IBM PureApplication System

Disaster recovery

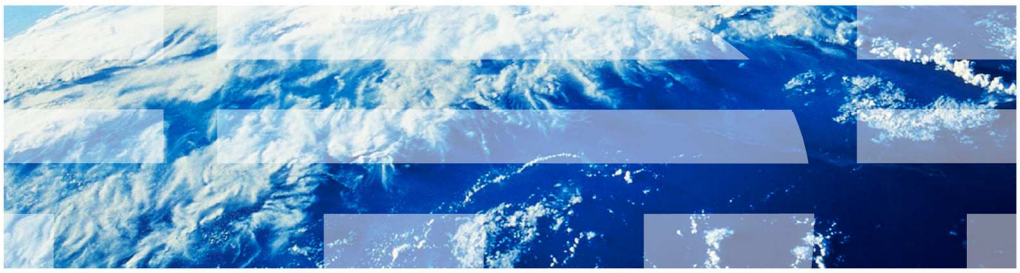

This presentation reviews disaster recovery in IBM PureApplication™ System.

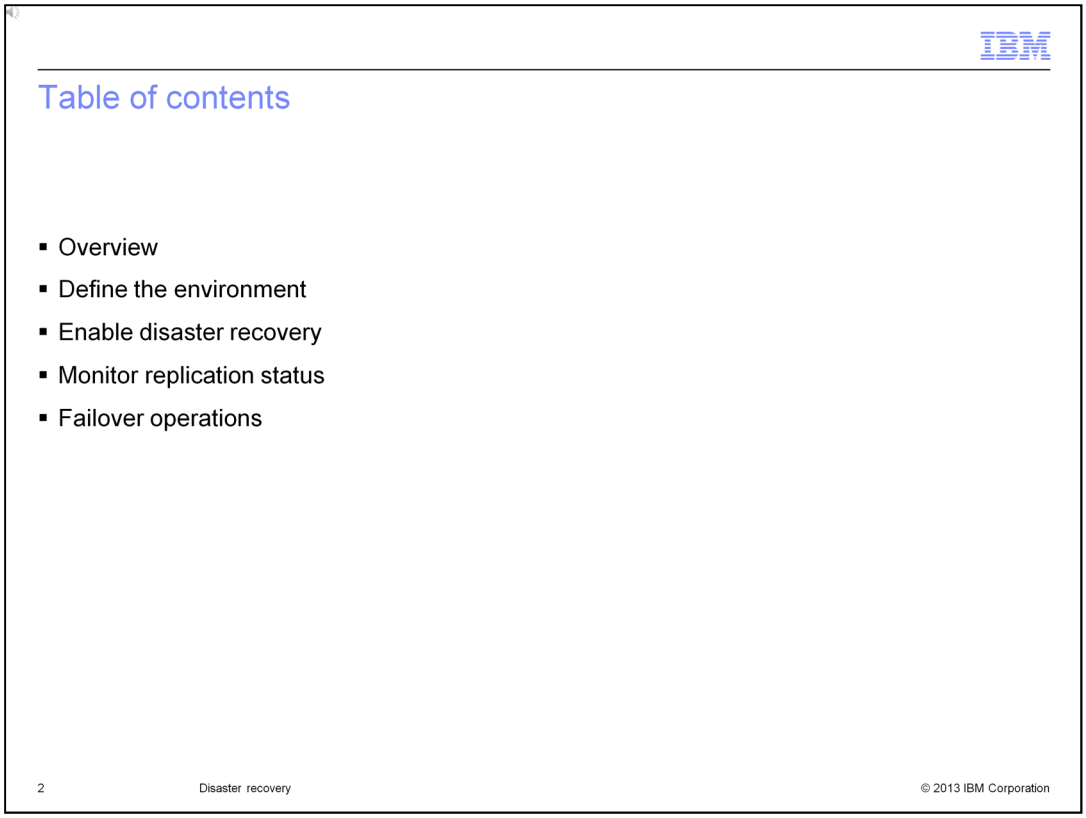

First, an overview of disaster recovery is presented, including capabilities, replication and life cycle. Then you see how to define and enable the disaster recovery environment. Once the environment is setup, you can use built-in monitoring functions to verify that the replications are successful. Finally, you review planned and unplanned failover scenarios.

Disaster recovery capabilities

- Address key business continuity scenarios including planned and unplanned failover
- Disaster recovery provides capability to failover from one rack to another
- Solution coverage for all application workloads
- Planned failover - Zero data loss
- Unplanned failover - Near zero data loss
- Time to recover – Measured in several hours to restart workloads on backup and update external network routers
- Distance between primary and backup racks – Up to 8000 kilometers

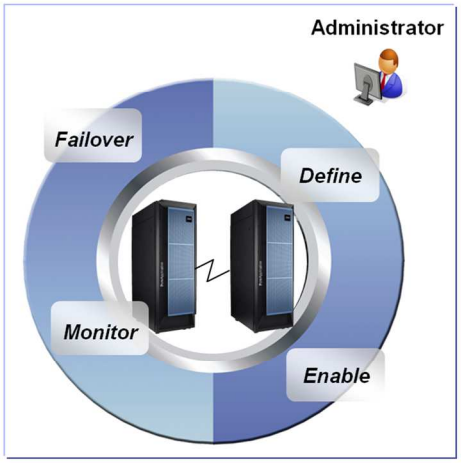

The disaster recovery capabilities in IBM PureApplication System cover both planned and unplanned scenarios. For planned failover, you can prove that the disaster recovery environment is setup properly by planning a failover and running for a period of time on the backup system. For unplanned failover, you can quickly recover from primary site failures or rack failures, by starting your workloads on a backup rack. The backup rack is continually updated with changes as they happen on the primary rack, and all workloads are covered in the environment.

For a planned failover, no data loss is experienced. For an unplanned failover, loss of data is limited to data that is in transit from an end-user of a deployed application or an unsaved administrative action. Disaster recovery includes a primary rack which uses asynchronous storage replication to keep a backup rack up-to-date with only seconds of data loss. When an unplanned failover happens, it should take only a few hours to initiate the failover on the backup rack, restart workloads and advertise IP addresses on the external routers.

Using global mirroring replication in the V7000 storage unit, the maximum distance between the primary rack and the backup rack is 8000 kilometers. You can achieve this using fiber channel extenders or storage area network routers.

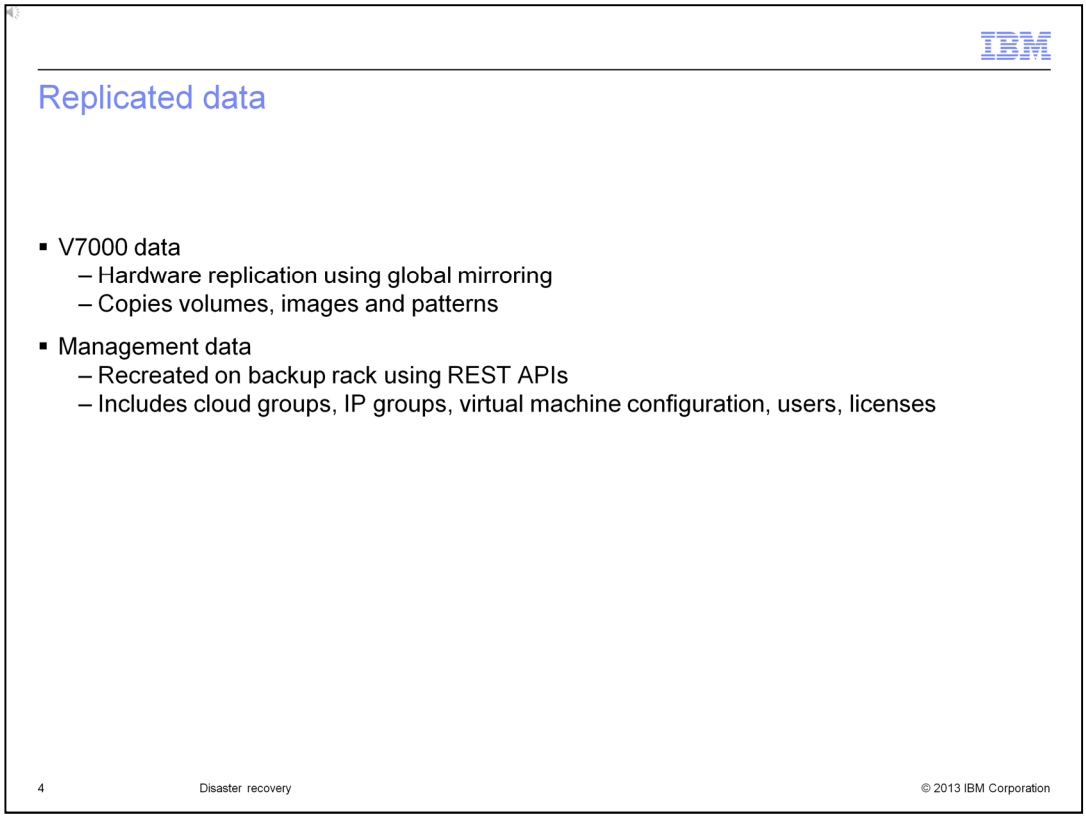

The V7000 storage unit uses global mirroring to make asynchronous copies of your disk. This means that the write is considered complete after it is complete at the local disk. It does not wait for the write to be confirmed at the remote cluster.

Management data is replicated through software using REST API's. This includes cloud groups, IP groups, virtual machine configuration, user information and licenses.

Note that after failover virtual machines are started up with the exact same IP addresses and VLAN's so you cannot bring up workloads on the primary rack.

