



InterSystems IRIS Release Notes

Version 2023.3
2024-05-16

InterSystems IRIS Release Notes

InterSystems IRIS Data Platform Version 2023.3 2024-05-16

Copyright © 2024 InterSystems Corporation

All rights reserved.

InterSystems®, HealthShare Care Community®, HealthShare Unified Care Record®, IntegratedML®, InterSystems Caché®, InterSystems Ensemble®, InterSystems HealthShare®, InterSystems IRIS®, and TrakCare are registered trademarks of InterSystems Corporation. HealthShare® CMS Solution Pack™ HealthShare® Health Connect Cloud™, InterSystems IRIS for Health™, InterSystems Supply Chain Orchestrator™, and InterSystems TotalView™ For Asset Management are trademarks of InterSystems Corporation. TrakCare is a registered trademark in Australia and the European Union.

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

1 New in InterSystems IRIS 2023.3	1
1.1 Enhancing Analytics and AI	1
1.1.1 Time Series Forecasting Support in IntegratedML	1
1.2 Enhancing Cloud and Operations	1
1.2.1 Journal Archiving	1
1.3 Enhancing Speed, Scale, and Security	2
1.3.1 Faster Global Size Estimates	2
1.3.2 Faster ObjectScript Runtime	2
1.3.3 TLS 1.3 Support Update for Windows (OpenSSL 3.1.1)	2
1.4 Platform Updates	2
1.4.1 OpenSSL 3.0 for AIX	2
1.4.2 CentOS	2
1.4.3 VxFS	2
1.5 Other Enhancements and Efficiency Improvements	3
1.5.1 Change to Automatic Freezing of SQL Query Plans	3
1.5.2 General Enhancements to InterSystems SQL	3
2 New in InterSystems IRIS 2023.2	5
2.1 Release Information for 2023.2	9
2.2 Private Web Server (PWS)	5
2.3 Enhancing Developer Experience	6
2.3.1 InterSystems SQL: Safe Query Cancellation	6
2.4 Enhancing Analytics and AI	6
2.4.1 Time Series Forecasting in IntegratedML (Experimental)	6
2.4.2 InterSystems Reports 2023.2	6
2.4.3 InterSystems Adaptive Analytics 2023.2	6
2.5 Enhancing Speed, Scale, and Security	7
2.5.1 SQL Performance Improvements	7
2.5.2 Kernel-Level Performance Improvements	7
2.6 Notable Platform Changes	7
2.6.1 End of Support for Windows Server 2012	7
3 New in InterSystems IRIS 2023.1	9
3.1 Release Information for 2023.1	9
3.2 Enhancing Analytics and AI	9
3.2.1 Columnar Storage	9
3.3 Enhancing Speed, Scale and Security	10
3.3.1 Foreign Tables	10
3.3.2 Memory Settings	10
3.3.3 Platform Scalability	10
3.4 Platform Updates	10
4 New in InterSystems IRIS 2022.3	11
4.1 Enhancing Developer Experience	11
4.1.1 SQL Development	11
4.2 Enhancing Analytics and AI	11
4.3 Enhancing Speed, Scale and Security	12
4.4 Platform Updates	12

5 New in InterSystems IRIS 2022.2	13
5.1 Enhancing Analytics and AI	13
5.1.1 Columnar Storage (Experimental Feature)	13
5.2 Enhancing Speed, Scale, and Security	13
5.2.1 SQL Enhancements	13
5.3 Platform Updates	14
6 New in InterSystems IRIS 2022.1	17
6.1 Release Information for 2022.1	17
6.2 Enhancing Developer Experience	17
6.2.1 Kafka Messaging Support	17
6.2.2 Embedded Python	17
6.2.3 Interoperability Productions in Python	18
6.2.4 Visual Studio Code ObjectScript Extension Pack Updates	18
6.2.5 Make SQL Queries With Minimal Code in Interoperability Productions	19
6.3 Enhancing Analytics and AI	19
6.3.1 SQL Loader	19
6.3.2 Adaptive Analytics Enhancements	19
6.4 Enhancing Cloud and Operations	19
6.4.1 Cloud Connectors	19
6.4.2 IKO Enhancements	19
6.5 Enhancing Speed, Scale, and Security	20
6.5.1 Online Shard Rebalancing	20
6.5.2 Adaptive SQL Optimizer	20
6.5.3 Saving on Storage	20
6.5.4 TLS 1.3 Support (OpenSSL 1.1.1)	21
6.5.5 New ^TRACE Utility	21
6.6 Other Enhancements and Efficiency Improvements	21
7 New in InterSystems IRIS 2021.2	23
7.1 Enhancing Developer Experience	23
7.1.1 Embedded Python	23
7.1.2 Interoperability Productions in Python	24
7.1.3 Visual Studio Code ObjectScript Extension Pack Updates	24
7.1.4 Make SQL Queries With Minimal Code in Interoperability Productions	24
7.2 Enhancing Analytics and AI	25
7.2.1 SQL Loader	25
7.2.2 Adaptive Analytics Enhancements	25
7.3 Enhancing Cloud and Operations	25
7.3.1 Cloud Connectors	25
7.3.2 IKO Enhancements	25
7.4 Enhancing Speed, Scale, and Security	26
7.4.1 Online Shard Rebalancing	26
7.4.2 Adaptive SQL Optimizer	26
7.4.3 Saving on Storage	26
7.4.4 TLS 1.3 Support (OpenSSL 1.1.1)	26
7.4.5 New ^TRACE Utility	27
7.5 Other Enhancements and Efficiency Improvements	27
8 New in InterSystems IRIS 2021.1	29
8.1 Release Information for 2021.1	29
8.2 Enhancing Analytics	29

8.3 Enhancing Developer Experience	29
8.4 IntegratedML Machine Learning	30
8.5 Enhancing Operations	30
8.6 Enhancing Interoperability	31
8.7 Block-level Compression Reduces the Overall Storage Footprint (Experimental Feature)	32
8.8 Other Enhancements and Efficiency Improvements	32
9 New in InterSystems IRIS 2020.1	35
9.1 Continuous Delivery Releases of InterSystems IRIS	36
9.2 API Management	37
9.2.1 InterSystem API Manager	37
9.2.2 Open API/Swagger Specification-First REST Development	37
9.3 In-Place Conversion from Caché and Ensemble	38
9.4 InterSystems Reports	38
9.5 Client Language Enhancements	38
9.5.1 InterSystems IRIS Native API for Python	38
9.5.2 InterSystems IRIS Native API for Node.js	39
9.5.3 Relational access for Node.js	39
9.5.4 Java and .NET Gateway Reentrancy	39
9.5.5 Native API for Java and .NET Enhancements	39
9.5.6 Execute TSQL Code via JDBC	39
9.6 New Look in the Management Portal	39
9.7 SQL Enhancements	40
9.7.1 Universal Query Cache	40
9.8 Interoperability Production Enhancements	40
9.8.1 New PEX Framework for Coding Production Components in Java and .NET	40
9.8.2 Port Authority for Monitoring Port Usage in Interoperability Productions	41
9.8.3 X12 Validation Enhancements	41
9.8.4 Enhanced DTL Support for X12	41
9.8.5 Import X12 Schemas from XSD Files	42
9.8.6 MQTT Adapters	42
9.9 Sharding Enhancements	42
9.9.1 Simplified Architecture	42
9.9.2 Flexible Sharded Schema Design and Objects Support	42
9.9.3 Unified Shard Queue Manager	43
9.10 Infrastructure and Cloud Deployment Improvements	43
9.11 New Automatic Configuration Customization	43
9.12 Analytics Enhancements	44
9.12.1 Selective Cube Build	44
9.12.2 PowerBI Connector	44
9.12.3 Pivot Table Preview	44
9.13 Natural Language Processing Enhancements	44
9.14 Improved Performance and Scalability of the Database	45
9.15 Other Enhancements and Efficiency Improvements	45
10 New in InterSystems IRIS 2019.1	47
10.1 Extended Maintenance and Continuous Delivery Releases of InterSystems IRIS	47
10.2 InterSystems Cloud Manager Enhancements	48
10.3 Client Languages Enhancements	48
10.4 Improved Scalability and Operations for Sharded Clusters	49
10.5 SQL Enhancements	49
10.6 Analytics Enhancements	50

10.7 Interoperability Enhancements	50
10.8 System Performance and Capabilities	51
10.9 New Features in 2019.1.1 Release	52
10.9.1 In-Place Conversion from Caché and Ensemble	52
10.9.2 InterSystems API Manager	53
10.9.3 X12 Element Validation in Interoperability Productions	53
11 Upgrading to This Release	55
11.1 Upgrading Containers	55
12 Known Issues and Notes	57
12.1 Unable to Open Documentation from the Launcher	57
12.2 AmazonCloudWatchLogs Classes are Deprecated	57
13 Deprecated and Discontinued Features	59
13.1 About Deprecated and Discontinued Technologies and Features	59
13.2 InterSystems Cloud Manager	59
13.3 InterSystems IRIS Natural Language Processing	60
13.4 Private Web Server (PWS)	60
13.5 Studio	60
13.6 Unstructured Information Management Architecture (UIMA) Integration	60
13.7 Spark Connector	61
13.8 Atelier	61
13.9 Zen	61
13.10 Zen Reports	61

1

New in InterSystems IRIS 2023.3

This page describes new features, enhancements, and other significant updates in the 2023.3 release of InterSystems IRIS® data platform, which is a continuous delivery (CD) release.

Planned features for InterSystems IRIS 2023.3 include:

- Support for base HL7® FHIR® version R5
- Integration with the IBM FHIR® validator
- Runtime performance improvements for InterSystems ObjectScript

Note: These features are expected but not assured to appear in the final release.

1.1 Enhancing Analytics and AI

1.1.1 Time Series Forecasting Support in IntegratedML

The capability to create time series models with IntegratedML, introduced with InterSystems IRIS 2023.2 as an experimental feature, is now fully supported for production use.

1.2 Enhancing Cloud and Operations

1.2.1 Journal Archiving

Starting with this release, system administrators can now configure an archive location for completed journal files. When configured, after a journal file switch, the completed journal file will first be compressed (using the [Journal Compression](#) feature) and then automatically moved to this archive location, which can be on a lower-cost storage tier, such as HDD drives, or cloud storage, such as Amazon S3. Archived journal files can then be automatically deleted from the local journal directory, reducing the overall footprint on the high-performance storage tier used for writing journal files and lowering the total cost of ownership for InterSystems IRIS deployments.

1.3 Enhancing Speed, Scale, and Security

1.3.1 Faster Global Size Estimates

This release introduces a new algorithm for fast estimation of the total size of a global that uses stochastic sampling rather than a block scan. The new algorithm provides accurate estimates (within five percent of the actual size) in seconds for globals that take up hundreds of gigabytes of space. It is available as an option in the `^%G$SIZE`, `%GlobalEdit`, and `%GlobalQuery` interfaces.

1.3.2 Faster ObjectScript Runtime

When using public variables in an ObjectScript routine, the InterSystems IRIS kernel needs them to be registered in an in-memory symbol table that is specific to the routine. Starting with InterSystems IRIS 2023.3, this symbol table is built lazily, only adding public variables upon their first reference. This means only variables actually used in the routine are added to the symbol table, which may significantly reduce the cost of building it. This is an entirely transparent change to internal memory management, but it may yield noticeable performance gains for ObjectScript code that makes significant use of public variables.

1.3.3 TLS 1.3 Support Update for Windows (OpenSSL 3.1.1)

With this version, the Windows kit now installs the OpenSSL 3.1.1 library.

1.4 Platform Updates

1.4.1 OpenSSL 3.0 for AIX

InterSystems IRIS now supports OpenSSL 3.0 on AIX 7.2 and 7.3 via the `aixopenssl30` kit. This kit has, as prerequisites, the latest Tech Level for either AIX 7.2 and 7.3.

1.4.2 CentOS

InterSystems no longer supports InterSystems IRIS on CentOS. InterSystems encourages developers who had been using CentOS to switch Red Hat's free developer program.

