command is issued.

The difference between using ENDDOMCRG (or Management Central or ENDCRG) and ENDDOMSVR in this regard is that ENDDOMCRG ends the server associated with the CRG first (just like ENDDOMSVR) and then ends the application CRG for the cluster. The server will be ended only if it had been started on the application CRG's current primary cluster node. ENDDOMSVR (or Operations Navigator "Domino" server interface) will just try to end the Domino server on the node where the command (or Operations Navigator "Domino" server interface) is issued. The application CRG will remain active in the cluster even when the server is ended with ENDDOMSVR. Ending the server on the current primary node will put the server in an *ENDED state. Also, if a Domino server was to fail and it was set to restart automatically, it would always restart on the same cluster node.

When a Domino server's application CRG is switched to another node using Management Central (or the CHGCRGPRI or CHGDOMPRI commands) the Domino server will always end on the cluster node where it is currently primary first (if the server is currently started on that node). Only when the Domino server has completely ended will the CRG switch over to the new current primary cluster node. At that time, the Domino server will attempt to start on the new primary node.

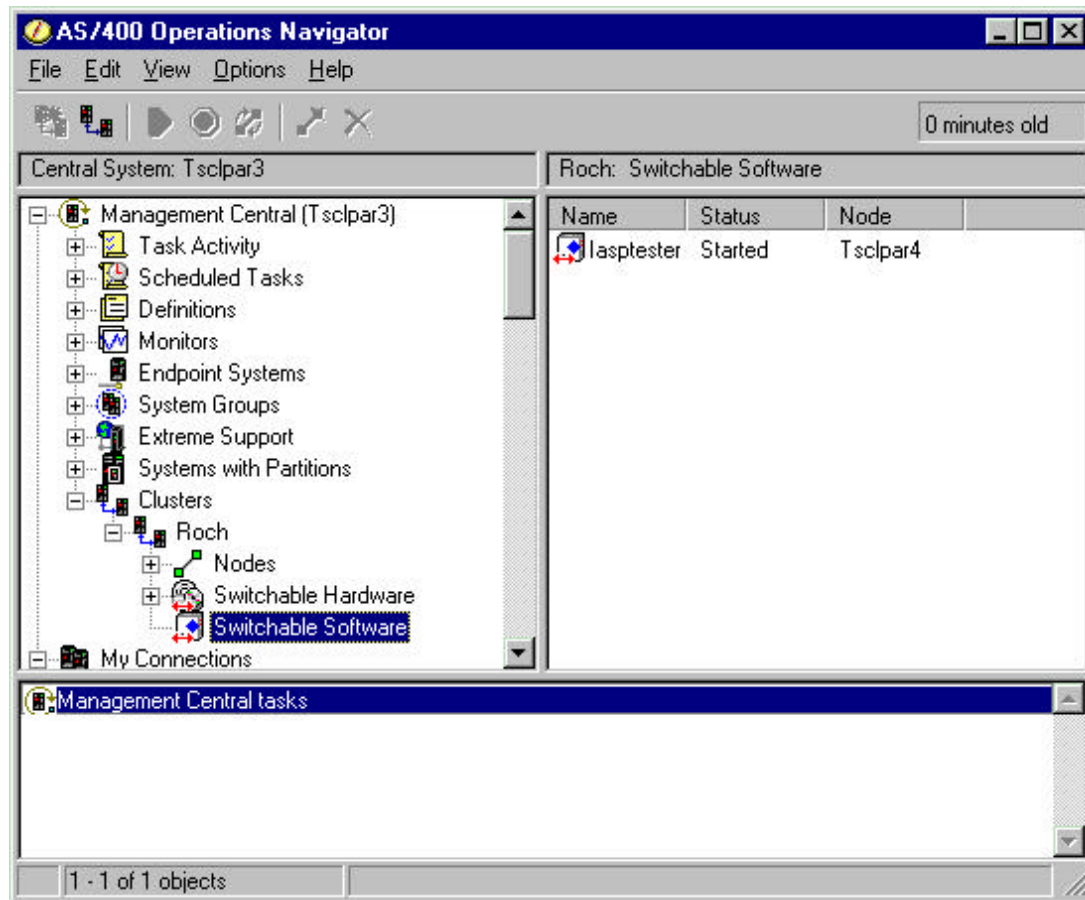
When the application CRG for a Domino server is ended or switched to another node, the server will also be ended and will be put in a *STANDBY state. Servers that are in a *STANDBY state cannot be started. It is also possible that a server in a *STANDBY state could be started at anytime if its application CRG is started or switched to the node where the server is in a *STANDBY state. Domino "stand alone" functions might work on a *STANDBY state server but there is no guarantee that the server's data directory will be able to be accessed from the same system where the server is already in *STANDBY state. Access to anything in the server's data directory would depend on the device CRG that managed the access to the device where the data directory was kept. A device CRG could be started and its device varied on for a node where the Domino server application CRG is not a current primary or where the Domino application CRG is not started at all.

Starting the Application CRG

To start a Domino server's application CRG the Management Central interface can be used. Note that when starting the application CRG for a server, the device CRG that server needs to be started and current primary system of the device CRG needs to match the current primary for the application CRG being started. Operators should check to see that the cluster node in the CRG's recovery Domains themselves have been started and that clustering is active on those cluster nodes. The device CRG should also be started and operators should make sure that current primary of the device CRG is the same as the Domino server application CRG's current primary. (More on this situation and how to resolve it is in the next part of this section.)

Using Operations Navigator Management Central to Manage an Application CRG

Once a Domino server's application CRG has been created it can be started using the Operations Navigator Management Central interface in the similar way that device CRG are able to be started. Domino application CRGs will be found under the "Switchable Software" icon instead. (Device CRGs are found under the "Switchable Hardware" icon.)



This shows the result of selecting the Switchable Software icon. In this example, IASPTESTER is the only application CRG. It was started and its current primary node is the cluster node TSCLPAR4. Note that there is no information except the CRG name that signifies this application CRG is associated with a Domino server, and if it is, which Domino server. If the CRG isn't recognizable by just its name, right clicking on the CRG name and selecting the "Properties" function will show some additional information about the application CRG, including its text description. Also looking at the exit program data of the properties of the CRG could help as if the CRG is associated with a Domino server, that server's name is in the first part of the CRG's exit data.

Note Application CRGs can be started, stopped, and switched using the Management Central functions or icons in the same manner as was described for device CRGs above.

Using the Cluster Management CL Commands

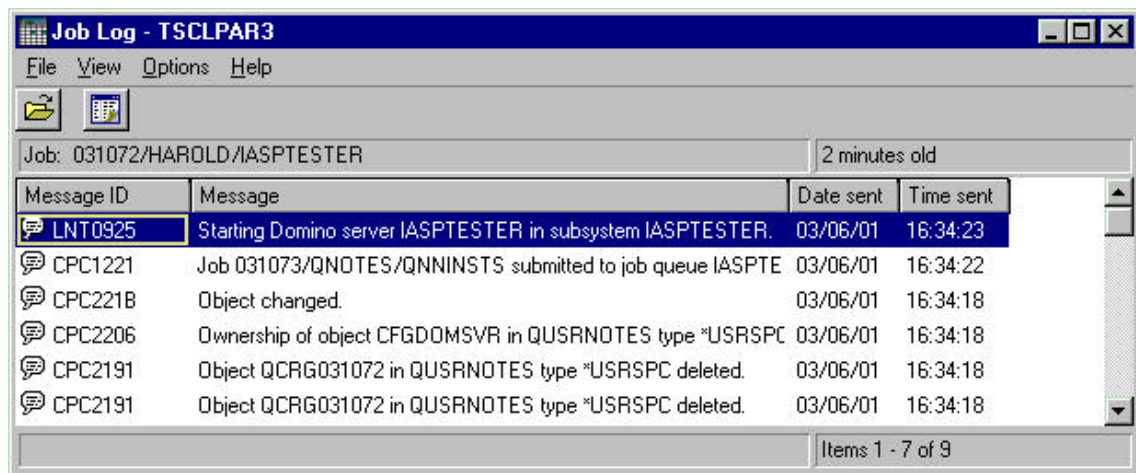
Alternatively, the Cluster Management CL commands (or the APIs) can be used to start, end, or switch a Domino application CRG. Given the cluster name and the CRG name, the STRCRG will start an application CRG on the active cluster nodes if not already started. ENDCRG will end an application CRG if not already ended. CHGCRGPRI can be used to change the primary node of a CRG to become primary on the next configured node in its recovery Domain.

Using the Domino CL Commands

Another alternative is the Domino provided interfaces to start, end, or switch a Domino server's application and/or the device CRG that manages access to its data directory. The three commands are STRDOMCRG, ENDDOMCRG, and CHGDOMPRI. These commands' CPPs build and invoke the Cluster Management APIs to do their function of managing CRGs. The difference is that they work off the Domino server's name. Instead of specifying the cluster and CRG names, the command parm is the name of the server that should be configured to be part of an application CRG. The CPPs for all these commands first find the configuration for the Domino server specified and then use internal configuration about that server to find the application CRG name and the device CRG name associated with that Domino server.