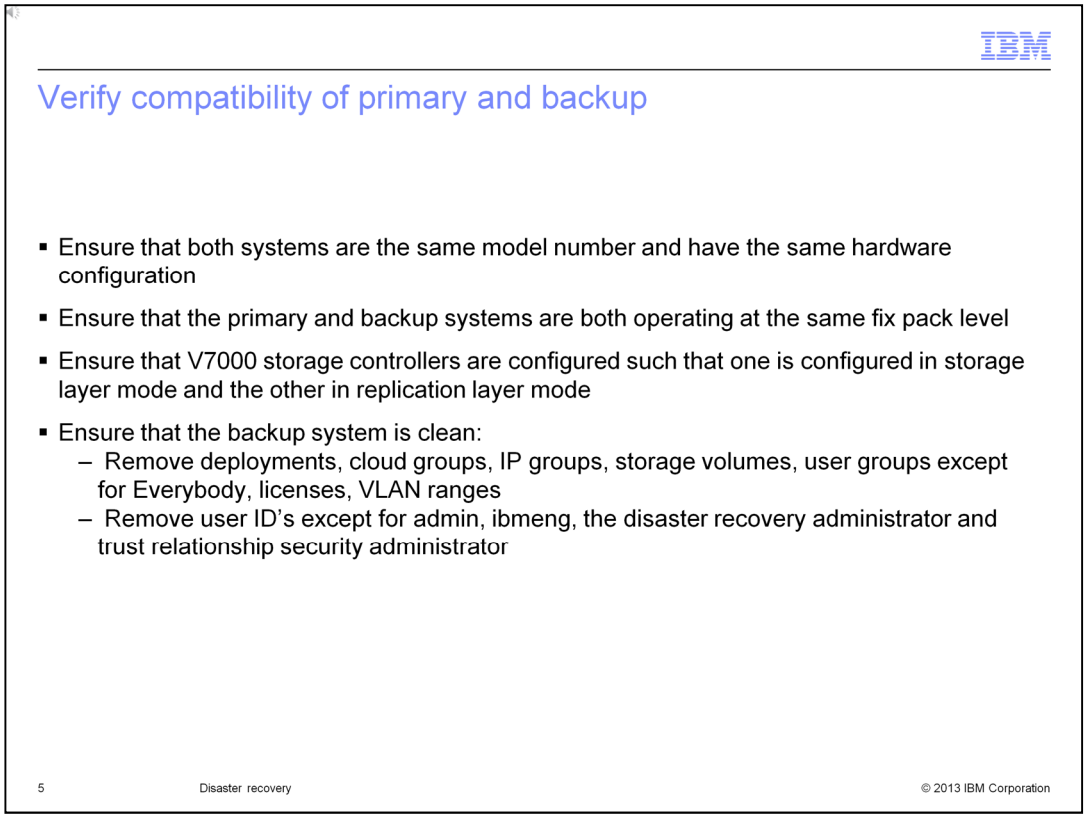

The primary rack and the backup rack must be compatible in a disaster recovery environment. There is a verification function in the user interface so you can perform these checks. Also, when you enable the disaster recovery profiles there is automatic verification with error messages for any failures.

You should ensure that the primary and backup systems are of the same hardware platform. For example, a W1500 system can only backup another W1500 system. A W1500 cannot backup a W1700 system. The racks should be the same model number with the same number of compute nodes. Also, both racks should have the same fix pack levels.

In previous releases, racks were shipped with both V7000 storage controllers configured in storage layer mode. To accommodate disk replication, one must be configured in replication layer mode while the other remains in storage layer mode.

The backup system should be a clean system without workload deployments and without configuration objects such as cloud groups, IP groups, volumes and user groups other than the 'Everybody' group, licenses and VLAN ranges. Also, remove user ID's except for admin, ibmeng, the disaster recovery administrator and the security administrator for the trust relationship between racks.

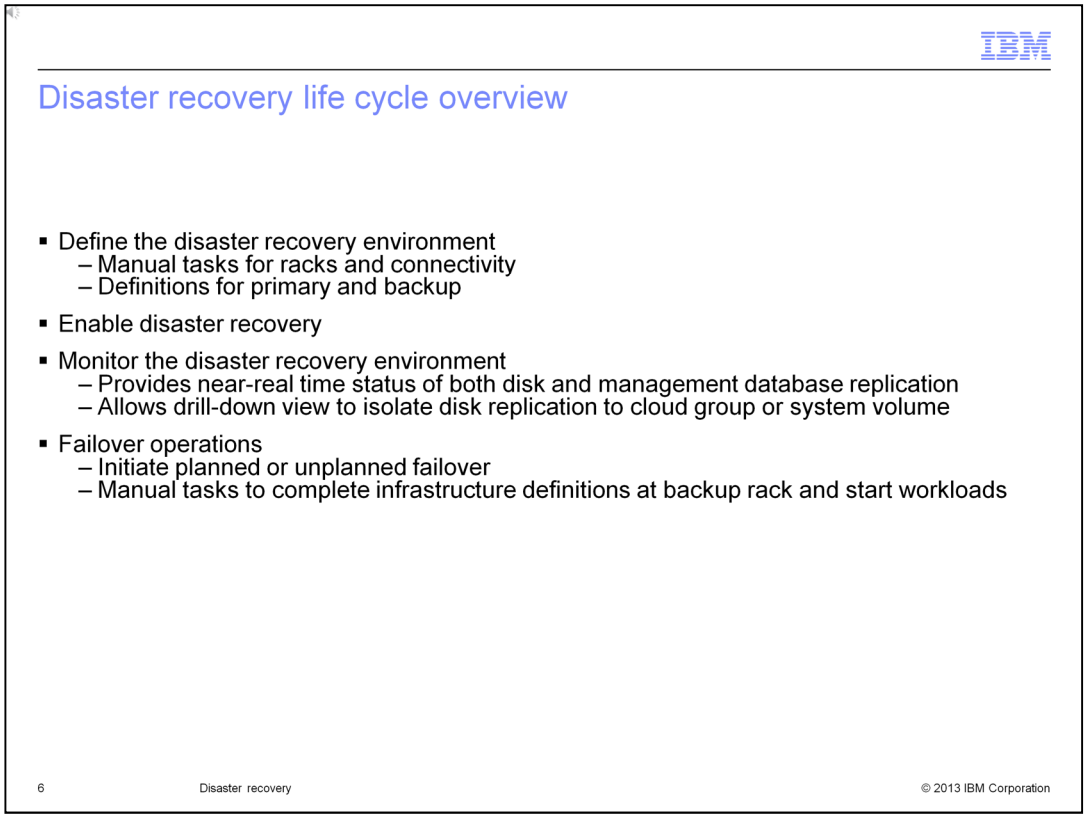

This slide shows the life cycle for the disaster recovery environment.

When defining the environment, there are a few manual tasks that are required to prepare the racks for disaster recovery and to setup fiber channel connectivity. Then, you create disaster recovery profiles, one on each rack in the disaster recovery relationship and you need to validate that the environment is configured properly before enabling disaster recovery.

When you enable disaster recovery, replication begins for the disks and management data. For a fully loaded V7000, it can take some time to do the initial copy, but after that, it is just delta copies. You enable the profile on the backup rack first, then you enable the profile on the primary rack.

Monitoring functions are provided to check the replication status between the racks. You can see the status of disk and management data, and drill down to check status by cloud group or specific volumes.

There are failover operations for both planned and unplanned failovers. For a planned failover, you start the failover on the primary profile, then follow that with a failover operation on the backup profile. For an unplanned failover, you initiate the failover operation on the backup profile. After failover, there are manual tasks required to start selected workloads and advertise IP addresses.

This section shows you how to define the disaster recovery environment.

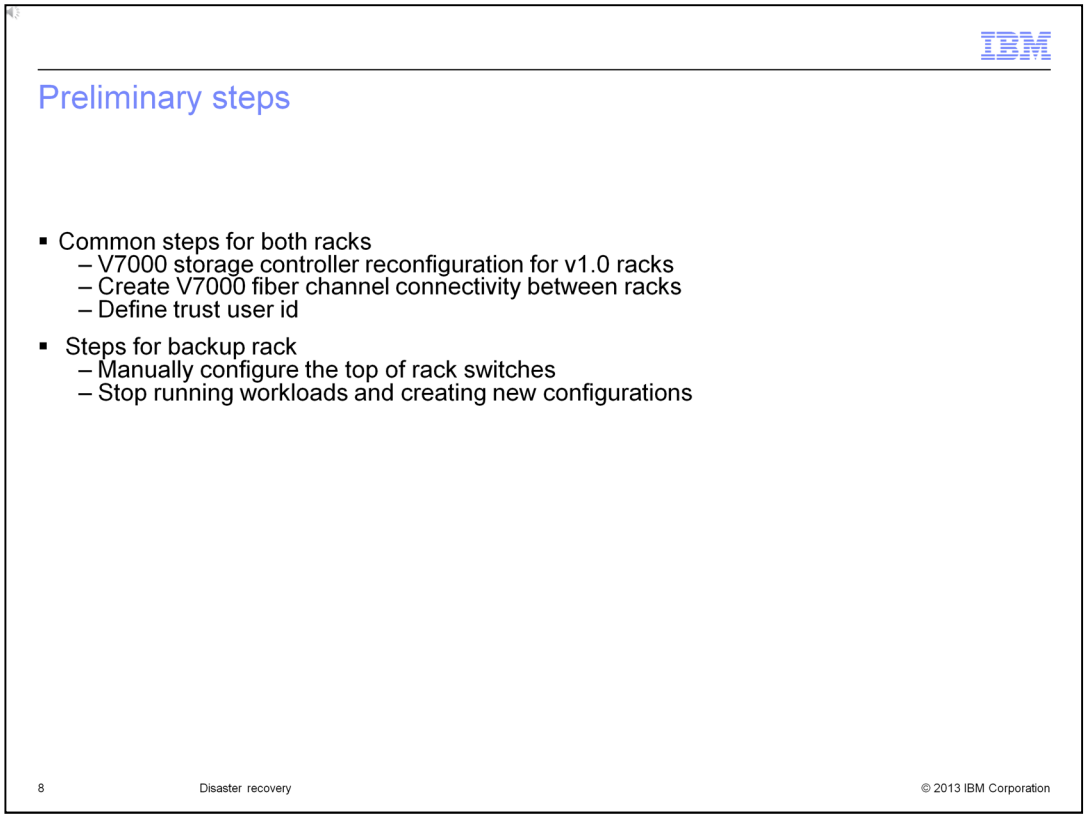

In previous releases, the racks were shipped with both V7000 storage controllers configured in storage layer mode. When configured in this manner, both V7000 storage controllers operate independently of each other and as a result, volumes managed by different V7000 storage controllers cannot be included in the same consistency group. To overcome this limitation, one of the V7000 storage controllers on a rack is configured in replication layer mode while the other remains in storage layer mode. New racks are already configured in this manner, however racks for previous versions need to be reconfigured.

A fiber channel connection to a device such as the Brocade Fiber Channel switch is required to enable the connection of storage controllers. Creating this connection involves creating a zone for the ports on the V7000 storage controllers. To create the zone, you need the world wide port number so this has been added to the user interface. You can access it in the menu: Hardware > Storage devices > Storage node.

When disaster recovery is enabled, the primary rack issues REST API calls to the backup system. A trust relationship simplifies exchanging data between primary and backup systems. You need to use the same user ID and password on both systems, and this user ID must have the security administrator role.

You need to manually configure the top of rack switches on the backup rack such that the VLAN assignments to external ports are equivalent to those of the primary rack. You should make sure that the switches are configured similarly from the beginning and that any changes to the primary rack switches are replicated in the backup rack' switches as they occur.

There is a restriction that the backup rack is dedicated for recovery. Workloads might not run on a rack configured as a backup. There should be no configuration performed on the rack other than the configuration required to be a backup. Any unexpected configuration on the backup rack prevents the backup role from being enabled.