1.4.3 VxFS

InterSystems no longer list VxFS as a supported filesystem. InterSystems recommends that SUSE customers on VxFS migrate to XFS or one of the other supported filesystems.

1.5 Other Enhancements and Efficiency Improvements

1.5.1 Change to Automatic Freezing of SQL Query Plans

Before this release, InterSystems IRIS would automatically freeze query plans upon upgrading to a new major version, in order to ensure existing plans and performance were maintained. This meant administrators had to manually unfreeze plans in order to benefit from IRIS SQL enhancements in the new release that would lead to a better plan. Starting with InterSystems IRIS 2023.3, when [Adaptive Mode](#) is enabled, query plans will no longer be frozen automatically, but invalidated upon upgrade so they get re-planned upon their first invocation after upgrading. For more details, see the [corresponding documentation](#).

1.5.2 General Enhancements to InterSystems SQL

This release introduces a number of general enhancements to IRIS SQL:

- Various performance improvements when querying tables using columnar storage, including operators such as IN and JOIN, various functions, and general data access patterns
- Various performance improvements for query statements involving date/time functions
- Several enhancements to the experimental [Foreign Tables](#) capability introduced in 2023.1
- Expanded support for Adaptive Parallel Execution to additional query types, increasing throughput for large queries
- The [Statement Index](#) now also tracks DDL statements, in addition to DML and query statements

2

New in InterSystems IRIS 2023.2

This page describes new features, enhancements, and other significant updates in the 2023.2 release of InterSystems IRIS® data platform, which is an continuous delivery (CD) release.

For other information you may wish to consider related to changes included in this release, see [Known Issues and Notes](#). For a more exhaustive list of the changes included in this release, refer to the [Upgrade Checklist](#).

2.1 Release Information for 2023.2

The current maintenance release is 2023.2.0. The posting for 2023.2.0 is build 2023.2.0.227.0.

2.2 Private Web Server (PWS)

Starting with this release:

- Local InterSystems IRIS installation will no longer install the private web server. [Access to the Management Portal and other built-in web applications](#) will require configuration of a connection to an external web server. Installation will include the option to [automatically configure](#) the Apache web server (on all platforms except Microsoft Windows) and the IIS web server (on Microsoft Windows), if already installed.
- InterSystems IRIS containers will no longer provide the private web server. Access to the Management Portal of containerized instances will require deployment of an associated Web Gateway container. Configuration examples will be available.

Note that the InterSystems IRIS Community Edition will continue to install a private web server. Remote configuration using the [Web Gateway](#) is still available.

Upgrades will not remove the private web server and existing instances may continue to use it; however, InterSystems strongly recommends the use of a web server that InterSystems IRIS supports. Automatic configuration during the installation (including upgrade installation) for Apache or IIS will assist with this change.

Users who do not need to serve InterSystems IRIS web applications can simply opt out of configuring a Web Gateway connection to a web server during installation.

Effective with the first EM release in 2026, the private web server will be discontinued; at that point, upgrades of existing InterSystems IRIS instances will remove the private web server.

An update of HealthShare Health Connect is not available with InterSystems IRIS 2023.2.

2.3 Enhancing Developer Experience

2.3.1 InterSystems SQL: Safe Query Cancellation

This release introduces a new command for InterSystems SQL to cancel running statements safely and reclaim any resources used by the statement, including memory and temporary space: **CANCEL QUERY**. The **CANCEL QUERY** command accepts the statement ID as an argument. This ID is available using the SQL Process View (`INFORMATION_SCHEMA.CURRENT_STATEMENTS`) which was introduced in a previous release. **CANCEL QUERY** ensures that participating worker threads are released without accidentally interfering with a different statement if the target one has already ended, which was a risk with manually terminating a query prior to the introduction of this capability. The InterSystems IRIS Management Portal has been updated to leverage this new functionality.

2.4 Enhancing Analytics and AI

2.4.1 Time Series Forecasting in IntegratedML (Experimental)

IntegratedML® can now be used for predicting future values for one or more time-dependent variables. The feature comprises new syntax for creating models (**CREATE TIME SERIES MODEL**), as well as for training models and consuming predictions (**SELECT WITH PREDICTIONS**). To use IntegratedML for time series forecasting, the source table or view must have a timestamp column that represents the collection time for the data. Time series models are only supported for the InterSystems AutoML and DataRobot providers; a future InterSystems IRIS release will extend this support to H2O.

This feature is currently experimental; customers who are interested should apply to take part in the accompanying Early Access Program (<https://www.intersystems.com/early-access-program/>).

2.4.2 InterSystems Reports 2023.2

InterSystems Reports 2023.2 incorporates [Logi Reports v23.2](#). New capabilities include:

- Increased functionality in Page Report Studio to enable customizations from Report Server.
- Additional Bookmark management.

This new component is available from [WRC Software Distribution](#) as **InterSystems Reports Designer** and **InterSystems Reports Server 23.2**.

2.4.3 InterSystems Adaptive Analytics 2023.2

InterSystems Adaptive Analytics 2023.2 incorporates [AtScale version 2023.2](#). New capabilities include:

- Token-based authentication to support Microsoft PowerBI in Azure environment (as requested by InterSystems)
- Scheduling of aggregate builds on a per-cube basis at convenient times for better performance management
- Usage Metrics Dashboard

This new component is available from [WRC Software Distribution](#) as **InterSystems Adaptive Analytics 2023.2**.

2.5 Enhancing Speed, Scale, and Security

2.5.1 SQL Performance Improvements

This release includes various performance enhancement across InterSystems SQL process which are 100% transparent to users:

- The Global Iterator is an internal module that shifts much of the work traversing indexes and data maps into the kernel, as opposed to leveraging ObjectScript. The use of the Global Iterator has been significantly expanded since the previous release of InterSystems IRIS.
- The Adaptive Parallel Execution framework (APE) was first introduced for queries involving columnar storage: rather than preparing and executing separate parallel subqueries, APE generates parallel worker code into the same cached query class as the main query. APE is now being expanded to an increasing number of query patterns which access regular (row) storage. This should lead to increased performance and manageability for your SQL-based applications.

For all of the preceding features, more enhancements are planned for future releases, gradually expanding the scope of these newer, more efficient query processing paradigms and frameworks.

2.5.2 Kernel-Level Performance Improvements

This release includes several kernel-level performance improvements that transparently benefit customer applications:

- Where possible, the global module now reads a set of physically contiguous big string blocks using a single, large IO operation. Big string blocks are used for storing long strings and columnar data.
- Mirroring now uses large asynchronous IO operations to read and write journal files through the agent channel, leading to a significant increase in mirror agent throughput.

2.6 Notable Platform Changes

2.6.1 End of Support for Windows Server 2012

Effective with this release, InterSystems IRIS no longer supports Windows Server 2012. To continue to run on a supported Microsoft server platform, please upgrade your operating system to Windows Server 2016, Windows Server 2019, or Windows Server 2022.

3

New in InterSystems IRIS 2023.1

This page describes the new and enhanced features in the 2023.1 release of InterSystems IRIS®, which is an extended maintenance (EM) release. Some of these features were also available in the continuous delivery (CD) releases since 2022.1, the previous EM release.

For a more exhaustive list of the changes included in this release, refer to the [Upgrade Checklist](#).

3.1 Release Information for 2023.1

The current maintenance release is 2023.1.4. The posting for 2023.1.4 is build 2023.1.4.580.0.

3.2 Enhancing Analytics and AI

3.2.1 Columnar Storage

Columnar Storage is a new storage option for InterSystems IRIS SQL tables. Columnar Storage offers analytical queries which are an order of magnitude faster than traditional row queries on InterSystems IRIS. Such queries typically aggregate data over very large tables and typically involve filters and groupings on one or more columns. By laying out the table data by column rather than by row (which works best for transactions on a handful of rows at a time), we can dramatically reduce the amount of I/O required to run such queries and exploit modern chipset-level optimizations called SIMD (Single Instruction Multiple Data) to further improve performance as part of vectorized query processing.

Note: This capability first became available in InterSystems IRIS 2022.2 as an experimental feature. It is now fully supported for production use in 2023.1, with the exception of using columnar storage for sharded tables. Support for this combination will be delivered in a future release. Customers who used the experimental version of this capability should reload all columnar table data after upgrading to 2023.1.

For more details, see [Choose an SQL Table Storage Layout](#).

3.3 Enhancing Speed, Scale and Security

3.3.1 Foreign Tables

This release introduces a new capability for leveraging external data in InterSystems IRIS. To any SQL queries you write in InterSystems IRIS, *Foreign Tables* present themselves as regular InterSystems IRIS tables. However, their data is not physically stored within the InterSystems IRIS server. They may be in remote files, third-party databases (on-prem or DBaaS), or a separate InterSystems IRIS server to which an [ECP](#) connection would not be practical. In other words, the data in these tables is not *managed* by the InterSystems IRIS instance, but it is *projected* to the InterSystems IRIS instance.

InterSystems IRIS 2023.1 includes support for projecting data from CSV files and JDBC data sources. The JDBC option leverages existing SQL Gateway infrastructure for managing connection details and credentials, and the syntax and capabilities for using file sources is fully aligned with the existing [LOAD DATA command](#).

Important: Foreign Tables are available in InterSystems IRIS 2023.1 as an Experimental Feature. This means they are not supported for production environments. However, the feature is well-tested and InterSystems believes it can add significant value to customers.

InterSystems is looking for feedback on this new capability based on customers' use in real-world environments. Please reach out on the Developer Community or contact the Worldwide Response Center (WRC) if you would like to share your experiences or you have questions.

For more details, see [Foreign Tables](#).

3.3.2 Memory Settings

New installations of InterSystems IRIS now use smarter defaults for shared memory and lock table size settings. The new defaults apply best practice configurations based on the configured global buffer size (which, in turn, considers available system memory if not set by the user) and work well for most workloads. As before, users may still override these defaults with specific values. Existing settings are not affected.

3.3.3 Platform Scalability

This release includes a number of scalability enhancements that enable large production deployments to meet highly demanding workloads. These enhancements include the asynchronous reading of journal files during de-journaling and changes to the infrastructure of the Enterprise Cache Protocol (ECP) which optimize resource usage and limit contention under very high load.

3.4 Platform Updates

This release adds support for the following new server platforms:

- macOS 13 (Ventura)

4

New in InterSystems IRIS 2022.3

This page describes the new and enhanced features in the 2022.3 release of InterSystems IRIS®, which is a Continuous Delivery (CD) release.

For a more exhaustive list of the changes included in this release, refer to the [Upgrade Checklist](#).

4.1 Enhancing Developer Experience

4.1.1 SQL Development

InterSystems SQL now offers a machine-readable format for query plans. Using this new option, the `$$SYSTEM.SQL.Explain()` method will now produce a JSON-based rendering of the query plan with significantly more detail on the plan itself as well as the tables and indexes it accesses. Where the previous XML-based format used simple English phrases to describe different steps, the new format is easier to consume for tools that wish to perform more thorough analysis or graphical rendering of query plans.

This release introduces an opt-in capability to sample runtime parameters for SQL statements. The Statement Index already records a rich set of metadata such as detailed runtime statistics and query plans for each statement, which typically has any literals substituted for placeholders that parameterize the cached query code. Now, the statement index can be extended with a sampling of the actual runtime values for those parameters. These can then be combined with the normalized statement text to, for example, build a representative workload that can be run against another deployment, for example to benchmark a new hardware environment or experiment with a different set of indexes for the schema.

InterSystems SQL now supports `CREATE SCHEMA` and `DROP SCHEMA` commands to include in scripts for setting up and tearing down your application environment.