When Does the Domino Server Actually Start?

Normally by just starting a Domino application CRG or switching it over to another cluster node the Domino server associated with that CRG on the primary node will be started automatically. The configured Domino CRG exit program determines if it is being run on the node that is configured as the current primary for the server's application CRG. If it is then the exit will eventually do a STRDOMSVR for the server on that cluster node. If all the conditions are met, especially that the server's data directory can be accessed from that node the exit does the STRDOMSVR (server will transition to the "**STARTING" state and eventually will be in the *STARTED state).

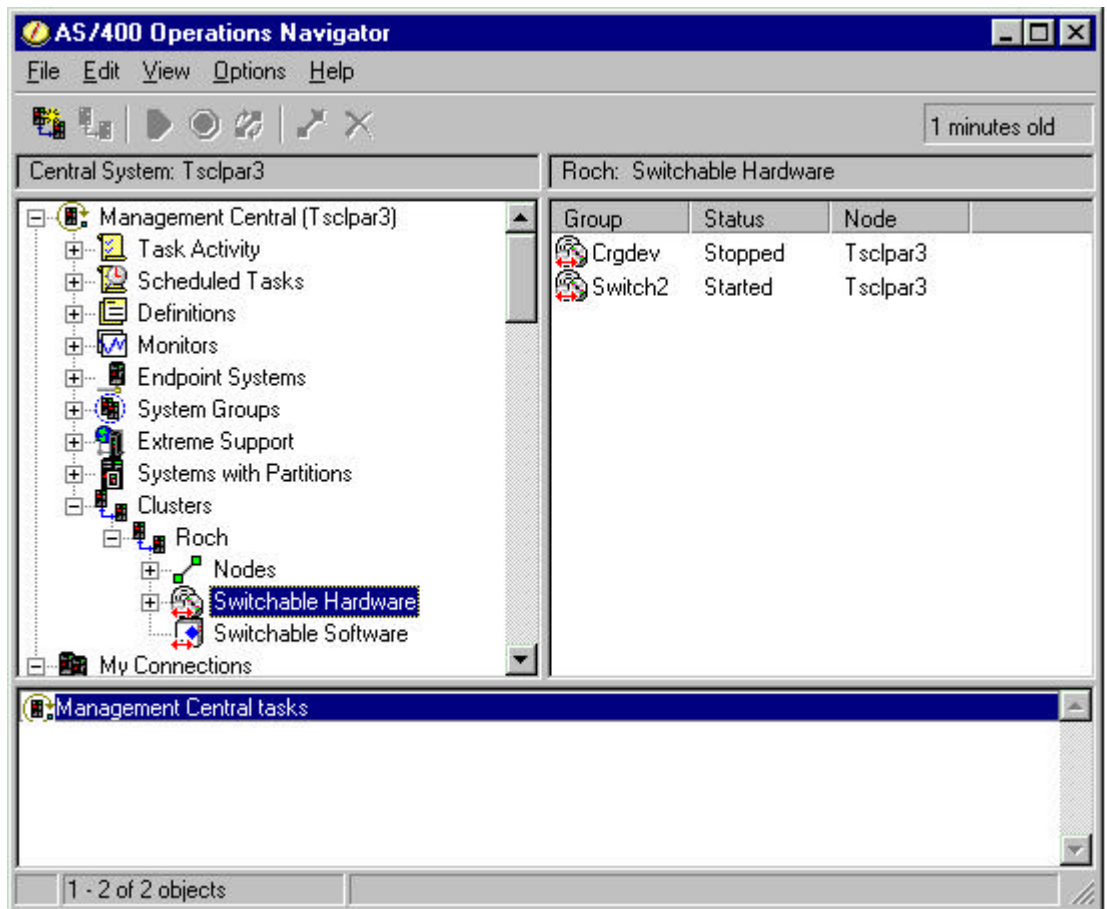


Message ID	Message	Date sent	Time sent
LNT0925	Starting Domino server IASPTESER in subsystem IASPTESER.	03/06/01	16:34:23
CPC1221	Job 031073/QNOTES/QNNINSTS submitted to job queue IASPTE	03/06/01	16:34:22
CPC221B	Object changed.	03/06/01	16:34:18
CPC2206	Ownership of object CFGDOMSVR in QUSRNOTES type *USRSPC	03/06/01	16:34:18
CPC2191	Object QCRG031072 in QUSRNOTES type *USRSPC deleted.	03/06/01	16:34:18
CPC2191	Object QCRG031072 in QUSRNOTES type *USRSPC deleted.	03/06/01	16:34:18

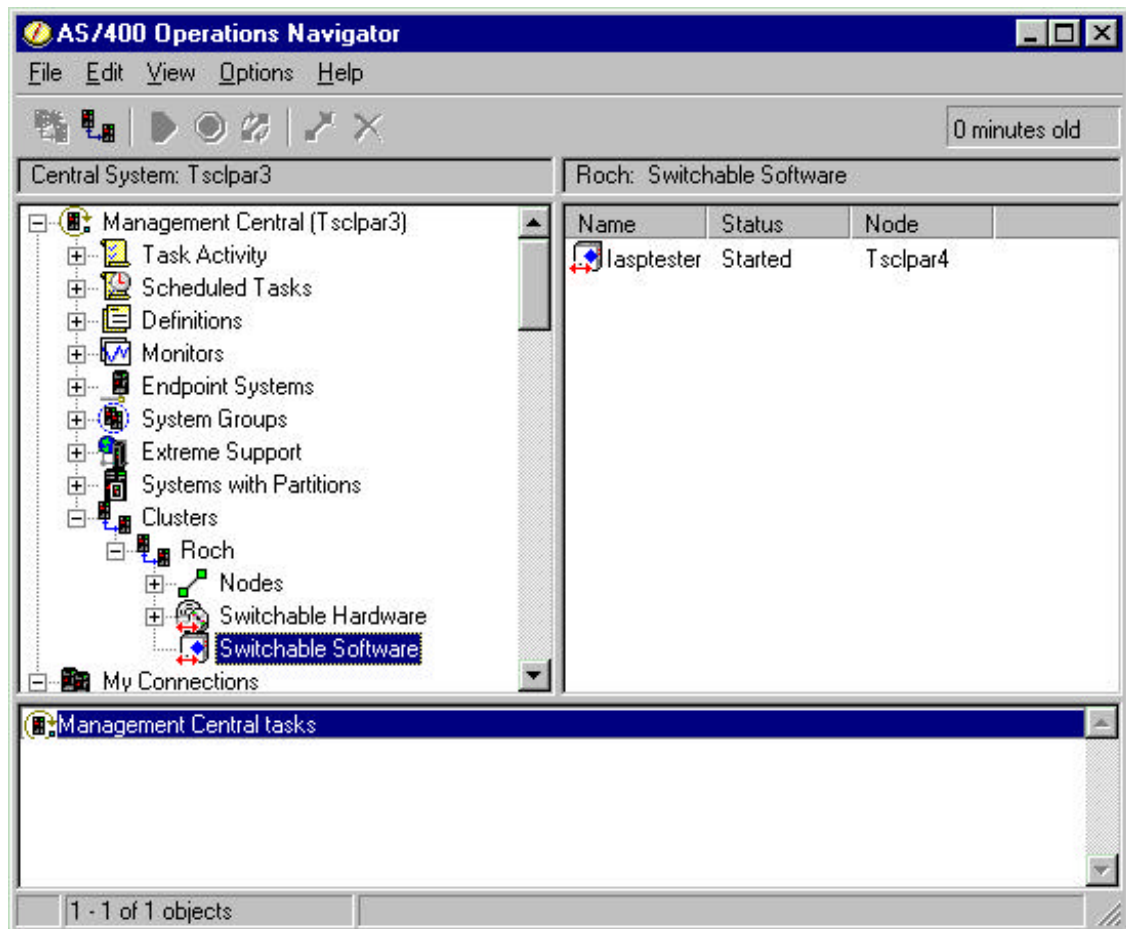
This shows the job log of the exit program's job. In this case, the job log shows that a Domino application CRG exit program has started a Domino server called IASPTESER (message LNT0925 in the job log). Note the exit program job for an application CRG is always named after the CRG (in this case the IASPTESER application CRG). The exit program will have been submitted to run as a batch job under the QBATCH subsystem by the Cluster Management function as a result of the "start" of the application CRG. In this case the application CRG was started on the TSCLPAR3 cluster node as it happened to be the current primary node for the application CRG at that moment. The exit program is submitted to run under the user profile, "Harold" that was configured with the application CRG (see the USRPRF on the CFGDOMCRG or CHGDOMCRG commands). In this example, the application CRG name and the Domino server name it is associated with happen to have the same name. They are not required to have the same name.

What Happens When the Domino Server Cannot Access its Directory on CRG Start Up?