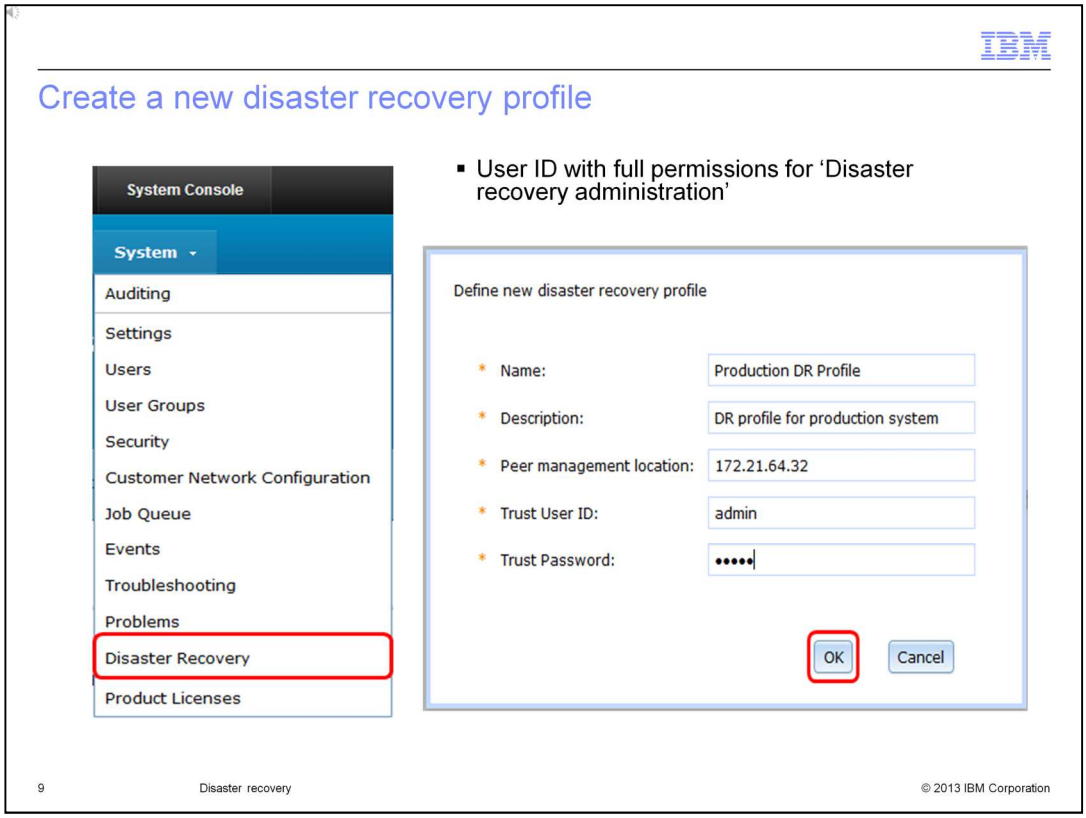

You need to create a disaster recovery profile on the primary rack and the backup rack. You can access the disaster recovery page in the system console within the System menu. To create a profile, your user ID should have full permissions for 'Disaster recovery administration'.

When creating the profile, you specify the peer management location which is the peer's management IP address. You also specify the user and password for establishing the trust relationship. This user should have full permissions for 'Security administration'.

Defined disaster recovery profile

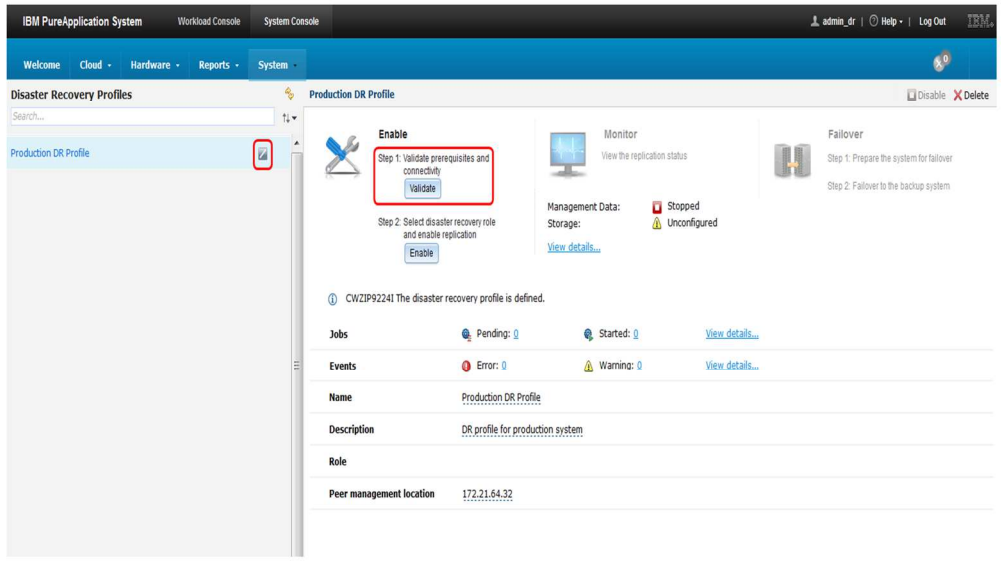

This shows the profile after it has been created. There are three sections on the profile which are Enable, Monitor and Failover. The latter two are disabled since the profile has not been enabled. But the next step is to validate that the rack is ready to participate in a disaster recovery relationship.

Validated disaster recovery profile

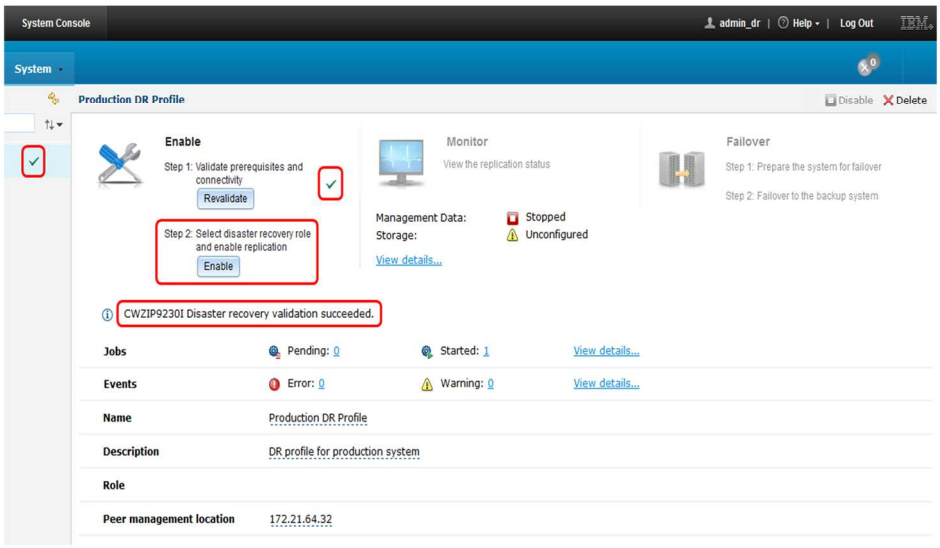

This slide shows a profile which has been validated, and a message indicates that the validation succeeded. Notice that the state of the profile changes to ‘Validated’ which is indicated by the green check mark. The next step is to enable the disaster recovery role.

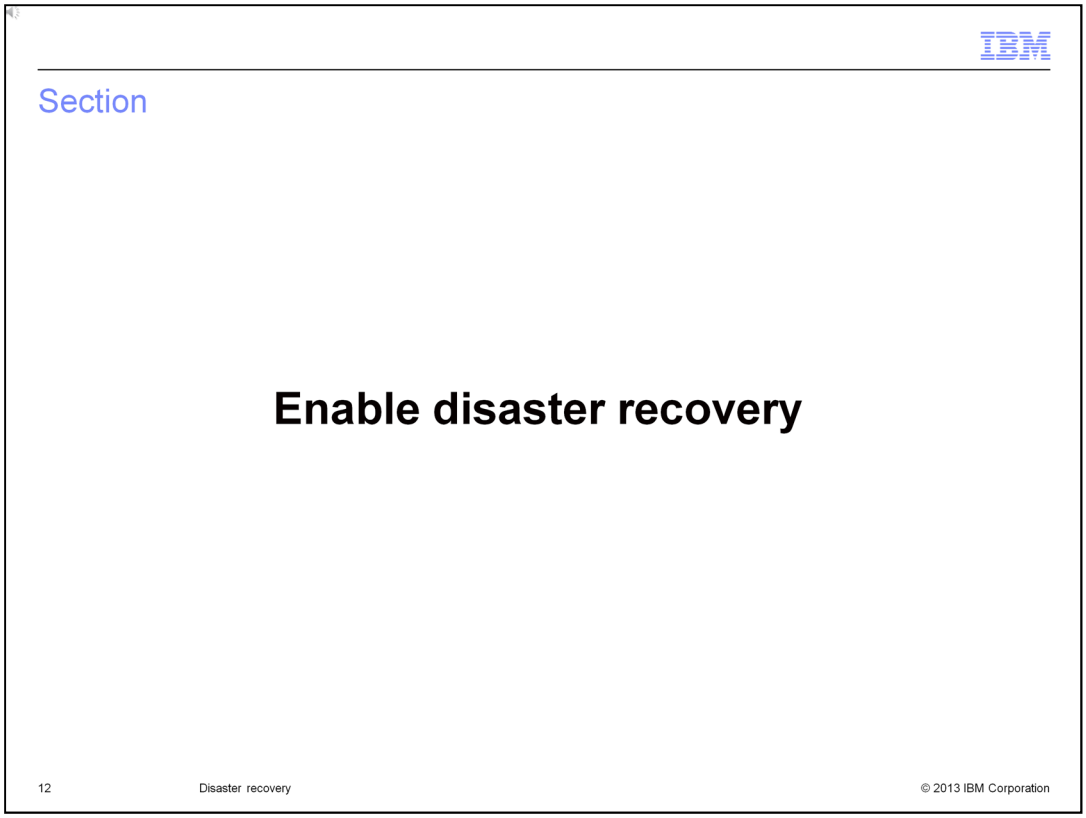

This section shows you how to enable the disaster recovery environment.

Enable disaster recovery

- Backup rack must be enabled before enabling the primary rack
- Once enablement is started on primary:
  - Initial copies for disk and management database replication from primary to backup system are started
  - Provides replication status information on both primary and backup system

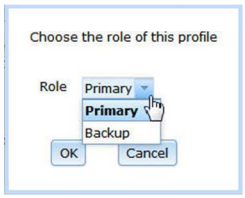

The backup rack must be enabled before the primary rack. If backup rack is not enabled, then the enable action on the primary rack fails.

