Bertrand Dufrasne
Christian Burns
Wenzel Kalabza
Sandor Lengyel
Patrick Schill
Christian Schoessler

IBM XIV Storage System Multi-Site Mirroring

Introduced with the IBM® XIV® Storage System software v11.5, the multi-site mirroring function, also referred to as the *3-way mirroring* in the XIV GUI, enables world-class high availability and disaster recovery capabilities. As such, the function also enables clients to comply with business, government, and industry-driven requirements and regulations.

Essentially, 3-way mirroring as implemented in XIV provides you with:

- Three concurrent copies of data
- Simple failover and failback mechanism, while keeping data mirrored to ensure business continuity and guarantees:
  - Minimized risk for data outage, service availability, and low business impact in general
  - Expedited failover and data restoration in the event of disaster
- Negligible performance impact, by building upon the ultra-efficient, field proven, XIV remote mirroring technology.

Note: While this technology provides more flexibility and maintainability in the disaster recovery, there are some limitations with the initial implementation delivered with the XIV Storage Software version 11.5. See “Boundaries” on page 10.

This paper describes the 3-way mirroring basic concepts, terminology, and practical usage. It contains several use cases (failover/failback scenarios) for disaster recovery.

For additional information about mirroring functions in XIV, refer to the IBM Redbooks® publication, *XIV Storage System Business Continuity Functions,* SG24-7759.

3-way mirroring overview

The 3-way mirroring, as the name indicates, includes three peers (sites) using different types of remote mirror relations. To be more specific, the solution combines methods from a synchronous replication and two asynchronous replications across three sites, in a concurrent topology configuration as shown in Figure 1. One of the asynchronous mirror couplings is in standby. The standby mirror can be defined either in advance, at the time of the 3-way mirror creation, or when needed for data recovery (it requires manual intervention).

The 3-way mirroring offers a higher level of data protection and more recovery capabilities than the common 2-way mirroring. It can definitely improve business continuity and consequently prevent against extensive downtime costs as well as improving the ability to accommodate some compliance-related requirements.

The 3-way mirror, as implemented in XIV, is a concurrent, multi-target topology, rather than a multi-hop, cascading topology.

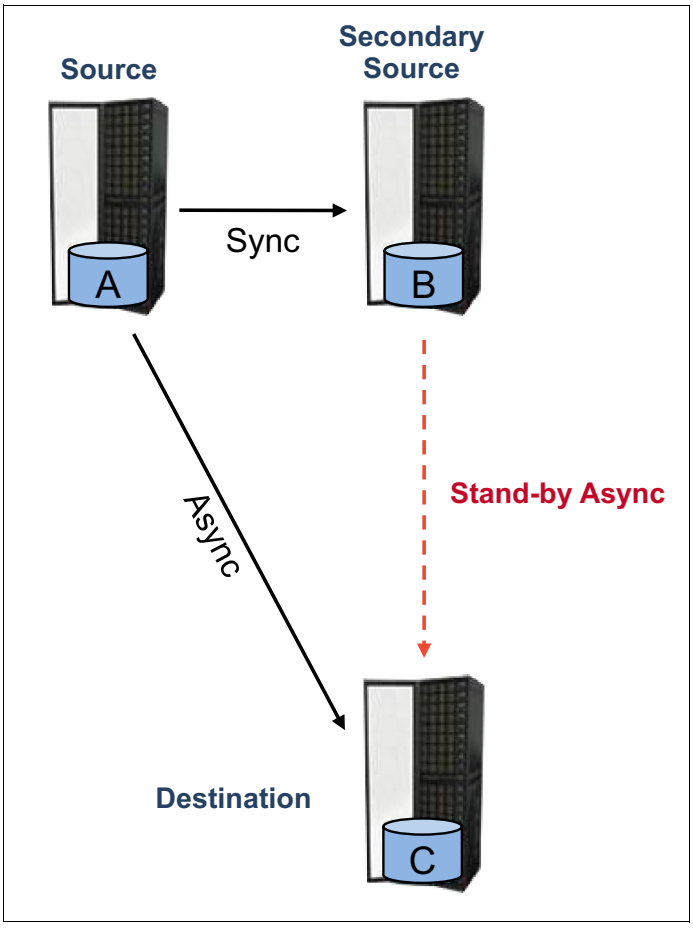

Figure 1 3-way mirroring configuration

Multiple 3-way mirroring can run concurrently on an XIV system, each with different mirror peers. Also, an XIV system can have different volumes with different roles in 3-way mirroring configurations.

3-way mirroring definition and terminology

A 3-way configuration is defined by extending 2-way mirror relationships. In other words, defining the 3-way mirroring assumes that at least there already exists, a fully initialized 2-way mirror relation that is based on either a synchronous or asynchronous mirror type.

The roles are defined as follows (refer to Figure 1 on page 2):

- **Source**

  The volume (volume A) that is mirrored.

  By association, the XIV Storage System that holds the source volume, is also referred to as the source system, or system A.

- **Secondary source**

  The secondary source (volume B) is synchronously mirrored with the source (volume A), and takes on the role of the source to the destination volume (volume C) when the source (volume A) becomes unavailable.

  By association, the XIV Storage System that holds the secondary source volume, is also referred to as the secondary source system, or system B.

- **Destination**

  The C volume that asynchronously mirrors the source.

  By association, the XIV Storage System that holds the destination volume, is also referred to as the destination system, or system C.

The 3-way mirror runs on the source system by activating the source to secondary-source and source to destination mirror relations. If one of these is already active, you do not need to change anything to that existing mirror relation.

Important: Running the same 3-way mirror configuration on the secondary source or on the destination is not allowed.

In a normal, steady state, the role of each volume is as indicated in Figure 1 on page 2. If the source (A) is unavailable, it is possible to establish a mirroring relationship between the remaining peers (B becomes the source of C).

Disconnecting or losing the target of one of the relations does not affect the other mirror relations in the solution.

Defining the standby mirroring relation in advance requires that the target connectivity between B and C (or at least its definitions) to be in place between all systems when the 3-way mirroring relation is configured. When defined, the B-C mirroring is in stand-by under normal conditions and only becomes active, by request, when a disaster occurs on system A, making volume A inaccessible (at which time volume B also must become the new source). This extended mirror relation has capability of enabling minimal data transfer and in this way it facilitates an expedited data consistency.

The 3-way view from each system is inherently different, as indicated in Table 1.

Table 1 Mirror relation view of each system in the 3-way mirroring solution

System volume	3-way role	Mirror relation 1	Mirror relation 2
A	source	A-B sync relation, A's role is source, the mirror is active	A-C async relation, A's role is source, the mirror is active
B	secondary source	A-B sync relation, B's role is destination, the mirror is active	B-C async relation, "stand-by" relation, B's role is the source
C	destination	A-C async relation, C's role is destination, the mirror is active	B-C async relation, "stand-by" relation, C's role is the destination

3-way mirroring states

The 3-way mirroring function in XIV, introduces new terms and defines new states as highlighted in Table 2. For a reminder and full overview of possible states in the simple two-peer mirroring relationship, refer to IBM Redbooks publication, *IBM XIV Storage System Business Continuity Functions*, SG24-7759.

While each individual mirroring definition has its own state, the 3-way mirroring definition has a global state, too. Among the possible global states, it is worth highlighting two new states denoted as Degraded and Compromised.

- **Compromised**

  Indicates that the 3-way mirroring relation is partially functioning. These are possible reasons for a compromised state:

  - Disconnection: The link is down for either A-B or A-C mirror coupling.
  - Resync: Either A-B or A-C are in resync and the secondary source did not yet take ownership.
  - Following a partial change of role: There was a role change on either A-B or A-C mirror coupling, so either of the 2-way mirrors are not active.

- **Degraded**

  When both A-B and A-C mirror couplings are active and A-C is in RPO lagging state, the global mirroring state appears as degraded.

Table 2 Three-site mirroring state

Source state	Condition
Inactive	Both mirror couplings (A-B, A-C) are inactive
Initializing	Copying all data from source to destination
Synchronized/Operational	Both couplings are synchronized and active → RPO OK
Degraded	Both couplings are synchronized and active → if A-C is RPO Lagging
Compromised	One coupling is synchronized and the other is in Resync or disconnected Following a partial change of role (role change on A-B or A-C)
Role Conflict (shown in GUI only, not XCLI)	Following a partial change of roles (two systems have a "source" role)

Source state	Condition
Secondary source state	**Condition**
Connected	The mirror with the source system is connected
Disconnected	The mirror with the source system is in disconnected state
Standby Mirror state	**Condition**
Up	Standby mirror is defined and connected
Down	Standby mirror is defined and disconnected
N/A	Standby mirror is not defined

3-way mirroring topology

There are two major existing 3-way topologies. The 3-way mirror as implemented in XIV is a concurrent, multi-target topology, rather than a multi-hop, cascading topology (refer to Figure 2).

- 1-to-n, also known as either “concurrent” or “multi-target”

  In this configuration, the source system replicates to two different destination systems, usually one replication type is asynchronous and the other is synchronous.

- Multi-hop also known as “Cascading”

  The source system has a synchronous replication to an intermediate system, which replicates an asynchronous relation to a far located system.

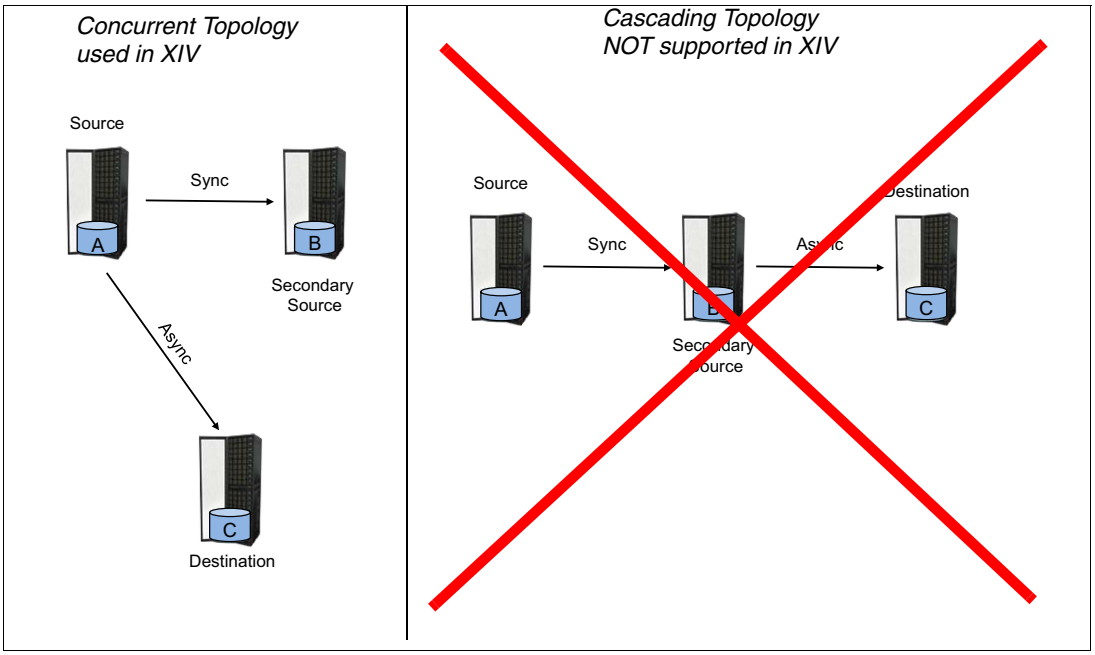

Figure 2 Concurrent topology (left) and Cascading topology

Important: The XIV Storage System supports the concurrent mirroring topology only.

For the XIV system architecture, the 3-way mirroring is established based on a concurrent topology in which A to B mirror coupling is synchronous (it means system A will not acknowledge the host before secondary source (B) is received and acknowledged, system A to C mirror coupling is asynchronous replication (it means source (A) starts a sync job for destination (C) at every async interval) while B to C mirror coupling is a stand-by asynchronous mirror relation. The latter one is optional. It may be configured in advance either after the 3-way mirror has been established or not at all. See Figure 3.

Note: The stand-by mirror relation consumes a mirror coupling from the predefined maximum number of 512.

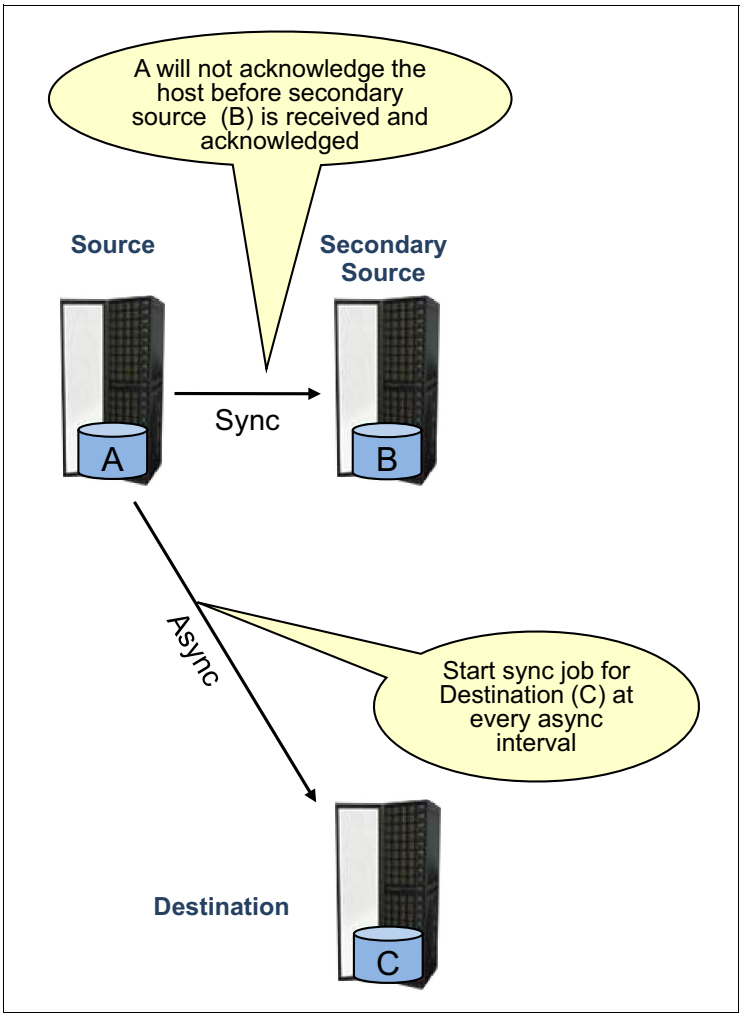

Figure 3 Concurrent topology including stand-by mirror relation

3-way mirroring characteristics

The 3-way mirroring relationships in XIV can be managed from the XIV GUI or the XCLI. The GUI automates some aspects of the 3-way mirroring processes.