4.2 Enhancing Analytics and AI

This release provides the following enhancements to analytics and AI:

- Updated version of InterSystems Reports (Logi Report 19.2): Key improvements are:
 - Bookmark feature - save parameters and filters on a web report
 - Report Studio available on the Report Server to provide additional report editing directly from Server

- For more details, see: <https://devnet.logianalytics.com/hc/en-us/articles/9898557594903-Logi-Report-v19-2-Release-Notes>
- Updated version of Adaptive Analytics (AtScale 2022.3) Key improvements are:
 - Support for the timeline feature in Microsoft Excel
 - Data Catalog API to expose the AtScale Semantic Layer to data catalog vendors
 - For more details, see <https://documentation.atscale.com/2022.3.1/release-notes/new-features-and-improvements> (requires login)

4.3 Enhancing Speed, Scale and Security

This release offers full elasticity for Sharded clusters. DBAs can now call an API method to mark a shard for removal. The method offloads data from the designated shard to other data nodes in the cluster and automatically disconnects the node when all data buckets have been successfully moved to other shards. This process leverages the same mechanism as Online Rebalancing, meaning users can continue to query sharded tables and ingest data into them while the data is being moved.

4.4 Platform Updates

This release adds support for the following new server platforms:

- Oracle Linux 9
- SUSE 15 SP4

5

New in InterSystems IRIS 2022.2

This page describes the new and enhanced features in the 2022.2 release of InterSystems IRIS®, which is a Continuous Delivery (CD) release.

5.1 Enhancing Analytics and AI

5.1.1 Columnar Storage (Experimental Feature)

Columnar Storage is a new storage option for IRIS SQL tables that offers an order-of-magnitude faster analytical queries compared to traditional row storage on IRIS. Such queries typically aggregate data over very large tables and typically involve filters and groupings on one or more columns. By laying out the table data by column rather than by row (which works best for transactions on a handful of rows at a time), we can dramatically reduce the amount of I/O required to run such queries, and exploit modern chipset-level optimizations called SIMD (Single Instruction Multiple Data) to further improve performance. This means the net performance gain will also depend on the chipset you run IRIS on.

Important: Columnar Storage is available as an Experimental Feature in release 2022.2. This means it is not supported for production environments. The feature is well tested and InterSystems has observed very significant performance benefits for analytical queries, though performance will vary depending on the query structure, the actual data, and the platform (chipset) you're running on.

InterSystems is looking for feedback around the performance benefits and potential space saving seen on customers' real systems — in test environments. Please contact the Worldwide Response Center (WRC) if you'd like to share your experiences or have questions.

For more details, see [Choose an SQL Table Storage Layout](#).

5.2 Enhancing Speed, Scale, and Security

5.2.1 SQL Enhancements

The 2022.2 release includes many enhancements to SQL.

5.2.1.1 SQL Process View

The SQL Process View offers a simple SQL-accessible view and corresponding page in the System Management Portal for consulting all the SQL statements that are currently running on the system. This enables administrators to quickly identify queries that may already have been running for an unusually long time and immediately compare this execution time to historical runtime statistics for that same query based on information captured in the Statement Index.

This change is fully integrated with our [System Alerting and Monitoring \(SAM\)](#) product, which leverages the same internal infrastructure to expose the corresponding metrics.

5.2.1.2 Distributing a Workload Across a Sharded Cluster

This release completes support for the object data model in sharded clusters. Where previously individual object-style access was already supported, such code, inherently procedural, would run on the node where it was invoked. With this release, we're introducing two easy-to-use API methods to invoke procedural code on all the nodes of a sharded cluster.

A **Broadcast()** method will simply invoke a routine or method once on each data and/or compute node, and can be used to kick off complex ObjectScript code across the cluster, for example to pull new data from a feed and insert it locally. A separate **Map()** method will invoke a routine or method once for every instance of a sharded class, on the data node where that instance is physically stored. These methods follow the same style of signatures as the Work Queue Manager, which distributes work across processes on a single server.

These two methods (in %SYSTEM.ShardWorkMgr) offer developers all the flexibility they need to take advantage of their sharded cluster.

5.2.1.3 Lateral JOIN Support

This release introduces support for *lateral joins*, a style of joining tables or subqueries where the different streams being joined together are not evaluated independently. A typical use case for this is when a subquery includes a reference to another table being joined, or a table-valued function taking a column value of another joined table as an input. LATERAL JOIN is a standard ANSI SQL construct.

5.2.1.4 Other

- Better schema management — This release introduces a number of convenience extensions to IRIS DDL such as the CREATE IF NOT EXISTS for tables and views and CREATE OR REPLACE for code artefacts such as procedures and functions. Also, the DDL export utility now covers more class definition features that can be expressed as DDL; see the **ExportDDL()** method of %SYSTEM.SQL.Schema.
- Projection of list collections — `List` of collection properties can now be projected to a child table, similar to how `Array` of collection properties were projected before. This means the physical storage option and SQL projection are now entirely independent.
- Performance improvements — When executing SQL queries, the use of JOIN, GROUP BY, ORDER BY and other constructs means a temporary data structure (referred to as "tempfile" internally) is built by one part of the query plan and then read by another part of the query plan. This change bundles a few changes to how internal constructs are used to store these temporary data structures to better leverage available memory and improve overall query performance. These changes take advantage of the new default for process-private memory (now defaulting to unlimited per the new default `bbsiz` setting).

5.3 Platform Updates

RedHat Enterprise Linux Updates: The 2022.2 release adds support for RHEL 9 and retires support for RHEL 7.

Ubuntu Updates: The 2022.2 release adds support for Ubuntu 22.04 and retires support for Ubuntu 18.04.

6

New in InterSystems IRIS 2022.1

This page describes the new and enhanced features in the 2022.1 release of InterSystems IRIS®, which is an Extended Maintenance (EM) release. Some of these features were also available in the continuous delivery (CD) releases since 2021.1, the previous EM release.

For a more exhaustive list of the changes included in this release, refer to the [Upgrade Checklist](#).

6.1 Release Information for 2022.1

The current maintenance release is 2022.1.5. The posting for 2022.1.5 is build 2022.1.5.940.0.

This release is the final maintenance release of version 2022.1.

6.2 Enhancing Developer Experience

6.2.1 Kafka Messaging Support

This release supports Apache Kafka <https://kafka.apache.org/>, an open-source distributed event streaming platform used for high-performance data pipelines, streaming analytics, data integration, and mission-critical applications. You can use Kafka in interoperability productions (see [Using Kafka Messaging](#)) or use the Common Messaging APIs outside of productions.

6.2.2 Embedded Python

InterSystems IRIS 2022.1 introduces Python fully integrated into the kernel, making Python a full peer to ObjectScript. Almost anything that you can do in ObjectScript, you can now also do in Python, including defining class methods. You can also interleave Python and ObjectScript, including directly calling Python libraries from ObjectScript without writing any Python code. Python provides access to many thousands of high-quality pre-built libraries, which can speed development and lower your maintenance cost. Python developers who are not familiar with ObjectScript can start developing without learning a new language. (first in 2021.2)

Any InterSystems IRIS object can be created and accessed with Embedded Python:

- Objects implemented in embedded Python are treated the same as objects implemented in ObjectScript.

- In Embedded Python, you have full and direct access from Python objects to ObjectScript objects and from ObjectScript objects to Python objects.
- Embedded Python has full access to globals, which are accessed as normal Python objects. You can use InterSystems IRIS persistence to store objects in the database, making them available in future sessions until the objects are explicitly deleted.

Embedded Python augments the InterSystems IRIS Python SDK, which includes client libraries and the external Python gateway.

For an introduction to embedded Python, see the [Embedded Python Overview](#).

Note: Embedded Python is designed to run with whatever version of Python you have installed on your machine.

If you are running Microsoft Windows and do not have Python already installed, the InterSystems IRIS installation kit installs it for you.

Many flavors of UNIX or Linux come with Python installed. If you need to install it, use the version recommended for your operating system by your package manager, for example:

- macOS: Install Python 3.9 using Homebrew (<https://formulae.brew.sh/formula/python@3.9>)
- Ubuntu: `apt-get install python3`
- Red Hat Enterprise Linux or Oracle Linux: `yum install python3`
- SUSE: `zypper install python3`

If you get an error that says “Failed to load python,” it means that you either don’t have Python installed or an unexpected version of Python is installed on your system. Install it or reinstall it using one of the above methods.

On a UNIX-based system, you may want to install Python packages with the `pip3` command. If you do not have `pip3` installed already, install the package `python3-pip` with your system’s package manager.

6.2.3 Interoperability Productions in Python

In this release, you can develop interoperability productions in Python using the Production EXtension (PEX) framework. This gives you the choice of developing productions in Python, Java, .NET, or ObjectScript. You can easily combine production components developed in different languages. You can develop in the language you are familiar with even if the other production components were developed in a different language. You can use Python with PEX to create new protocol adapter, perform complex analysis or calculations, and to create persistent messaging and long-running business processes. For more information, see [Developing Production Components with External Languages](#). (first in 2021.2)

6.2.4 Visual Studio Code ObjectScript Extension Pack Updates

The Visual Studio Code ObjectScript Extension Pack is available from the [Visual Studio Code download page](#) and has the following enhancements that make developing code faster and easier:

- Integrated documentation — hover-over in-line documentation, browse class hierarchies, and preview custom class documentation.
- Server-side source improvements — search and support for many client-side web application workflows.
- Debugging — inspect properties of objects and improved reliability.

The extension pack includes the ObjectScript extension and the Language Server extension. For more information, see the [VSCoDe ObjectScript Extension documentation https://intersystems-community.github.io/vscode-objectsript/](https://intersystems-community.github.io/vscode-objectsript/). (first in 2021.2)

6.2.5 Make SQL Queries With Minimal Code in Interoperability Productions

In this release, new SQL business services and operations make it easy to perform SQL queries in a production. See [Using SQL Business Services and Operations](#) for details. (first in 2021.2)

6.3 Enhancing Analytics and AI

6.3.1 SQL Loader

The SQL `LOAD DATA` command loads data from a CSV file or JDBC source into an SQL table efficiently. This allows you to easily populate a table with well-validated data. You can refine the command with `COLUMNS` and `VALUES` clauses similar to an `INSERT` statement and override default behavior with the `USING` clause, similar to how it's used in IntegratedML. For details, see [LOAD DATA](#) in the SQL Reference. (first in 2021.2)

6.3.2 Adaptive Analytics Enhancements

In this release, Adaptive Analytics added two additional features:

- Validation of InterSystems Reports as a client of Adaptive Analytics — now our customers can provide reports using InterSystems Reports with the same data model that is shared with PowerBI, Tableau, and other business intelligence tools. (first in 2021.2)
- Import of InterSystems IRIS Business Intelligence (BI) cubes — we have the ability to export an InterSystems IRIS BI cube definition and import it as a virtual cube in Adaptive Analytics. Please note that there are caveats with this — some things cannot be exported from a cube such as Cube Relationships and cubes based on data connectors — see the [Adaptive Analytics documentation](#) for more information. (first in 2021.2)

6.4 Enhancing Cloud and Operations

6.4.1 Cloud Connectors

This release contains cloud connectors that make it easier for you to manage InterSystems IRIS applications in Amazon Web Services and use connectors to access services. This release has the following adapters:

- Inbound and outbound adapters for S3 (Amazon Simple Storage Service) (first in 2021.2)
- Outbound adapter for Cloudwatch (Amazon monitoring service) (first in 2021.2)
- Outbound adapter for SNS (Amazon Simple Notification Service) for messaging (first in 2021.2)

6.4.2 IKO Enhancements

This release makes it easier to deploy and manage InterSystems IRIS in Kubernetes with the following new InterSystem Kubernetes Operator (IKO) features:

- IKO can deploy and manage InterSystems System Alert and Monitoring (SAM) and the InterSystems API Manager (IAM) with your InterSystems IRIS cluster. This makes it easier to administer and scale your system. (first in 2021.2)

- IKO can deploy locked down InterSystems IRIS and InterSystems Web Gateway containers. (first in 2021.2)
- IKO can deploy InterSystems Web Gateway containers with Nginx as well as Apache web servers. (first in 2021.2)
- IKO can include ephemeral as well as persistent volumes in deployments. (first in 2021.2)