On the profile page, select the button labeled ‘Enable’, then you are prompted to choose the role, either ‘Primary’ or ‘Backup’.

After enablement is complete for both racks then initial copies begin for the disks and management data. This process can take a long time to complete depending on the network speed and size of the data. After the initial copies are complete, then the changes are replicated as they occur. You can monitor the status of the replication at any time on either of the racks.

Backup role enablement in progress

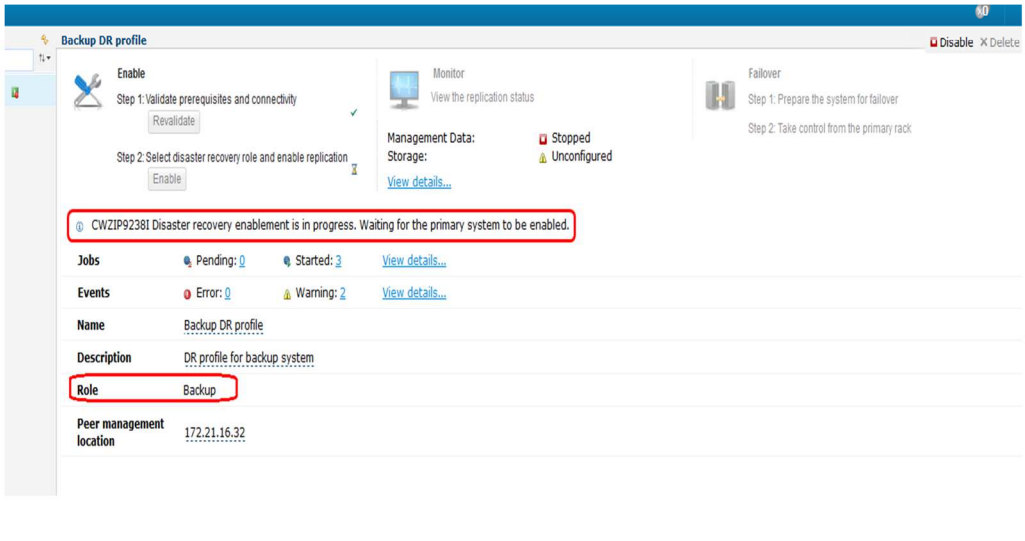

On this slide, you can see the backup profile after enable has been selected. A message indicates that enablement in progress, because the primary rack has not been enabled yet. The role field shows that it is in the backup role.

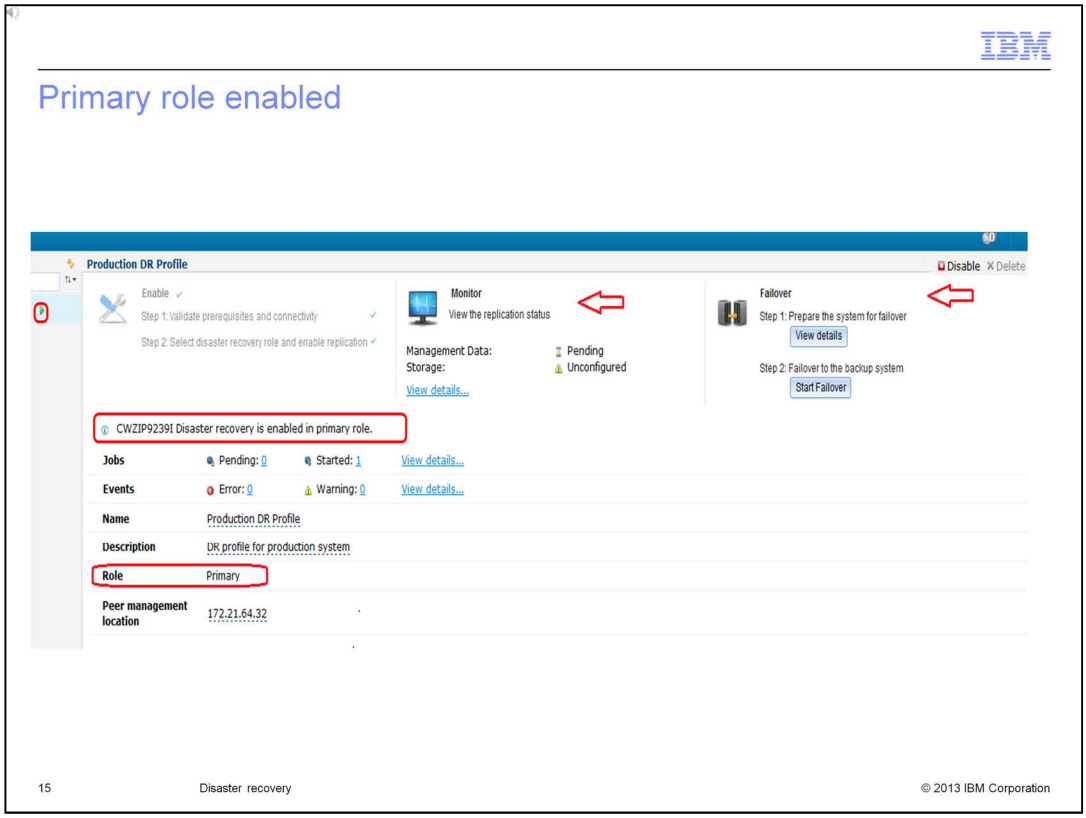

Here you can see the primary profile after it has been enabled. A message indicates that disaster recovery is enabled in the primary role. The role field also shows that it is in the primary role. Notice also that the Monitor and Failover sections are now available for your use.

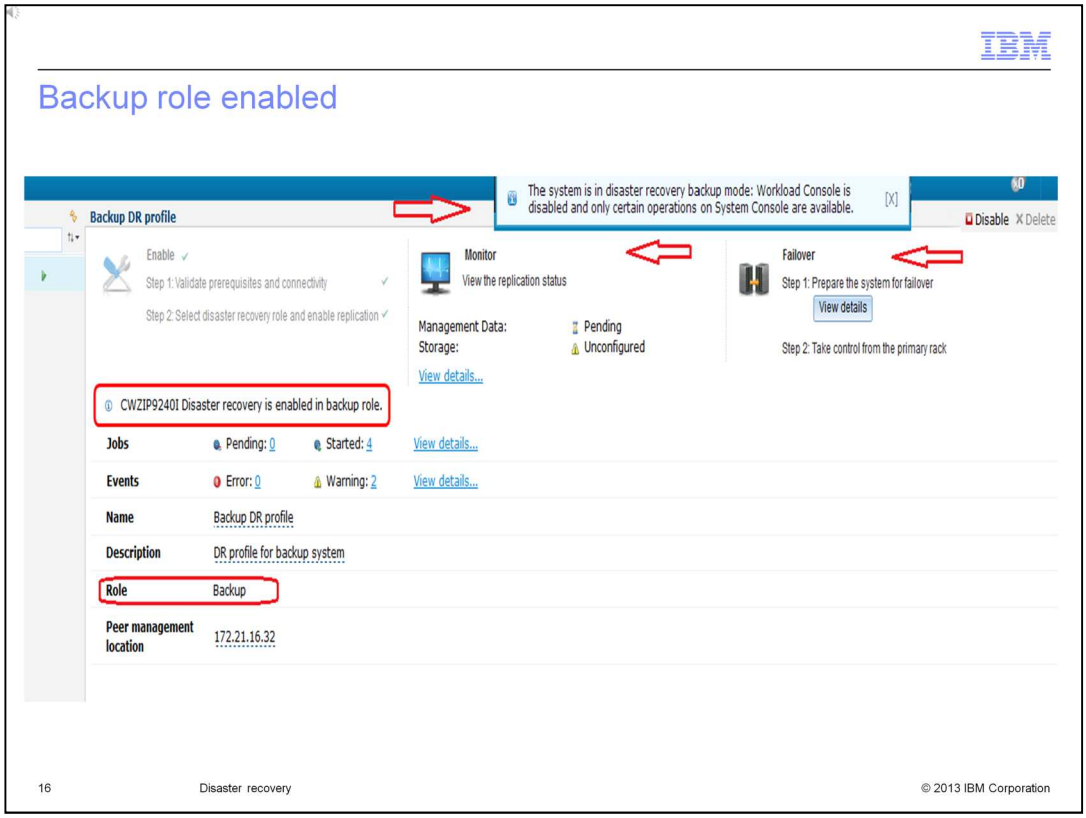

On this slide, you can see the backup profile after the primary profile has been enabled. A message indicates that backup rack is now enabled. The role field also shows that it is in the backup role. Notice also that the Monitor and Failover sections are now available for your use and the Workload Console is disabled.

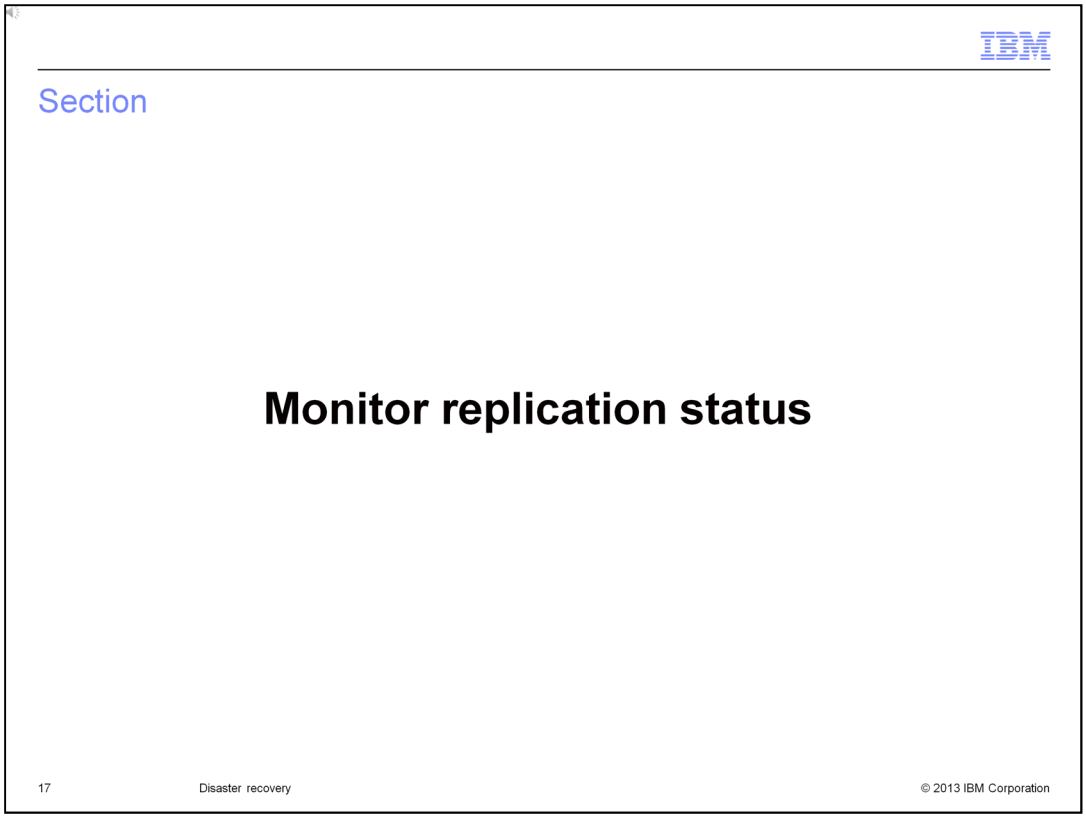

This section shows you the monitoring capability for checking disaster recovery status.