The 3-way mirror is internally represented in the XIV storage software as a new object, called *xmirror*. The management of the 3-way mirror relies on functions that directly act upon the

xmirror object. However, the design and implementation are such that it remains possible to manage and monitor each of the participating replication processes as independent mirroring relationships. As depicted in Figure 4, the creation of a 3-way mirror relation can be expedited through offline initialization, either for adding a synchronous replication on-top of an existing asynchronous copy or vice versa.

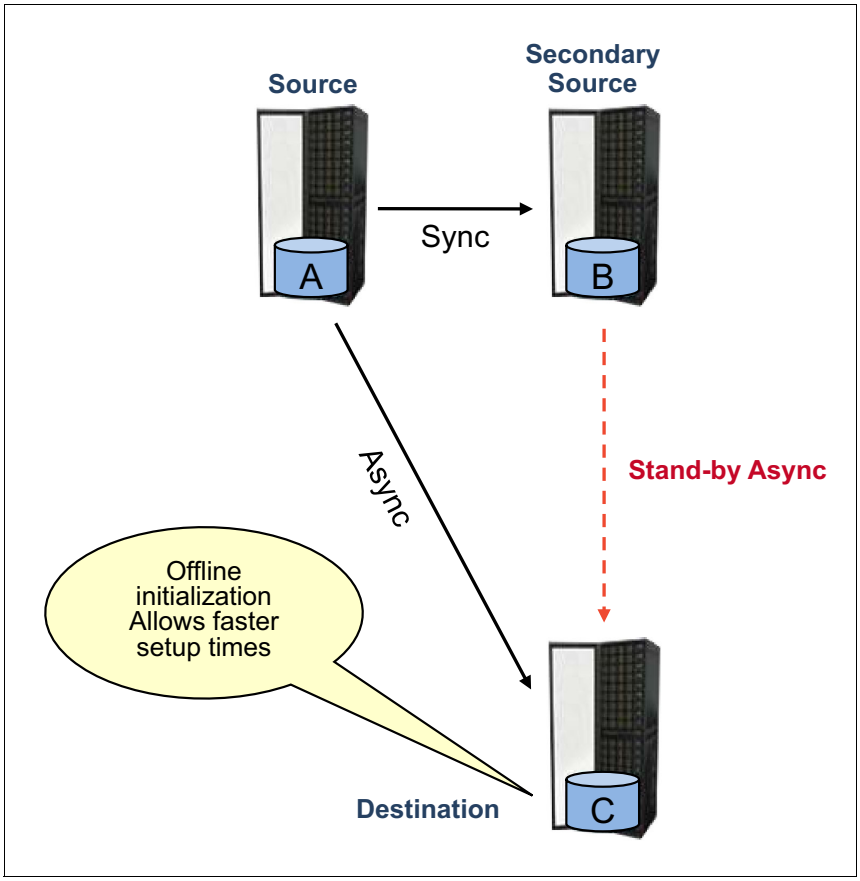

Figure 4 Offline initialization

The connectivity between XIV systems is supported over Fibre Channel (FC) or iSCSI links. Furthermore, it is also possible to define heterogeneous connectivity types, that is, typically, using an FC connection for the synchronous connection between A and B and an iSCSI protocol for the asynchronous connection between A and C, or also the stand-by connection between B and C. Indeed, each individual mirror is created over a separate source or target.

It is important as well to remember that hosts can only access the source volume.

Tip: It is essential to configure a Network Time Protocol (NTP) server for each of the XIV storage systems and make sure that the NTP server is indeed reachable from any of the locations where the systems are located.

Advantages of 3-way mirroring

As implemented in XIV, the 3-way mirroring solution offers the following advantages, as discussed hereafter.

Simplicity

The XIV system 3-way mirroring technology provides ease-of use implementation. The configuration is very simple and straightforward since the 3-way configuration simply builds upon the existing 2-way mirror relations, already familiar to users. Assigning roles to each volume in a 3-way mirror can be achieved through simple XCLI commands or GUI actions. The XIV GUI also provides support for adding automatically the standby (third) mirror coupling and subsequently activating the 3-way mirror relationship. Those two actions are optional.

High performance

Concurrent topology offers the best protection against local and regional disasters. Although, it can be seen as having a higher impact on production volumes when compared to the cascading topology. Indeed in the cascading approach, the source (volume A) peer is synchronized with secondary source only (volume B) while in concurrent topology, the source (volume A) must synchronize data with both the secondary source (volume B) and destination (volume C). However, thanks to the unique architecture of XIV, adding another mirror relation between A and C has barely an impact on performance. Indeed, each individual mirror is created with separate targets (B and C), communicating over different interface modules. Refer to Figure 5.

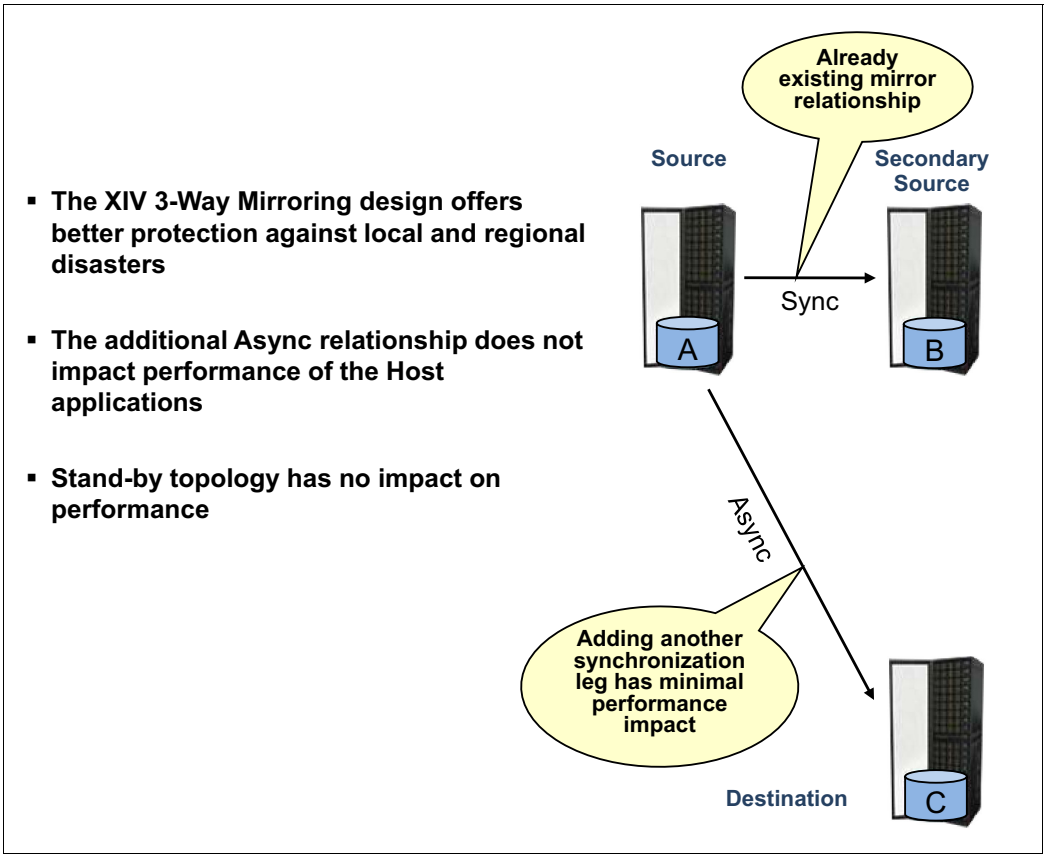

Figure 5 Minimal performance impact

Helps to accelerate synchronization

The 3-way mirroring as implemented in XIV, offers an incredible setup speed to establish the connection between secondary source (B) and destination (C) in case of data recovery requirement. Both the source (A) and secondary source (B) hold and maintain the asynchronous related snapshots; that is the Most Recent Snapshot (MRS) and the Last Replicated Snapshot (LRS) as depicted in Figure 6.

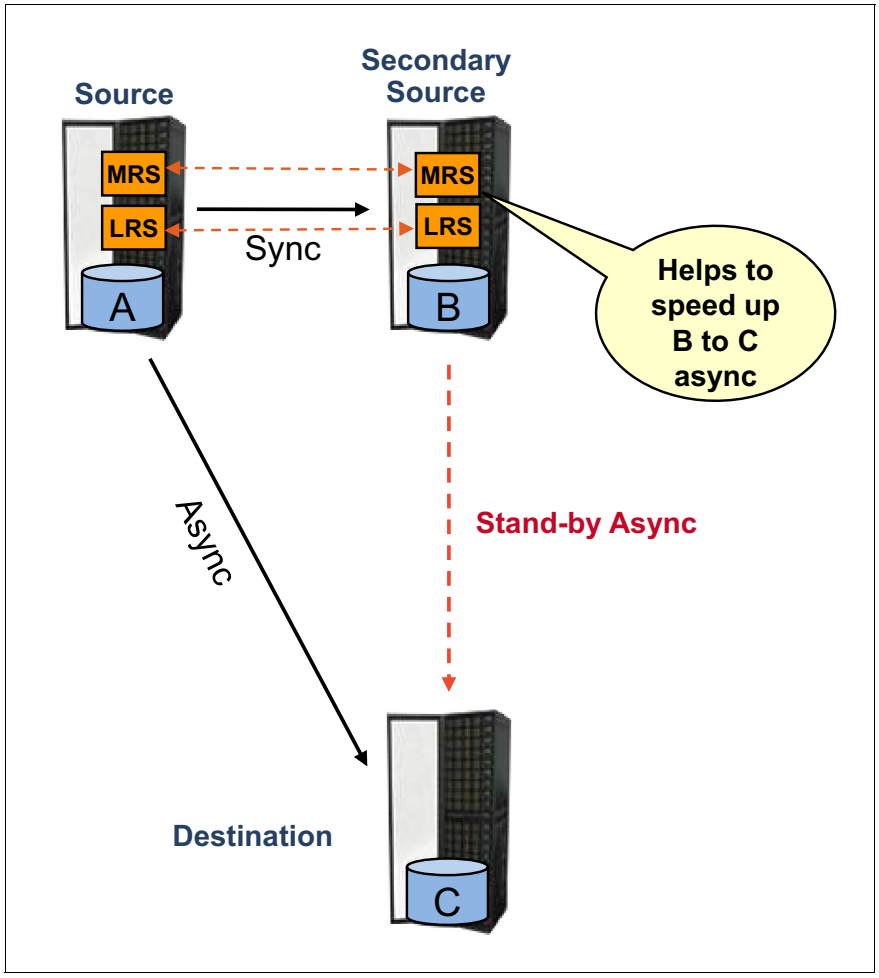

Figure 6 Internal asynchronous snapshots help to accelerate setup speed

As a reminder, the MRS is a snapshot of the source that the mirroring process uses to calculate the data that needs to be replicated to the destination on the next sync job. This snapshot captures that content of the sync job before it is initiated. The MRS exists only on the source.

The LRS is the latest snapshot that is known to be fully replicated by the destination (this snapshot captures the content of the sync job after it successfully completes). Both the source and the destination have a copy of the LRS.

Those snapshots are used by the system to know precisely the minimum data that needs to be replicated to the destination on the next sync job, which helps to speed up synchronization to the destination.

In a 3-way mirror, the source (A) notifies the secondary source (B) when it creates the MRS for the other destination (C). The secondary source maintains this MRS and changes the previous MRS to LRS (previous LRS gets deleted). The solution requires some extra capacity

on system B and also costs in processing the various snapshot copies required on system B. However, this process minimizes data transfer.

The MRS and LRS on volume B are used for fast recovery in the scenario where system A fails and B needs to take over, becoming the new source and activating the (stand-by) asynchronous relation with C. This way, the amount of data that needs to be synchronized is at most the size of the last sync job.

More than that, in some recovery cases, resynchronization between the new source (volume B) and destination (volume C) is not needed at all since the new source (B) might already be fully synchronized with volume C.

Flexibility

Multiple 3-way mirroring configurations can run concurrently per system. Any given system can be represented in several 3-way configurations, each referencing different systems.

A system can host mirroring peers with different roles in different multi-site configurations as shown in Figure 7.

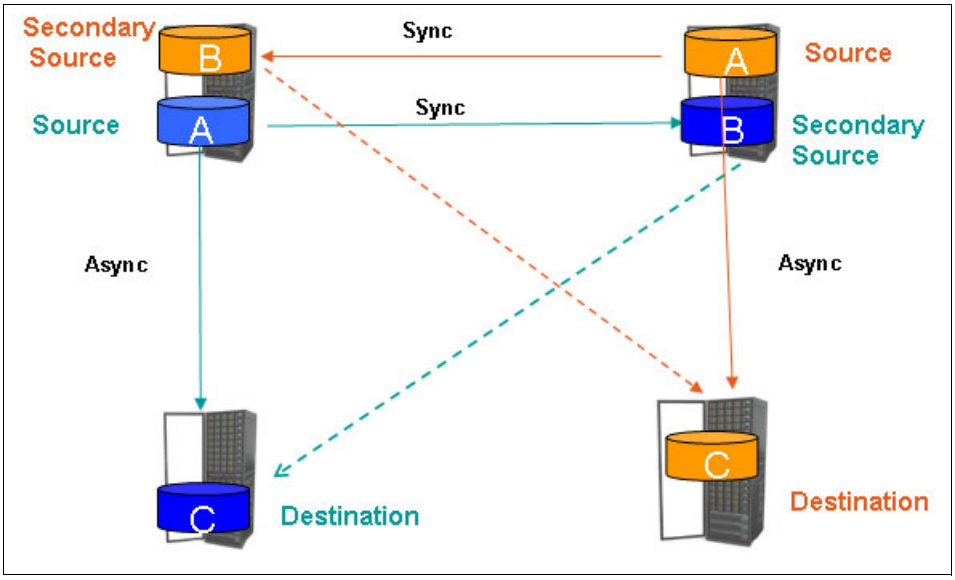

Figure 7 Flexibility of 3-way mirroring

Boundaries

With its initial implementation in the XIV Storage software, version 11.5, the 3-way mirroring has the following boundary conditions:

- A to B relation is always synchronous while A to C and B to C relations are always asynchronous
- Role switching is not supported in 3-way mirror relationships
- Online Volume Migration (OLVM) is not possible while maintaining the 3-way mirror relation of the migrated volume
- You cannot create any 3-way mirror Consistency Groups (CGs)
- There is a maximum of 256 cross mirror pairs

- No support for disaster recovery scenario where C becomes the real source of the 3-way replication. The Volume C role can be changed to become the source; however, it cannot be activated when it is part of a 3-way mirror setup.
- No ad hoc sync job support
  - Remote snapshot (ad hoc sync job) cannot be created for asynchronous mirroring that is part of the 3-way relation
  - Ad hoc snapshot can be done on the sync coupling only (A->B)

Setting up 3-way mirroring

To set up 3-way mirroring, use the XIV Storage System GUI or XCLI session.