6.5 Enhancing Speed, Scale, and Security

6.5.1 Online Shard Rebalancing

InterSystems IRIS [Sharding](#) distributes data and its associated workload across multiple nodes. Even data distributions offer near linear scalability for analytical query workloads. Therefore, if the data volume or workload increases, you may add additional data nodes to meet your performance goals. After adding a data node, a [rebalancing operation](#) can be used to redistribute the older data across all available nodes, which enhances the elasticity of sharded clusters. Starting in this release, this rebalancing happens online: users can continue to query and update the data while rebalancing is running. (first in 2021.2)

6.5.2 Adaptive SQL Optimizer

The InterSystems IRIS SQL Optimizer leverages table statistics to derive the best query plan for each user-submitted statement and uses an efficient query cache to reuse the generated code. When those statements include parameters, the values submitted at runtime may provide opportunities for faster execution using an alternative query plan. The new Run-Time Plan Choice (RTPC) infrastructure introduced with this release ensures InterSystems IRIS SQL takes advantage of such opportunities efficiently. RTPC scans for the use of outlier values and efficiently estimates the selectivity of range conditions based on more detailed table statistics. This leads to more adaptive query planning and significant savings in execution time and I/O for many real-world datasets. (first in 2021.2)

In addition, InterSystems IRIS now uses block-level sampling rather than full or row-based scanning to gather the table statistics used by the optimizer. This efficient algorithm enables gathering statistics (such as by using the [TUNE TABLE command](#)) for even the largest tables with billions of rows within seconds. Also, InterSystems IRIS SQL will now gather table statistics on-the-fly when a table has none to ensure appropriate query plans. (first in 2021.2)

6.5.3 Saving on Storage

In this release, stream and journal compression can significantly reduce storage needed for your InterSystems IRIS deployment:

- **Stream compression** – is now on by default for all globals-based stream classes, with no application change required. Existing data remains readable and will be compressed upon the next write. Experiments with real-world data have indicated compression ratios ranging from 30% for short texts to 80% and more for XML and other document types. (first in 2021.2)
- **Journal compression** — compresses inactive journal files immediately after journal switch. Rollback and roll forward are executed directly from the compressed format. This significantly reduces the storage requirements for this vital part of InterSystems IRIS data integrity strategy. See [Journaling Best Practices](#) for more information. (first in 2021.2)

6.5.4 TLS 1.3 Support (OpenSSL 1.1.1)

With this version, InterSystems IRIS includes support for OpenSSL 1.1.1 and fully supports TLS 1.3. With TLS 1.3, users will see faster performance among other improvements, such as cutting the encryption latency in half. This is accomplished by eliminating an entire round trip from the handshake process. (first in 2021.2)

Beginning with InterSystems IRIS 2021.2, we will no longer ship OpenSSL libraries on UNIX but depend on the operating system to provide those. One benefit of this change is that updates to the OpenSSL library no longer requires a new installation of InterSystems IRIS but can be performed with the usual operating system updates. For more information on this new approach, see [Relationship of TLS Version to Operating System and Its Version](#). Because InterSystems products require access to the operating-system-provided OpenSSL library, the product will now perform a check during the installation and the startup of an instance. The call can also be manually invoked. See [Installing the Required Dependencies](#) for details.

The change to not ship OpenSSL libraries also triggered an adjustment we needed to make for kits. Every kit is specific to a major version of OpenSSL (OpenSSL 1.1.1 is the major version, minor versions are indicated by a letter following the major version, such as OpenSSL 1.1.1f).

On Windows, the kit does install the OpenSSL library.

Note: If the correct version of OpenSSL is not installed on your UNIX system, the installation will not succeed. You must install OpenSSL and then reinstall InterSystems IRIS. For example, on MacOS, you can install OpenSSL using Homebrew, see <https://formulae.brew.sh/formula/openssl@1.1>.

6.5.5 New ^TRACE Utility

This release introduces a new tool for tracing raw events from one or more processes. Existing utilities such as %SYS.MONLBL and PERFMON track mostly the same event types but immediately generate a report formatted for a specific type of analysis. The new ^TRACE tool captures these events in a more generic file format and allows interactive navigation and summarization of the captured information through a command-line interface or API. Supported event types include, but are not limited to global sets and kills, physical writes, network requests, cache hits and reads, and various journal events. Information captured for these events includes the routine line and call stack, as well as the full global reference where applicable. This offers a single interface for a broader set of performance analysis tasks. (first in 2021.2)

6.6 Other Enhancements and Efficiency Improvements

In each release, InterSystems makes many efficiency improvements and minor enhancements. This release includes:

- DataMove is enhanced for general robustness and for recoverability around mirror failover.
- Separate kits are provided for installing InterSystems IRIS with IntegratedML or without, making the installation process more efficient.
- Compact double support for external clients, including JDBC, .NET clients, Python, and IRISNative. (first in 2021.2)
- In this release, the security tables now have an embedded version number, which allows finer access over allowable imports. You can export security tables from version 2021.1 and then import them to this version. For details, see [^SECURITY](#). (first in 2021.2)
- This release updates the Log4j library to version 2.17.0. (first in 2021.2)
- This release updates the node.js library to version 14. (first in 2021.2)

7

New in InterSystems IRIS 2021.2

This page describes the new and enhanced features in the 2021.2 release of InterSystems IRIS®, which is a [Continuous Delivery \(CD\) release](#). This is the first CD release to be [available on server platforms](#), in addition to an [OCCI-compliant container format](#).

7.1 Enhancing Developer Experience

7.1.1 Embedded Python

InterSystems IRIS 2021.2 introduces Python fully integrated into the kernel, making Python a full peer to ObjectScript. Almost anything that you can do in ObjectScript, you can now also do in Python, including defining class methods. You can also interleave Python and ObjectScript, including directly calling Python libraries from ObjectScript without writing any Python code. Python provides access to many thousands of high-quality pre-built libraries, which can speed development and lower your maintenance cost. Python developers who are not familiar with ObjectScript can start developing without learning a new language.

Any InterSystems IRIS object can be created and accessed with Embedded Python:

- Objects implemented in embedded Python are treated the same as objects implemented in ObjectScript.
- In Embedded Python you have full and direct access from Python objects to ObjectScript objects and from ObjectScript objects to Python objects.
- Embedded Python has full access to globals, which are accessed as normal Python objects. You can use InterSystems IRIS persistence to store objects in the database, making them available in future sessions until the objects are explicitly deleted.

Embedded Python augments the InterSystems IRIS Python SDK, which includes client libraries and the external Python gateway.

For an introduction to embedded Python, see the [Embedded Python Overview](#).

Note: Embedded Python is designed to run with whatever version of Python you have installed on your machine.

If you are running Microsoft Windows and do not have Python already installed, the InterSystems IRIS installation kit installs it for you.

Many flavors of UNIX or Linux come with Python installed. If you need to install it, use the version recommended for your operating system by your package manager, for example:

- macOS: Install Python 3.9 using Homebrew (<https://formulae.brew.sh/formula/python@3.9>)
- Ubuntu: `apt-get install python3`
- Red Hat Enterprise Linux or Oracle Linux: `yum install python3`
- SUSE: `zypper install python3`

If you get an error that says “Failed to load python,” it means that you either don’t have Python installed or an unexpected version of Python is installed on your system. Install it or reinstall it using one of the above methods.

On a UNIX-based system, you may want to install Python packages with the pip3 command. If you do not have pip3 installed already, install the package python3-pip with your system’s package manager.

7.1.2 Interoperability Productions in Python

In this release, you can develop interoperability productions in Python using the Production EXtension (PEX) framework. This gives you the choice of developing productions in Python, Java, .NET, or ObjectScript. You can easily combine production components developed in different languages. You can develop in the language you are familiar with even if the other production components were developed in a different language. You can use Python with PEX to create new protocol adapter, perform complex analysis or calculations, and to create persistent messaging and long-running business processes. For more information, see [Developing Production Components with External Languages](#).

7.1.3 Visual Studio Code ObjectScript Extension Pack Updates

The Visual Studio Code ObjectScript Extension Pack is available from the [Visual Studio Code download page](#) and has the following enhancements that make developing code faster and easier:

- Integrated documentation — hover-over in-line documentation, browse class hierarchies, and preview custom class documentation.
- Server-side source improvements — search and support for many client-side web application workflows.
- Debugging — inspect properties of objects and improved reliability.

The extension pack includes the ObjectScript extension and the Language Server extension. For more information, see the [VSCode ObjectScript Extension documentation](#) at <https://intersystems-community.github.io/vscode-objectsript/>.

7.1.4 Make SQL Queries With Minimal Code in Interoperability Productions

In this release, new SQL business services and operations make it easy to perform SQL queries in a production. See [Using SQL Business Services and Operations](#) for details.

7.2 Enhancing Analytics and AI

7.2.1 SQL Loader

The SQL `LOAD DATA` command loads data from a CSV file or JDBC source into an SQL table efficiently. This allows you to easily populate a table with well-validated data. You can refine the command with `COLUMNS` and `VALUES` clauses similar to an `INSERT` statement and override default behavior with the `USING` clause, similar to how it's used in IntegratedML. For details see [LOAD DATA](#) in the SQL Reference.

7.2.2 Adaptive Analytics Enhancements

In this release, Adaptive Analytics added two additional features:

- Validation of InterSystems Reports as a client of Adaptive Analytics — now our customers can provide reports using InterSystems Reports with the same data model that is shared with PowerBI, Tableau and other Business Intelligence tools.
- Import of InterSystems IRIS Business Intelligence (BI) cubes — we have the ability to export an InterSystems IRIS BI cube definition and import it as a virtual cube in Adaptive Analytics. Please note that there are caveats with this — some things cannot be exported from a cube such as Cube Relationships and cubes based on data connectors — see the [Adaptive Analytics documentation](#) for more information.

7.3 Enhancing Cloud and Operations

7.3.1 Cloud Connectors

This release contains cloud connectors that make it easier for you to manage InterSystems IRIS applications in Amazon Web Services and use connectors to access services. This release has the following adapters:

- Inbound and outbound adapters for S3 (Amazon Simple Storage Service)
- Outbound adapter for Cloudwatch (Amazon monitoring service)
- Outbound adapter for SNS (Amazon Simple Notification Service) for messaging

7.3.2 IKO Enhancements

This release makes it easier to deploy and manage InterSystems IRIS in Kubernetes with the following new InterSystem Kubernetes Operator (IKO) features:

- IKO can deploy and manage InterSystems System Alert and Monitoring (SAM) and the InterSystems API Manager (IAM) with your InterSystems IRIS cluster. This makes it easier to administer and scale your system.
- IKO can deploy locked down InterSystems IRIS and InterSystems Web Gateway containers.
- IKO can deploy InterSystems Web Gateway containers with Nginx as well as Apache web servers.
- IKO can include ephemeral as well as persistent volumes in deployments.

7.4 Enhancing Speed, Scale, and Security

7.4.1 Online Shard Rebalancing

InterSystems IRIS [Sharding](#) distributes data and its associated workload across multiple nodes. Even data distributions offer near linear scalability for analytical query workloads. Therefore, if the data volume or workload increases, you may add additional data nodes to meet your performance goals. After adding a data node, a [rebalancing operation](#) can be used to re-distribute the older data across all available nodes, which enhances the elasticity of sharded clusters. Starting in this release, this rebalancing happens online: users can continue to query and update the data while rebalancing is running.

7.4.2 Adaptive SQL Optimizer

The InterSystems IRIS SQL Optimizer leverages table statistics to derive the best query plan for each user-submitted statement and uses an efficient query cache to reuse the generated code. When those statements include parameters, the values submitted at runtime may provide opportunities for faster execution using an alternative query plan. The new Runtime Plan Choice (RTPC) infrastructure introduced with this release ensures InterSystems IRIS SQL takes advantage of such opportunities efficiently. RTPC scans for the use of outlier values and efficiently estimates the selectivity of range conditions based on more detailed table statistics. This leads to more adaptive query planning and significant savings in execution time and I/O for many real-world datasets.

