



First Look: Exploring Data Resiliency with InterSystems IRIS

Version 2018.1
2018-01-31

First Look: Exploring Data Resiliency with InterSystems IRIS

InterSystems Version 2018.1 2018-01-31

Copyright © 2018 InterSystems Corporation

All rights reserved.



InterSystems, InterSystems Caché, InterSystems Ensemble, InterSystems HealthShare, HealthShare, InterSystems TrakCare, TrakCare, InterSystems DeepSee, and DeepSee are registered trademarks of InterSystems Corporation.



InterSystems IRIS Data Platform, InterSystems iKnow, Zen, and Caché Server Pages are trademarks of InterSystems Corporation.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Memorial Drive, Cambridge, MA 02142, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

InterSystems Worldwide Response Center (WRC)

Tel: +1-617-621-0700

Tel: +44 (0) 844 854 2917

Email: support@InterSystems.com

Table of Contents

First Look: Exploring Data Resiliency with InterSystems IRIS	1
1 Data Resiliency: What InterSystems IRIS Provides	1
2 Data Resiliency: How InterSystems IRIS Provides It	2
2.1 Write Image Journaling: Protection Against Physical Data Corruption and Loss	3
2.2 Journaling: Protection Against Logical Data Corruption and Loss	3
2.3 Mirroring: High Availability and Disaster Recovery Solutions	3
3 Mirroring a Database with InterSystems IRIS: Try It	4
4 InterSystems Data Resiliency: Where to Look Next	5

First Look: Exploring Data Resiliency with InterSystems IRIS

In this document, we will introduce you to all the features of InterSystems IRIS™ data resiliency and disaster recovery. Data resiliency consists of three goals – crash recovery, high availability, and disaster recovery – that are achieved by several InterSystems IRIS features.

One of the key features of InterSystems IRIS Data Platform™ is the capability to provide continuous and undisrupted access to your data by utilizing logical data replication. Data can be replicated synchronously to allow automatic failover with no data loss under a broad range of outage scenarios, either planned (such as software upgrade) and unplanned (such as hardware failure).

Synchronously replicating data requires low latency between two nodes and therefore is not always suitable for disaster recovery (DR) scenarios where one would like to transfer data across country. For those scenarios, InterSystems IRIS Data Platform provides built-in asynchronous data replication.

This document will introduce you to all the features of InterSystems IRIS data platform resiliency and disaster recovery; you will create a primary and failover mirror member, bring both members online, and make changes to the primary member. These changes will automatically be replicated to the failover member.

These activities will only use the default settings, so you can acquaint yourself with the fundamentals of these features. For the full documentation, see the [InterSystems IRIS Data Integrity Guide](#).

1 Data Resiliency: What InterSystems IRIS Provides

Data resiliency covers multiple topics, but they all revolve around one principle: once data is recorded in the data platform, it is accessible, no matter what happens.

In order to achieve this goal, there are a few different areas that need to be addressed:

- Guaranteeing data validity even in the event of errors or power failures
- Crash recovery from a local failure of the InterSystems IRIS server, be it physical failure or software related
- Protection against a full site disaster (loss of power, network issues, etc.)

To address these requirements, InterSystems IRIS provides:

- Structural database integrity (the contents of the database blocks on disk) and protection against internal integrity failures (the data represented within the database)
- Logical database integrity through transaction processing, locking, and automatic rollback
- An economical high availability solution with automatic failover
- Logical data replication that minimizes risks of carry-forward physical corruption and has no dependency on shared resources
- A solution for both planned and unplanned downtime
- Business continuity benefits via a geographically dispersed disaster recovery configuration

InterSystems IRIS data resiliency protects the data on your production systems 24/7.

2 Data Resiliency: How InterSystems IRIS Provides It

The cost of system downtime can range from thousands to millions of dollars, depending on the type and length of outage and the type of system affected. Not only is fast recovery from downtime important but also the ability to recover data and ensure your data is protected from loss and corruption. InterSystems IRIS write image journaling technology provides structural database integrity (the contents of the database blocks on disk) and protects against internal integrity failures (the data represented within the database) due to system crashes. InterSystems IRIS backup and journaling systems provide rapid recovery from physical integrity failures. Data resiliency ensures logical database integrity through transaction processing, locking, and automatic rollback. In addition to the features provided with journaling, InterSystems IRIS also allows you to mirror your databases to provide fast, efficient data replication and disaster recovery.