Using the GUI for 3-way mirroring

We first look at how to create and manage the 3-way mirror with the XIV GUI. It offers some advantages over the XCLI by automating some of the tasks.

Create a 3-way mirroring

To establish a 3-way mirroring relation, the two-way mirroring relations (sync relation between A and B and async relation between A and C, and optionally between B and C) must first be created.

In the GUI, select the source XIV system and click **Remote** → **Mirroring**, as shown in Figure 8. Complete the following steps.

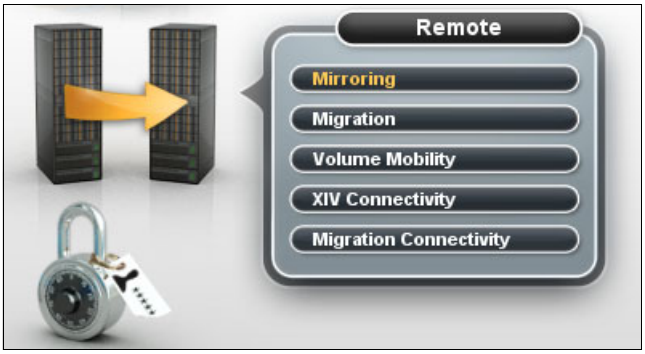

Figure 8 Selecting Mirroring

1. Right-click an existing 2-way mirrored volume coupling and select **Extend to 3-way,** as illustrated in Figure 9 on page 12.

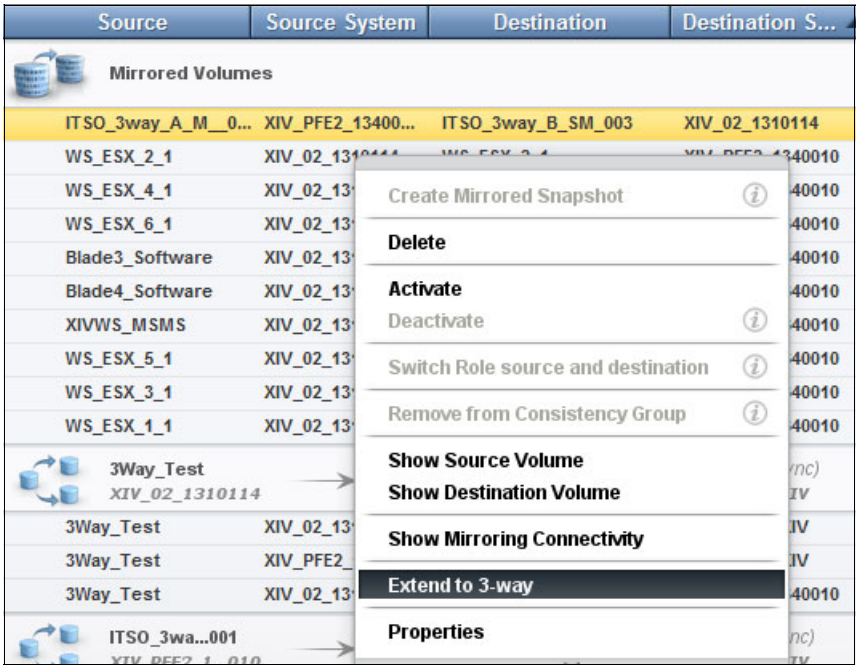

Figure 9 Selecting Extend to 3-way Mirror from dialog window

If the mirroring relation has the source and target connectivity (or at least its definitions) in place between all of the systems, a dialog similar to the one shown in Figure 10 is displayed. In this illustration, we have synchronous connectivity set up between XIV_PFE2_1340010 (source, volume A) and XIV_02_1310114 (secondary source, volume B) and asynchronous connectivity between XIV_PFE2_1340010 (source, volume A) and vvol_Demo_XIV (destination, volume C). In addition, the asynchronous (stand-by) relation between XIV_02_1310114 (secondary source, volume B) and vvol_Demo_XIV (Destination, volume C) is also already set up.

Figure 10 Extending to 3-way Mirror (when the mirror relation connectivity is in place)

If connectivity of all individual mirroring relations is not defined, the Extend to 3-way Mirror dialog box opens, as shown in Figure 11 on page 13.

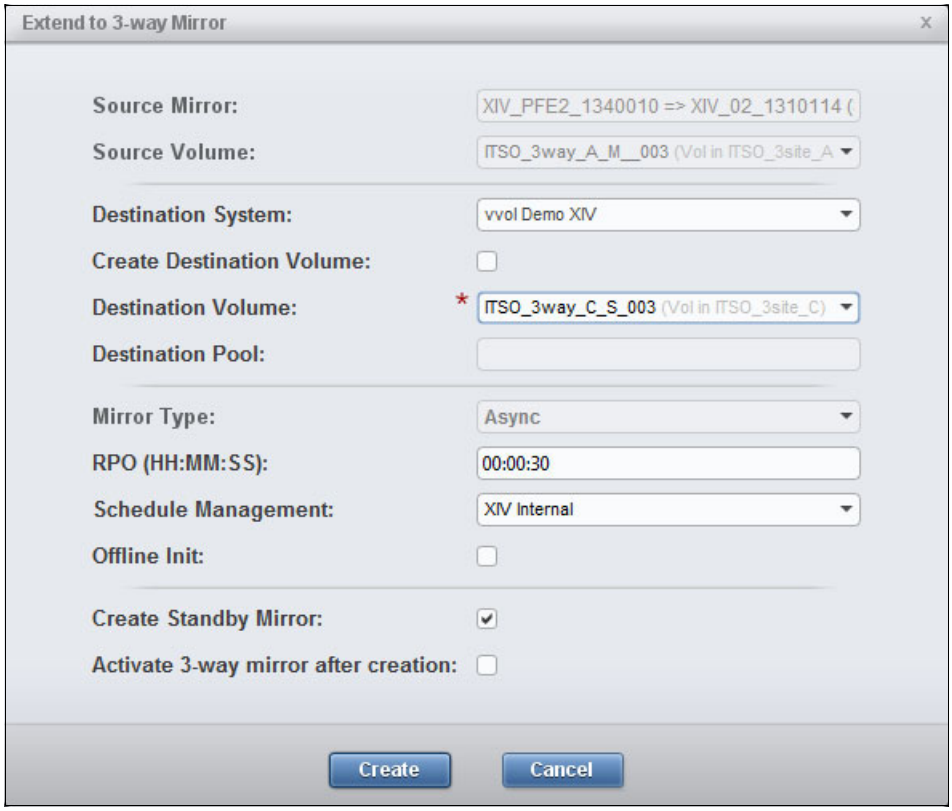

Figure 11 Extend to 3-way Mirror input panel

2. Make the appropriate selection and entries, as follows:
   - Source Mirror

     This is the existing primary mirror. The mirror type can be either synchronous (mirror coupling A-B) or asynchronous (mirror coupling A-C) to be extended to 3-way mirroring.
   - Source Volume

     This is the volume at the source site to be mirrored. Select the volume from the pull-down list.
   - Destination System

     This is the IBM XIV at the third site called destination (site C) that contains the destination volume. Select the destination system (C) from the list of known targets.
   - Create Destination Volume

     If selected, a destination volume is automatically created in the selected pool. If not selected, the user must specify the volume manually. By default, the volume name and size are taken from the source volume.
   - Destination Volume

     This is a name of the destination volume. If the **Create Destination Volume** option was selected, a destination volume of the same size and name as the source will be automatically created in the chosen pool. The name can be changed. If the Create Destination Volume option was not selected, the destination volume must exist and can be selected from the list. It must have the same size as the source volume and needs to be formatted, except if Offline Init (see below) is chosen.

– Destination Pool

This is a storage pool on the destination IBM XIV that contains the destination volume. The pool must already exist. This option is available only if **Create Destination Volume** is selected.

– Mirror Type

The mirror type is automatically selected in accordance with the source mirror. If source 2-way mirror is established on the basis of a synchronous link, the mirror type will be set as asynchronous. It cannot be changed.

– RPO (HH:MM:SS)

This option is disabled if Mirror Type is Sync. RPO stands for Recovery Point Objective and is only relevant for asynchronous mirroring. RPO time designation is a maximum time interval at which the mirrored volume can lag behind the source volume. Defined replication interval (the schedule) helps to achieve the requested target RPO.

– Schedule Management

This option is disabled if Mirror Type is Sync. Schedule Management is relevant only for asynchronous mirroring. Set the Schedule Management field to *XIV Internal* to create automatic synchronization using scheduled sync jobs. The *External* option means that no sync jobs are scheduled by the system. Therefore, this setting requires the user to run an ad hoc mirror snapshot to initiate a sync job.

Note: In the current XIV software v 11.5, remote snapshot (ad hoc sync job) cannot be created for the asynchronous mirroring connection that is part of the 3-way relation.

– Offline Init

This field is only available for selection if the **Create Destination Volume** option is not selected. Offline initialization is also known as *truck mode*. Upon activation of the mirror, only the differences between the source and the destination volume need to be transmitted across the mirror link. This option is useful if the amount of data is huge and synchronization might take more time than a manual transport because of the available bandwidth.

– Create Standby Mirror

This option enables the client to establish a third mirror of the 3-way mirroring optionally at the point of 3-way mirror creation. If this option is selected, an asynchronous mirror coupling is automatically configured in to standby state so that to avoid having further manual intervention to define. It becomes operational only by request in case of disaster recovery. This extended mirror relation has the capability of enabling minimal data transfer and in this way it facilitates an expedited data consistency. If it is not selected, you can configure it after the 3-way mirror has been created. Note that the stand-by mirror relation consumes a mirror coupling from the predefined maximum number.

– Activate 3-way mirror after creation

This option activates the 3-way mirroring immediately after its creation and therefore reduces the number of clicks when compared to doing a manual activation afterward.

Note: In our example, the **Activate 3-way mirror after creation** option is cleared. You may select this option at the creation phase to avoid having further manual intervention to activate 3-way mirror relation. Note that the state of the existing mirror is not changed and is independent of this option.

3. After all the appropriate entries have been specified, click **Create**.

   A 3-way mirror relation is created and is in Inactive mode as shown Figure 12. In this state, data is not yet copied from the source to the target volumes. The global status of 3-way mirroring, highlighted in amber background, provides its overall status in one line by listing the names of all three involved volume peers and name of the source, secondary, and destination source system as well as the link type in between the mirror relations.

   Drill down to see the related mirror relations created and to check for the status of each individual coupling. The corresponding mirror couplings also are automatically created on the secondary source and the destination XIV system where the same global and individual status of the 3-way mirroring can be seen. They are also in Inactive state accordingly.

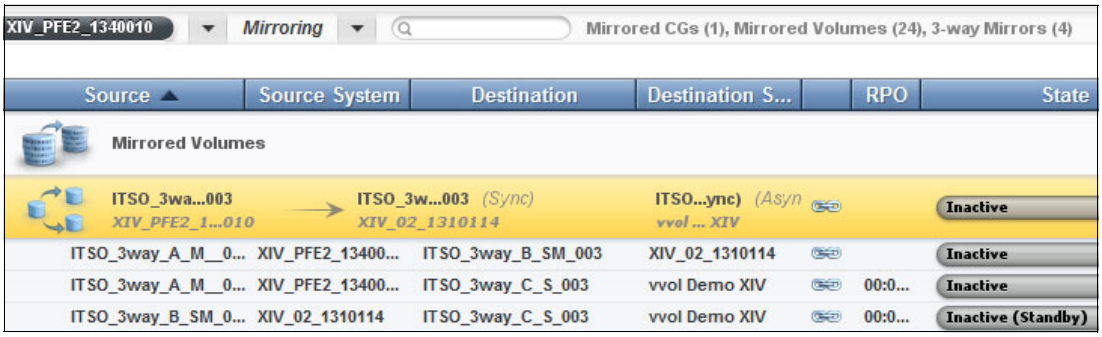

Figure 12 3-way Mirror coupling in Inactive state at source side (Source XIV system)

4. Repeat steps 1- 3 to create additional 3-way mirror relations.

3-way mirror activation

To activate the 3-way mirror couplings, proceed as follows:

1. On the source system, go to **Remote** → **Mirroring**. Highlight all the 3-way mirror relations that you want to activate, right-click, and select **Activate**, as shown in Figure 13.