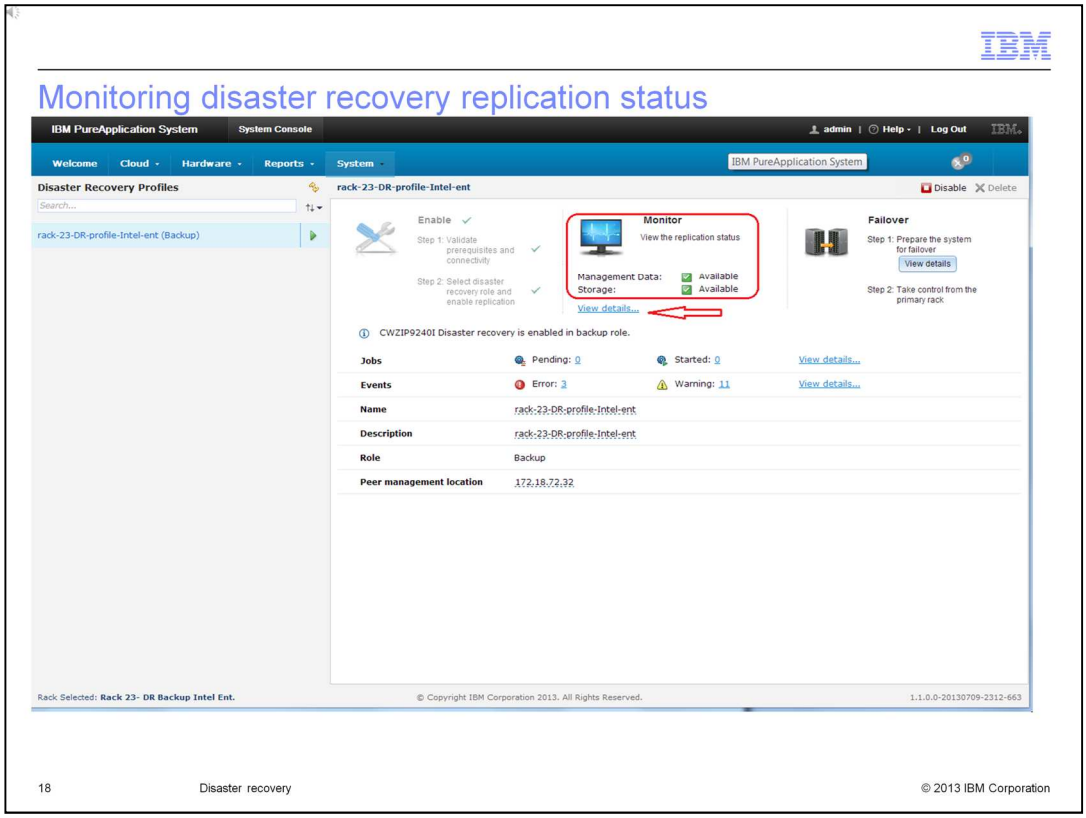

Monitoring is provided on both the primary and backup racks and the information displayed is similar. A summary of the replication monitoring status is displayed on the disaster recovery profile panel. There are two types of replication monitoring: management data monitoring and storage monitoring. Detailed monitoring information for both is obtained by clicking on the “view details” link.

Full visibility to real-time replication status

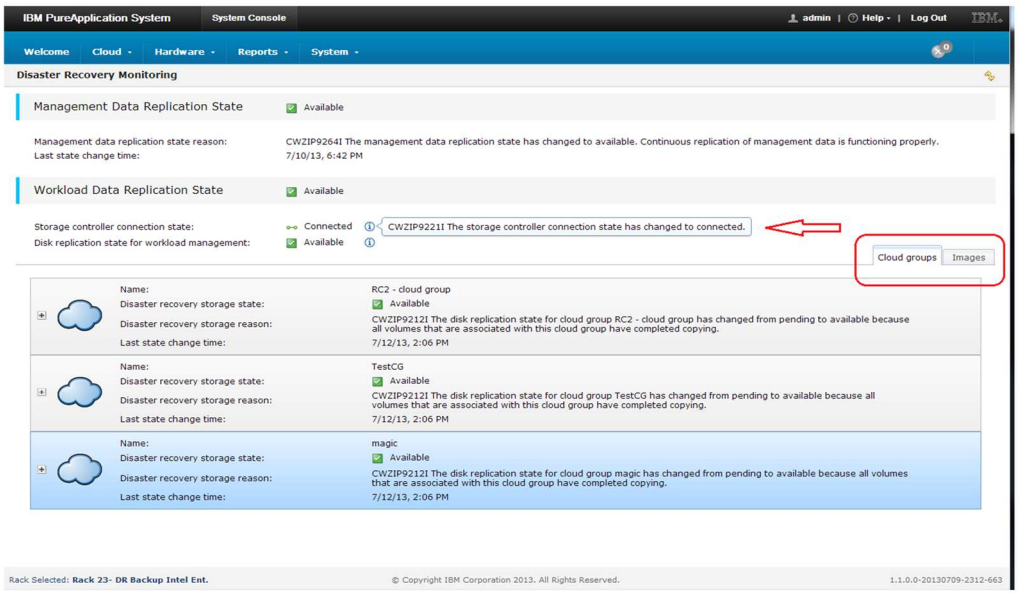

You can see the replication state of the management database and disk. For disk replication, you can drill-down to volumes by cloud group, and you can view replication by images. There are two tabs, one for cloud groups and one for images.

You can obtain additional information about the storage controller connection state and the disk replication state for workload management by hovering the cursor over the information icon.

Drill down to see detailed status

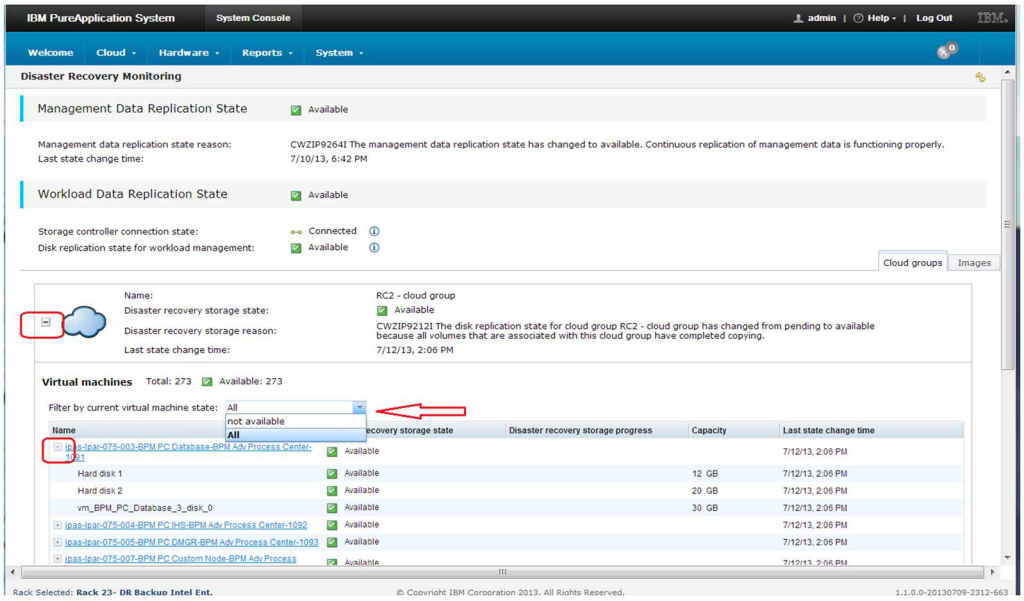

Here you see the cloud group tab which shows the status by cloud group. One of the cloud groups has been expanded so you can see more detailed information. In the drill down for this cloud group, you see a list of virtual machines. One of the virtual machines has been expanded so you can see detailed information about its disks. If an initial copy of a disk was being performed, the percentage complete would be shown as well.

By default, only the virtual machines that are not available are shown. To see all virtual machines, select “All” for the drop down box “Filter by current virtual machine state”.

Image status

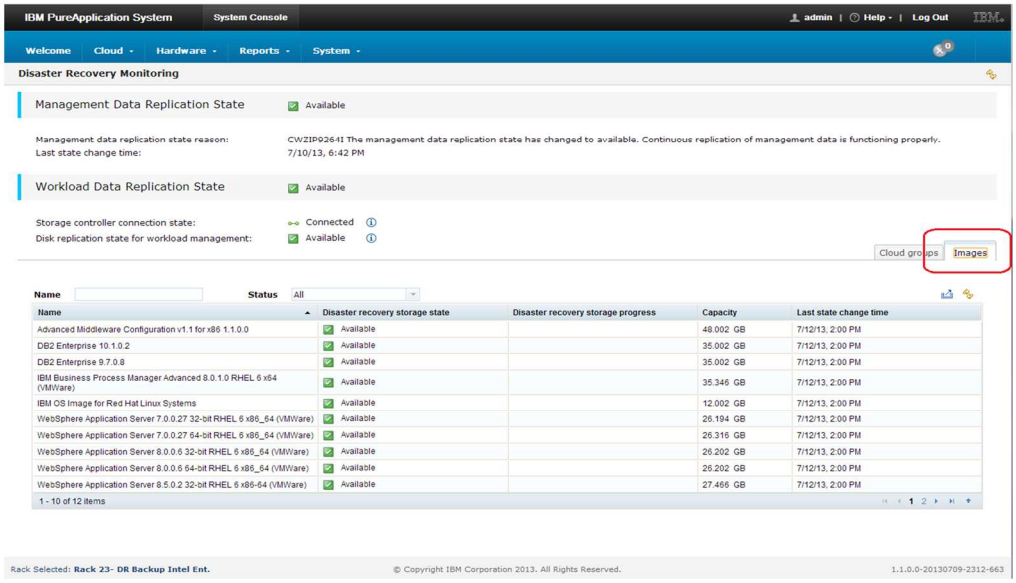

On this slide, you see the images page where the images are listed along with disaster recovery storage state, storage progress, capacity and last state change time.