The Domino CRG exit program will always make sure the device CRG that manages access to the Domino server is set to the same current primary node as where the exit is attempting to start the Domino server. There is no reason to actually start a Domino server on that node if its data directory is not also accessible to that node. A Domino server's application CRGs can be independently managed from the device CRG on which it depends. The situation of current primary "mismatch" was described above as the state where the Domino server's application CRG is active and primary on a different cluster node than its device CRG (or the device CRG might not be active at all). In this case the application CRG is actually "started" on its current primary cluster node as far as Cluster Management is aware. The exit cannot start the server however. This section explains this situation and how to resolve it.



In this example the CRG device CRGDEV is not started. If it was to be started, its current primary node would be TSCLPAR3.



This shows the application CRG IASPTESTER as Started. Its current primary node is TSCLPAR4. This server depends on the device CRG CRGDEV to be active and is attempting to verify that the current primary node of the device CRG is the same as the current primary node for the application CRG. Only then will the server's data directory be accessible and the Domino server started on that node by the Domino application exit program. (Until that happens the server will remain in a *STANDBY state even though the application CRG shows that it is "Started" on TSCLPAR4.)

This is an illustration of the possible "mismatch" of current primary states that was discussed above. In this case the device CRG hasn't even been started. The idea is for the operator to resolve this situation in order to get the Domino server started. In this example there are two possible approaches the operator can take:

1. The Domino application CRG can be switched over to be the same primary node (TSCLPAR3) as the device CRG on which it depends. (The device CRG would also have to be started in this case.) Once the application CRG is switched and the device CRG on which it depends has started (and its device varies on), the application CRG exit program will detect that the device is available on TSCLPAR3 and start the Domino server managed by application CRG IASPTESTER on that node.
2. The device CRG CRGDEV can be started (on node TSCLPAR3) and then it can be switched over to TSCLPAR4. The Domino application CRG exit program will detect that the device is available on TSCLPAR4 and start the Domino server managed by application CRG IASPTESTER on that node.

Either approach to resolving the mismatched primary states of the two CRGs ends up with a Domino server started and the current primary nodes of the server's application CRG and device CRG "matched" to the same cluster node. The difference of course is that in one case the resolution was to make the current primary for both nodes TSCLPAR3 (1) and in the other case TSCLPAR4 (2). Which method of resolving the mismatched primaries is more correct depends on which cluster node the operator intended to end up with the Domino server started. Note that in either case, yet more switches of both CRGs can be done to (re)start the Domino server on the other cluster node. Note that the Domino application CRG can be switched to a new primary node even if it never actually had the chance to start the Domino server.

Note There is another possible error condition that produces these same symptoms. It is possible under some circumstances that the device CRG is active and currently primary on the same current node as the Domino server's CRG however the device associated with the device CRG did not get varied on. If the device CRG is active and the current primary node for the device CRG is matched up with the server's application CRG, the operator should check the specific device for the device CRG (in this example it is "SWITCHABLE") on the current node and vary on the device on that node.

When the exit program job is attempting to start a Domino server it makes sure the device CRG that controls access to the server's data directory is active and has the same current primary node. It will continuously do this for approximately five minutes. If the device CRG still has not been started or is not switched to the same node where the exit program job is trying to start the Domino server, the exit program will issue a message, LNT1221, to the QSYSOPR message queue.

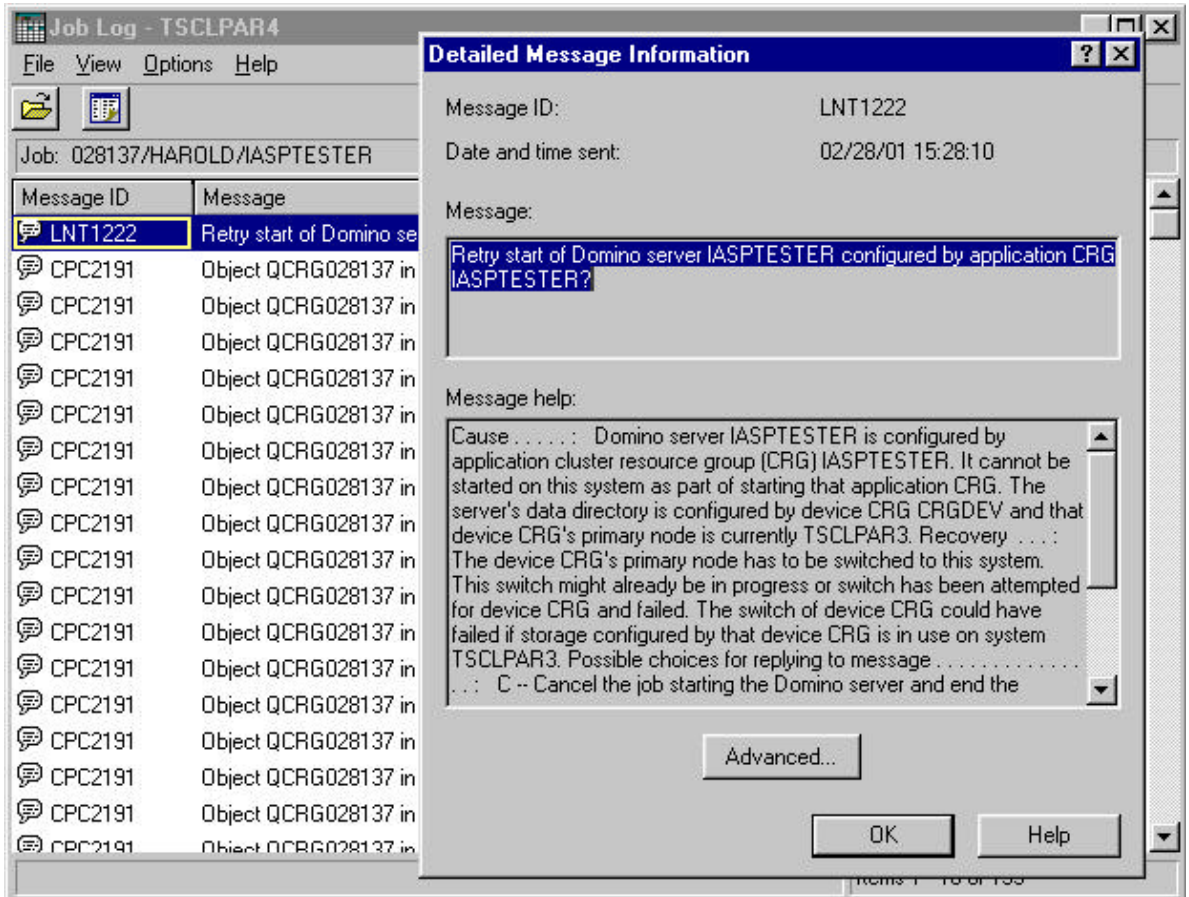
Message ID	Message	Date sent	Time sent
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:46
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:31
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:16
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:01
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:01
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:05:01
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:04:46
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:04:31
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:04:16
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:04:00
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:03:46
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:03:31
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:03:15
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:03:01
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:02:45
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:02:30
CPC2191	Object QCRG030068 in QUSRNOTES type *USRSPC deleted.	03/12/01	09:02:15

This shows the job log of an application CRG exit program. In this case the job is attempting to start the Iasptester Domino server. The job is running in the QBATCH subsystem and was submitted by Cluster Management as a result of starting the application CRG on that cluster node. The job is named after the application CRG (in this case it happens to have the same name as the Domino server it is supposed to start). This job has spent approximately five minutes trying to verify that the device CRG CRGDEV is active and is primary on the same cluster node and has failed to verify this has happened. A clue to this is the appearance of the CPC2191 message over and over in the job log. This message appears every time the job uses a Cluster Management API to get information about the device CRG's current primary node. Normally this will be done a few times when the job was started and the device CRG was verified to have the same current primary node.

This job has issued the LNT1221 to the QSYSOPR message queue as a means to record the reason why it has not started the Domino server. The idea is that there will be an operator who can monitor for this message and realize there is a primary node mismatch and take some appropriate action to resolve the situation. In the meantime the exit program will continuously try to verify that the device CRG's primary node is switched to

match its own node so that it can start the Domino server. Approximately every five minutes it will again issue the LNT1221 message if it cannot start the Domino server due to the CRG primary node mismatch situation.

Since it is quite possible the device CRG is already in the process of being switched over or is simply being delayed as it is varied off its current primary system, the exit program will keep trying for approximately 25 more minutes to determine if the device CRG's current primary has been switched to match up with the exit CRG's node where it is attempting to start the Domino server. If the primary mismatch is not resolved within this time, the exit program will stop trying to determine the current primary node of the device CRG. It will instead halt and issue an inquiry message, LNT1222, to its own job log and to the QSYSOPR message queue.