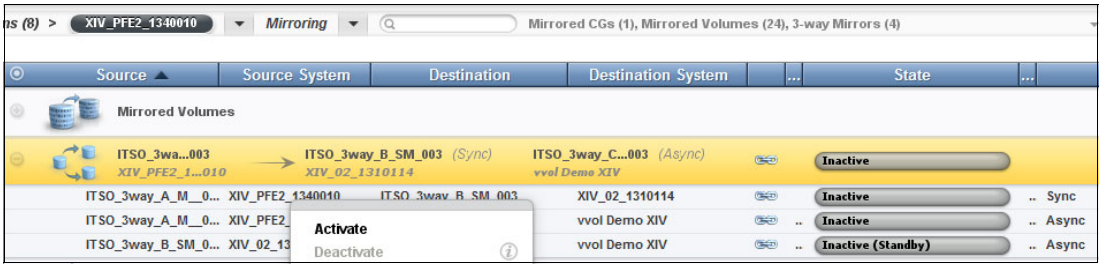

Figure 13 3-way mirror activation

2. The global state of the 3-way mirror then enters an *Initialization* state as shown in Figure 14 on page 16. The same state can be observed on each site. In this phase, all data is copied from the source volume peer to the destination volume. Therefore, its individual state shows *Initialization* too. Getting to the synchronized state (for this mirror relation) might take some time depending on the amount of data to be transferred and the bandwidth of the link. Instead of over-the-wire initialization, you can do an offline initialization (also known as “truck” mode), whereas a backup of the source can be shipped to the remote site.

   Then, once being restored the system can be instructed to establish the mirroring after first comparing the source with the restored destination replica. This can shorten the initialization process considerably and save expensive bandwidth. The volume peers

between the source (A) and secondary source (B) get synchronized faster because they were fully initialized before. The third asynchronous mirror coupling (between secondary source and destination volume) stays in Inactive (standby) state and it becomes operational only by request in case of disaster recovery.

Figure 14 3-way mirror initialization phase

3. After the *Initialization* phase is complete, the global state of 3-way mirror is *Synchronized* and the state of the mirror relation between source and destination system is *RPO OK*, as shown in Figure 15.

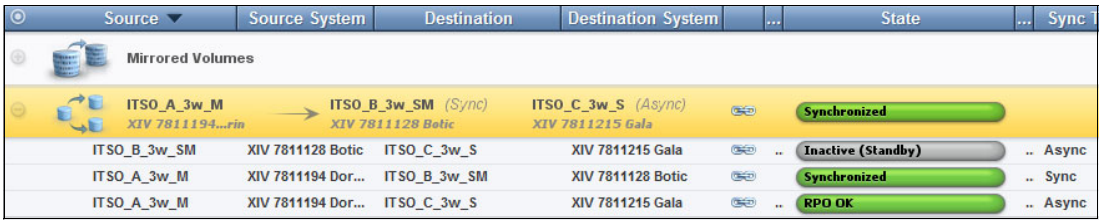

Figure 15 3-way mirror synchronized

Adding standby mirror to 3-way mirroring later

If there is no standby mirror defined at the point of the 3-way mirror creation, you can add a standby mirror at a later time:

1. In the XIV GUI, select any XIV system involved in 3-way mirroring and from the main window, select the function icon **Remote** → **Mirroring**. In the mirrored volumes list, right-click the relevant 3-way mirror relation, as shown in Figure 16.

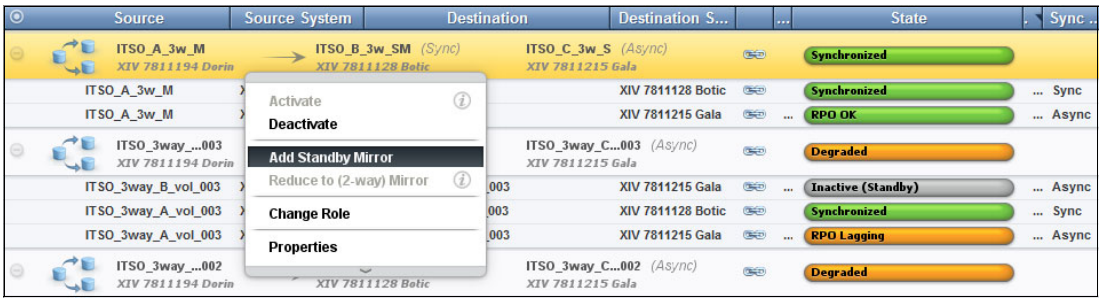

Figure 16 Adding a standby mirror to 3-way mirroring

2. From the pop-up menu, simply select **Add Standby Mirror**. An additional row displaying the standby mirror as inactive is added to the mirroring view, as shown in Figure 17.

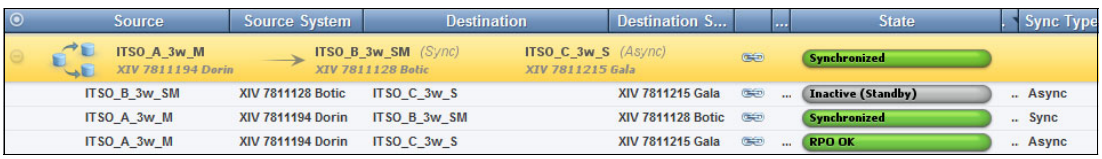

Figure 17 Standby mirror is added posteriorly to 3-way mirroring

Reducing from a 3-way to 2-way mirror relation

The XIV system 3-way mirroring feature allows you as well to subsequently reduce the number of mirror couplings, thereby reverting to a 2-way mirror solution.

Deletion of 3-way relation is a *disband* operation where the client instructs which mirror coupling should be kept alive. The other individual mirrors under the 3-way mirror relationships get deleted automatically when using the XIV GUI, whereas in the XCLI session the individual mirror relations are not deleted automatically and just go into a standby state. You must then manually delete the standby relations or decide to define them again in a 3-way mirror setup.

Reverting to a 2-way mirror can be necessary in a disaster recovery situation. Suppose that the source system (A) and secondary source (B) fail, then hosts operations will have to be routed to the destination (system C). System C can have its role changed and become the source so that it can now serve I/Os from applications; however it cannot be activated when it is part of a 3-way mirror configuration. In this case, you first must remove the mirror relation with the secondary source, making the relation a regular asynchronous relation. Then you can activate the mirror coupling (C-A).

The 3-way mirror relation must be placed into an *Inactive* state before reducing to 2-way mirror coupling. In the XIV GUI, select any XIV system involved in 3-way mirroring, and from the left pane choose the function icon **Remote** → **Mirroring**. Then, right-click the relevant 3-way mirror relation and select **Deactivate**, as shown in Figure 18.

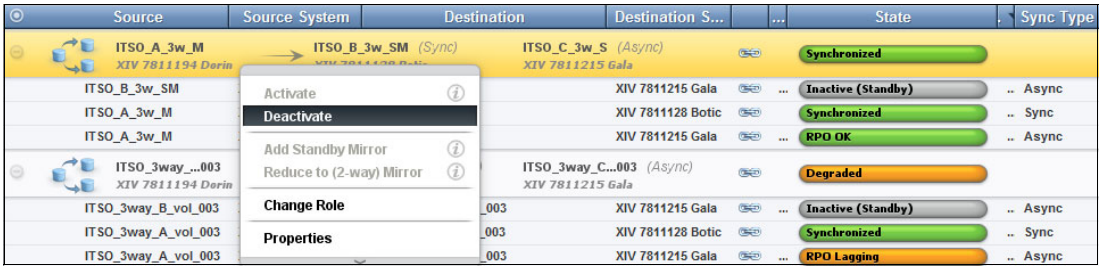

Figure 18 Deactivating 3-way mirror relation

3. Wait for 3-way mirroring to be deactivated (its state turns to *Inactive*). Subsequently, right-click the now deactivated 3-way mirroring relation and select **Reduce to (2-way) Mirror** from the pop-up menu, as shown in Figure 19.

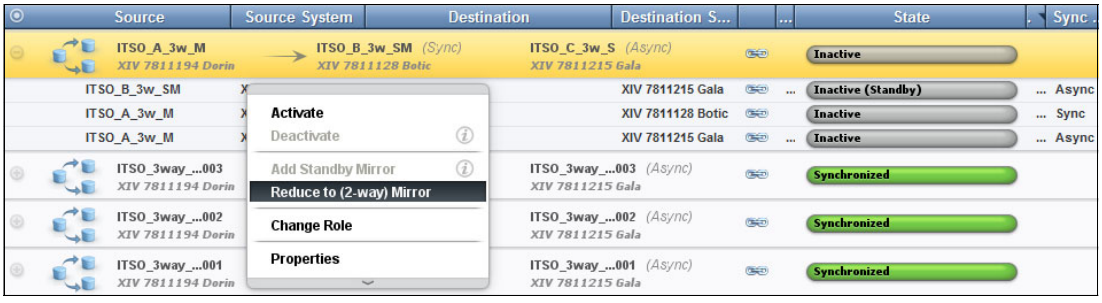

Figure 19 Reducing from 3-way to 2-way mirror relation

4. The Reduce to 2-way Mirror window displays. Select the mirroring relation to keep from the drop-down list and click **OK**. Refer to Figure 20 on page 18.

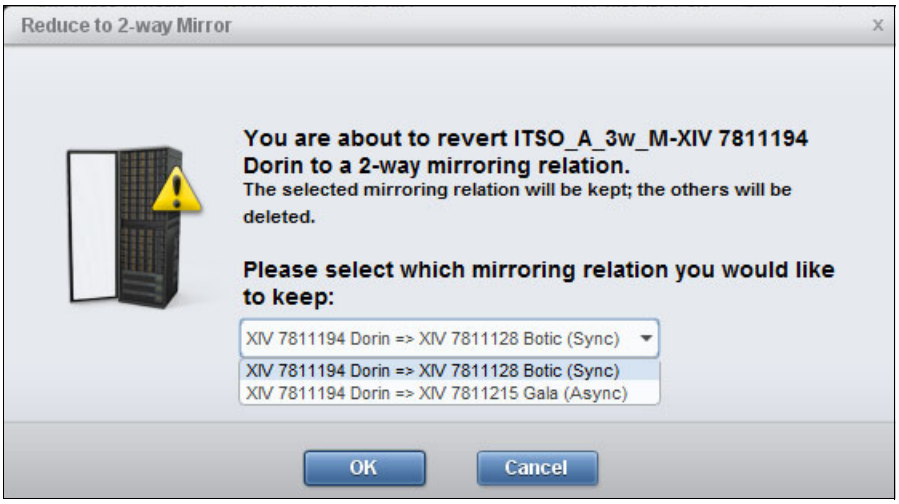

Figure 20 Reduce to 2-way Mirror window

Using XCLI for 3-way mirroring setup

For completeness, our illustration describes all the setup steps to perform when using the XCLI to create a 3-way mirror relation including the synchronous link as well. We also assume that all the volumes involved in the 3-way mirroring already exist on the XIV systems and are not in any mirroring relation.

Note that in XCLI commands, the master, smaster, and slave parameters respectively refer to the source, secondary source, and destination volumes.

Creating the 3-way mirror relation

1. Open an XCLI session for the source XIV system and define the other two systems as targets using the **`target_define`** command as shown in Example 1. (Note that target definition must be done on each system.)

Example 1 Define targets on source system

```
XIV 7811194 Dorin>>target_define target="XIV 7811128 Botic" protocol=FC
Command executed successfully.
XIV 7811194 Dorin>>target_define target="XIV 7811215 Gala" protocol=FC
Command executed successfully.
```

2. Run **`target_mirroring_allow`** with the other two systems successively specified as the target= parameter (refer to Example 2). This command is performed on a local storage system in order to allow the target storage system permission to read, write, view, create volumes, and define existing volumes as destinations. (Note that this command must be run on each system.)

Example 2 Run target mirroring allow on source system

```
XIV 7811194 Dorin>>target_mirroring_allow target="XIV 7811128 Botic"
Command executed successfully.
XIV 7811194 Dorin>>target_mirroring_allow target="XIV 7811215 Gala"
Command executed successfully.
```

3. Define connectivity to the other two systems using the **target_connectivity_define** command as shown in Example 3. The *fcaddress* parameter represents the Fibre Channel (FC) address of the port on the remote target. The l*ocal_port* parameter is self-explanatory and pertains to an FC port only. Note that connectivity must be defined on each system and if there is already a defined connectivity, use it. Only one connection must be in place between two peers at a given time.

Example 3 Define target connectivity on source system

```
XIV 7811194 Dorin>>target_connectivity_define target="XIV 7811128 Botic"
fcaddress=500173802B780140 local_port=1:FC_Port:4:4
Command executed successfully.
XIV 7811194 Dorin>>target_connectivity_define target="XIV 7811128 Botic"
fcaddress=500173802B780150 local_port=1:FC_Port:5:4
Command executed successfully.
XIV 7811194 Dorin>>target_connectivity_define target="XIV 7811215 Gala"
fcaddress=500173802BCF0140 local_port=1:FC_Port:4:4
Command executed successfully.
XIV 7811194 Dorin>>target_connectivity_define target="XIV 7811215 Gala"
fcaddress=500173802BCF0150 local_port=1:FC_Port:6:4
Command executed successfully.
```

4. Create a 2-way synchronous mirror relation between the source (A) and secondary source (B) volume using the **mirror_create** command on the source system, as shown in Example 4. The mirror type is *SYNC _BEST_EFFORT* in accordance with synchronous mirror.

Example 4 Create 2-way sync mirror relation on source system

```
XIV 7811194 Dorin>>mirror_create target="XIV 7811128 Botic" vol="ITSO_A_3w_M"
slave_vol="ITSO_B_3w_SM" type=SYNC_BEST_EFFORT
Command executed successfully.
```

5. Create a 2-way asynchronous mirror relation between source (A) and destination (C) volumes using the **mirror_create** command on the source system, as shown in Example 5. Set the mirror type to *ASYNC_INTERVAL* on the basis of async mirroring. As usually advised, keep the value for the schedule interval at about one third of the recovery point objective (RPO).

Example 5 Create 2-way async mirror relation on source system

```
XIV 7811194 Dorin>>mirror_create remote_schedule="min_interval" vol="ITSO_A_3w_M"
schedule="min_interval" part_of_xmirror="yes" slave_vol="ITSO_C_3w_S" rpo="60"
target="XIV 7811215 Gala" remote_rpo="60" type="ASYNC_INTERVAL" create_slave="No"
Command executed successfully.
```

6. Go to the secondary source (B) XIV system and open an XCLI session. Create a 2-way asynchronous mirror relation between the secondary source (B) and destination (C) volume using the **mirror_create** command as shown in Example 6.