In addition, InterSystems IRIS now uses block-level sampling rather than full or row-based scanning to gather the table statistics used by the optimizer. This efficient algorithm enables gathering statistics (such as by using the [TUNE TABLE command](#)) for even the largest tables with billions of rows within seconds. Also, InterSystems IRIS SQL will now gather table statistics on-the-fly when a table has none to ensure appropriate query plans.

7.4.3 Saving on Storage

In this release, stream and journal compression can significantly reduce storage needed for your InterSystems IRIS deployment:

- **Stream compression** – is now on by default for all globals-based stream classes, with no application change required. Existing data remains readable and will be compressed upon the next write. Experiments with real-world data have indicated compression ratios ranging from 30% for short texts to 80% and more for XML and other document types.
- **Journal compression** — compresses inactive journal files immediately after journal switch. Rollback and roll forward are executed directly from the compressed format. This significantly reduces the storage requirements for this vital part of InterSystems IRIS data integrity strategy. See [Journaling Best Practices](#) for more information.

7.4.4 TLS 1.3 Support (OpenSSL 1.1.1)

With this version, InterSystems IRIS includes support for OpenSSL 1.1.1 and fully supports TLS 1.3. With TLS 1.3 users will see faster performance among other improvements, such as cutting the encryptions latency in half. This is accomplished by eliminating an entire round trip from the handshake process.

Beginning with InterSystems IRIS 2021.2 we will no longer ship OpenSSL libraries on UNIX but depend on the OS to provide those. One benefit of this change is that updates to the OpenSSL library no longer requires a new installation of InterSystems IRIS but can be performed with the usual OS updates. For more information on this new approach, see [Relationship of TLS Version to Operating System and Its Version](#). Because InterSystems products require access to the OS provided OpenSSL library, the product will now perform a check during the installation and the startup of an instance. The call can also be manually invoked. See [Installing the Required Dependencies](#) for details.

The change to not ship OpenSSL libraries also triggered an adjustment we needed to make for kits. Every kit is specific to a major version of OpenSSL (OpenSSL 1.1.1 is the major version, minor versions are indicated by a letter following the major version, such as OpenSSL 1.1.1f).

On Windows, the kit does install the OpenSSL library.

Note: If the correct version of OpenSSL is not installed on your UNIX system, the installation will not succeed. You must install OpenSSL and then reinstall InterSystems IRIS. For example, on MacOS, you can install OpenSSL using Homebrew, see <https://formulae.brew.sh/formula/openssl@1.1>.

7.4.5 New ^TRACE Utility

This release introduces a new tool for tracing raw events from one or more processes. Existing utilities such as %SYS.MONLBL and PERFMON track mostly the same event types but immediately generate a report formatted for a specific type of analysis. The new ^TRACE tool captures these events in a more generic file format and allows interactive navigation and summarization of the captured information through a command-line interface or API. Supported event types include, but are not limited to global sets and kills, physical writes, network requests, cache hits and reads, and various journal events. Information captured for these events includes the routine line and call stack, as well as the full global reference where applicable. This offers a single interface for a broader set of performance analysis tasks.

7.5 Other Enhancements and Efficiency Improvements

In each release, InterSystems makes many efficiency improvements and minor enhancements. This release includes:

- Compact double support for external clients, including JDBC, .NET clients, Python, and IRISNative.
- In this release, the security tables now have an embedded version number, which allows finer access over allowable imports. You can export security tables from version 2021.1 and then import them to this version. For details, see [^SECURITY](#).
- This release updates the Log4j library to version 2.17.0.
- This release updates the node.js library to version 14.

8

New in InterSystems IRIS 2021.1

This page describes the new and enhanced features in the 2021.1 release of InterSystems IRIS®, which is an extended maintenance (EM) release. Some of these features were also available in the continuous delivery (CD) releases since 2020.1, the previous EM release.

8.1 Release Information for 2021.1

The current maintenance release is 2021.1.3. The posting for 2021.1.3 is build 2021.1.3.389.0.

8.2 Enhancing Analytics

With InterSystems IRIS 2021.1, customers can deploy InterSystems IRIS [Adaptive Analytics](#), an add-on product that extends InterSystems IRIS to deliver greater ease of use, flexibility, scalability, and efficiency to analytics end users regardless of their business intelligence (BI) tools of choice. It enables defining an analytics-friendly business model and transparently accelerates analytic query workloads that run against this model by autonomously building and maintaining interim data structures in the background.

Other enhancements for Analytics use cases include:

- SQL users querying InterSystems IRIS directly now have access to SQL standard [Window Functions](#) for easily expressing complex aggregations in a single query. This enables pushing data-intensive calculations typical for reporting and BI use cases closer to the data and improves both net performance and simplicity.
- InterSystems IRIS embedded BI capability will experience measurable performance improvements thanks to enhancements to its native MDX query engine.

8.3 Enhancing Developer Experience

InterSystems IRIS 2021.1 introduces a new Python and R Gateway to run code in those languages out-of-process or on a different server as needed. The Python Gateway supports virtual environments, meaning that each Gateway can use its own version of Python to maximize developer flexibility. All gateways now support starting automatically on first use, as well as simple and secure re-entrant connections. This means, for example, that external code can transparently reach back into

InterSystems IRIS for data access without having to explicitly open a new connection and provide user credentials. The .NET Gateway now supports .NET Core 2.1 (first in 2020.3).

New in InterSystems IRIS 2021.1 is the ability to define external stored procedures in SQL. This enables developers to leverage code written in Java, Python or .NET from SQL, using simple SQL syntax. The external code is invoked transparently through the corresponding Gateways. This release also adds SQL user-defined aggregate functions with [CREATE AGGREGATE](#) and [DROP AGGREGATE](#).

This release also adds a number of significant enhancements and extensions to the client APIs, including a fully native Python client SDK that runs on all platforms supporting Python, and support for \$list and \$order operations in all four Native SDK languages (Python, Java, C# and Node.JS). The ODBC driver has been enhanced with additional T-SQL support. XEP adds support for deferred indexing and indexes can be built as a background process (first in 2020.3).

Java developers can now take advantage of Java SE 11 LTS. InterSystems tests and supports both the Oracle OpenJDK and AdoptOpenJDK implementations of the standard for all of its Java-based components. See the corresponding section in the Supported Technologies list for more detail (first in 2020.4).

The InterSystems IRIS JDBC driver now fully supports Connection Pooling for efficient managing of database connections from your Java applications. See the Connection Pooling section in the JDBC documentation for more detail (first in 2020.4).

For ObjectScript developers, the VSCode-ObjectScript Version 1.0 is available. [VSCode-ObjectScript](#) is an open source extension for the VSCode IDE to enable practical development of ObjectScript applications for InterSystems IRIS. For details, see the [VS Code InterSystems ObjectScript documentation](#).

This release has support for Spark 2.4.4 (first in 2020.3).

8.4 IntegratedML Machine Learning

This release includes IntegratedML, a new feature that brings best of breed machine learning to analysts and developers via simple and intuitive SQL syntax (first in 2020.3). Developers can now easily train and deploy powerful predictive models from within InterSystems IRIS, right where their data lives. For details, see [Using IntegratedML](#) and [Learn IntegratedML in InterSystems IRIS](#).

For this release:

- Standard and Community Edition containers are available from the [InterSystems Container Registry](#) (ICR). See [Using the InterSystems Container Registry](#) for information on the container registry.
- Community Edition containers are also available from [Docker Hub](#).
- Kits (and container tarballs) are available from the [WRC Software Distribution site](#).

Note: Full installation kits are provided for a subset of server platforms on the WRC. When using the installer, you must specify a custom install and select the IntegratedML option to install it on your system.

8.5 Enhancing Operations

This release provides the following enhancements to the deployment and operations experience, both in the cloud and on-premises:

- The [InterSystems Kubernetes Operator](#) (IKO) packages InterSystems IRIS-specific knowledge and best practices into an easy-to-use, automated tool for provisioning and operating dynamic clusters. Starting with 2021.1, IKO also supports deploying InterSystems [System Alerting & Monitoring](#) (SAM).

- The InterSystems Cloud Manager (ICM) adds support for [InterSystems API Manager](#) (first in 2020.3) and [SAM](#) (first in 2020.4) deployments.
- This release include asynchronous mirroring support for sharded clusters (first in 2020.3). Users can now configure mirroring (synchronous or asynchronous) on an existing cluster, or fail over the entire cluster to the set of asynchronous mirror members in another data center in Disaster Recovery scenarios. See the corresponding section in the [Scalability Guide](#) for more details (first in 2020.4).
- The [InterSystems SQL syntax](#) has been extended with a set of new commands for managing and configuring your database from a SQL prompt. This enables users with just JDBC or ODBC access to perform most administrative tasks without requiring access to the System Management Portal or an ObjectScript terminal prompt. It includes common tasks such as building indexes and managing frozen plans. For details, see [BUILD INDEX](#), [FREEZE PLANS](#), [PURGE CACHED QUERIES](#), [CREATE INDEX](#), and new options in [SET OPTION](#) (first in 2020.4).
- You can now manage Work Queues from the System Management Portal (first in 2020.3).
- The newly available iris-locked-down container is a security-hardened container image that implements many security best practices, offering peace of mind for customers deploying sensitive applications in complex environments. Users of the Web Gateway container will be pleased to see improvements to its default configuration.
- Starting with 2021.1, InterSystems IRIS is now available for ARM platforms, both as full kits and pre-packaged containers. This enables customers to deploy their applications to cost-efficient hardware platforms, both physical and in the cloud. For more information, refer to the Supported Platforms guide.
- This release simplifies the deployment of InterSystems Reports, the new reporting capability for InterSystems IRIS (first in 2020.4). As part of a closer integration, InterSystems Reports now uses the same user accounts as InterSystems IRIS for managing, building and executing reports. In addition, all configuration and management data for InterSystems Reports uses InterSystems IRIS if the setup scripting is used. A script to complete the initial configuration of InterSystems IRIS Report Server for on-prem deployments and a docker-compose file for Docker deployments of the Reports Server are both available as part of this release.

8.6 Enhancing Interoperability

With InterSystems IRIS 2021.1, customers can deploy [InterSystems API Manager](#) (IAM) 2.3, which includes many enhancements broadening the reach of this crucial component in a modern API-centric environments.

There are the following other interoperability enhancements:

- New SOAP Business Service and Business Operation, `EnsLib.EDI.X12.Service` and `EnsLib.EDI.X12.Operation` that allow you to use SOAP to receive and send X12 messages (First in 2020.2).
- Improved X12 error handling (First in 2020.2).
- This release adds support for a new "foreach" action, which can be used within Routing Rules used for segmented virtual documents (ASTM, EDIFACT and X12). The foreach action is supported in the Rule Type "Segmented Virtual Document Message Routing Rule". The foreach action can loop over repeating segments in the virtual document and nested loops are supported. This enables developers to build rules that match certain conditions regardless of the position of a segment within a repeating group. For details, see [About Actions](#) (first in 2020.4).

8.7 Block-level Compression Reduces the Overall Storage Footprint (Experimental Feature)

Block-level compression reduces the overall storage footprint (amount of disk required). Depending on the data cardinality (extent of repetition in data values), compression can reduce storage consumption significantly. Sparsely populated data (lots of zeros and spaces in the data) compress much better.

Block-level compression is an *experimental* feature in InterSystems IRIS 2020.2. This means that it is not supported for production. The feature is well tested, and compression can provide substantial space savings; however there may be an impact on performance. In some cases, performance may increase as less data needs to be moved from disk to memory; in others it could decrease because of the computation to do compression and decompression. Space savings and performance impact will depend on the application and platform.

InterSystems is looking for feedback around the space saving and performance overhead seen on customers' real systems — in test environments. Please contact the [Worldwide Response Center](#) (WRC) if you are measuring this in your environment.

Currently, the three compression types provided are zlib, zstd, and lz4. Only data and big string blocks are compressed, although this could be extended to other block types in the future. A given database block is compressed only if the compression will allow freeing at least one 4KB chunk of space on disk. Block compression is only supported on Linux systems that support sparse files.