Fiber connection broken – Monitoring summary

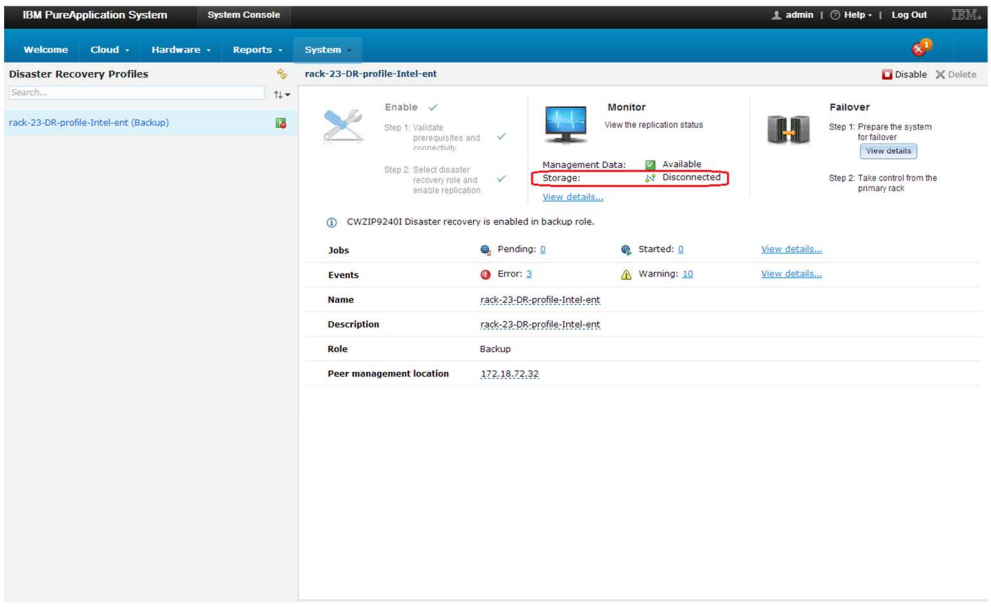

22 Disaster recovery

When a storage replication error occurs, the status in the storage monitoring summary on the disaster recovery profile page changes from ‘Available’ to ‘Disconnected’, ‘Failed’, or ‘Stopped’. In this case, the fiber connection is broken, which creates a storage replication error and the status changes to disconnected.

Fiber connection broken – Events issued

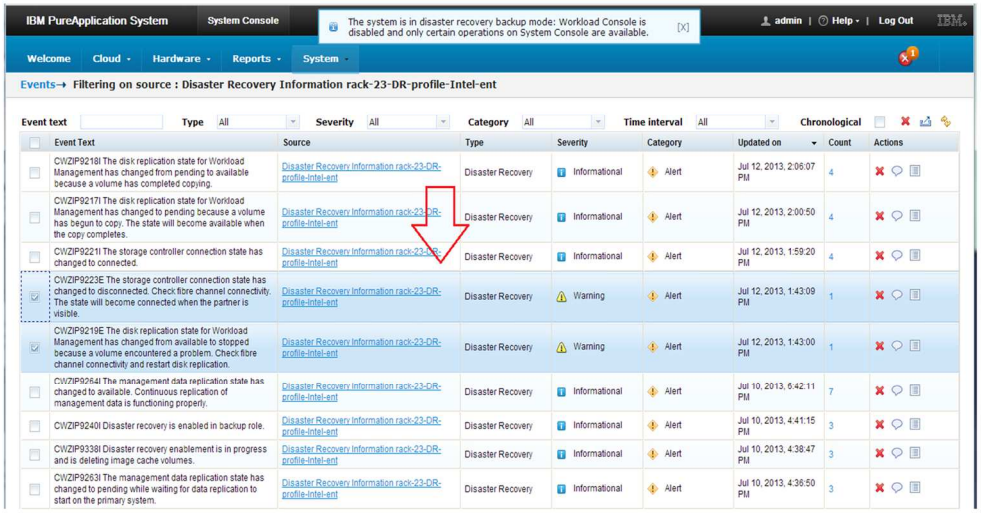

23 Disaster recovery

Most likely, you are not monitoring the disaster recovery profile looking for a change in the storage monitoring summary. So an event is triggered when a disk replication error is encountered. As shown in this slide, two events were issued when the fiber connection was broken.

Fiber connection broken – Monitoring details

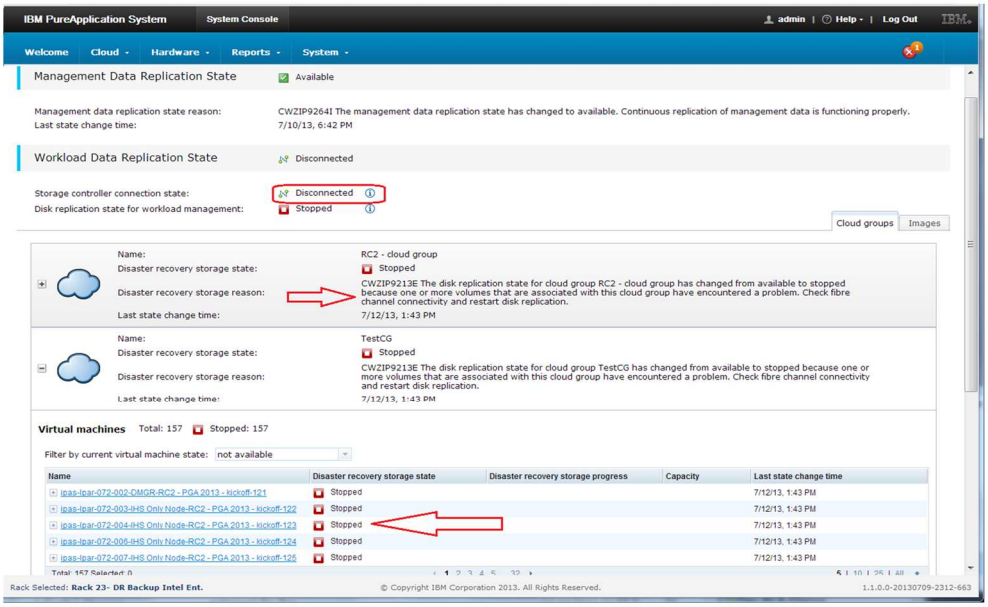

Additional diagnostic information may be required when a disk replication error event is triggered. The monitoring details page provides an explanation of disk replication issues along with diagnosis and potential corrective actions. In this case, since the storage area network connectivity was lost the storage controller connection state is disconnected. Since the storage controller state is disconnected, the status of all disk replication is also stopped. Since the status of all disk replication is stopped, the status of all virtual machines is stopped.

Fiber connection restored – Monitoring details

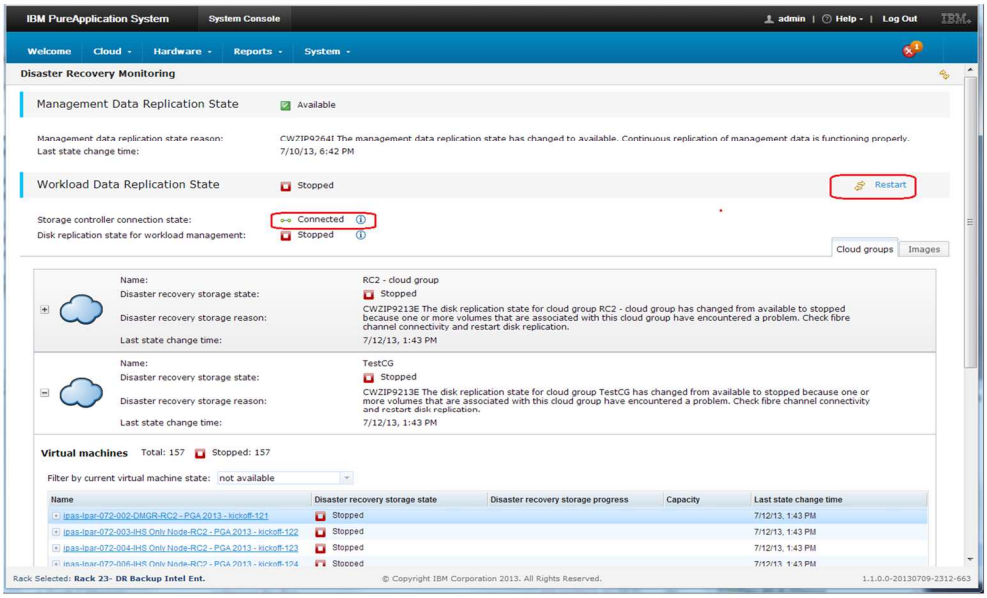

Once the fiber connection is restored the storage controller connection state becomes connected. The states of the disks and virtual machines remain stopped. An option to restart disk replication opens. While in the stopped state, the disks are in a consistent state. Restarting disk replication temporarily puts the disks into an inconsistent state while it catches things up. A planned failover cannot occur while disks are in an inconsistent state.

Disk replication restarted – monitoring details

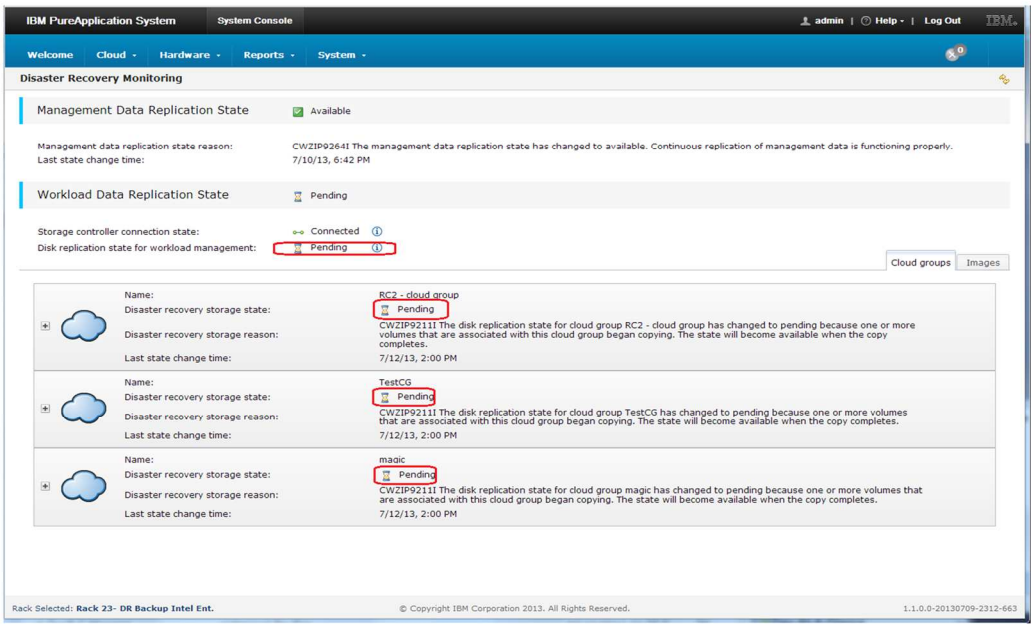

Once restart is selected, the disk replication state moves to pending until all remote copies catch up. Pending indicates that a disk is in an inconsistent state. The more common reason for a disk to be in an inconsistent state is that the initial copy of the disk is being performed.