Example 6 Create 2-way async mirror relation on secondary source system

```
XIV 7811128 Botic>>mirror_create remote_schedule="min_interval" vol="ITSO_B_3w_SM"
schedule="min_interval" part_of_xmirror="yes" slave_vol="ITSO_C_3w_S" rpo="60"
target="XIV 7811215 Gala" remote_rpo="60" type="ASYNC_INTERVAL" create_slave="No"
Command executed successfully.
```

7. Go to the source (A) XIV system. Create a 3-way mirror relation using the **xmirror_define** command as shown in Example 7. Any name can be given to the xmirror object but it must be unique in the system.

Example 7 Create a 3-way mirror relation on source system

```
XIV 7811194 Dorin>>xmirror_define vol="ITSO_A_3w_M" slave_target="XIV 7811215
Gala" xmirror="ITSO_A_3w_M-XIV 7811194 Dorin" smaster_target="XIV 7811128 Botic"
Command executed successfully.
```

8. Activate the 3-way mirror relation using the **xmirror_activate** command on the source system as illustrated in Example 8.

Example 8 Activate 3-way mirror relation on source system

```
XIV 7811194 Dorin>>xmirror_activate xmirror="ITSO_A_3w_M-XIV 7811194 Dorin"
Command executed successfully.
```

Monitoring the 3-way mirror

The **xmirror_list** command shows all of the existing 3-way mirror relations. This xcli command can be executed for monitoring purpose on any system that is part of a 3-way mirroring, as shown successively in Example 9, Example 10, and Example 11 on page 21. Note that there is no standby mirror coupling created (B-C) at the first listed 3-way mirror relation. Its standby mirror state is shown up accordingly through each site.

Example 9 Lists 3-way mirror relations on the source system

```
XIV 7811194 Dorin>>xmirror_list
Name                                    Xmirror ID         local volume name
ITSO_3way_A_vol_001-XIV 7811194 Dorin   00A138002BBA3AA6   ITSO_3way_A_vol_001
ITSO_3way_A_vol_002-XIV 7811194 Dorin   00F938002BBA3AA7   ITSO_3way_A_vol_002
ITSO_A_3w_M-XIV 7811194 Dorin           00A638002BBA003A   ITSO_A_3w_M

Local Xmirror Role   Xmirror State   Standby Mirror   Master   SMaster
Master               Operational     NA               Local    XIV 7811128 Botic
Master               Degraded        Up               Local    XIV 7811128 Botic
Master               Operational     Up               Local    XIV 7811128 Botic

Slave
XIV 7811215 Gala
XIV 7811215 Gala
XIV 7811215 Gala
```

Example 10 Lists 3-way mirror relations on the secondary source (SMaster) system

```
XIV 7811128 Botic>>xmirror_list
Name                                    Xmirror ID         local volume name
ITSO_3way_A_vol_001-XIV 7811194 Dorin   005E38002BBA3AA6   ITSO_3way_B_vol_001
ITSO_3way_A_vol_002-XIV 7811194 Dorin   00F938002BBA3AA7   ITSO_3way_B_vol_002
ITSO_A_3w_M-XIV 7811194 Dorin           00A638002BBA003A   ITSO_B_3w_SM

Local Xmirror Role   Xmirror State   Standby Mirror   Master              SMaster
SMaster              Connected       Down             XIV 7811194 Dorin   Local
SMaster              Connected       Up               XIV 7811194 Dorin   Local
SMaster              Connected       Up               XIV 7811194 Dorin   Local

Slave
```

```
NA
XIV 7811215 Gala
XIV 7811215 Gala
```

Example 11 Lists 3-way mirror relations on the destination system

```
XIV 7811215 Gala>>xmirror_list
Name                                    Xmirror ID         local volume name
ITSO_3way_A_vol_001-XIV 7811194 Dorin   005E38002BBA3AA6   ITSO_3way_C_vol_001
ITSO_3way_A_vol_002-XIV 7811194 Dorin   00F938002BBA3AA7   ITSO_3way_C_vol_002
ITSO_3way_A_vol_003-XIV 7811194 Dorin   00CA38002BBA3AA8   ITSO_3way_C_vol_003
ITSO_A_3w_M-XIV 7811194 Dorin           00A638002BBA003A   ITSO_C_3w_S

Local Xmirror Role   Xmirror State   Standby Mirror   Master
Slave                Connected       Down             XIV 7811194 Dorin
Slave                Connected       Up               XIV 7811194 Dorin
Slave                Connected       Up               XIV 7811194 Dorin

SMaster             Slave
NA                  Local
XIV 7811128 Botic   Local
XIV 7811128 Botic   Local
```

Deleting the 3-way mirror relationship

Since the XIV storage software has an explicit *xmirror* object to represent 3-way mirror relations, destroying the 3-way relation is done by simply deleting the relevant xmirror object.

However, before the xmirror object can be deleted, the master must no longer have two active mirror relations (they must be deactivated). Then, the deletion of the xmirror object is done on the source system.

1. Deactivate the xmirror object using the **xmirror_define** command on any affected system. A warning message prompts you to confirm the deactivation, as shown in Example 12.

Example 12 Deactivate 3-way mirror relation

```
XIV 7811194 Dorin>>xmirror_deactivate xmirror="ITSO_A_3w_M-XIV 7811194 Dorin"

Warning:   Are you sure you want to deactivate mirroring? y/n: y
Command executed successfully.
```

2. Go to source system and delete the 3-way mirror relation by issuing the **xmirror_delete** command. The force (force=yes/no) is an optional parameter that deletes the xmirror on the local system only. After the xmirror object is deleted, delete the asynchronous mirror relation between the source and destination system (A-C), using **mirror_delete**. See outputs of the executed commands in Example 13.

Example 13 Delete 3-way mirror relation on source system

```
XIV 7811194 Dorin>>xmirror_delete xmirror="ITSO_A_3w_M-XIV 7811194 Dorin"
Command executed successfully.

XIV 7811194 Dorin>>mirror_delete vol="ITSO_A_3w_M" target="XIV 7811215 Gala"

Warning:   Are you sure you want to delete this mirroring relationship? y/n: y
Command executed successfully.
```

3. Go to the secondary source system. The relation between secondary source (B) and destination (C) is not deleted automatically by the system since it was not created automatically by the system. Proceed to delete the standby mirror relation using the **mirror_delete** command as shown in Example 14.

Example 14 Delete standby mirror relation on secondary source system

```
XIV 7811128 Botic>>mirror_delete vol="ITSO_B_3w_SM" target="XIV 7811215 Gala"

Warning:   Are you sure you want to delete this mirroring relationship? y/n: y
Command executed successfully.
```

Note: The GUI is doing some operations for the user (for ease-of-use) like creating and deleting mirror between the secondary source (B) and the destination system (C).

Disaster recovery scenarios with 3-way mirroring

In this section, we consider three general categories of disaster situations:

- The primary storage site, where the source of the 3-way mirror resides, is destroyed.
- The secondary source is destroyed.
- The destination system is destroyed.

The (simpler) recovery cases when only the connection/mirror links between sites becomes interrupted or are broken, can be addressed using re-synchronization steps (see IBM Redbooks publication, *IBM XIV Storage System Business Continuity Functions*, SG24-7759).

Figure 21 on page 23 shows the normal working state of a 3-way mirror setup, which is our baseline for the scenarios illustrated in this section.

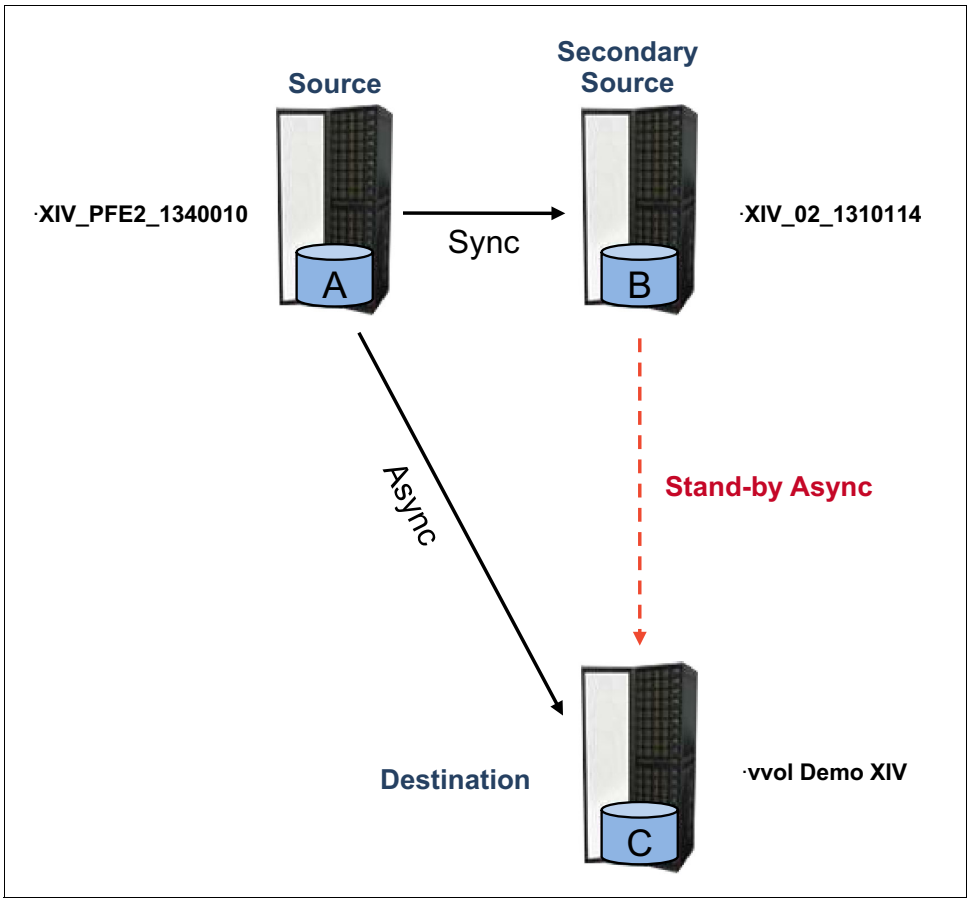

Figure 21 3-way mirror: normal working state

Our scenarios assume a single volume being mirrored to keep the GUI views compact. We use the following XIV storage systems:

- XIV_PFE2_1340010 as source storage system A (with volume A)
- XIV_02_1310114 as secondary source storage system B (with volume B)
- vvol Demo XIV destination storage system C (with volume C)

The mirror relation from A to B is a synchronous mirror over Fibre Channel. A to C and B to C are asynchronous mirror relations over iSCSI (for long distance).

The representation in the XIV GUI, for this normal, active 3-way relationship appears as shown in Figure 22. The expanded GUI view shows the synchronous mirror and asynchronous mirror relations listed with their actual states (refer to the State column). The state is shown for each volume involved in this particular 3-way mirror.

Figure 22 3-way mirror: normal state in GUI view

Source system (site A) failure scenario

In this scenario, the source system is affected by a disaster and can no longer send or receive I/Os, that is, it can no longer communicate with hosts and the mirror links are inactive. This

also means that A can also no longer communicate updates to the secondary source (B), which will now act as a source, and destination volume (C), as depicted in Figure 23.

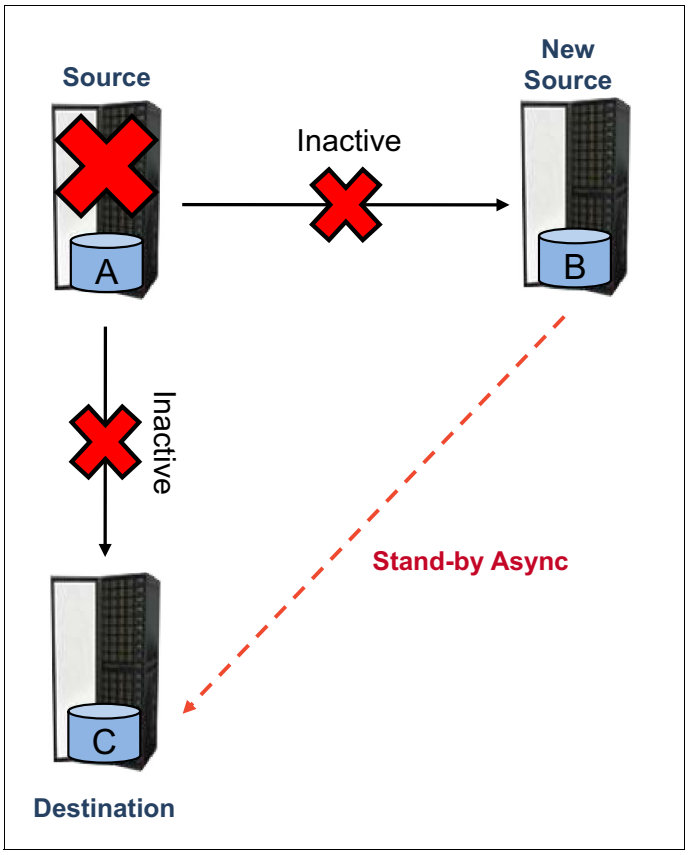

Figure 23 3-way mirror - site A failure scenario

After site A fails, the GUI view shown in Figure 24 indicates that all mirror links were affected and went into an inactive state. The synchronous connection (A to B) is in an unsynchronized state and the asynchronous connection between A and C is in an RPO lagging state.

Figure 24 3-way mirror - site A failure scenario, expanded view with details

With site A (source) down, a manual intervention is required (if there is no recovery software in place) to move applications toward site B (secondary source). There are usually two options:

- If servers/applications at site A are still intact, they can be redirected over the SAN to point to storage system at site B.
- If a set of servers/applications are also maintained as a backup at site B, they can simply be started at site B.

Use the following steps to get volume B (secondary source) operational as source:

1. Change volume B role from “secondary source” to “source” by right-clicking the empty area (red dot) of the global state of the mirror relation (highlighted in orange). From the pop-up menu, select **Change Role**.