8.8 Other Enhancements and Efficiency Improvements

In each release, InterSystems makes many efficiency improvements and minor enhancements. This release includes these improvements:

- You can now use [Proof Key for Code Exchange \(PKCE\)](#) with OAuth authentication. PKCE enables you to securely perform the OAuth exchange from public clients and mitigates the threat of having the authorization code intercepted. PKCE is supported in both the [OAuth clients and servers](#).
- The configuration and utility functions in `$$SYSTEM.SQL` have now been organized thematically in subclasses such as `%SYSTEM.SQL.Functions`, `%SYSTEM.SQL.Schema` and `%SYSTEM.SQL.Stats.Runtime`, making their signature and behavior more consistent across the board. The old entry points in `%SYSTEM.SQL` have been deprecated but are still available for backwards compatibility (first in 2020.3).
- You can now use Transact-SQL through JDBC. Please see the [Transact-SQL Migration Guide](#) for more on hosting Transact-SQL applications on InterSystems IRIS (first in 2020.3).
- Node.js Native SDK now includes the List class. See [Native SDK Quick Reference for Node.js](#) (first in 2020.3).
- Java Messaging Service (JMS) adapter is able to connect to a broader range of servers. (first in 2020.3)
- InterSystems IRIS on Linux has been enhanced to use Asynchronous I/O for writes to database files, as it always has on all UNIX® and Windows platforms (first in 2020.3). This is coupled with automatic use of direct I/O instead of buffered I/O. This change optimizes the disk I/O characteristics for database files in the following ways:
 - Improves application responsiveness at higher scaling levels by more fairly sharing I/O bandwidth with database reads and journal writes.
 - Improves integrity check performance by allowing integrity check to read multiple blocks asynchronously.
 - Improves effectiveness of asynchronous reads performed by `$prefetchon`.

- InterSystems SQL saw a number of performance enhancements that will speed up many different types of queries when upgrading to this release. A complex customer benchmark composed of millions of queries ran 6% faster on 2020.4 and 2021.1 compared to earlier releases.
- This release introduces a new algorithm for estimating field selectivity when gathering table statistics. This improves the ability of the SQL optimizer to choose the fastest query plan for any given SQL statement (first in 2020.4).
- For a number of applicable scenarios, the SQL engine will now use a kernel-level iterator to read through temporary results. This simplifies generated code and can speed up certain steps in a query plan, such as sorting, by up to 40% (first in 2020.4).
- Recent updates to browser security have changed handling of third-party cookies. These updates use the SameSite attribute to reduce the risk of cross-site request forgery (CSRF) attacks, unauthorized access to data, and other possible security issues. Chrome (starting with v.84) enforces stricter rules for SameSite behavior, and these rules can cause issues with existing websites and web applications. These issues may include login problems, the login page being displayed repeatedly, and page elements not displaying properly. In previous versions, you could not modify the SameSite attribute; hence web applications running on these versions may have such issues.

This release (first in 2020.4) solves these problems by setting the SameSite attribute for cookies and by allowing you to set change the default setting; however, you may need to modify your code to customize values setting the session cookie scope and the user cookie scope. Additionally, if you are using “SameSite=None”, you must ensure that your web application can support secure HTTPS connections. For details, see [About the SameSite Attribute](#).

- This release improves performance on newly installed systems where the database cache size has not been configured. Under most circumstances, you should carefully configure cache sizes and Configure Huge and Large Pages for optimal system performance. Configuring cache sizes is especially important for live production systems, systems with heavy loads, and systems with multiple instances. (first in 2020.3)
- This release improves security of the command-line history by not recording it if the user is ‘root’ or has administrative privileges. Usually command line history is written to ~/.iris_history, where ~ expands to the value of \$HOME (the user's home directory). If you scroll before the first command in the current session, the command history from the log is used. When the user is ‘root’ command history is not written to the log or read from previous sessions so as not to expose any commands executed as superuser.

9

New in InterSystems IRIS 2020.1

This document describes the new and enhanced features in the 2020.1 release of InterSystems IRIS®. It addresses the new features in 2020.1 that were not present in the 2019.1.0 version (see [New and Enhanced Features for InterSystems IRIS 2019.1](#)). Some of these features were first introduced in a 2019.1 maintenance release or in a 2019.2, 2019.3, or 2019.4 continuous delivery release. These features are identified in the descriptions. The following sections describes the 2020.1 release and its new capabilities and enhancements:

- [Continuous Delivery releases of InterSystems IRIS](#)
- [API Management](#)
 - [InterSystems API Manager](#)
 - [Open API/Swagger Specification-First REST Development](#)
- [In-Place Conversion from Caché and Ensemble](#)
- [InterSystems Reports](#)
- [Client Language Enhancements](#)
 - [InterSystems IRIS Native API for Python](#)
 - [InterSystems IRIS Native API for Node.js](#)
 - [Relational access for Node.js](#)
 - [Java and .NET Gateway Reentrancy](#)
 - [Native API for Java and .NET Enhancements](#)
 - [Execute TSQL Code via JDBC](#)
- [New look in the Management Portal](#)
- [SQL Enhancements](#)
 - [Universal Query Cache](#)
- [Interoperability Production Enhancements](#)
 - [New PEX Framework for Coding Production Components in Java and .NET](#)
 - [Port Authority for Monitoring Port Usage in Interoperability Productions](#)
 - [X12 Validation Enhancements](#)
 - [Enhanced DTL Support for X12](#)

- [Import X12 Schemas from XSD Files](#)
- [MQTT Adapters](#)
- [Sharding Enhancements](#)
 - [Simplified Architecture](#)
 - [Flexible Sharded Schema Design and Objects Support](#)
 - [Unified Shard Queue Manager](#)
- [Infrastructure and Cloud Deployment Improvements](#)
- [New Automatic Configuration Customization](#)
- [Analytics Enhancements](#)
 - [Selective Cube Build](#)
 - [PowerBI Connector](#)
 - [Pivot Table Preview](#)
- [Natural Language Processing Enhancements](#)
- [Improved Performance and Scalability of the Database](#)
- [Other Enhancements and Efficiency Improvements](#)

9.1 Continuous Delivery Releases of InterSystems IRIS

InterSystems IRIS 2020.1 is both an extended maintenance release and a continuous delivery release of InterSystems IRIS. There are now two streams of InterSystems IRIS releases:

- **Continuous delivery releases** — These releases provide access to new features and are ideal for developing and deploying applications in the cloud or in local Docker containers.
- **Extended maintenance releases** — These releases are less frequent than the continuous delivery releases but provide the increased stability of maintenance releases. These releases are ideal for large enterprise applications where the ease of getting fixes in maintenance releases is more important than getting early access to new features.

Continuous delivery releases are provided in container format and are available on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, Docker Hub, and the InterSystems WRC download site. You can run a continuous delivery release on any of these cloud platforms or a local system using Docker container. InterSystems does not provide maintenance releases for continuous delivery releases, but instead fixes issues in subsequent continuous delivery releases.

The initial major extended maintenance release is provided on all InterSystems IRIS Supported Platforms, including UNIX, Windows, the cloud platforms, and the Docker container. Following maintenance releases are provided on all server and cloud platforms in the *InterSystems IRIS Supported Platforms*, but are not provided on the Docker container. If you are on a Docker container, you can upgrade to a continuous delivery release.

If your application runs on a non-container platform, you can only use an extended maintenance release for that application but can consider using the continuous delivery releases for:

- Evaluating new features and testing your custom code — this will reduce your upgrade costs when you upgrade to the next extended maintenance major release.
- Using it for new projects that can be deployed in the cloud or in local containers.

In addition to providing fully supported releases, InterSystems provides access to prerelease software for developers who want to get an early look at new features.

9.2 API Management

This release includes two new API Management features:

- [InterSystems API Manager](#)
- [Open API/Swagger Specification-First REST Development](#)

9.2.1 InterSystems API Manager

This release includes the InterSystems API Manager (IAM) enabling you to monitor and control traffic to and from your web-based APIs. The API Manager was released with the maintenance release 2019.1.1 and the continuous delivery release 2019.2 (an early version of 2019.2 did not include the API Manager).

If you are building service-oriented application layers, you are very likely to find the number of APIs you are using quickly rise. The more distributed your environment the more critical it becomes to properly govern and monitor your API traffic. The API Manager enables you easily route all your traffic through a centralized gateway and forward API request to appropriate target nodes. This enables you to:

- Monitor all your API traffic in a central spot.
- Plan, document, and update the list of APIs you are using and the servers that provide them.
- Identify issues before they become critical.
- Control API traffic by throttling throughput, configuring allowed payload sizes, whitelist and blacklist IP addresses and domains, and quickly taking an endpoint into maintenance mode.
- Onboard internal and external developers by providing interactive API documentation through a dedicated and customizable developer portal.
- Secure your API's in a central place.

The API Manager is interoperable, reliant, intuitive, and scalable. You can perform all configuration using a simple web-based user interface, but can also configure the API Manager using API calls, which makes it easy to perform remote deployments,

The API Manager is released in its own container. You can configure the API Manager as a cluster of multiple nodes, but even a single node can handle the load of multiple tens of thousands of requests per second.

For more information, see [InterSystems API Manager](#).

9.2.2 Open API/Swagger Specification-First REST Development

This release enhances the API Management service so that it can generate the ObjectScript code for REST services from OpenAPI 2.0 specifications. This generated code handles the incoming REST call and you only have to write custom code to perform the specific function performed by the service. If you are implementing a service that is already defined in an OpenAPI 2.0 specification, your work is significantly reduced. Even if there is no existing OpenAPI 2.0 specification, it is much easier to create a new specification than to write the custom code required to define the REST API and the specification also provides documentation and aids anyone developing client code for the service. For details see [Creating REST Services](#). (first released in 2019.2)

9.3 In-Place Conversion from Caché and Ensemble

This release of InterSystems IRIS allows you to convert an existing instance of Caché or Ensemble to InterSystems IRIS. The conversion process may require some changes to application code, configuration scripts, and other procedures, but will be relatively easy for the majority of cases. As with any major upgrade, you should thoroughly test your custom code, including any production business services, processes, and operations, in a test environment before deploying to a live production environment.

Before performing an in-place conversion, it is important that you read the *IRIS In-Place Conversion Guide* and the *InterSystems IRIS Adoption Guide* for background information on the differences between Caché or Ensemble and InterSystems IRIS. You can download these documents from the InterSystems Worldwide Response Center [documents distribution page](#).

- Important:** Only the HealthShare Health Connect and InterSystems IRIS for Health products support the HL7 and DICOM features that are available in Ensemble. The InterSystems IRIS product does not support these features. Consequently, if your Ensemble productions use the Ensemble health care features, you should not perform an in-place conversion to InterSystems IRIS, but should rather convert to InterSystems IRIS for Health or HealthShare Health Connect.
- If you are using Ensemble as a general-purpose data platform for health care development or plan to develop health care applications in the future, you should convert to InterSystems IRIS for Health 2020.1 or later.
 - If you are using Ensemble as an integration engine, you should convert to HealthShare Health Connect 2020.1 or later.

For details, see the *InterSystems IRIS for Health Release Notes* or the *HealthShare Health Connect Release Notes* and the *IRIS In-Place Conversion Guide* and the *InterSystems IRIS Adoption Guide*.

9.4 InterSystems Reports

InterSystems Reports can be used with InterSystems IRIS and InterSystems IRIS for Health. InterSystems Reports is a repackaging of Logi Report (formerly named JReport®), a product of Logi Analytics®. We plan to provide documentation on how to use InterSystems Reports with our database products, but for customers who want to start using InterSystems Reports now, you can start with the documentation at <https://www.jinfont.com/documentation/>. If you have questions about using InterSystems Reports, contact the [Worldwide Response Center](#) (WRC).