Disk replication returns to available – monitoring details

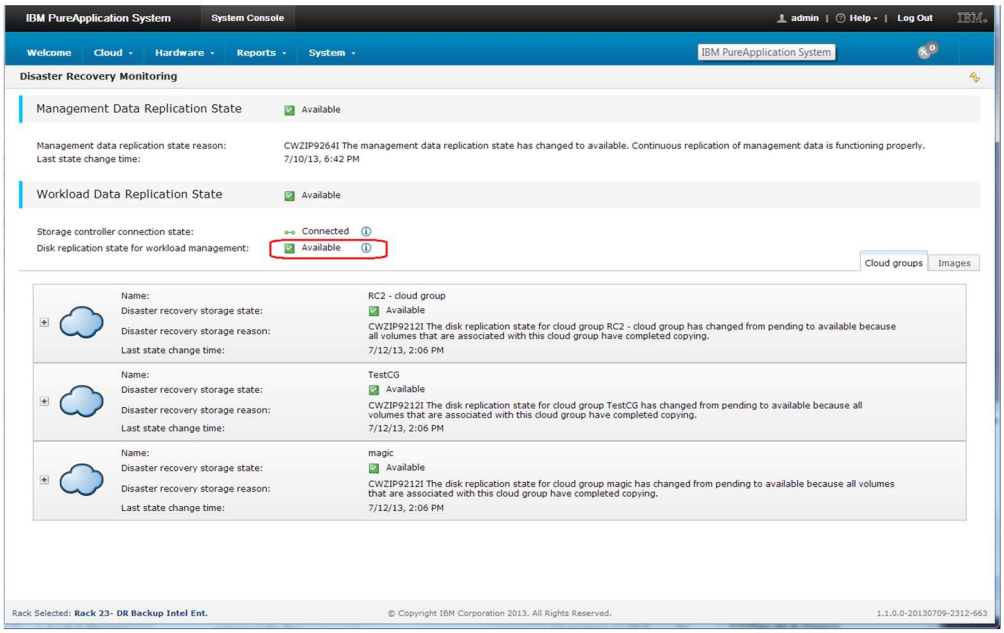

Once things catch up, the disk replication state returns to available.

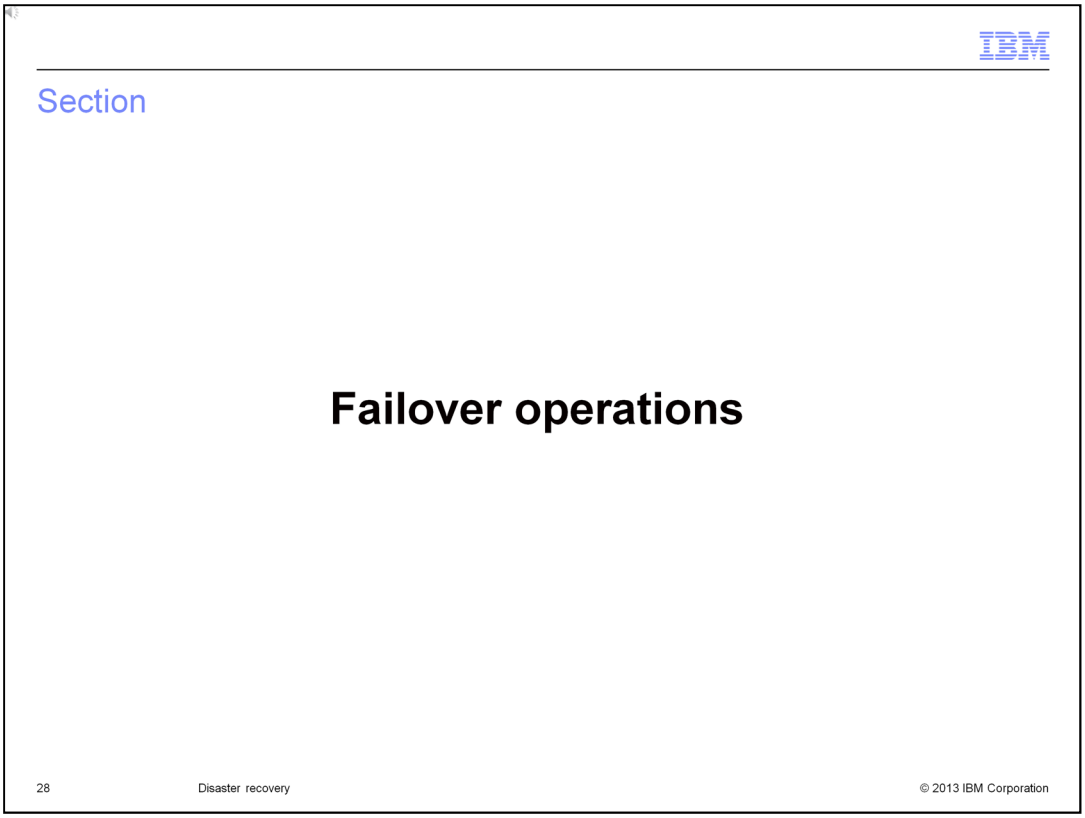

This section covers failover operations for planned and unplanned scenarios.

Planned – Prepare to failover on primary

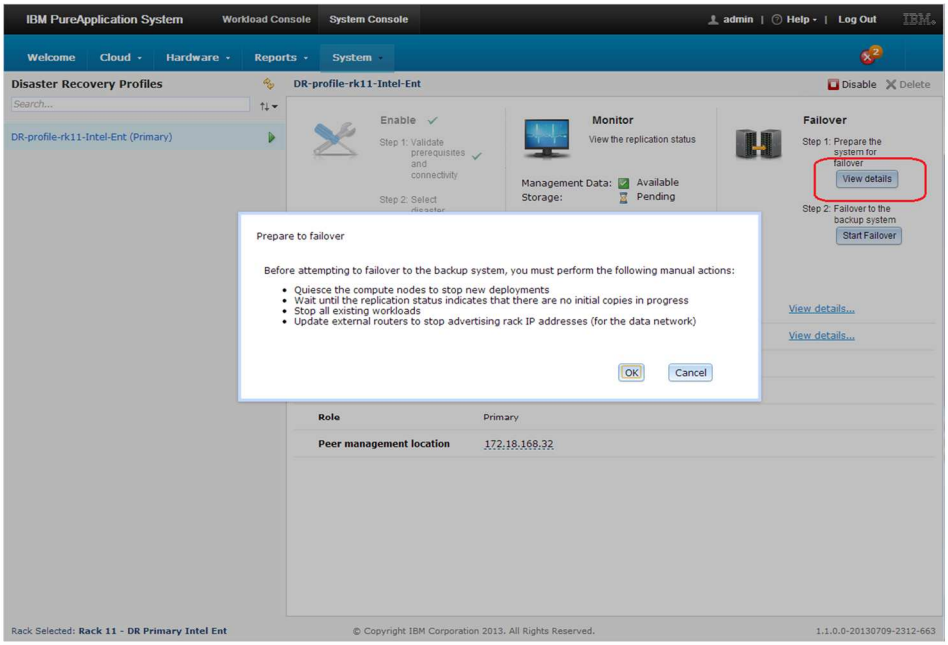

For a planned failover, you initiate it on the primary disaster recovery profile, then follow up with the actual failover on the backup rack. On the primary profile in the first step, click ‘View Details’ and a popup reminds you of the manual actions that need to be performed. You should quiesce the compute nodes to ensure no new deployments take place. Also, verify that the replication status shows that the initial copies are complete. You should stop the workloads on the system and update the external routers to stop advertising the rack IP addresses.

Planned – Failover on primary

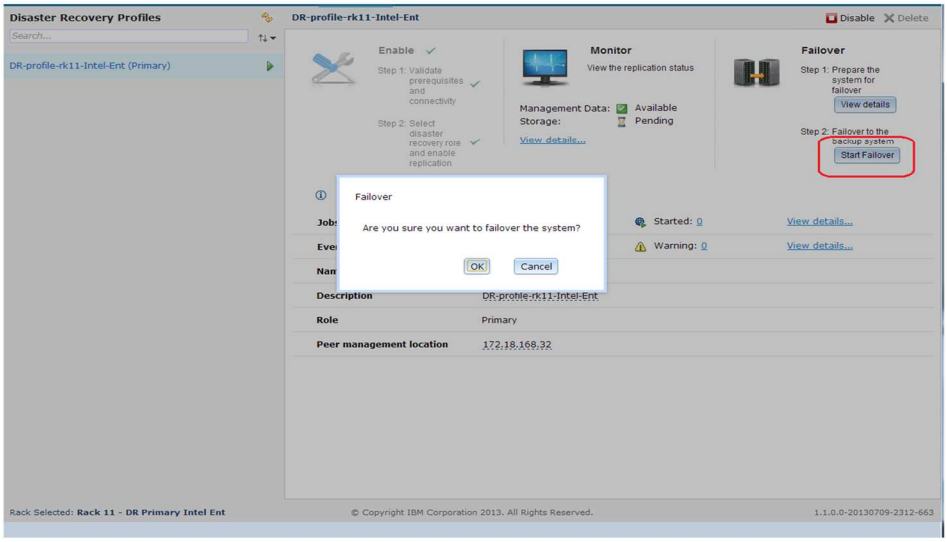

- After failover, role shows '(none)' and state shows 'Validated'
- Need to complete the failover on the backup

When you are ready to perform a planned failover, on the primary profile click 'Start Failover' and then click 'OK'. After it is finished, the operation the profile role shows '(none)' and the profile state shows 'Validated'. At this point, a failover on the backup must be performed before workloads can be started on the backup rack.