This again shows the job log of a the same application CRG exit program. In this case the job has continuously attempted (for approximately thirty minutes since it first was started) to verify that the device CRG CRGDEV is active and is primary on the same cluster node and has failed to verify this has happened . It has issued the LNT1222 inquiry message to its own job log and to the QSYSOPR message queue. The exit has halted and will wait for a response to that inquiry message. The idea of halting and posting a query message is that the resolution of the current primary mismatch is not likely going to happen without operator intervention. Instead of continuing to detect when it does happen, the exit will wait for an operator to let the exit know that the mismatch has been resolved and that it should again attempt to verify the device CRG's primary node (or if the exit should just stop attempting to start the Domino server).

Note Even when it has halted in this state and is waiting for an answer to the LNT1222 query message, the application CRG can still be ended or switched to another primary node using the methods available to manage a CRG primary node, as previously documented.

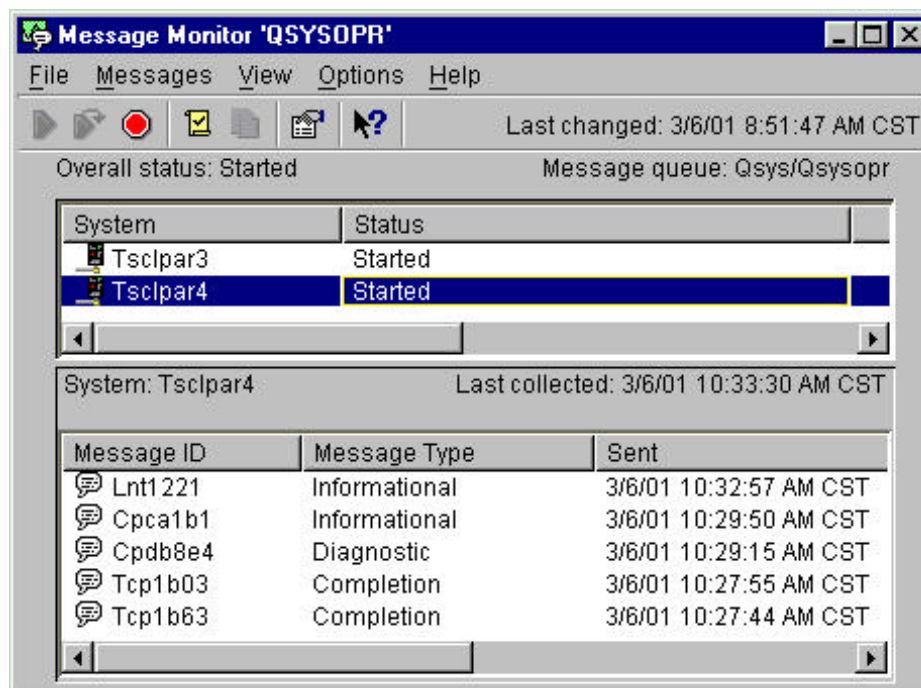
Alternatively, the operator can respond to the LNT1222 query message in one of two ways:

- "R" response to indicate the exit should retry to verify the current primary node of the device CRG. It will do so once and if the device CRG primary node is still not matched up with the exit program's node it will again issue the LNT1222 query message. The appropriate use of this response is when the operator has managed to get the device CRG's current primary switched over and wants the Domino exit program to verify this and then start the Domino server.
- "C" response to indicate the exit should cancel or halt attempting to start the Domino server. This will end the application CRG in the cluster (just as if an ENDDOMCRG or ENDCRG had been issued against the application CRG).

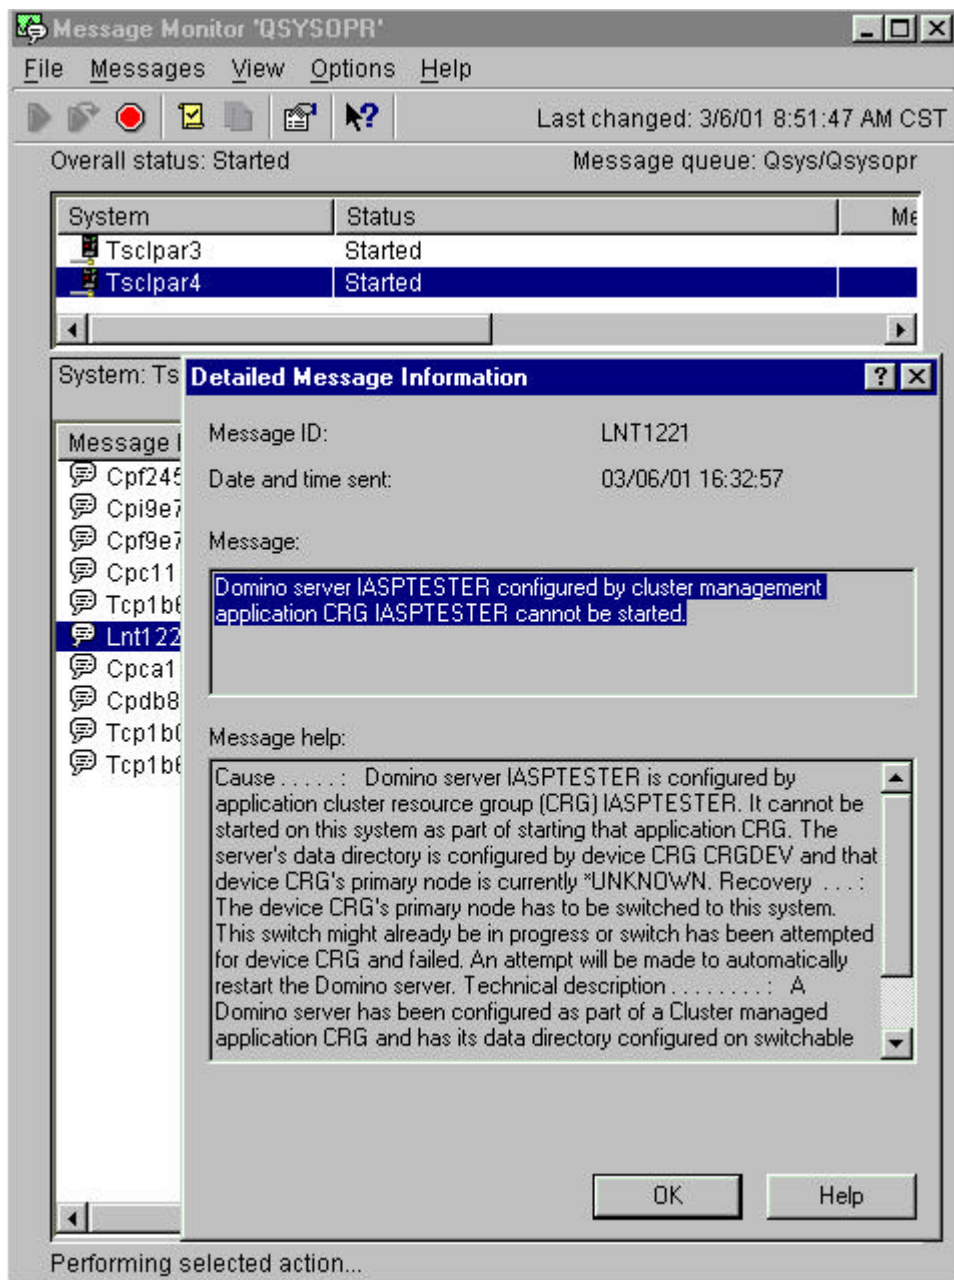
Using Operations Navigator Message Monitors to help Manage Domino Application CRGs

As explained previously, the Domino server's application CRG and the device CRG on which it depends are two different cluster managed objects. They can be started and stopped and possibly switched independently. Even when an operator is skilled at managing applications and device CRGs, it is likely the situation of mismatched current primary states can occur. Besides knowing how to resolve the situation as explained above, it is important to be able to detect when it might be occurring so that the resolution can be attempted.

A means to help an operator detect and manage this situation for Domino application CRGs is via the use of Operations Navigator Message Monitors. (Information on how to configure and manage Message Monitors appears at the end of this document.)

