


Article

[Bob Binstock](#) · Sep 7, 2016  6m read

# Mirror Outage Procedures: How to handle planned and unplanned Caché mirror outages

## [Mirror Outage Procedures](#)

Caché mirroring is a reliable, inexpensive and easy to implement high availability and disaster recovery solution for Caché and Ensemble-based applications. This article provides an overview of recommended procedures for dealing with a variety of planned and unplanned mirror outage scenarios. (For detailed information about mirroring and a wide range of mirror-related procedures, see [Mirroring 101](#).)

A Caché mirror typically consists of two Caché instances on physically independent hosts, called [failover members](#). (A mirror can also be configured with a single failover member). The mirror automatically assigns the role of primary to one failover member, while the other becomes the backup. Applications update the primary's databases, while the mirror keeps the backup's databases synchronized with the primary's.

When the primary fails or becomes unavailable, the backup automatically takes over as primary, and application connections are redirected to the new primary's databases. When the primary instance is restored to operation, it automatically becomes the backup.

Two primary factors [determine whether the backup can take over automatically](#):

- The backup must confirm that the primary is down, or at least can no longer operate as primary. When direct communication between the failover members is interrupted, the backup gets help confirming this from a third system, the [arbiter](#), which maintains independent contact with both failover members.
- The backup must either be [active](#), which means that it has the latest journal data from the primary, or be able to obtain this data through the [ISCAgent](#), a process running independently of the Caché instance on the primary's host.

Assuming an arbiter is configured, almost all [unplanned primary outages](#) with an active backup are covered by automatic failover.

Operator-initiated failover can also be used to maintain availability during planned outages for maintenance or upgrades.

The procedures referred to in this article cover planned and unplanned Caché instance outages or host system outages of one or both failover members.

Important: This article is not sufficient to guide you in performing these procedures; you must follow the links to the documentation to obtain the actual steps and details required.

## Procedures List

[Planned Outage Procedures](#)

[Maintenance of the Backup Failover Member](#)

[Maintenance of the Primary Failover Member](#)

[Upgrade of Caché Instances in the Mirror](#)

[Unplanned Outage Procedures](#)

[Unplanned Outage of the Backup Failover Member](#)

[Unplanned Outage of the Primary Failover Member With Automatic Failover](#)

[Unplanned Outage of the Primary Failover Member When Automatic Failover Does Not Occur](#)

[Unplanned Outage of Both Failover Members](#)

[Disaster Recovery Procedures](#)

[Manual Failover to a Promoted DR Async During a Disaster](#)

[Planned Failover to a Promoted DR Async](#)

[Temporary Replacement of a Failover Member with a Promoted DR Async](#)

## Planned Outage Procedures

To perform planned maintenance, you typically need to temporarily shut down the Caché instance on one of the failover members (possibly prior to shutting down the system hosting it). This is done by gracefully shutting down the failover member using the [ccontrol stop command](#).

Planned outage procedures include the following:

[Maintenance of the Backup Failover Member](#)

When you need to perform maintenance on the backup, use this procedure to ensure that a possible restart of the primary while the backup is down does not affect the mirror's operation.

[Maintenance of the Primary Failover Member](#)

When you need to fail over to keep the mirror operating while you perform maintenance on the member that is currently the primary, [verify that the backup is active](#), then use this procedure. (You can also shut down the primary [without triggering failover](#) if you wish).

## [Upgrade of Caché Instances in the Mirror](#)

All failover and DR async members of a mirror must be of the same Caché version, and can differ only for the duration of a mirror upgrade. There are several procedures to choose from, depending on whether it is a major version or maintenance release upgrade and whether you are making changes to mirrored databases.

## Unplanned Outage Procedures

When a failover member unexpectedly fails or becomes unavailable, the appropriate procedures depend on which Caché instance has failed, the [failover mode](#) the mirror was in, the status of the other failover member instance, the availability of both failover member's ISCAgents, and the mirror's settings. Before using these procedures, you may want to review [the mirror's response to various outage scenarios](#).

### [Unplanned Outage of the Backup Failover Member](#)

When the backup's Caché instance or host system fails, applications may [experience a brief pause](#) but otherwise continue functioning without incident. When the failed member returns to service, failover capability will be restored automatically.

### [Unplanned Outage of the Primary Failover Member With Automatic Failover](#)

[Automatic failover can occur](#) under almost all [primary outage scenarios](#), assuming an arbiter is configured. When the failed member returns to service, automatic failover capability is automatically restored.

### [Unplanned Outage of the Primary Failover Member When Automatic Failover Does Not Occur](#)

When the backup is unable to automatically take over from an unresponsive primary, solutions may include

- Restarting both the primary's host system and the primary Caché instance, allowing the failover members to negotiate until one becomes primary.
- Restarting the primary's host system but not the primary Caché instance, allowing the backup to obtain the needed journal data from the primary's ISCAgent.
- [Manually forcing the backup to become primary](#), which may involve copying journal files from the primary's host system.
- In the case of [unplanned network isolation of the failover members](#), pursuing one of several courses of action depending on the circumstances.

### [Unplanned Outage of Both Failover Members](#)

When both failover members unexpectedly fail, the appropriate procedure depends on whether you can restart either or both of the failover members within the limits of your availability requirements. The longer the mirror can be out of operation, the more options you are likely to have.

## Disaster Recovery Procedures

A [disaster recovery \(DR\) async](#) mirror member maintains read-only copies of all of the primary's mirrored databases, making it possible for the DR async to be [promoted to failover member](#) should the need arise. There are three scenarios in which you can use DR async promotion:

### [Manual Failover to a Promoted DR Async During a Disaster](#)

When the mirror is left without a functioning failover member, for example when a data center-wide failure occurs,

---

you can manually fail over to a promoted DR async. There are three scenarios under which this is an option:

- In a true disaster recovery scenario, in which the host systems of both failover members are down and their journal files are inaccessible, you can [promote the DR async to primary without obtaining the most recent journal data](#) from the former primary. This is likely to result in some data loss.
- If the primary's host system is running, but the Caché instance is not and cannot be restarted, you can [update the promoted DR async with the most recent journal data through the primary's ISCAgent](#).
- If the host systems of both the primary and the backup are down but you have access to the primary's journal files, or to the backup's journal files and console log, you can [update the DR async with the most recent journal data before promoting](#) it.

### [Planned Failover to a Promoted DR Async](#)

If you have included one or more DR asyncs in a mirror to provide disaster recovery capability, it is a good idea to regularly test this capability through a planned failover to each DR async. You may also want to fail over to a DR async for other reasons, such as a planned power outage in the data center containing the failover members.

### [Temporary Replacement of a Failover Member with a Promoted DR Async](#)

Some of the procedures reviewed in [Planned Outage Procedures](#) and [Unplanned Outage Procedures](#) involve temporary operation of the mirror with only one failover member—that is, with no backup. While it is not necessary to maintain a running backup, you can temporarily promote a DR async member to backup failover member if desired, protecting you from interruptions to database access and potential data loss should a primary failure occur.

[#Databases](#) [#System Administration](#) [#High Availability](#) [#Mirroring](#) [#Failover](#) [#Tips & Tricks](#) [#Caché](#)

Source URL: <https://community.intersystems.com/post/mirror-outage-procedures-how-handle-planned-and-unplanned-cach%C3%A9-mirror-outages>