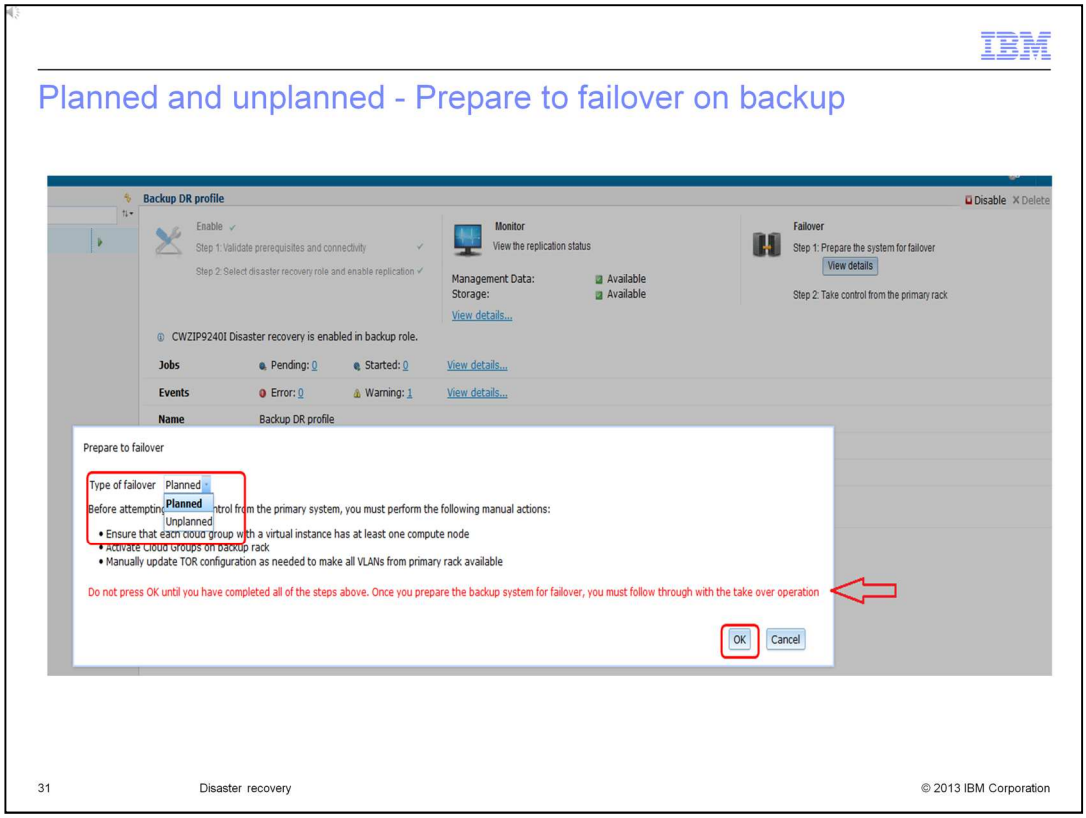

For both planned and unplanned scenarios, it is a two step process to initiate the failover on the backup rack. In the backup profile click ‘View Details’, then in the dialog select the type of failover, either Planned or Unplanned. In the dialog, it reminds you of the manual actions to perform on the backup rack, including adding compute nodes to the cloud groups and configuring the top of rack switches. Once these manual tasks are complete, click OK. Adding a compute node to a cloud group will activate the cloud group.

Planned and unplanned - Failover on backup

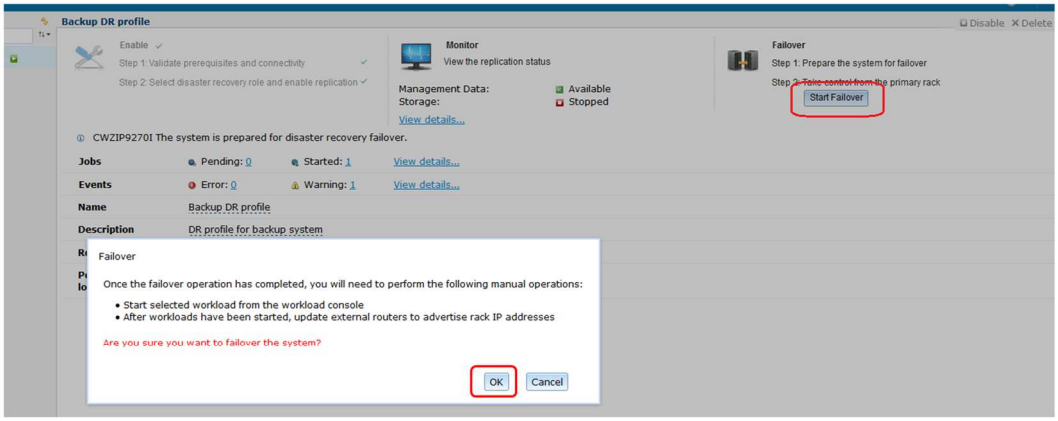

- After failover, role shows ‘(none)’ and state shows ‘Validated’

When you are ready to perform the failover, on the backup profile click ‘Start Failover’ and then click 'OK'. After completing the operation, the profile role shows ‘(none)’ and the profile state shows ‘Validated’. After the operation completes, you can start your workloads and update the routers to advertise rack IP addresses.

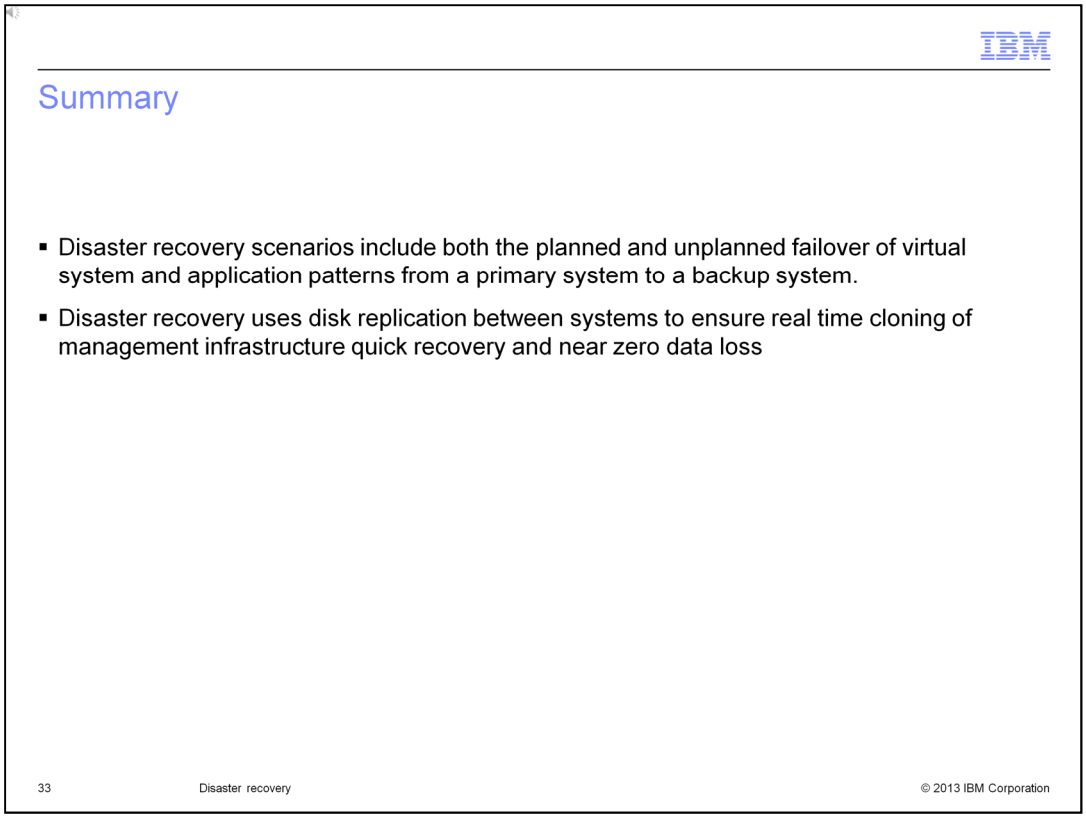

In this presentation, you have reviewed the disaster recovery scenarios including planned and unplanned failovers. You have learned about replication for disk and management data, monitoring replication state, and failover operations.

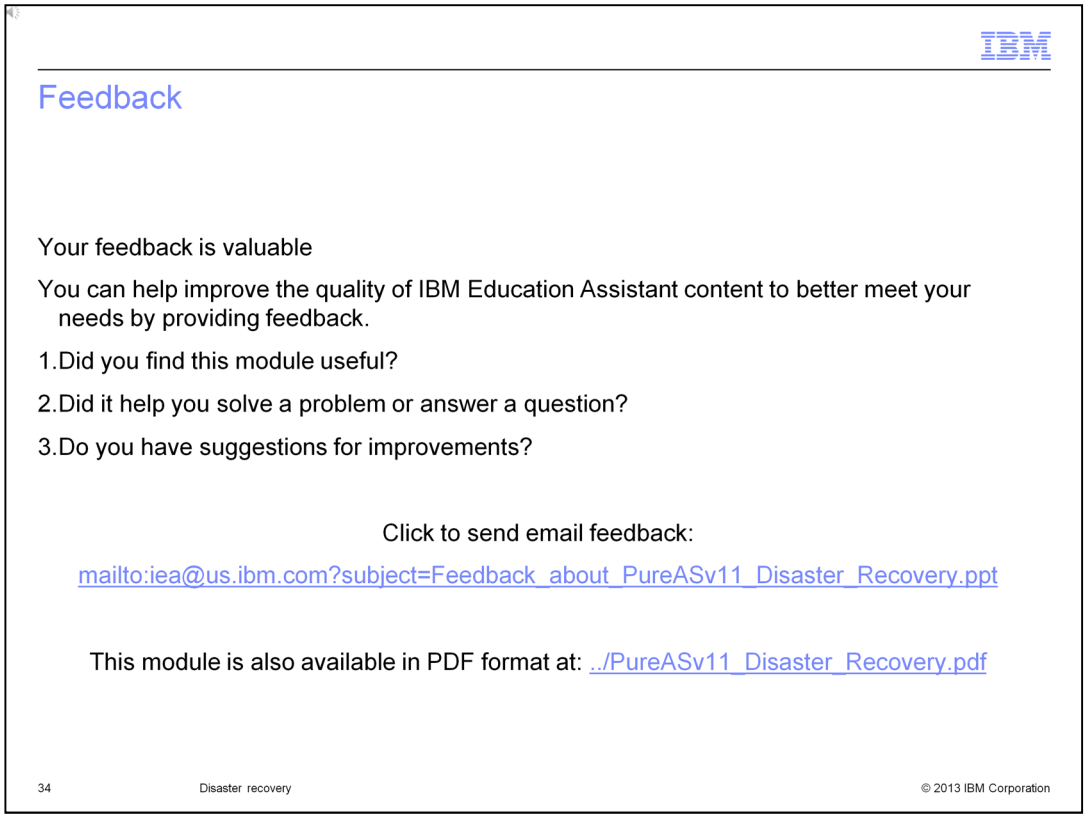

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, disclaimer, and copyright information