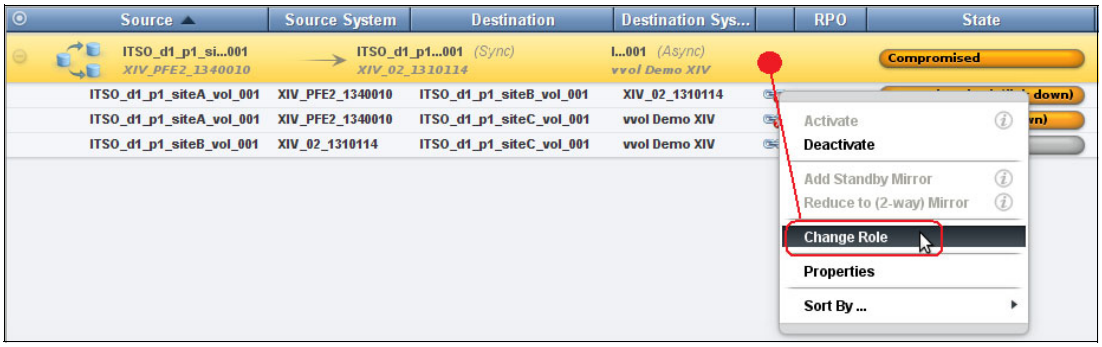

Figure 25 3-way mirror - site A failure, preparing site B for changing role to source - step 1

2. In the Change Role window that opens, select (from the pull-down menu) the system for which the role needs to be changed (site B). When you press **OK**, volume B becomes source.

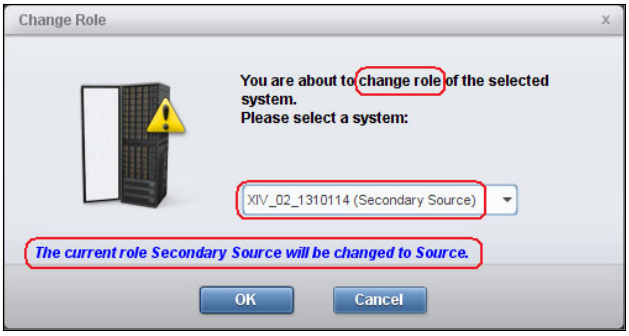

Figure 26 3-way mirror: change role menu

As a result of this step, there is a role conflict (also known as *split brain*) situation because there are now two source volumes (A and B) in one mirror relation defined. As depicted in Figure 27, the role conflict is shown three times: first in the global mirror relation state, second and third on the site A and B mirror relations.

Figure 27 3-way mirror: site A failure, role conflict after the role change of B to source

3. Now we activate volume B as the source for the stand-by asynchronous connection with volume C. Again, right-click the empty area of the global mirror row (as indicated by the red dot in the orange background, in Figure 28. From the pop-up menu, select ‘Activate on XIV <site B>’ (in our example, XIV_02_1310114) to get B activated in a new source role for the mirror relation with C.

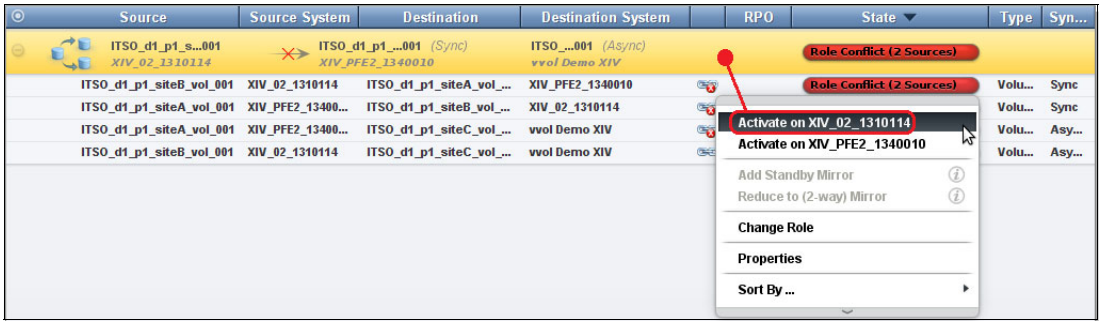

Figure 28 3-way mirror: site A failure, site B activation after role change of B to source

Figure 29 shows the GUI view following site B activation as a source.

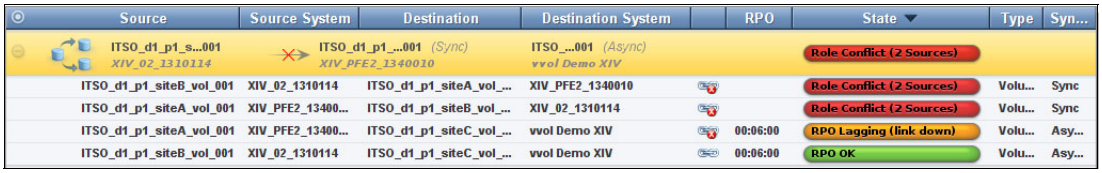

Figure 29 3-way mirror: site A failure, site B activated as source

With B as the source of the asynchronous mirror connection with C:

– Volume C blocks out any I/Os it might still receive from A (this will not normally happen since site A was destroyed in this scenario).

– Before activating the B-C mirror, the system examines the secondary source snapshots to determine which snapshot to use for resync. The following three situations are possible:

  • A failure after the last A-C sync job

    The system compares the time stamps of MRS on B to LRS on C and finds them identical. No recovery is required. B becomes the new source master and B-C becomes alive.

  • A failure during the last A-C sync job

    The time stamps of LRS on B and LRS on C are identical to the last A-C completed sync job. This indicates that A failed during the last sync job, and B sends the difference between its LRS and MRS to C during the next sync job.

  • A failure after updating C and before updating B

    This scenario is least likely to occur. For LRS on C to be more recent than MRS on B, the A-B mirror must fail first, and the A-C mirror must be up long enough for an asynchronous snapshot to be taken and passed to C, only to fail as well. If C’s LRS has a more recent time stamp than B’s MRS (that is, A continued to update C while the B-A mirror was disconnected, and C is more up-to-date than B), then reinitialization of this mirror is required.

– Site C restores from its LRS, overriding any writes that site A may have written to site C before site B became the source.

4. Finally, perform the necessary SAN zoning, volume mapping XIV site B volumes to the hosts servers (or their backup) server and restart the server/applications connected to site B (new source). In short, start the production backup site.

Figure 30 illustrates the new active setup with site B changed to source.

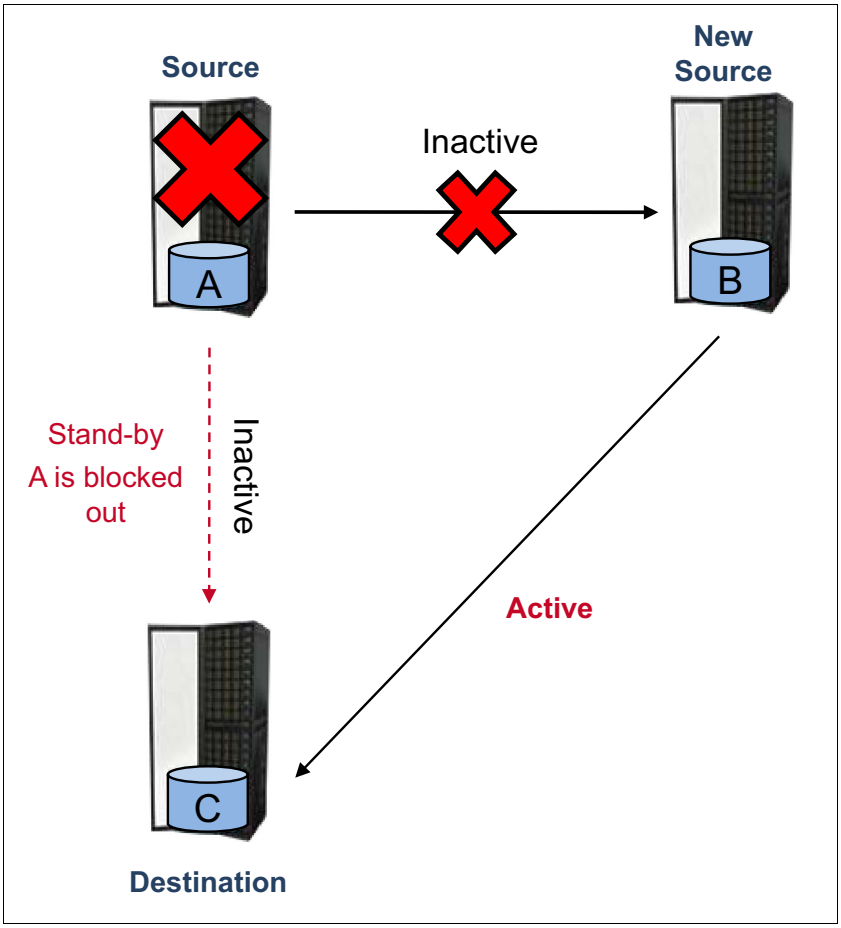

Figure 30 3-way mirror: site B changed to source

This setup/situation will remain until site A is recovered and the mirror connectivity from site A to C and site A to B is restored.

Failback to site A as data source

After site A is fully recovered and operational and the mirror links to both sites B and C are also again in working order, you can prepare for and execute a failback to the normal production site (site A).

Once the mirror links were recovered, they can be activated again from site A to C and between site A and B. The asynchronous connection (Site A to C) changes its state to “Configuration Error” (change from amber to red color in the GUI). Note that the state could also be “inactive standby”, depending on the failure scenario. Refer to Figure 31 on page 28.

The sequence of steps is detailed below:

1. Site A needs first to be changed to secondary source. This change is required to synchronize site A with the data updates that took place at the backup production site (site B) and site C, while A was out of service. The synchronization with site B and C will run in parallel.

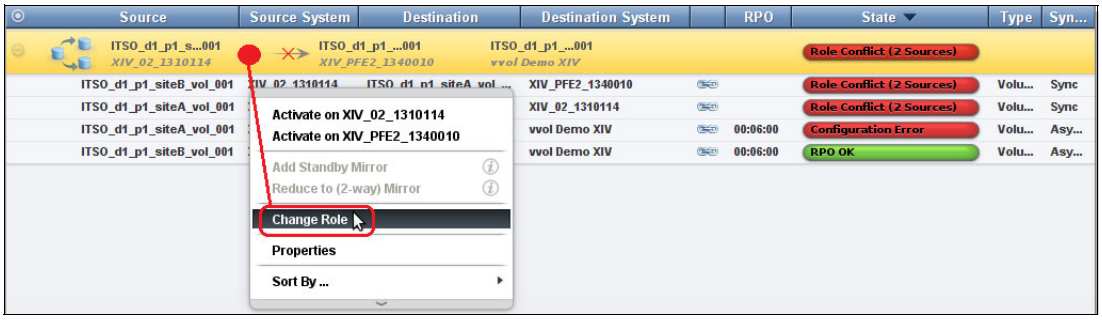

Figure 31 3-way mirror: site A failure recovery, site A change to secondary source - step 1

To change the role of Site A, right-click in the amber zone as shown in Figure 31 and select **Change Role** from the pop-up menu.

2. The Change Role dialog is displayed. Select the system at site A from the pull-down menu to have it change to a secondary source, as shown in Figure 32. Press **OK**.

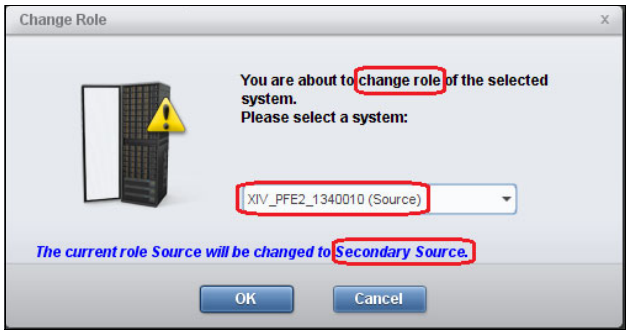

Figure 32 3-way mirror - site A failure recovery - site A change to secondary source - step 2

Site A is now the secondary source. The 3-way mirror state and the global state of the 3-way mirror displays as Compromised, as can be seen in Figure 33.

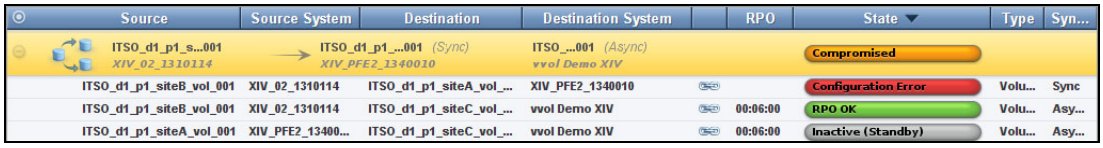

Figure 33 3-way mirror - site A failure recovery - site A changed to secondary source

3. Now the synchronization from B to A must be activated, by selecting Activate from the pop-up menu, as illustrated in Figure 34.

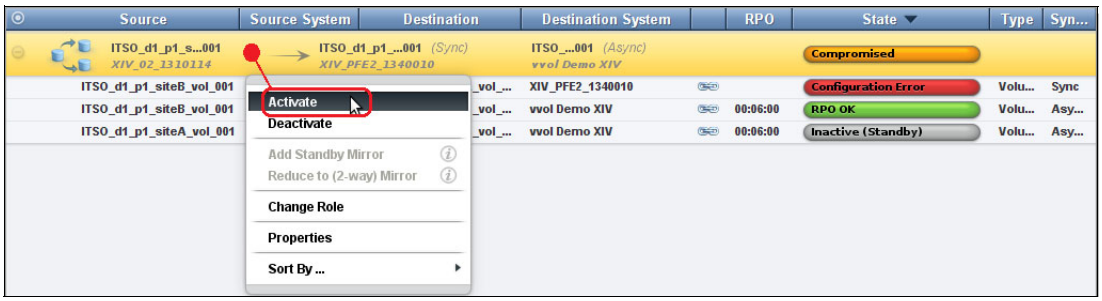

Figure 34 3-way mirror - site A failure recovery - site B to site A synchronization activation

Once the site B to site A synchronization is activated, the state changes from Configuration Error to Unsynchronized, as shown in Figure 35.

Figure 35 3-way mirror - site A failure recovery - site B to site A synchronization started

4. As the synchronization takes place, its progress is reflected in the GUI until site A is fully synchronized, as shown in Figure 36. In parallel, the production on the customer's backup site is ongoing.

