Content Restoration and Fail-Over

Some common uses of Content Router are disaster recovery, backup, and archiving. In all these cases, it will sometimes be necessary to recover lost content data from a remote cluster. The exact nature of such a recovery process, and the ease and speed with which it happens, will depend somewhat on the topology of the cluster network.

Administrative Disaster Recovery

If the topology involves a one-way connection between the primary and DR clusters, then recovering from a disaster that causes loss of data, perhaps because one or more nodes in the primary cluster have been destroyed, must involve an administrative process to reverse the data flow. That's because the primary cluster in such a topology is running only a Publisher service, which moves data from the primary to the DR. In order to move data in the other direction for recovery purposes, the primary cluster (and perhaps also the DR cluster) will need to be reconfigured to run one or more Replicators, allowing the DR site to reconstitute lost data in the primary. Any lost data will be unavailable until the loss is discovered and corrected by the two Content Router nodes.

> ***Note:*** *Administrators who temporarily install a Replicator or Publisher service on a Content Router node as part of a DR event should plan to uninstall the temporary service when the recovery is complete. Failure to do so may result in unintentional creation of a mirrored configuration when the cluster returns to normal since both services will be installed on the same server.*

Content Mirroring

In a mirrored architecture where the primary and DR clusters are both running paired Publishers and Subscribers, the Content Router nodes will ensure that all appropriate data exists in both clusters at all times. If there is a disaster in the primary cluster, the data can be repopulated using the “Republish” function in the DR Publisher admin console.

Application-Assisted Fail-Over

In both the above scenarios, the loss of data can be corrected, but will be unavailable to applications in the primary cluster until the Content Router nodes do their jobs and recover the data. If these applications are made aware of the presence and address of the remote DR cluster, and if they have visibility and access to the remote, then the applications can automatically fail-over to the remote cluster in the event a stream is not currently accessible in the primary.