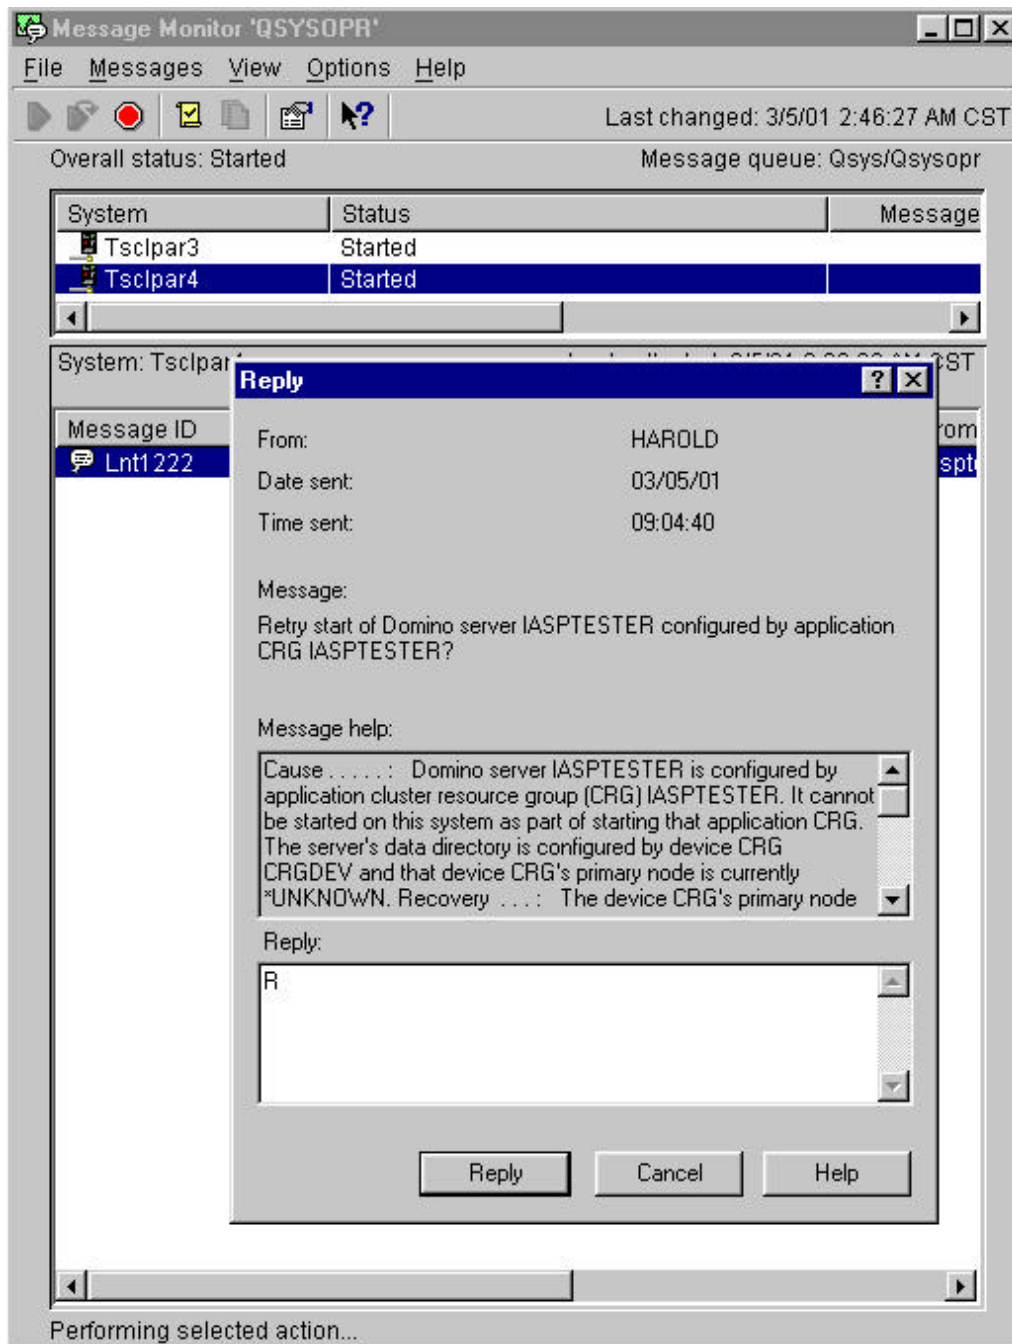


This shows a Message Monitor named "QSYSOPR" that has been set up to monitor all messages that would appear on the QSYSOPR message queue for either cluster node (TSCLPAR3 or TSCLPAR4). The monitor shows an LNT1221 message has been issued to the QSYSOPR message queue on system TSCLPAR4. In this case the monitor has been configured to automatically sample the QSYSOPR message queue for new messages every 15 seconds. Any new messages posted to QSYSOPR since the last time it sampled are automatically added to the monitor and appear in the monitor window.

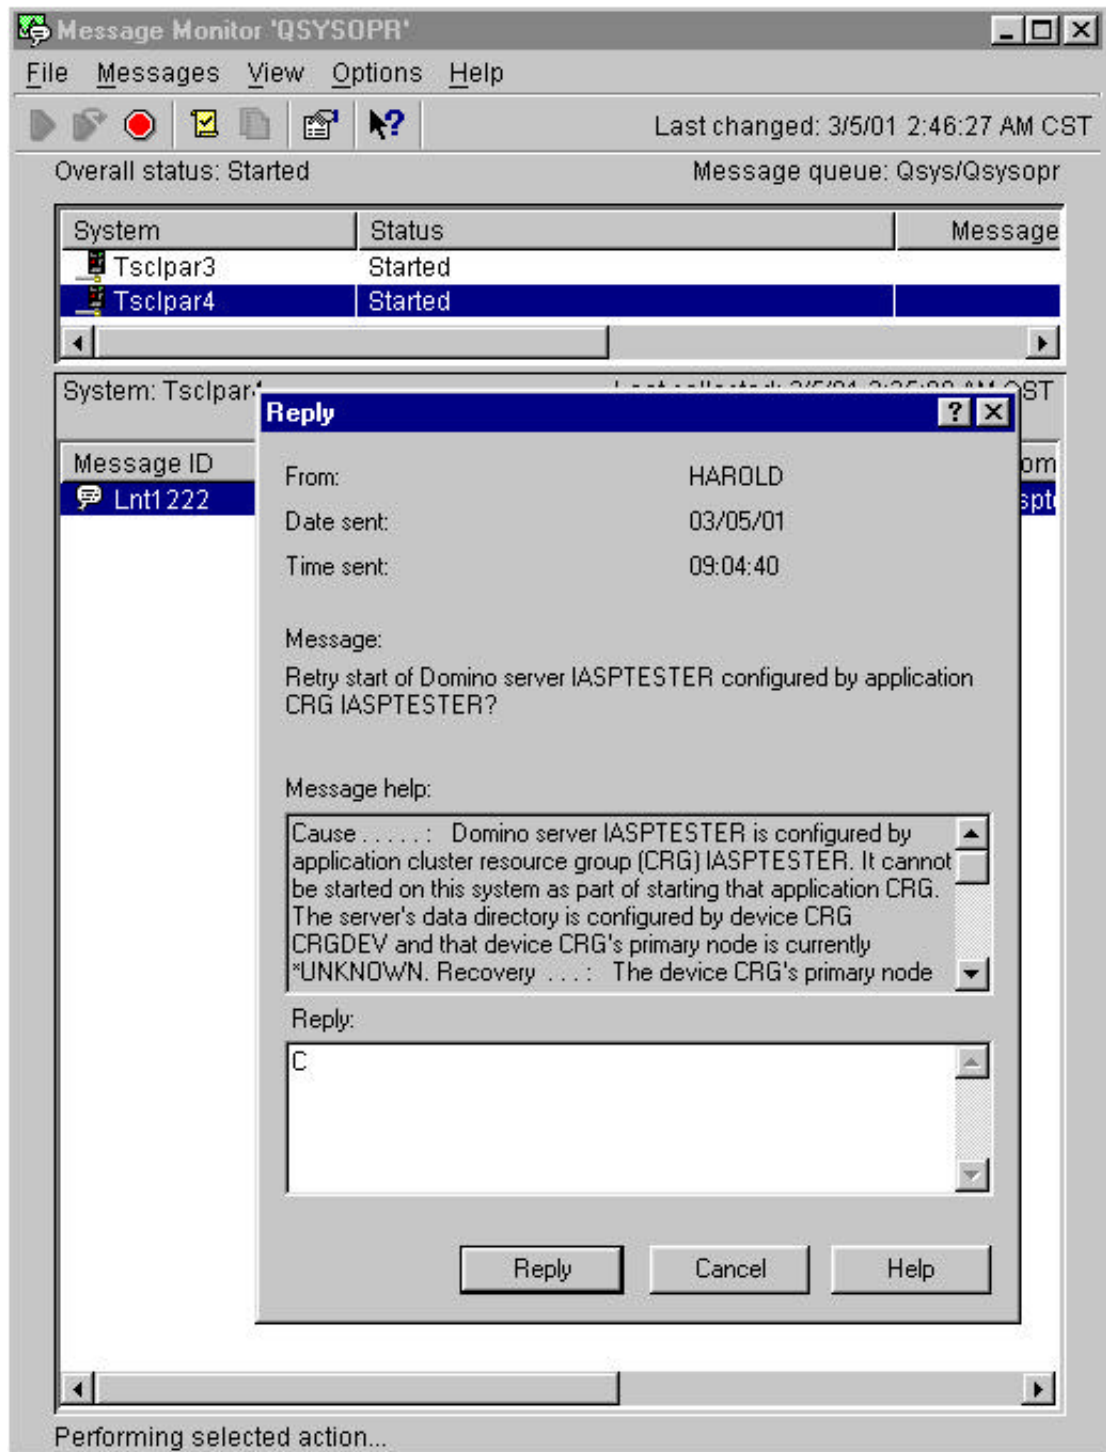


By selecting a specific message in the monitor window, the operator can determine that there is a problem with a Domino application CRG exit program attempting to start a Domino server. The LNT1221 message will contain the name of the application CRG, the Domino server that cannot be started, and the device CRG name that the exit needs to be switched over by the operator to resolve the situation so the Domino server can start.