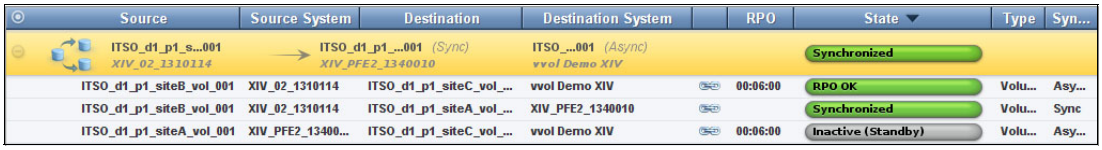

Figure 36 3-way mirror - site A failure recovery - site B and site A synchronized

5. Up to now, the production is still running on the backup site (site B). Now comes the time to return to production at the regular site (site A). All three sites are synchronized or consistent; all hardware is in healthy state. Next steps will bring site A back as the source again.
6. Production on backup site B must be terminated by stopping host I/Os to the XIV system at site B and B needs to become again the secondary source.

Important: At this stage, neither site A or site B can accept host I/Os. Production is temporarily stopped.

The 3-way mirror that was set on site B must be deactivated, as shown in Figure 37.

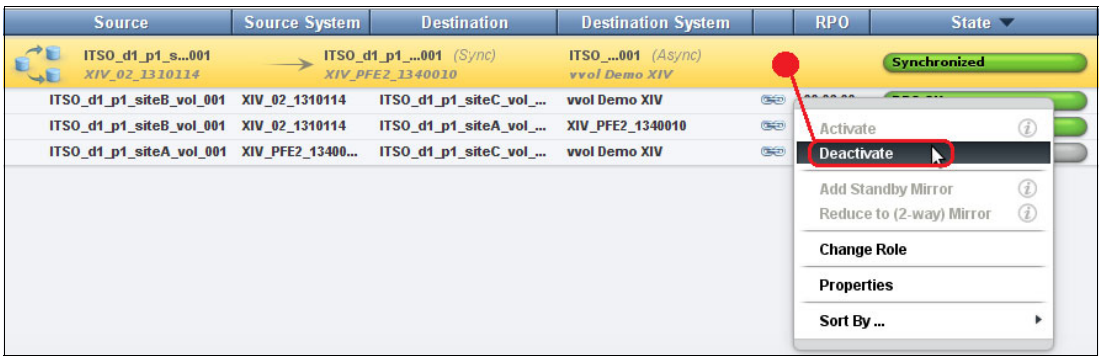

Figure 37 3-way mirror - site A failure recovery - site B mirror deactivation step 1

The mirror deactivation will set the mirror states to inactive, like expected and can be observed in Figure 38.

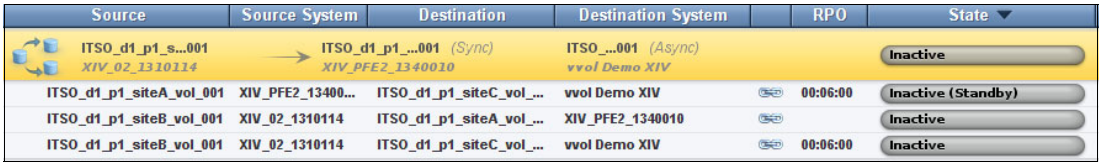

Figure 38 3-way mirror - site A failure recovery - mirror deactivation on site B

7. Now again, a change role needs to take place to set volume B back as secondary source. Follow the illustrations given in Figure 39 and Figure 40.

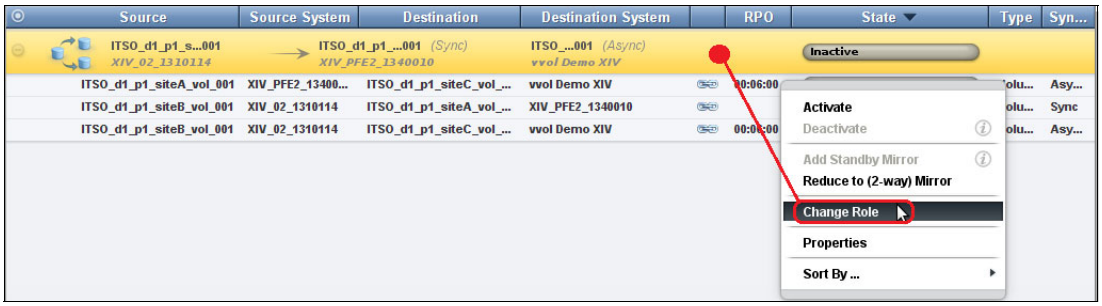

Figure 39 3-way mirror - site A failure recovery - site B preparation to become secondary source

Make sure to pick the correct XIV Storage System that needs to become the secondary source from the pull-down menu shown in Figure 40. Then, click **OK**.

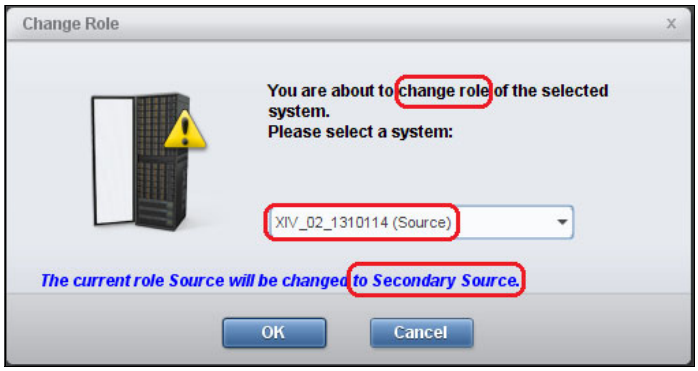

Figure 40 3-way mirror - site A failure recovery - site B choose as new secondary source

8. The new setting of site B as the secondary source again results into a role conflict like shown in Figure 41. Indeed, there are now two secondary sources.

   Note that up to this stage, neither site A or site B can accept host application write I/Os.

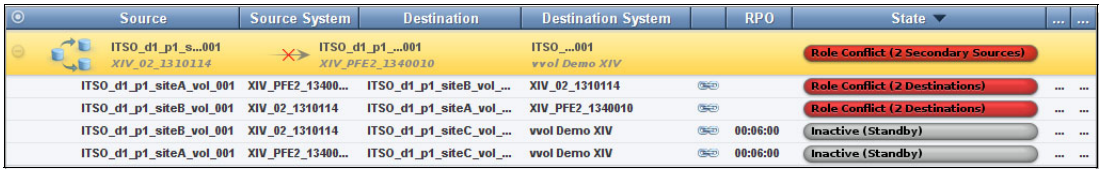

Figure 41 3-way mirror - site A failure recovery - role conflict

9. Now, we can finally operate a role change for site A to become source. Refer to Figure 42 and Figure 43.

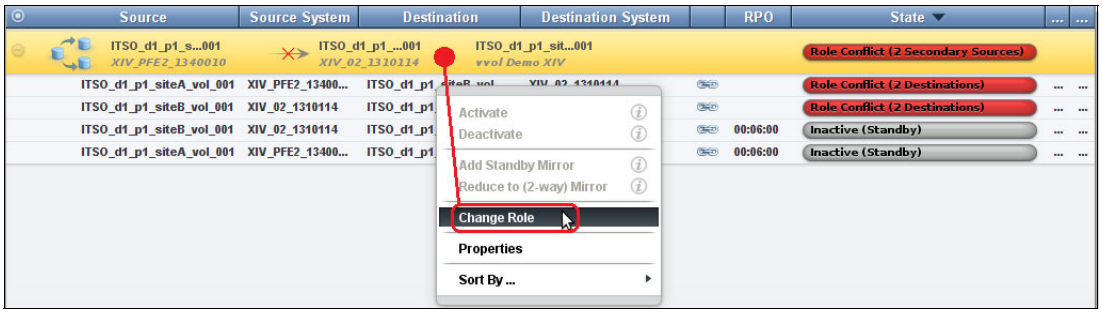

Figure 42 3-way mirror - site A failure recovery - site A switch back to source - step 1

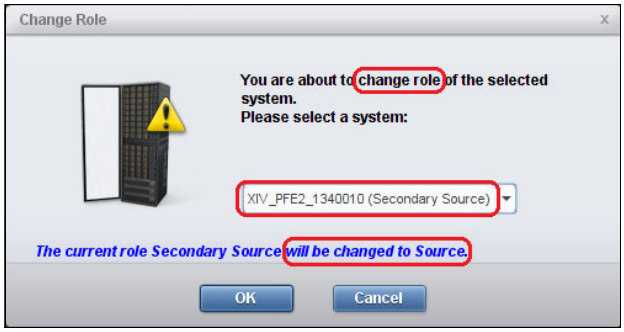

Figure 43 3-way mirror - site A failure recovery - site A switch back to source - step 2

Although Site A is now set as the source, the mirror is still inactive as can be seen in Figure 44.

Figure 44 3-way mirror - site A failure recovery - site A is back to source - mirror inactive

The storage system at site A can now accept host I/Os and production is restored.

Note that you can also delay restoring the production until after the 3-way mirror has been reactivated and data synchronized. In this case, perform step 10 on page 32 and 11 on page 32 first.

10. The 3-way mirror can be activated again (see Figure 45), which will activate the site A to site B synchronous mirroring and site A to site C asynchronous mirroring; Site B to site C will be again in asynchronous standby state.

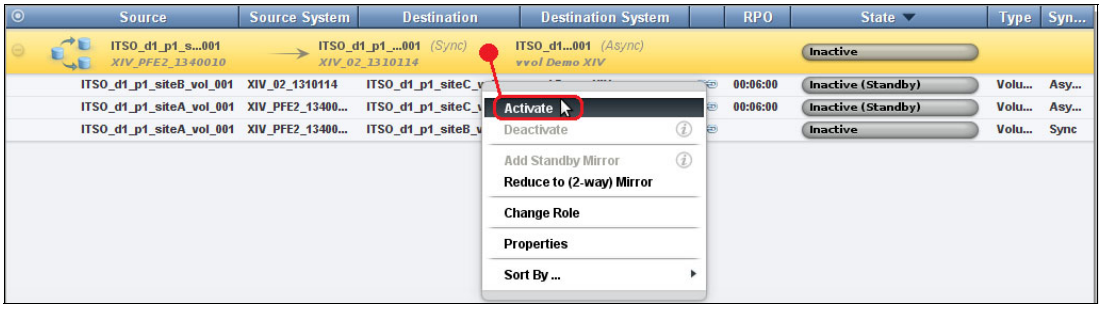

Figure 45 3-way mirror - site A failure recovery - 3-way mirror reactivation

11. Activating the 3-way mirror on site A will initially leave the mirrors in compromised, unsynchronized, and RPO Lagging states, as seen in Figure 46.

Figure 46 3-way mirror - site A failure recovery - 3-way mirror is synchronizing

After the data between sites is verified and synchronized, the mirrors will come back to normal states like shown in Figure 47.

Figure 47 3-way mirror - site A failure recovery - back to normal production

As previously noted, you can also resume production at site A at this stage.

Secondary source recovery test (validating site B)

This scenario shows how to simulate a site A failure without impacting the normal production on site A, but with the caveat that the synchronous mirror data copy (between site A and B) will not be available during that test scenario time frame.

The goal of this scenario is to verify that B site is a valid disaster recovery site without impacting normal production. For this purpose, the backup production (or a simulated production) will be activated on site B, while normal production still takes place on site A and asynchronous replication between A and C remains active.

Based on a working 3-way mirror setup, hosts are writing and reading data to and from the storage system at site A. Also, bidirectional writes and reads between site A and site B are possible (but obviously on different volumes for site A and B).

The XIV storage systems in use and their roles are as described in “Source system (site A) failure scenario” on page 23.

Follow these steps:

1. Deactivate the A to B mirror relation from the XIV system at site, as illustrated in Figure 48.

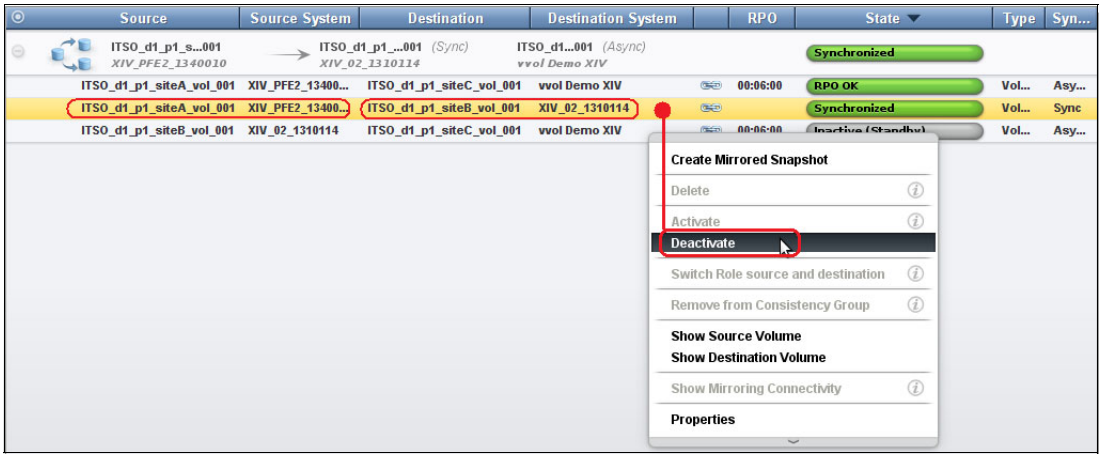

Figure 48 3-way mirror: validating site B; deactivation of the site A-B mirror relation on volume level

2. The 3-way mirror state changes to Compromised and the A to B synchronous mirror state goes to Inactive, as shown in Figure 49.

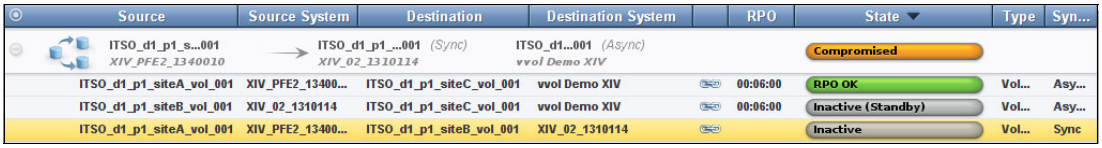

Figure 49 3-way mirror: validating site B; site A-B mirror relation deactivated

3. Now B is decoupled from A (mirror inactive) and C (mirror inactive (standby)).

   Volume B needs to become writable for the further test steps. To set site B volumes now to a writable state from hosts, a role change for volume B from secondary source to source is required. Initiate a Change Role as shown in Figure 50.

