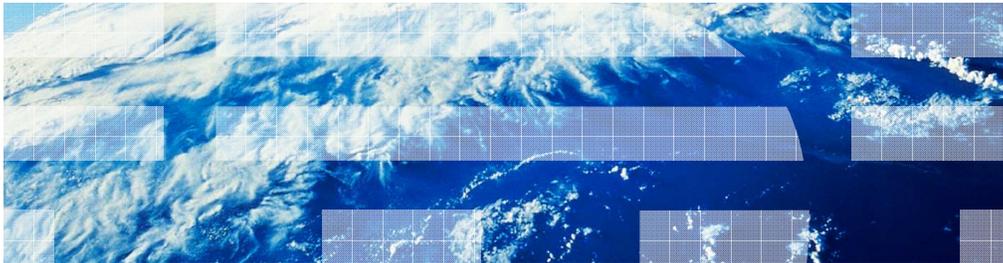


# IBM iCluster for I

## Switching concepts for IBM iCluster for i



This presentation is designed to provide you a head start in developing your own switch routine.

## Switching concepts with IBM iCluster for i

- Switching
  - Configuration actions and user exit routines created to ensure users can switch between production database and recovery database
- iCluster® allows configure switching manually or automatically
- Final step using iCluster for I
- Switch concepts and information
- Highlight aspects of switch over process

Switching describes the configuration actions and user exit routines that have to be created that will ensure that your users can switch between the production database and the recovery database. iCluster for i contains features that allow you to configure switching manually or automatically. This presentation defines switching concepts and provides information to get you started as you work towards a switch test in your environment.

Switching with iCluster is the final step in using iCluster for i as a means to use your backup system as production in a disaster recovery situation. This presentation describes at a high level how iCluster will provide the switch over for your replicated objects that are configured in your environment. This presentation also highlights other aspects of a switch over process that need to be considered and will compliment the iCluster software switch.

## Switching iCluster or switching environment

- Aspects of a switch over to consider
  - Actual iCluster switch – considered a software switch for replicated objects setup in iCluster by groups defined
  - Production system switch – other manual/automated procedures required to ‘quiet’ production server and change backup into production
    - Batch jobs
    - Communications
    - Scheduled jobs
    - User connectivity

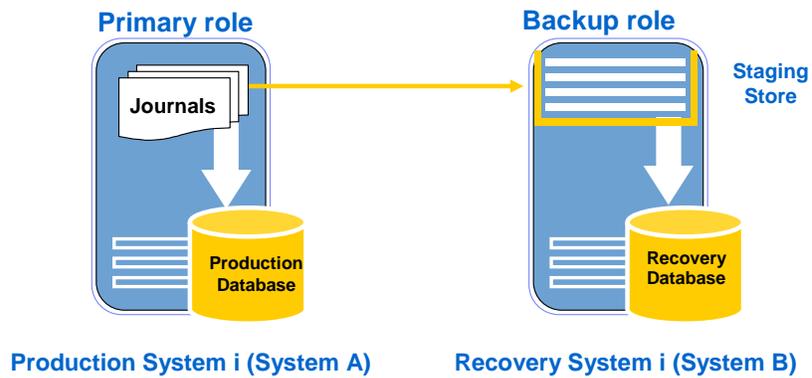
One way to think of switching is to somewhat separate the iCluster software switch, the act of switching your replication groups and the actual switch over of your environment.

The software switch performed within iCluster is quite simple and though there are different verifications required before, during and after the switch; the software switch really comes down to one command – DMCHGROLE.

The actual Production System switch requires more time spent on documentation and discussion as the switch over plan is being put together. Examples of items to discuss include: How will the users connect to the backup system when it is production? What items need to be shutdown to stop transactions from being generated on production? How to limit user access and authorization of applications?

These items and more require discussion so you can be sure the current production system is restricted in such a way that it looks more like the backup server. These same items need to be reviewed and prepared on the backup system. They start when the switch over is completed and the backup system becomes the new production.