9.5 Client Language Enhancements

9.5.1 InterSystems IRIS Native API for Python

This release introduces the Native API for Python, which is a lightweight Python interface to the native multidimensional storage data structures that underlie the InterSystems IRIS object and SQL interfaces. The Native API allows you to implement your own data structures by providing direct access to global arrays, the tree-based sparse arrays that form the basis of the multidimensional storage model. By providing direct access to global arrays, the Native API for Python allows you to define very efficient storage structures in Python and, consequently, very efficient applications. For details see [Using the Native API for Python](#) and [Native API for Python Reference](#). (first released in 2019.2)

9.5.2 InterSystems IRIS Native API for Node.js

This release introduces the Native API for Node.js, which is a lightweight Node.js interface to the native multidimensional storage data structures that underlie the InterSystems IRIS® object and SQL interfaces. The Native API allows you to implement your own data structures by providing direct access to global arrays, the tree-based sparse arrays that form the basis of the multidimensional storage model. By providing direct access to global arrays, the Native API for Node.js allows you to define very efficient storage structures in Node.js and, consequently, very efficient applications. For details see [Using the Native API for Node.js](#) and [Native API for Node.js Reference](#). (first released in 2019.2)

9.5.3 Relational access for Node.js

This release provides ODBC access to InterSystems IRIS databases to Node.js developers. (first released in 2019.2)

9.5.4 Java and .NET Gateway Reentrancy

This release allows you to use reentry on the Java and .NET Gateway with forward and reverse proxy. Forward and reverse proxy allow freer use of objects. Proxy objects, or object originally created in InterSystems IRIS to represent objects in Java and .NET, will be created in more situations involving method invocation. You can create objects without concern for the proxy objects. (first released in 2019.3)

9.5.5 Native API for Java and .NET Enhancements

The IRIS Native API, which allows you to access InterSystems IRIS data using globals, has been extended to include \$LIST and pass by reference:

- \$LIST allows you to easily iterate through data structures without needing to parse them in detail. This supports scenarios such as developing Java and .NET applications that access existing global structures from existing applications, as well as simplifying development and speeding performance. (first released in 2019.4)
- Passing parameters by reference allows you to create applications that use less memory, using cleaner code, and have improved performance. Passing by reference provides an effective mechanism for two-way communication between the referencing routine and the function. Any change the function makes to a variable in its formal list is also made to the corresponding by-reference variable in the actual list.

9.5.6 Execute TSQL Code via JDBC

This release provides the ability to execute Transact-SQL (TSQL) code directly from JDBC.

9.6 New Look in the Management Portal

This release represents the beginning of a new, more modern look for the Management Portal. In this first phase, the menus and buttons have a new look but the functionality is unchanged. This new implementation provides the basis for future streamlining and improvements to the user interface. (first released in 2019.2)

9.7 SQL Enhancements

As with every release, InterSystems IRIS includes a number of enhancements to its SQL engine, based on advancements in the underlying software and continuous benchmarking against industry-standard and customer workloads. Customers are likely to observe measurable increases in query throughput for high-load scenarios compared to the 2019.1 release and are encouraged to share their experiences with InterSystems in case there is an opportunity for extending our benchmarking to include specific new use cases.

Note that when you upgrade to a new major version, existing Query Plans are automatically frozen. This ensures that a major software upgrade will never degrade the performance of an existing query. For performance-critical queries, you should test if you can achieve improved performance.

This release includes the following enhancements:

- [Universal Query Cache](#)
- Improvements to our parallelization engine that enable more types of queries and DML to be parallelized (automatically) and make more efficient use of CPU capacity. (first released 2019.4)
- Sharded queries can now use implicit joins using `->` syntax. (first released 2019.4)
- Queries issued from the SQL explorer page in the System Management Portal are executed in the background. While this enables query cancellation and avoids web request timeouts, it also means certain legacy stored procedures that depend on foreground execution and write to the current device may no longer display this logging information properly. (first released 2019.3)

9.7.1 Universal Query Cache

This release introduces a Universal Query Cache, which enables every query (including embedded and class queries) to be saved as a cached query. Previously, the use of embedded SQL meant application code needed to be recompiled in order to pick up current table statistics or newly available indexes. Now, all query plans are managed in a single cache and can be purged (per query, table or namespace) when appropriate. This significantly improves the ability for applications to adapt to actual data characteristics when deployed to multiple sites.

Also, all query types can now equally take advantage of more efficient data access implemented in generated query code.

9.8 Interoperability Production Enhancements

9.8.1 New PEX Framework for Coding Production Components in Java and .NET

This release includes the Production EXtension (PEX) framework that provides you with a choice of implementation languages when you are developing interoperability productions. In this release you can use Java and .NET to develop business services, processes, and operations and, also, inbound and outbound adapters. In previous releases, you could only code business services and operations, could only code in Java, and had to use a special code generator wizards in the Management Portal. The PEX framework provides the flexible plumbing that connects your Java and .NET code to the interoperability production components. You can connect your Java and .NET code using PEX with minimal or no ObjectScript coding. The PEX package includes the following classes:

- `EnLib.PEX.BusinessService`

- EnsLib.PEX.BusinessProcess
- EnsLib.PEX.BusinessOperation
- EnsLib.PEX.InboundAdapter
- EnsLib.PEX.OutboundAdapter
- EnsLib.PEX.Message

For details, see PEX: Developing Productions with Java and .NET.

9.8.2 Port Authority for Monitoring Port Usage in Interoperability Productions

The Port Authority utility allows you to monitor how ports are used in your interchange systems. The Port Authority examines the business services and business operations in multiple productions and instances to determine which ports are being used on each system. You can determine which ports are free for new services and operations and reserve ports for specific uses. For details, see Managing Port Usage.(first released in 2019.3)

9.8.3 X12 Validation Enhancements

This release provides two kinds of enhanced X12 validation:

- SNIP levels 1 and 2 validation — validates the X12 message according to the standards developed by the Workgroup for Electronic Data Exchange (WEDI) Strategic National Implementation Process (SNIP).
- X12 element validation — (first released in 2019.1.1 and 2019.2)

In previous releases, you could not use SNIP validation and could only validate the overall segment structure. There was no mechanism to validate the contents of the segment.

SNIP allow you to validate that:

- SNIP level 1 — segments are valid , segment order is valid, element attributes are valid, numeric data elements have numeric values, and message conforms to X12 rules.
- SNIP level 2 — meets HIPAA requirements, such as presence of required elements, non-use of elements marked as not used, and values conforming to the code tables.

X12 element validation enables you to validate that:

- Required fields are present and that all fields are allowed by the schema.
- Number of fields within a segment and whether they are repeated as allowed by the schema.
- Data types for fields and components are correct.
- Field values conform to the code tables specified.
- Field and components conform to length restrictions.

For details, see X12 Validation.

9.8.4 Enhanced DTL Support for X12

In this release you can define data transformations for an entire X12 batch including schemas for the interchange envelope, functional groups, and transaction sets. This allows you to process X12 batch messages using a single data transformation without having to use subtransformations. This release also improves the user interfaces and also provides convenience functions that make it easier to handle repeating elements. For details, see Creating an X12 Data Transformation.

9.8.5 Import X12 Schemas from XSD Files

In previous versions, you could only import X12 schemas from SEF files or InterSystems proprietary XML format. In this release, you can also import X12 schemas from the newer XSD schema files. For details, see [Loading X12 Schemas](#).

9.8.6 MQTT Adapters

This release includes MQTT adapters that support Message Queuing Telemetry Transport (MQTT) protocol, which is often used in Internet of Things (IoT) applications. For details, see [Using MQTT Adapters in Productions](#).

9.9 Sharding Enhancements

9.9.1 Simplified Architecture

This release introduces a simple and straightforward blueprint for sharded clusters, the node-level architecture, which can be configured through the new `%SYSTEM.Cluster` API. In this cluster architecture, we've implemented some best practices on laying out the different foundational elements introduced with sharding in the initial InterSystems IRIS 2018.1 release that will make it significantly easier to deploy and expand your cluster. As the node-level architecture is essentially a smart way of leveraging the existing infrastructure, it is fully transparent to application code and does not require any changes to existing deployments. For details, see [Elements of Sharding](#) and [Sharding APIs](#) in the *Scalability Guide* and see the `%SYSTEM.Cluster` class documentation in the *InterSystems Class Reference*. (first released in 2019.2)

9.9.2 Flexible Sharded Schema Design and Objects Support

This release introduces further improvements to how sharding supports designing your application's schema in the following ways:

- You can now coshard any two sharded tables. Before, only tables with a common user-defined shard key could be cosharded (that is, explicitly defining the shard key for an Order and OrderLine table to be OrderID, the field on which they are joined). With this release, you can use `COSHARD WITH` syntax in DDL or the `CoShardWith` index keyword to coshard a new table with an existing table that has a system-assigned shard key. This significantly increases the flexibility for designing your application's schema, preserving the operational benefits of using system-assigned shard keys. For details, see [Create the Tables](#) in the *Scalability Guide* and [Defining a Sharded Table](#).
- Where previously sharded schema design could only happen through DDL, you can now mark a persistent class (table) as sharded through its class definition, using the new "Sharded" class keyword. The class compiler has been extended to warn against using class definition features incompatible with sharding, such as customized storage definitions, at compile time. For details, see [Defining a Sharded Table by Creating a Persistent Class](#).
- You can use the object paradigm to interact with sharded classes. This means you can create new or open existing sharded class instances with `%New()` or `%OpenId()` methods and save them with `%Save()`. The sharding infrastructure will make sure new instances are appropriately distributed across the cluster, fully transparent to the application. Note that this object code will still be executed on the machine your client is connected to.

These sharding enhancements were first released in 2019.2.

9.9.3 Unified Shard Queue Manager

In this release, the Unified Shard Queue Manager improves sharding efficiency when a sharded cluster is being queried by large numbers of clients by queuing up sharded query work rather than spawn individual processes for each shard-local query and potentially flood the system. An efficient algorithm then ensures an appropriate number of worker processes is used to process work from the queue, based on available hardware and system load.

9.10 Infrastructure and Cloud Deployment Improvements

This release contains improvements to the infrastructure and cloud deployment, including the following:

- Tencent Cloud Support — InterSystems Cloud Manager (ICM) now provides end-to-end cloud provisioning and deployment for applications based on InterSystems IRIS and running on Tencent Cloud. (first released in 2019.4)
- Support for Docker named volumes in addition to bind mounts. (first released in 2019.4)
- InterSystems Cloud Manager (ICM) support for elastic scaling — Existing configurations can now be scaled, that is, reprovisioned and redeployed with more or fewer nodes. For details, see [Reprovisioning the Infrastructure](#) and [Redeploying Services](#) in the *InterSystems Cloud Manager Guide*. (first released in 2019.2; scale out DATA nodes and scale in/out COMPUTE nodes in node-level architecture first released in 2019.4)
- Improved user experience when packaging your own container. (first released in 2019.3)
- InterSystems Cloud Manager (ICM) support for node-level sharding. (first released in 2019.3)
- Containers use non-root default user, which is a container best practice and improves security. (first released in 2019.3)
- ICM support for creating and deploying on a private network, in which bastion servers connect your private network to the public network and provide improved security protection from denial of service attacks. (first released in 2019.3)
- Support for service discovery with secure RPC communication. (first released in 2019.3)
- ICM support for multi-region deployments, which can provide high availability even if an entire region stops functioning. (first released in 2019.3)
- Ability to upgrade ICM and retain knowledge of deployed systems. (first released in 2019.3)
- Containerless Mode — The following were previously restricted but can now be performed: deploying sharded configurations on Google Cloud Platform using containerless mode and deploying the Web Gateway on Ubuntu or SUSE nodes using containerless mode. (first released in 2019.2)

9.11 New Automatic Configuration Customization

A new InterSystems IRIS configuration feature enables customization of the configuration parameter file (CPF) of an InterSystems IRIS instance prior to startup, upon which the custom configuration is automatically implemented. This feature greatly simplifies automation and supports the use of configuration management tools such as Kubernetes with InterSystems IRIS, and is also included in ICM in this version. Automatic configuration customization is an important new capability that will be expanded in future versions. (first released in 2019.4)

9.12 Analytics Enhancements

This release contains the following analytics enhancements:

- [Selective Cube Build](#)
- [PowerBI Connector](#)
- [Pivot Table Preview](#)

9.12.1 Selective Cube Build

This release provides Selective Cube Build, a feature of InterSystems IRIS Business Intelligence, that allows you to select the measures and dimensions to be built individually. You can make changes and selectively rebuild without taking the full cube out of service. The user interface also automatically flags the dimensions or measures that have been added or changed so that you know what to rebuild.