2.1 Write Image Journaling: Protection Against Physical Data Corruption and Loss

Due to the sequential nature of disk access, any sudden, unexpected interruption of disk or computer operation can halt the update of multiple database blocks after the first block has been written but before the last block has been updated. The consequences could be as severe as a database that is totally unusable, with all data irretrievable by normal means. The InterSystems IRIS write image journaling technology protects against this kind of data corruption. It prevents an incomplete update from leading to an inconsistent database structure.

Write image journaling safeguards database updates by using a two-phase approach. The InterSystems IRIS instance first writes updates from memory to a transitional journal, **CACHE.WIJ**, and then to the database. If the system crashes during the second phase, you can reapply updates upon recovery.

The **Write** daemon creates the write image journal (WIJ) file when InterSystems IRIS starts and records database updates in the WIJ before writing them to the InterSystems IRIS database. Once it enters all updates to the WIJ, it sets a flag in the file and the second phase begins where the **Write** daemon writes the same set of blocks recorded in the WIJ to the database on disk. When this second phase completes, the **Write** daemon sets a flag in the WIJ to indicate it is deleted. When InterSystems IRIS starts, it automatically checks the WIJ and runs a recovery procedure if it detects that an abnormal shutdown occurred. When the procedure completes successfully, the internal integrity of the database is restored. InterSystems IRIS also runs WIJ recovery following a shutdown as a safety precaution to ensure that data can be safely backed up. WIJ recovery is necessary if a system crash or other major system malfunction occurs. When InterSystems IRIS starts, it automatically checks the WIJ. If it detects that an abnormal shutdown occurred, it runs a recovery procedure. Depending on where the WIJ is in the two-phase write protocol process, recovery does the following:

- If the crash occurred after the last update to the WIJ was completed but before completion of the corresponding update to the databases, the WIJ is restored.
- If the crash occurred after the last WIJ update was durably written to the databases, a block comparison is done between the most recent WIJ updates and the affected databases.

If the recovery procedure runs, then when it completes successfully, the internal integrity of the databases has been restored.

Write image journaling provides a record of all database modifications that have been made in memory but not yet written to the database. When a system crash occurs, the system automatically writes the contents of the write image journal to the database when it restarts.

2.2 Journaling: Protection Against Logical Data Corruption and Loss

Journaling is a feature that you can enable on a per-database basis that provides a complete record of all database modifications. In the event that some modifications are lost, for example because they occurred after the most recent backup of a recovered database, you restore them to the database by restoring the contents of the journal file.

The InterSystems IRIS recovery process affords maximal protection by using the “roll forward” approach. If a system crash occurs, the recovery mechanism completes the updates that were in progress. It also protects the sequence of updates; if an update is present in the database following recovery, then all preceding updates are also present. This protects the incremental backup file structures, as well as the database. You can run a valid incremental backup following recovery from a crash.

Journaling is the basis for the third feature of data resiliency: mirroring.

2.3 Mirroring: High Availability and Disaster Recovery Solutions

InterSystems IRIS Database Mirroring (mirroring) falls within the automatic failover category of system availability strategies. InterSystems IRIS provides mirroring at a fraction of the cost of other database technologies, thus providing an economical, comprehensive, reliable, and robust enterprise solution for database availability. Traditional availability solutions that rely on shared resources (such as shared disk) are often susceptible to a single point of failure with respect to that shared resource. Mirroring reduces that risk by maintaining independent resources on the primary and backup mirror members. Further, by utilizing logical data replication, mirroring avoids the risks associated with physical replication technologies such as SAN-based replication, including out-of-order updates and carry-forward corruption.

InterSystems IRIS provides two options for high availability mirroring: Mirroring and Virtualization.

2.3.1 Mirroring

InterSystems IRIS mirroring with automatic failover relies on logical data replication between fully independent systems. This avoids the risk of having a single point of failure when using shared storage for multiple systems. It also ensures that a production system can immediately fail over to an alternate InterSystems IRIS instance in almost all failure scenarios — system, storage, and network.

In an InterSystems IRIS mirror, one InterSystems IRIS instance, called the primary failover member, provides access to the production databases. Another instance on a separate host, called the backup failover member, communicates synchronously with the primary, retrieving its journal records, acknowledging their receipt, and applying them to its own copies of the same databases. In this way, both the primary and the backup always know whether the backup has the most recent journal files from the primary, and can therefore precisely synchronize its databases with those on the primary.