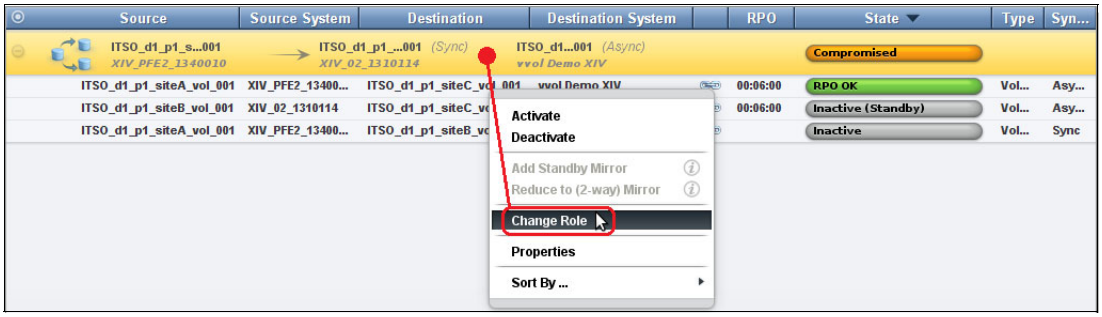

Figure 50 3-way mirror: validating site B; site B change role to source preparation

4. In the Change Role dialog, verify that you selected the right XIV system (at Site B) and click **OK**. Refer to Figure 51 on page 34.

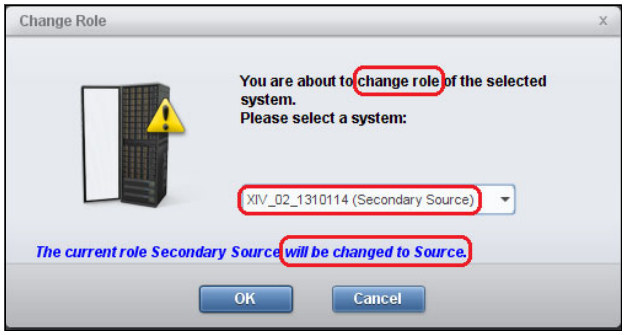

Figure 51 3-way mirror: validating site B; site B change role to source menu

5. Now two source volumes are defined, one on site A and the other on site B, which is also reflected in the GUI view (role conflict indicated in red), as illustrated in Figure 52. Volume A to C asynchronous mirror is ongoing (RPO OK - in green).

Figure 52 3-way mirror: validating site B; site B changed to source - site A to C mirror active

For the purpose of validating site B, the backup server can now be mapped to the related volumes at site B, if not already done. The backup production (or similar test setup) can be started on the XIV Storage System at site B. Additional verification tests and related aspects of the disaster recovery plan can be documented.

Site B failback after recovery test (site A will update site B)

After the site B validation test has completed, volume at site B must go back to its role as secondary source.

While tests were being performed at site B, actual production data changes were taking place on site A. To restore the normal situation, the data changes must now be applied at site B.

Follow these steps:

1. Stop I/Os (related to the site B validation test volumes) from the backup servers at site B and unmap the test volumes. In other words, restore site B as it was before the site B validation test.
2. Next, site B requires a role change to be set back as the secondary source (site B relevant volumes becomes read only (RO) from host side). Invoke the change role menu as shown in Figure 53.

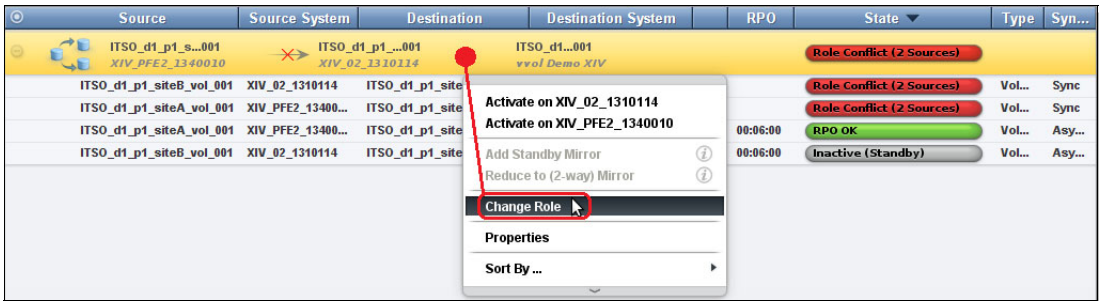

Figure 53 3-way mirror: validating site B; site B role change back to secondary source - step 1

3. In the Change Role dialog, specify the XIV Storage System (at site B), as shown in Figure 54. Click **OK**.

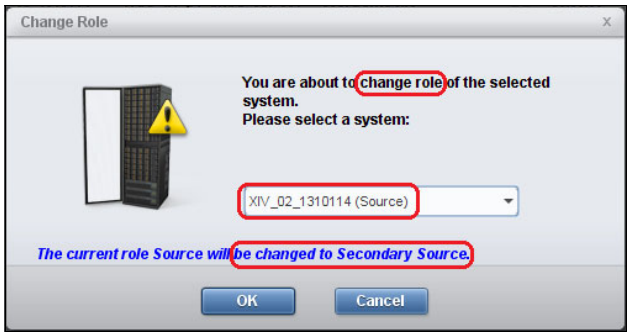

Figure 54 3-way mirror: validating site B; site B change role back to secondary source - menu

4. The role conflicts that were showing in Figure 53 on page 34 disappear. The 3-way mirror setup is almost back to normal state, except that the A to B mirroring is still in inactive state, as shown in Figure 55.

Figure 55 3-way mirror: validating site B; site B role changed back to secondary source

5. Now we activate the 3-way mirror (refer to Figure 56).

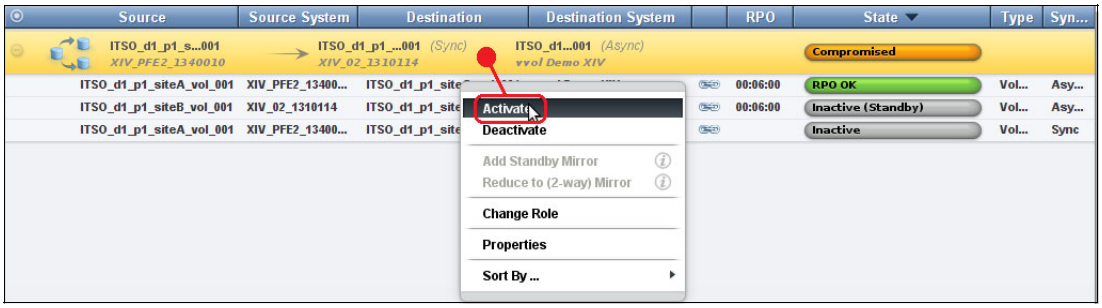

Figure 56 3-way mirror: validating site B; 3-way mirror reactivation

6. As soon as the 3-way mirror is reactivated, synchronization of volume B from A starts and changes that took place on volume A during the site B validation test, are applied to B. The A to B mirror is temporarily unsynchronized as shown in Figure 57.

Figure 57 3-way mirror: validating site B; 3-way mirror reactivated, site A to B unsynchronized

When volume A to B synchronization has completed, the 3-way mirror is back to the normal state depicted in Figure 58 on page 36. Direct host writes to B (during the site B validation) have been overridden by A.

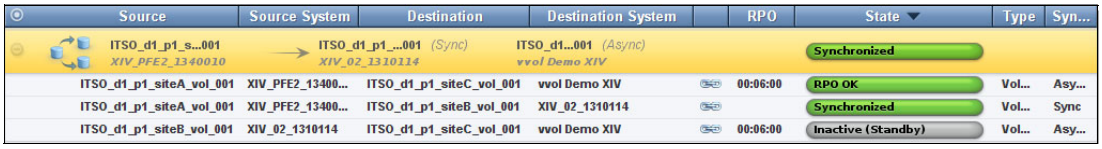

Figure 58 3-way mirror: validating site B; 3-way mirror synchronized again - final

At this point, the site B validation test is completed.

Secondary source storage system (B) failure scenario

The secondary source storage system failure scenario is very similar to the site B validation test. The difference is only that until site B is repaired and completely back in operation, no I/Os will be served from site B.

Destination storage system (C) failure scenario

We simulate the site C destination storage system failure scenario by deactivating the mirror links from site A to site C. In doing so, site C will not receive updates from site A, and also not from site B because of the inactive standby links. In this scenario, the synchronous mirror relation between sites A and B remains fully functional.

When the mirror links between site A and C are deactivated, the GUI shows the 3-way mirror in compromised state, site A to site B remains synchronized, site A to site C is RPO lagging with mirror links down and the state of the mirror relation between site B and C, is still in Inactive (Standby), as shown in Figure 59.

Figure 59 3-way mirror: site C failure

This example applies also for an XIV Storage System disaster case where site C is completely down.

Because the site A to site B mirroring remains fully functional, no actions are required until site C is fully recovered (unless you have specific applications running at site C).

Destination volume (C) failback scenario

Site C recovery is easy to do because site C is just a target (destination) storage system.

Once site C is recovered and the mirror links reactivated, the only action is to synchronize site C with the data changes that occurred at site A, while site C was down, and then being repaired. Data changes are all contained in site A most recent snapshots. Until synchronization from the most recent snapshot is not fully completed, the site A to C mirror relation remains in RPO Lagging state, as shown in Figure 60.

Figure 60 3-way mirror: site C failure recovery in RPO lagging state

After a while, synchronization is complete, and the state of the A to C mirror changes to RPO_OK as shown in Figure 61.

Figure 61 3-way mirror: site C failure recovery completed

Source (A) and secondary source (B) failure scenario

This scenario is applicable when for instance the source and secondary source are located at the same site and there is a disaster at that site that destroys both systems.

In this scenario, C must take over as the source.

Source A and secondary source (B) failback scenario

When the storage systems at site A and B have been recovered, proceed as follows to redefine the 3-way mirror as it was before the disaster:

1. If the 3-way mirror is still defined on A, B, or C, it must be removed.
2. The changes that occurred on volume C while it became the new source, must be replicated to volume A. You can also use an offline initialization to speed up the synchronization:
   - To synchronize A with C, simply activate a *synchronous* mirror from C to A, if the distance between A and C is appropriate for a synchronous connection (usually less than 100 km). It is advantageous to use a synchronous relation if possible, so that you can do a role switch next, without requiring to stop I/Os.

     After A and C are synchronized, switch roles between C and A, making A the source again and C a destination volume. Proceed to step 3.
   - If the distance between A and C does not allow synchronous mirroring, you must use an asynchronous mirroring.

     After A and C are synchronized, do a switch role. Keep in mind that with async mirroring, a switch role is not possible while there is an on-going sync job. The switch role is also not possible if C still receives host I/Os and is more updated than A. In that case, you must stop I/Os on C, wait for the sync job to complete, and then only do the switch role. Proceed to step 3.
3. Reroute all host traffic from C to A and remove the synchronous mirror between A and C.
4. Finally, re-create the 3-way mirror.

Authors

This paper was produced by a team of specialists from around the world working at the IBM International Technical Support Organization, San Jose Center.

Bertrand Dufrasne is an IBM Certified Consulting I/T Specialist and Project Leader for IBM System Storage® disk products at the International Technical Support Organization, San Jose Center. He has worked in a variety of I/T areas at IBM. Bertrand has authored many IBM Redbooks publications and has also developed and taught technical workshops. Before joining the ITSO, he was an Application Architect in IBM Global Services. He holds a Master of Electrical Engineering degree.

Christian Burns is an IBM Storage Solution Architect based in New Jersey. As a member of the Storage Solutions Engineering team in Littleton, MA, he works with clients, IBM Business Partners, and IBMers worldwide, designing and implementing storage solutions that include various IBM products and technologies. Christian's areas of expertise include IBM Real-time Compression™, SAN Volume Controller, XIV, and IBM FlashSystem™. Before joining IBM, Christian was the Director of Sales Engineering at IBM Storwize® before its becoming IBM Storwize. He brings over a decade of industry experience in the areas of sales engineering, solution design, and software development. Christian holds a BA degree in Physics and Computer Science from Rutgers College.

Wenzel Kalabza is a certified XIV Product Field Engineer (PFE) based in the IBM storage competence center in Mainz, Germany. Wenzel joined IBM in 1998 as a customer quality engineer for IBM disk drive failure and performance analysis. He joined the Back Office for the high end storage system (ESS) in June 2002. In 2005, Wenzel started a PFE role for the IBM disk storage DS6000™. In June 2008, he became a PFE for the XIV storage product. Wenzel holds a degree in Electrical Engineering and Power Economy, and several storage-related certifications.

Sandor Lengyel is an IBM Technical Advisor in the EMEA region for the XIV Storage System, based in Vac, Hungary. He joined IBM in 1998 as a Test Engineer and worked at IBM Vac Manufacturing for IBM System Storage disk products. In 2010, Sandor joined the Storage Solution Competence Center as a Technical Advisor and certified for XIV Administration and Specialist. Until 2011, he also supported the ProtecTier, Scale Out Network Attached Storage, and FlashSystem storage products. He holds a degree in Information Technology Engineering.

Patrick Schill is a Senior Certified Storage IT specialist within IBM GBS Federal Sector. He has 12 years of experience designing high performance, turnkey solutions supporting various complex host platforms, and application workloads. Patrick’s expertise and focus surrounds x86 Virtualization: Hypervisors, System Storage, Microsoft messaging solutions, and a variety of high transactional I/O database designs that support Cloud, and analytical solutions. Patrick holds a Bachelor of Science degree from DePaul University.

Christian Schoessler is an IT Specialist in the IBM ATS Disk Solution Europe Team in Mainz, Germany. He is working as a consultant for IBM for 13 years on different storage products. For the last three years, his role as technical sales support with ATS, is mainly focused in XIV. Other areas of expertise include Flash, Performance, and Copy Services calculations. He holds a degree in Physics from the Technical University of Darmstadt.

Special thanks to Oded Kellner, Ramy Buechler, Osnat Shasha for their help and advice.

Now you can become a published author, too!

Here’s an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  https://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

This document REDP-5129-00 was created or updated on September 15, 2014.

Send us your comments in one of the following ways:
- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Trademarks