It is possible that the operator takes no action on resolving this situation and the device CRG is not switched over to have the same primary as the Domino server's application CRG. After approximately 30 minutes of continuously trying to see if the device CRG primary matches its own, the exit program will issue an LNT1222 inquiry message to the QSYSOPR message queue, and it will halt and wait for an operator response.



This again shows the same "QSYSOPR" monitor for messages posted to the QSYSOPR message queue on TSCLPAR4. In this case the exit program has continuously attempted (for approximately thirty minutes since it first was started) to verify that the device CRG, CRGDEV in this example, is active and is primary on the same cluster node and has failed to verify this has happened . It has issued the LNT1222 inquiry message to the QSYSOPR message queue and it was posted in the monitor. The exit has halted and is waiting for a response to the LNT1222 inquiry message. In this example the operator is responding to the message with an "R" response to indicate to the application CRG exit program that it should again retry to verify that the device CRG's current primary matches the application CRG's for the exit's node.



This again shows the same "QSYSOPR" monitor. The exit program issued the LNT1222 inquiry message to the QSYSOPR message queue and it was posted in the monitor. The exit has halted and is waiting for a response to the LNT1222 inquiry message. In this example the operator is responding to the message with an "C" response to indicate to the application CRG exit program that it should cancel the attempt to start the Domino server. The exit's application CRG will be ended.

Note Another possible operator response can be to use Management Central or Cluster Management CL commands (or the Domino provided CRG management commands) to switch the application CRG over to the cluster node where the device CRG was currently started and primary. The operator could also use a similar interface to end the application CRG and stop attempts to start a Domino server altogether.

Changing a Domino Application CRG

Some attributes of a Domino server's application CRG can be changed after the application CRG has been created. Interfaces to change CRGs are provided through Operations Navigator Management Central or using the Cluster Management API or CHGCRG command. There is a CHGDOMCRG command that allows some attributes of a Domino application CRG to be changed. This is the most "controlled" way to change a Domino application CRG and is the only recommended way to change an existing Domino server's application CRG. The other interfaces provided will "work" but also allow attributes about the application CRG to be changed that could disable the application CRG from being used to manage the Domino server it was associated with when it was created with the CFGDOMCRG command.

There are some important restrictions on changing Domino application CRGs that make the rules for changing them slightly different than any other application CRG. Note that these attributes can be changed using some of the interfaces provided change CRGs as nothing stops the interface or API from altering information about the application CRG. The restrictions are in place for the following information in that if it is ever changed for a Domino application CRG, the CRG will likely fail to start again and may, in fact, result in the Domino server being started without first deleting the application CRG.

The following bullets describe details about specific Domino server application CRG information that should NOT be changed:

- The application CRG exit program (and exit program library) should never be changed. These are set to QNNHACRG for the program name and QNOTES for the library name when the Domino server's application CRG is created with CFGDOMCRG. This is the only CRG exit program that will work for Domino server application CRGs
- The application CRG exit program data. This is set to a specific value when the Domino server's application CRG is created with the CFGDOMCRG command. The exit data contains some information about the Domino application CRG, most importantly the name of the Domino server that is associated with the application CRG and the device CRG that the server will rely on to manage access to its data directory. While viewing an application CRG's exit data will help an operator distinguish one Domino application CRG from another, the exit data has a specific format the exit program relies upon. A Domino application CRG should never be modified in an attempt to get it to start a different Domino server. This will create a conflict with other internal information that configures both servers and the result is that neither server could be started. It is better to delete the CRG entirely and then recreate the CRG specifying the new server's name.
- The Domino servers application CRG's recovery domain should never be changed directly. This is set to be the same recovery domain as the device CRG identified by the DCRG command parameter when the CFGDOMCRG command is used to create the application CRG. The only time the recovery domain should be changed is when the recovery domain of the device CRG is itself changed. The only recommended way to change a Domino server's application CRG to match the new recovery domain of the device CRG is use the CHGDOMCRG command. This command's CPP will determine the device CRG that was originally used to set the recovery domain for the server's application CRG. It will then find the current (now changed) recovery domain for that device CRG and again, automatically, change the application CRG's recovery domain to match.

Most clustering interfaces that allow a CRG to be changed support the ability to use "*SAME" for the application CRG attributes that should not be altered when some specific attribute of the CRG is changed. Caution should be used when using any interface to change an existing Domino server application CRG that changing one specific attribute, such as the text description, will not indirectly alter (or reset to "blanks") the exit program or exit program data that is configured for that application CRG. (Using the CHGDOMCRG command to change the Domino application CRG will protect the CRG attributes that should not be changed.)

Adding Nodes to an Existing Application CRG

Whenever a new node is added to the cluster or to a recovery domain, the pre-check steps mentioned above that should be done for any cluster node before creating a new Domino server application CRG should be performed on the new cluster node. CRGs have to be ended to add new nodes to their recovery domains. As fits the usual pattern of changing a recovery domain of a Domino application CRG, the recovery domain of the device CRG should be changed to add the new cluster node to its recovery domain, then the device CRG

should be tested to use this new node, then a CHGDOMCRG command is used to reset the Domino servers application CRG's recovery domain so that it matches the new recovery domain of the device CRG.

Note The Operations Navigator Management Central GUI only supports two nodes in the cluster it manages. Adding a third node will disable this interface. In addition, CHGDOMCRG does not allow a recovery domain to be specified. That information will always be taken from the recovery domain of the device CRG that was originally specified when the application CRG was created with CFGDOMCRG.

Removing Nodes from an Existing Application CRG

Cluster nodes can be removed from an existing application CRG. The CRGs have to be ended whenever their recovery domain is changed. As fits the usual pattern of changing a recovery domain of a Domino application CRG, the recovery domain of the device CRG should be changed to remove the node from its recovery domain, then tested to make sure it still works without the node. A CHGDOMCRG command should then be used to reset the Domino servers application CRG's recovery domain so that it matches the new recovery domain of the device CRG.

In the case of a Domino application CRG, the node where the Domino server was originally configured should never be removed from that CRG's recovery domain. In the case where this original node has to be removed (for example the node is being removed from the cluster entirely) all Domino server application CRGs that had that node as their originally configured node should be deleted first. (See the "Deleting a Domino Application CRG" below.) Then the Domino servers themselves should be moved to a remaining cluster node. A new application CRG should be configured with CFGDOMCRG on that cluster node to where the servers have been moved.

Changing the Configured IP Address of a Domino Server Associated with an Application CRG Node

Whenever a new Domino application CRG is configured, the IP address that is configured with the server's application CRG is derived from the Domino server's configuration by reading it in the server's "NOTES.INI" file. This happens in the CPP of the CFGDOMCRG when the application CRG is first configured. Should the IP address of a Domino server associated with an application CRG ever have to be changed, the IP address configured for the server's application CRG must also be changed.

To change the application CRG's IP address, the Domino server must be ended first. The application CRG should also be ended but the device CRG should be left active and primary on the same node. Then change the Domino server's IP address (it can only be changed on the system where the device CRG is still primary since this changes the server's configuration in its data directory). When the Domino server's IP address has been changed, the CHGDOMCRG command can then be run again for that server. The CPP for the CHGDOMCRG command will retrieve the NEW IP address for the Domino server from the server's "NOTES.INI" file and change its application CRG to use the new IP address.

An alternate way to change the CRG is to use the Cluster Management CHGCRG CL command with the CRGTYPE(*APP) and the new IP address for the CRG specified with the TKVINTNETA parameter. The IP address specified must be the one for the server associated with the application CRG being changed. The application CRG should be ended when its IP address is changed. When the CRG's IP address has been changed, the Domino server's application CRG can be restarted so that it restarts the Domino server and will use the new IP address.

Note The pre-checks and configurations that were documented above for an IP address to be used by an application CRG still apply to a new IP address.

Whenever the IP address is changed for a server (and its application CRG) it is a good idea to always test the ability of the Cluster Management CRG switch primary function to be able to end and start the new IP address on every cluster node in the application CRG's recovery domain. Do this by switching the application CRG to all these nodes and determine if the new IP address is activated on each node as the application CRG is started on that node.

Deleting a Domino Application CRG

A domino server's application CRGs can be deleted using the Management Central interface, the Cluster Management API or CL command DLTCRGACT, or the Domino provided interface DLTDOMCRG. The server definition should be removed from all the notes in the application CRG's recovery domain when the CRG is deleted from the cluster except the node where it was originally configured when the application CRG was created. Deleting the application CRG does NOT delete the Domino server from the cluster node where it was originally configured. A Domino server should not be removed with the CFGDOMSVR OPTION(*REMOVE) if it is configured as part of an application CRG. However, once the application CRG has been deleted, the Domino server can then be removed using CFGDOMSVR OPTION(*REMOVE).