Another great feature of InterSystems IRIS mirroring allows for configuration of a special async member, which can receive updates from multiple mirrors across the enterprise. You can configure an async member for disaster recovery of a single mirror, which allows it to seamlessly take the place of one of the failover members should the need arise. A single mirror can include up to 16 members, which allows you to configure numerous geographically dispersed DR async members. This model provides a robust framework for distributed data replication, thus ensuring business continuity benefits to the organization.

As an added benefit, mirroring with an async member allows a single system to act as a comprehensive enterprise data warehouse. This gives you enterprise-wide data mining and business intelligence.

2.3.2 Virtualization Platforms

Virtualization platforms generally provide HA capabilities, which typically monitor the status of both the guest operating system and the hardware it is running on. The use of mirroring in a virtualized environment, in which the InterSystems IRIS instances constituting a mirror are installed on virtual hosts, creates a hybrid high availability solution combining the benefits of mirroring with those of virtualization. While the mirror provides the immediate response to planned or unplanned

outages through automatic failover, virtualization HA software automatically restarts the virtual machine hosting a mirror member following an unplanned machine or OS outage. This allows the failed member to quickly rejoin the mirror to act as backup (or to take over as primary if necessary). On the failure of either, the virtualization platform automatically restarts the failed virtual machine, on alternate hardware as needed. When the InterSystems IRIS instance restarts, it automatically performs the normal startup recovery, maintaining structural and logical integrity as if you restarted InterSystems IRIS on a physical server.

3 Mirroring a Database with InterSystems IRIS: Try It

Now that you have some background knowledge of InterSystems IRIS Data Resiliency and what it offers, let's see a part of it in action through mirroring.

It's easy to set up a mirrored environment with InterSystems IRIS. This simple procedure walks you through the basic steps of setting up a mirror.

Creating a mirror involves configuring the primary failover member. After you create the mirror, you can add databases. This exercise will walk you through setting up a mirrored environment using Amazon Web Services.

To deploy an InterSystems IRIS mirror in the Amazon Web Services public cloud using the InterSystems Cloud Manager (ICM), follow the procedure in [Try It! Deploy InterSystems IRIS in the Cloud with ICM](#) in [First Look: InterSystems Cloud Manager](#), and in [Customize the Sample Configuration Files](#), make the following changes:

- In the defaults.json file, change "Mirror": "False", which appears at the end of the file, to "Mirror": "True".
- Customize the definitions.json file:
 - in the DM node definition, change "Count": "1" to "Count": "2"
 - remove the DS node definition
 - add an AR node definition:

```
{
  "Role": "AR",
  "Count": "1",
}
```

ICM walks you through setting up the primary and failover. Once that is done, you can try out the features by setting some globals on the primary and then seeing the updates on the failover! We'll be remaining in the ICM container. To do this:

1. Open the Terminal on the primary mirror member by running the command:

```
icm terminal -machine node-name
```

Nodes are names `Label-Role-Tag.NNNN` (Role is the ICM node type, in this case DM). For example, if you set TRYIT and TEST, the nodes will be called TRYIT-DM-TEST-0001 (primary) and TRYIT-DM-TEST-0002 (backup).

Note: You are in the DB namespace specified in ICM. A database already exists.

To see the node names you have, type:

```
icm inventory
```

2. Enter the following command in the ICM terminal:

```
Set ^testglobal="testglobal"
```

3. Press **Enter**.

4. Enter the following command in the ICM terminal:

```
write ^testglobal
```

5. Press **Enter**.

You will see the global that you have just set.

On the failover instance in the ICM terminal:

1. Enter the following command in the ICM terminal:

```
write ^testglobal
```

2. Press **Enter**.

You will see that same global returned!

Now, if you try to change the value of the testglobal global directly on the failover database you will see that the failover database is read-only. However, if you were to power down the primary and force the failover to take over, you would see that you could change the global and it would then change on the primary when it is brought back up. The primary and failover mirror members are immediately ready to take over for one another and protect your data from corruption and loss.

4 InterSystems Data Resiliency: Where to Look Next

InterSystems has lots of resources to help you learn more about InterSystems IRIS Data Resiliency:

Further reading:

- [The InterSystems IRIS High Availability Guide](#)
- [The InterSystems IRIS Data Integrity Guide](#)
- [The InterSystems IRIS Cloud Manager Guide](#)

Hands-on learning:

- [The High Availability and Disaster Recover Resource Guide](#)
- [InterSystems Cloud Manager Introduction](#)
- [Getting Started with ICM](#)