## Normal mode



4

Switching concepts for IBM iCluster for i

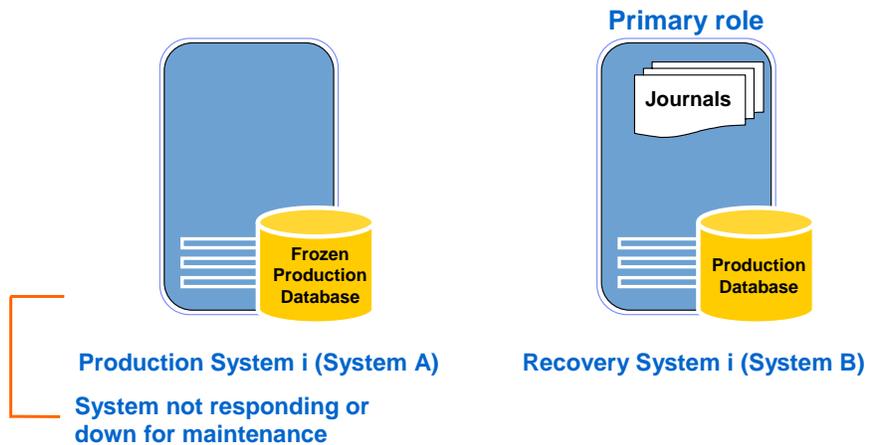
© 2010 IBM Corporation

This slide shows Normal mode replication.

The production system journals are used by a scrape process that replicates the changes to the backup system.

All changes in normal mode are received by the recovery server and put within the staging library where they are then applied by the backup apply process. This is the process that continues as your day to day transactions are being processed.

## Failover mode



5

Switching concepts for IBM iCluster for i

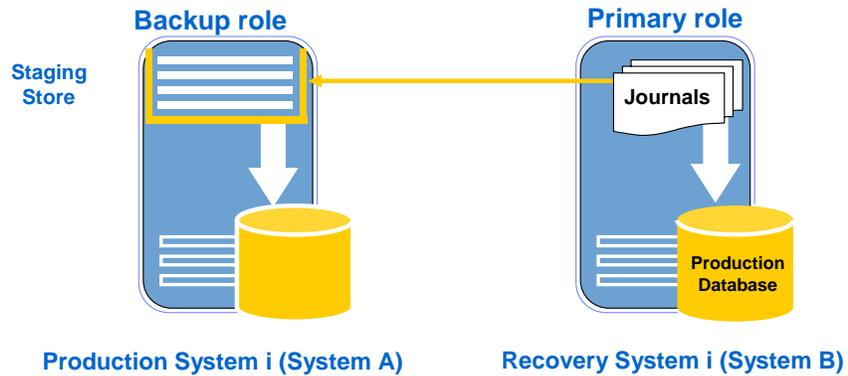
© 2010 IBM Corporation

Failover mode is when your production system is either lost or you have taken actions to 'quiet' the production system and perform the iCluster switch over; which changes the primary role to the backup recovery server.

In this mode you can see that since the users are now on the recovery system, any new changes are written to the journals on the recovery server. This server is now considered production.

This mode can be used for the time the original production machine is down for service or repairs. While the original production machine is down, all transactions are being staged to the journal receivers. It is important that no journal receivers are deleted until the resynchronization process has been completed or data is lost.

## Re-sync mode



Re-sync mode is performed when the original production server is again available. Before it can assume the production role again, it needs the changes stored on the recovery server.

For this reason the recovery server will continue to act as the production server until all transactions are sent and applied to the original production server. Once all transactions are applied then you can determine an appropriate time to switch back and allow users to again access the original production server.

## Tips for getting started (1 of 2)

- **Items to review and document**
- Use CHGJRN command to generate new receivers
- Use separate IP address for iCluster and general user connectivity
- Review and modify while switched to backup server
  - Startup routine after IPL
  - Journal cleanup
  - Sync check process
  - Review job scheduler as well to ensure the proper system is setup and backup is on hold
- Generate a Switching Checklist that guides users through the process

Some tips to help you get started and review what it will take to switch over in your environment include the use of the CHGJRN command to generate new receivers.

Some customers find it beneficial after the switch completes and before users gain access to generate a new journal receiver and reset the sequence number to '1'. This way you can easily see the first transactions generated by the user community as well you can easily verify the starting point for replication.

Other items to consider when reviewing your systems is how iCluster communicates and how the users communicate. If this can be separated, you can keep iCluster nodes communicating while keeping the ability to control when users can access the server and when they cannot. Similar to allowing users to connect, any items normally scheduled or configured for productions need to be listed as they will now be performed on the backup system when it is primary.