Recommended Procedure for Deleting a Domino Server's Application CRG

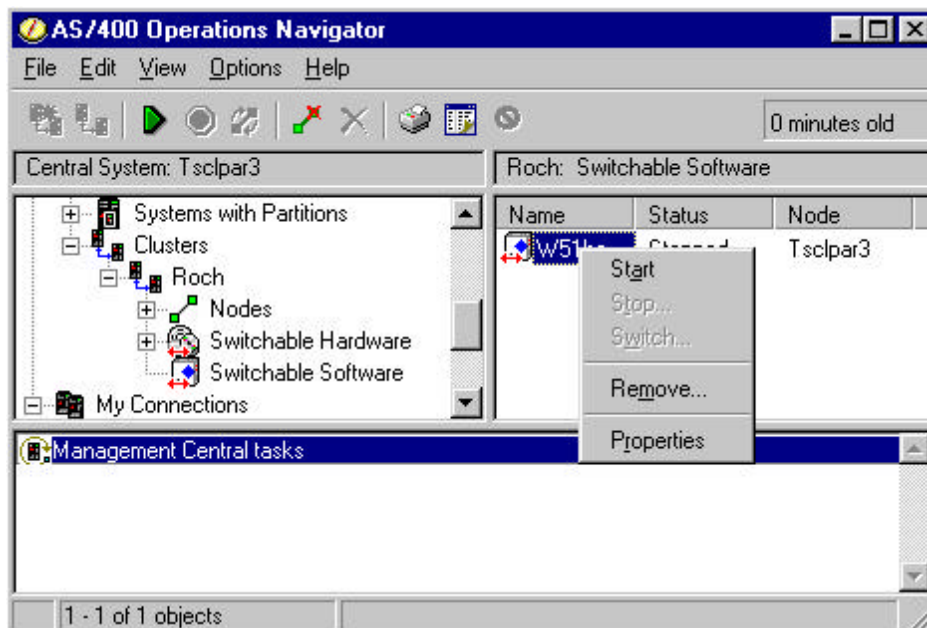
Before deleting an application CRG using any of the interfaces provided, follow these steps:


1. The application CRG must be ended. An active CRG cannot be deleted. The Domino server will be ended when its application CRG is to be ended. The Domino server that was associated with the application CRG is ended with ENDDOMSVR and will go into a *STANDBY state.
2. All the nodes in the application CRG's recovery domain should be active in the cluster if possible so that the CRG can be removed from all the nodes as it is deleted. The node in the application CRG's recovery domain where the Domino server was originally configured when the CFGDOMCRG command was used MUST be active in the cluster. The CRG could be deleted even if other nodes in the recovery domain are not active (or are in a "Failed" state). The result is that the CRG will still be configured on the failed nodes. If the CRG is configured there, the Domino server is still configured there. For some amount of time, until those nodes are started, the cluster is not in a "consistent" state regarding the deleted application CRG. Eventually the inactive nodes should be restarted, and will rejoin the cluster and the Domino server's application CRG and the server configuration will be deleted from that node at that time.
3. If possible the device CRG should be active. Access to the device CRG should be switched so that its current primary is switched to the node where the Domino server was originally configured when the CFGDOMCRG command was issued.

When these conditions have been met, the Domino server's application CRG can be deleted.

Deleting the application CRG using Management Central

The application CRG configured for a cluster can be removed by selecting the "Switchable Software" icon under the cluster and then right clicking on the CRG to be removed



This shows an application CRG and the "Remove" menu item seen when right clicking on the application CRG. Alternatively just selecting the application CRG and then clicking on the  icon will remove the CRG. Note that an active CRG cannot be removed. It has to be ended ("End" is "Stop" in Management Central) or the remove menu item and icon will not be activated and CRG cannot be removed.

Deleting the application CRG using the Cluster Management CL command

The Cluster Management CL command DLTCRGACT can be used. Given the cluster name and the name of the application CRG, the command will delete the CRG from all the active nodes in the cluster. This command can be used from any active node in the cluster to delete Domino server application CRGs.

Deleting the application CRG using the DLTDOMCRG command

This command's CPP does some error checking and then internally builds the parameters for the QcstDeleteClusterResourceGroup system API. This command is provided as an alternate way to delete a Domino server's application CRG, given the Domino server's name. The command can be used from any active cluster node in the application CRG's recovery domain where the Domino server specified is configured.

After the Domino Server's Application CRG Has Been Deleted

The Domino server configuration should only remain on the original node where it was configured when the CFGDOMCRG command was used to create the application CRG. It should still be able function as a "normal" server but only on that node. The Domino server could also be configured as part of another new application CRG at this point. Note that when the CRG has been deleted from the cluster, the Domino server should automatically go to the *ENDED state. It is currently ended and is no longer configured as part of an application CRG so it should no longer be in a *STANDBY state.

The Domino server still may not be able to be started. Deleting the server's application CRG does NOT change how access to it's data directory is controlled through the device CRG on which it stills depends. Note that if this device CRG was not active or its current primary node is not set to allow the server's data directory to be accessed, the Domino server cannot be started. The Domino server could be removed with the CFGDOMSVR OPTION(*REMOVE) but if the server's data directory was not accessible from the node where the CFGDOMSVR OPTION(*REMOVE) command was done, the data directory associated with that server would not be deleted. The server's data directory would have to be removed later by using IFS directory commands or Operations Navigator's "Integrated File Systems" interface on the cluster node where the device CRG is primary.

Once the Domino server's application CRG has been successfully deleted, it should function as a "normal" Domino server. To truly remove the server's dependency entirely from the device CRG and from Cluster Management, the same set of steps that are described above for copying a server's data directory can be used.

Caution Before proceeding with the following instructions, the existing server's data directory and system configuration should be backed up. The Domino server should not be active and no "stand alone" Domino functions should be run against the server while this procedure is being followed.

A variation on these same steps described above is repeated here, with the idea being to reverse the outcome this time and move the server's data directory off the disk storage managed by the device CRG.

1. Make sure the device CRG is started and that it is switched so that its current primary is the node where the data directory is to be copied from the device's UDFS.
2. Copy the existing Domino server's data directory from the device CRG's IASP managed device to a directory managed by the system's ASP. In this example the UDFS name was "Switchable". The domino data directory was simply copied back to the root IFS file directory so that the directory name remained the same, except that "Switchable" was dropped from the path name (i.e. from /Switchable/DominoData/Iasptester copied to /DominoData/Iasptester).

The copy command for this example is:

```
CPY OBJ('/Switchable/DominoData/Iasptester/*') SYMLNK(*YES)
TODIR('/DominoData/Iasptester/') SUBTREE(*ALL) OWNER(*KEEP)
```

Note SYMLNK(*YES), SUBTREE(*ALL), and OWNER(*KEEP) must be specified on the CPY command as all the symbolic links and sub directories are to be copied. The owner of all the data directory objects must remain as QNOTES in the copy.

3. Change the server's internal OS/400 configuration path names to reference the now copied data directory on the system.

For example to change the data directory of the server, IaspTester, to /DominoData/IaspTester call the program QNOTES/QNNHACHIR as follows:

```
CALL QNOTES/QNNHACHDIR PARM(IaspTester '/DominoData/IaspTester')
```

*ALLOBJ authority is needed by the caller of the program. Calling this program will result in changing where the Domino administration commands find the server's data directory. It will not change where the Domino server itself finds the data directory. This will require modifying the notes.ini, located in the new data directory, and changing every reference to the old data directory path name to the new data directory path name in the notes.ini.

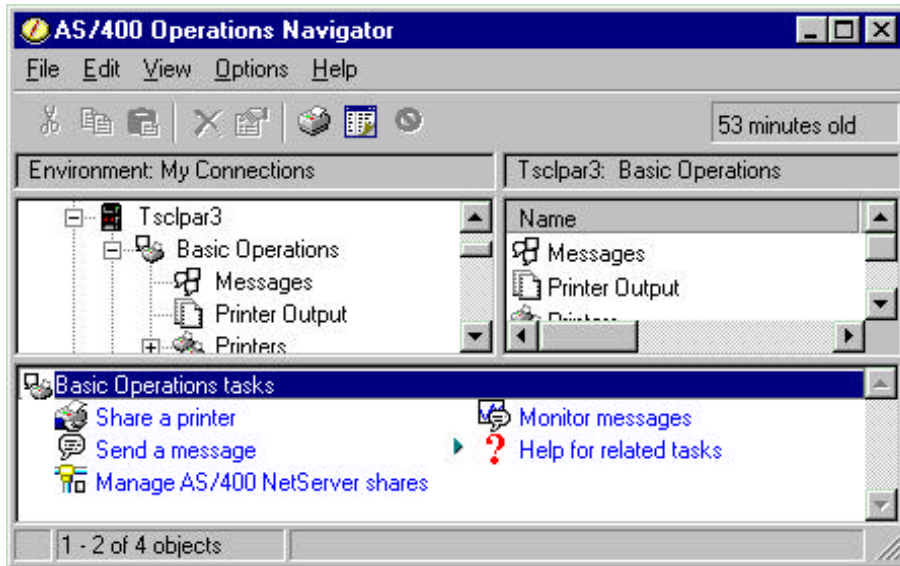
To verify that the program call changed the data directory, use the WRKDOMSVR command. Select PF11 on the desired server to view the data directory. Verify that the directory displayed is the new directory name entered on the program call to QNOTES/QNNHACHDIR.

4. Change the server's notes.ini references to any data directory object to use the new configuration path name that does NOT include the UDFS name in the path. EVERY reference that used to refer to a path with the UDFS in the path name has to be changed.
5. Start the Domino server to verify that it can start and use the copied data directory and the path names to that directory have been correctly revised.
6. <Optionally> Remove the device CRG managed copy of the server's data directory from the UDFS as it should no longer be used.

When these steps have been completed, the result should be that the server and its data directory are no longer associated or controlled by CRGs or affected by Cluster Management functions.

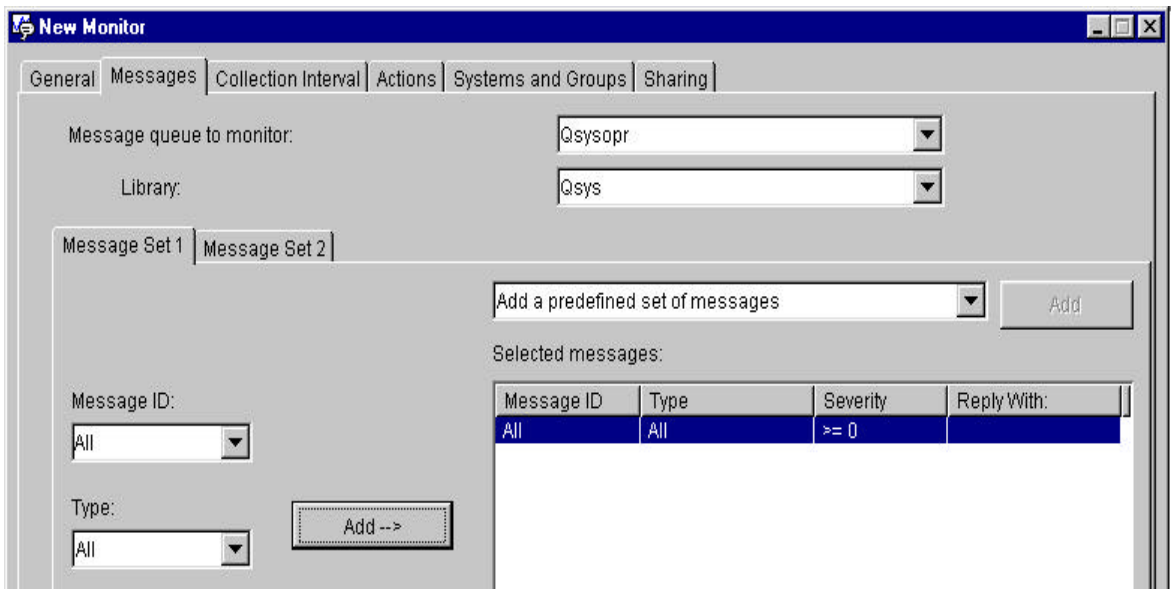
Setting up a Message Monitor with Operations Navigator

Operations Navigator can help manage Domino server application CRGs in a cluster by allowing the creation of a message Monitor. A new Monitor can be created by selecting "Basic Operations" using Operations Navigator for one of the cluster nodes, and then selecting "Monitor messages" under the Basic Operations task list.

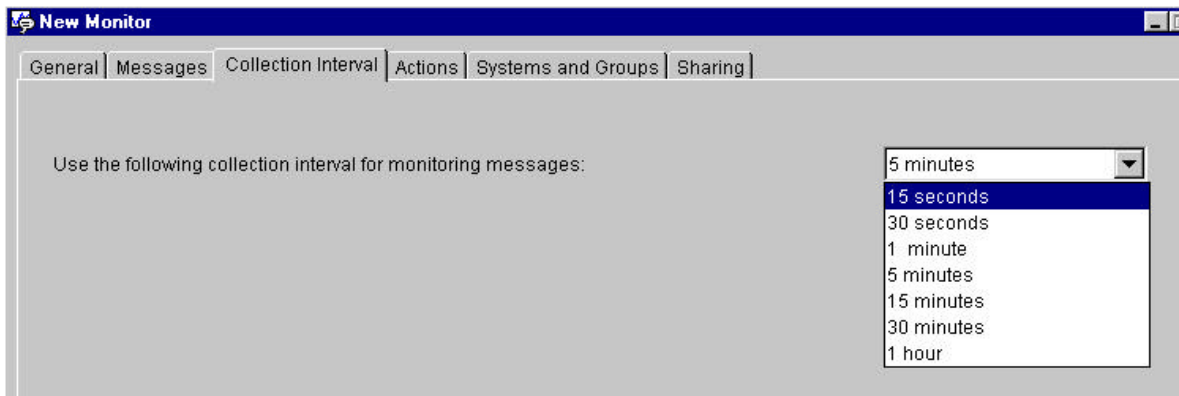


This is the window that will allow a new monitor to be created.

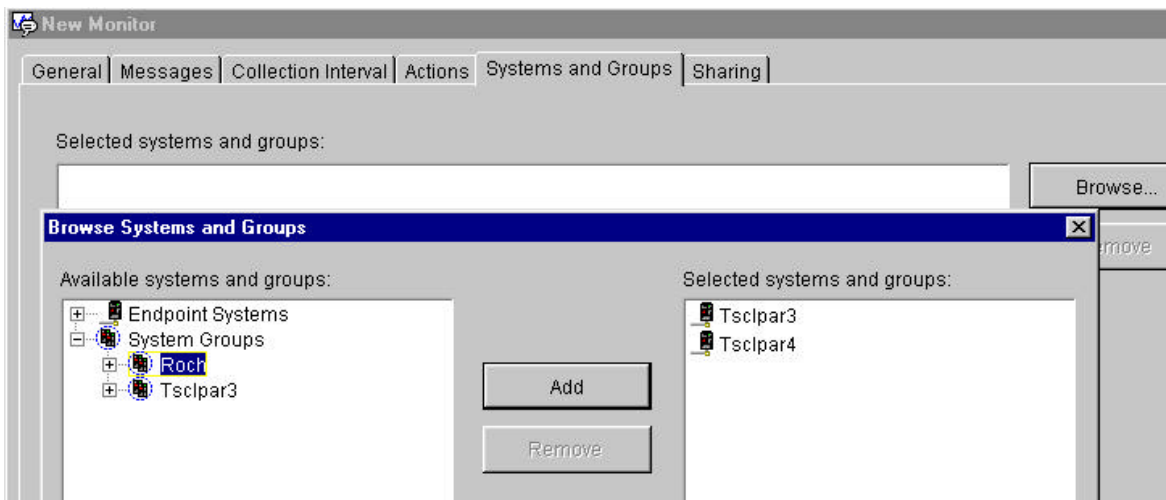
The monitor must first be given a name under the "General" tab (not shown).



Selecting the "Messages" tab allows the message queue monitor (in this case use QSYSOPR) to be configured. The message IDs and types to be monitored also must be configured (in this case "All" IDs and types have been added with the "Add -->" button).



Selecting the "Collection Interval" tab can allow adjustments to the collection interval (how often the Monitor will test the message queue for new messages posted on that queue). The default is every 5 minutes (in this case the interval was set to every 15 seconds).



A powerful capability of message Monitors is that they can be created to operate against multiple systems. This capability is ideal for clustering as all the cluster nodes can be defined for the new monitor by selecting the "Systems and Groups" tab. The "Browse..." button will show another window and the cluster nodes to be monitored can be selected and added to the new Monitor with the "Add" button. In this example both cluster nodes TSCLPAR3 and TSCLPAR4 from cluster ROCH have been added.

Selecting the "Ok" button (not shown) will then create the Monitor.

The message monitor created in this example is set up to Monitor all messages posted to QSYSOPR message queue for both of the cluster node systems. This kind of Operations Navigator function is invaluable to an operator trying to determine the occurrence of the primary state mismatch that is preventing a Domino server from starting on a cluster node because the device CRG on which it depends is either not started, or has its current primary switched to another cluster node.

To locate and start or stop an existing message monitor select Management Central "Monitors" and then select "Message" under that. Monitors may be added and removed using this Management central function.



An IBM Company

© Copyright 2001 Lotus Development Corporation. All rights reserved.

Not for reproduction or other use without express written consent of Lotus Development Corporation