Since the list can be large, it is important to document the steps it takes to perform the switch. Document the actions required to end processes on production and then starting the same processes on the recovery server while it is acting as production.

## Tips for getting started (2 of 2)

- **Before you initiate a switch**
- Use DSPJRN on primary to view changes before issuing DMCHGROLE
- If transactions found, check user and job that submitted the transactions
  - Remove ability for new transaction or wait until completed before doing switch
- **After you complete the switch**
- Ensure DM\_PRIMPRE is checked before users are allowed to connect to new production

One most important step within any switch over is verifying if all transactions are processed and that there are no new transactions being generated. The easiest way to do this is to monitor the journals with the DSPJRN (Display Journal) command. If any transactions are found, you can easily see who or what job is generating them.

Similarly, once the switch over completes, review the iCluster software switch by reviewing the DM\_PRIMPRE job log which is generated for each replication group. This job is discussed in more detail in a later slide.

## Basic switching steps for a planned switch

### **On original production node (primary node)**

1. Quiet production system to ensure users are not connected and no new transactions are occurring
2. Verify all transactions have been sent and applied
3. End mirroring

### **On original backup node**

4. Process DMCHGROLE command to perform the switch
5. Verify iCluster software switch by reviewing DM\_PRIMPRE job log
6. Allow users and applications to function on NEW production machine (original backup)

The steps displayed on this slide are straight forward. They are basic steps to highlight critical points of a switch over.

These basic steps are for a planned switch. This is when you have access to both systems and are able to 'quiet' the production system at your own pace and review all system processes. The switch process is slightly different for an unplanned switch.

As you review and document your own switch routine, add more steps that perform tasks or document manual procedures that are required to complete the switch over in your environment.

## Basic switching steps for an unplanned switch

### On original production node (primary node)

1. Quiet production system to ensure users are not connected and no new transactions are occurring
2. Verify all transactions have been sent and applied
3. End mirroring

### On original backup node

4. Process DMCHGROLE command to perform the switch
5. Verify iCluster software switch by reviewing DM\_PRIMPRE job log
6. Allow users and applications to function on NEW production machine (original backup)

This slide displays the steps for an unplanned switch. The steps here in red because in an unplanned switch, the original production is not available so you move to step 4. As the production system is down there are no actions to consider to 'quiet' the production system.

Since the current production is not available full focus is now on the backup server and the steps that are required to make it a production system.

At this point, the focus is on making the backup system production but do not forget that when the original production system is available again, it will attempt to start as production. If you have turned your backup server into production, you need to bring up the original production server in a restricted state. The only items started are items required when the server is in the backup recovery mode.

Certain batch should be on hold, users should not be able to connect to the original production until the switchback is completed. These items and scenarios need to be discussed to ensure that users and batch are running where they should be and that your environment's connectivity is controlled.

## DMCHGROLE command

- Functions of DMCHGROLE
  - Submits JOB DM\_PRIMPRE
    - Starts apply job with force drain option
    - Enable triggers
    - Enable constraints
    - Starts journaling for objects (if option is specified)
    - Works with identity columns
    - Identifies marked position in journal used for start-up
    - Switches groups nodes so original backup node listed as primary node
- Review DM\_PRIMPRE job log

The DMCHGROLE command is used to do the switch over for each replication group. The command submits a job for each group called DM\_PRIMPRE. This job is submitted by the DMCLUSTER user and can be easily found on the backup server using the WRKSPLF command.

Be aware of the items processed by the DM\_PRIMPRE. The staging store that holds transactions on the backup server is drained and any remaining transactions are applied.

As normal replication disables triggers and constraints for replicated files on the backup server, the switch over process will now enable triggers and constraints. You can use this job log to verify that all triggers and constraints are enabled as part of the switch over.

A marked position is captured and used as the starting point for replication when the re-sync mode is started later. This is one main reason that it is critical for this DM\_PRIMPRE job to complete and be verified before users are allowed on the system. All user or batch transactions should be listed after the marked position point so no transactions are missed when data is resynchronized back to the original production server.

The final act of the software switch is to change the role of the groups now showing the backup node as the 'primary' node for each group.

Best practices for verifying the switch procedure in iCluster is based around review of the job log.

## When to switch

- All types of failures are critical to mirroring
  - TCP is ended
  - Interface failed ended
  - Other communication issues
  - AS/400® is lost – production failure
- Switch necessary or only mirroring affected
- Planned switch for maintenance

There can be many reasons to use the switch routine, such as the TCP is ended, the interface failed ended, other communication issues and a production failure. Also, any planned maintenance you have that sets a date and time that you do a planned switch over. Any of these issues will change the node to a failed status and the replication will end.

One item to consider for the switch over and when to do it is the severity of the issue. If there is a connection issue that causes iCluster to fail yet the users are still on production and normal daily transactions can be completed, a switch over might not be required.

If the connection issue affects some users and more can be connected on the backup being production, then a decision can be made to perform the switch over or not. Each production issue should be discussed to understand the severity of the issue and what is limiting production. From that discussion, the focus is on one of two issues. Either on the recovery from the issue and maintain the current production or, first the switch over to the backup server so normal production processes can be maintained while the disaster on the original production is being resolved.

## Planning for initial switch testing

- Plan for switching
  - Review production, users, applications, batch processing and connections
  - Review how to enable or start list from backup node while in production
  - Document approach in checklist form
  - Test checklist and switching of groups during off hours to verify connections, user access and applications
  - Perform quarterly switch tests

This slide highlights some of the items noted in previous slides. It is meant as a starting point to get you thinking about different aspects of switching and your environment variables that need to be considered when doing a switch over. Document your approach in a checklist form that contains the iCluster software switch and items required by your environment.

Test your checklist and the switching of your groups during off hours to verify your connections, user access and applications to ensure you are ready for a disaster.

Different approaches can be done to complete the initial switch tests. Additional review and testing should be completed before a full switch is initiated and confirm production users are processing on the backup node.

## Initial switch test - A

- Performed during business hours
- Common first step
  - End mirroring control
  - Perform full system save on backup node
  - Allow small number of users to connect and test
- No production changes performed
- Verify and note issues
- Restore backup system
- Production not affected

The initial switch test A is a common first step of switching, which is really not a switch at all. It can be done during business hours. One common test is to end mirroring control. Perform a full system save on the backup node, then allow a small number of users to connect and test the applications and connectivity on the backup node. Note that production changes should not be performed; only testing transactions, applications and connectivity.

This first step gives you a sense of how you can use the recovery system as production. Verify and note any issues users have so they can be resolved and not affect a proper switch.

Completion of this test is to restore the backup system as it was at the 'pretesting' stage and then restart mirroring and pick up where you left off. Also note the entire duration of this test. Users and batch are still on the current production machine. All changes are queued up in the journal receivers on production, so again production is not affected by this test.

## Initial switch test - B

- Actual switch over
  - Follow switch checklist
  - Minimal changes made
  - Limited users allowed access
    - Move production users to backup node
  - Document issues
- Main focus
  - Test checklist
  - Complete iCluster software switch
  - Successful user connectivity

Initial switch test – B is an actual switch over. Follow a switching checklist that you generated or obtained from Professional Services. If this is your first switch, do it in a controlled manner and only switching over for a short time. Minimal changes are made and limited number of users are allowed access. This is advised so you can control the environment while switched and if there are any issues, you can easily determine the changes made and decide any actions required before or after the switch back.

This switch test requires downtime for production. Move production users to the backup node to simulate a disaster. Document any issues in your switch over routine so you can fine tune your documentation and procedures. Your main focus is that your checklist is tested, the iCluster software switch completes and user connectivity is successful. Each test you do will allow your team to be more comfortable with the procedure and make your document or checklist more efficient.



## Need assistance?

- Contact IBM iCluster Support Team for questions regarding switching or command usage
- IBM Professional Services are available to help design, test and document a switching procedure

Contact the IBM iCluster Support Team for questions regarding switching or command usage. Also, IBM Professional Services are available to help you design, test and document a switching procedure for your environment.



## Feedback

Your feedback is valuable

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

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

Click to send email feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_iCluster\\_SwitchingConcepts.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_iCluster_SwitchingConcepts.ppt)

This module is also available in PDF format at: [./iCluster\\_SwitchingConcepts.pdf](http://iCluster_SwitchingConcepts.pdf)

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



## Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, AS/400, iCluster, and System i are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

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

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