9.12.2 PowerBI Connector

InterSystems customers can now use Microsoft Power BI to access tabular and cube data stored on InterSystems IRIS. This allows combining the data visualization capabilities offered by Power BI with the high-performance data management and querying capabilities offered by InterSystems IRIS. While the connector leverages ODBC, it will also allow customers to access InterSystems IRIS BI cubes directly from Power BI when connecting to InterSystems 2019.2 or above. The connector ships as part of Power BI starting with its April 2019 release. For details, see [InterSystems IRIS Connector for Power BI](#).

9.12.3 Pivot Table Preview

This release contains the Analytics Pivot Table Preview, a new mode for the Analyzer that presents a representative pivot table based on a truncated data set. This will allow previewing a pivot table much more quickly than analyzing the complete result set. A **Show All** button is also presented when in Preview mode to indicate that the result set is not complete. Selecting the **Show All** button automatically turns off Preview mode. (first released in 2019.2)

9.13 Natural Language Processing Enhancements

This release contains the following natural language processing enhancements.

- InterSystems IRIS Natural Language Processing (NLP) will now extract value and unit information for measurement attributes in the English language model. This information can be visualized through highlighting or retrieved through the query and REST APIs. (first released in 2019.3)
- This release introduces support for the Czech language with InterSystems IRIS Natural Language Processing. The embedded NLP engine will now also identify concepts and their context for natural language text written in Czech as it does for the 10 other languages supported previously. (first released in 2019.2)

9.14 Improved Performance and Scalability of the Database

This release has significant optimizations in the database engine. This is especially important for very large systems and significantly increases the ability to scale systems to handle very heavy loads.

One of the efficiency changes in this release improves efficiency when traversing globals. For a database block in memory that is accessed frequently but not modified often, the system may automatically build an optimization structure, called a node table, to speed up searches for nodes within the block. This speeds up global accesses, particularly when access to nodes are sparsely or randomly distributed, or for patterns that access the nodes in reverse collation order (including reverse \$order / \$query). The memory for this comes from the database cache itself, a small fraction typically less than one percent.

9.15 Other Enhancements and Efficiency Improvements

In each release, InterSystems makes many efficiency improvements and minor enhancements. In this release these improvements include:

- Journal performance enhancements.
- Easier configuration for mirrored environments.
- Support for new versions of Apache Spark version 2.3 and 2.4.
- Support for WebSocket client in addition to the existing WebSocket server support in the Web Gateway.
- Source Control for Productions — source control hooks have been added to allow check-in and check-out of a production as an entity, simplifying change tracking and configuration management. (first released in 2019.4)
- Whitelists to Support Penetration Testing — customers performing their own security penetration testing can reduce or eliminate false positives related to CSP, Zen, and REST. (first released in 2019.4)
- Upgrades .NET support to .NET Core 2.1. (first released in 2019.3)
- Improved efficiency for ODBC database access. (first released in 2019.3)
- Structured logging to improve access to log messages. (first released in 2019.3)
- Improved API to fetch alerts. (first released in 2019.3)
- Compiler now tests for global names that are too long and reports an error. Previously, these global names were silently truncated. (first released in 2019.3)

Maintenance release 2019.1.1, continuous delivery release 2019.3, and subsequent releases include the set of changes identified as JournalingGroup2019, which corrects issues associated with journaling and mirroring. The changes associated with these issues are SML2776, SML2781, SML2782, SML2783, SML2785, JO2990, JO3117, JO3137, JO3140, JO3141, RJF391, RJF392, HYY2362, HYY2364, and HYY2373.

10

New in InterSystems IRIS 2019.1

This page describes the 2019.1 release of InterSystems IRIS®.

This release includes new capabilities and enhancements in the following areas:

- [Extended Maintenance and Continuous Delivery Releases of InterSystems IRIS](#)
- [InterSystems Cloud Manager](#)
- [Client languages](#)
- [Improved scalability and operations for sharded clusters](#)
- [SQL performance](#)
- [Analytics](#)
- [Interoperability](#)
- [System Performance and Capabilities](#)
- [New Features in 2019.1.1 Release:](#)
 - [In-Place Conversion from Caché and Ensemble](#)
 - [InterSystems API Manager](#)
 - [X12 Element Validation in Interoperability Productions](#)

10.1 Extended Maintenance and Continuous Delivery Releases of InterSystems IRIS

InterSystems IRIS 2019.1 is an extended maintenance delivery release of InterSystems IRIS in contrast with InterSystems IRIS 2019.2, which is a continuous delivery release. There are now two streams of InterSystems IRIS releases:

- **Continuous delivery releases** — These releases provide access to new features and are ideal for developing and deploying applications in the cloud or in local Docker containers.
- **Extended maintenance releases** — These releases are less frequent than the continuous delivery releases but provide the increased stability of maintenance releases. These releases are ideal for large enterprise applications where the ease of getting fixes in maintenance releases is more important than getting early access to new features.

Continuous delivery releases are provided in container format and are available on Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, Docker Hub, and the InterSystems WRC download site. You can run a continuous delivery release on any of these cloud platforms or a local system using Docker container. InterSystems does not provide maintenance releases for continuous delivery releases, but instead fixes issues in subsequent continuous delivery releases.

The initial major extended maintenance release is provided on all InterSystems IRIS Supported Platforms, including UNIX, Windows, the cloud platforms, and the Docker container. Following maintenance releases are provided on all server and cloud platforms in the InterSystems IRIS Supported Platforms, but are not provided on the Docker container. If you are on a Docker container, you can upgrade to a continuous delivery release.

If your application runs on a non-container platform, you can only use an extended maintenance release for that application but can consider using the continuous delivery releases for:

- Evaluating new features and testing your custom code — this will reduce your upgrade costs when you upgrade to the next major release.
- Using it for new projects that can be deployed in the cloud or in local containers.

In addition to providing fully-supported releases, InterSystems provides access to preview software for developers who want to get an early look at new features.

10.2 InterSystems Cloud Manager Enhancements

InterSystems Cloud Manager (ICM) provides you with a simple, intuitive way to provision cloud infrastructure and deploy services on it. In this release ICM has the following enhancements:

- *Availability Zone Support* — This enhancement allows you to span multiple zones within a given region with cloud providers that provide this facility. For more information, see [Deploying Across Multiple Zones](#) in the *InterSystems Cloud Manager Guide*.
- *Asynch Mirror Support* — This enhancements lets you configure asynch mirror members. For more information, see [Mirrored Configuration Requirements](#) in the *InterSystems Cloud Manager Guide*.
- *Containerless Support* — This enhancement allows you to use ICM to deploy noncontainerized InterSystems IRIS instances from installation kits. For more information, see [Containerless Deployment](#) in the *InterSystems Cloud Manager Guide*.
- *Service Discovery* — Service Discovery mode gives multiple users in any networked locations management access to a single ICM deployment. For more information, see [Sharing ICM Deployments](#) in the *InterSystems Cloud Manager Guide*.

10.3 Client Languages Enhancements

This release contains the following enhancements and performance boosts to access InterSystems IRIS using client languages:

- *IRIS Native API for .NET* — Provides low-level access to the underlying global storage from .NET applications.
- *Relational access for Python*.
- *Dynamic Java Gateway*.

- *Shared memory support for Java Gateway* — Support for shared memory connections has been extended to Java Gateway. For more information, see [Using Shared Memory Connections in *Using Java with the InterSystems JDBC Driver*](#).
- *Hibernate* — This release is compatible with Hibernate 5.2 or 5.3. For more information, see [Hibernate Support in the *InterSystems Implementation Reference for Java Third Party APIs*](#).
- *Bulk loader in Java* — The bulk loader is a new utility that can be used for massive data transfer from one data source to another.

10.4 Improved Scalability and Operations for Sharded Clusters

An InterSystems IRIS® sharded cluster partitions both data storage and caching across a number of servers, providing flexible, inexpensive performance scaling for queries and data ingestion while maximizing infrastructure value through highly efficient resource utilization. This release provides improved scalability and operations for sharded clusters including the following:

- *Scalability Enhancements* — Support a broader set of scalability scenarios for SQL. Data nodes can now be added to a sharded cluster at all times, irrespective of the database schema and shard keys used. Furthermore, after nodes are added, data can be rebalanced across the available nodes to ensure an even distribution of data and, subsequently, work to improve overall cluster performance. For more information, see [Rebalance Sharded Data Across Additional Shard Data Servers](#) in the *Scalability Guide*.
- *Management Portal Enhancements* — New page to review and configure the sharded cluster's layout.
- *API for Backups* — A new API for coordinating the creation of backups of a sharded cluster's data. For details, see [Coordinated Backup and Restore of Sharded Clusters](#) in the *Scalability Guide*.
- *Bulk Loader Support* — New Bulk Loader client utility also optimizes ingestion of large datasets into a sharded cluster.

10.5 SQL Enhancements

This release contains significant enhancements to SQL usability and performance including the following:

- *Auto-parallel queries* — This release provides improved efficiency by automatically using parallel queries where appropriate, significantly improving throughput for machines with many CPU cores. For more information, see [System-Wide Parallel Query Processing](#) in the *InterSystems SQL Optimization Guide*.
- *SQL Usability Enhancements* — The new TUNE TABLE command tunes a table based on the data currently in the table and is available from the SQL shell. For more information, see [TUNE TABLE](#) in the *InterSystems SQL Reference*.

In addition, this release includes several other SQL shell enhancements, such as the ability to browse the schemas, tables, and views defined in or accessible from the current namespace. For more information, see [Using the SQL Shell Interface](#) in *Using InterSystems SQL*.

ShowPlan function and EXPLAIN command — Now display sub-plans for composite plans such as parallel and/or sharded queries. For details, see [Show Plan](#).

- *Comment Options* — This release supports Comment Options specified in the SQL code that cause the Optimizer to override a system-wide compile option for that query. For more information, see [Comment Options](#) in the *InterSystems SQL Optimization Guide*.
- *General performance enhancements* — With each release, InterSystems includes various enhancements to its SQL engine that are 100% transparent from the application's perspective. For 2019.1, an especially broad array of improvements went into the query optimizer and subsequent code generation that define your SQL query performance. Combined with the now automated parallel query execution, InterSystems IRIS SQL users should see a noticeable to significant improvement in throughput, depending on their query set.

10.6 Analytics Enhancements

This release contains the following analytics enhancements:

- *Partial Date type in Business Intelligence* — Partial dates allow you to specify incomplete dates, such as just a year or a year and a month. For more information, see Partial Dates in *Defining Models for InterSystems IRIS Business Intelligence*.
- *%SQLRESTRICT dimension for cubes* — This new cube dimension allows run-time restrictions on an MDX query via a SQL SELECT statement or WHERE clause. For more information, see %FILTER Clause in the *InterSystems MDX Reference*.
- *Pivot Table Headers* — When a large pivot table requires scrolling to see all columns and/or rows, the header columns and rows remain in place so that the labels continue to be visible as you scroll.
- *Work Queue Manager replaces Agents* — The Work Queue Manager is used in InterSystems IRIS Business Intelligence to distribute work to multiple concurrent processes in order to manage performance. The Work Queue Manager is a standard component of InterSystems IRIS. For more information on the Work Queue Manager, see [Using the Work Queue Manager](#).

10.7 Interoperability Enhancements

This release contains new interoperability capabilities that speed configuring and troubleshooting of productions. These include the following:

- *Interface Maps* — Users can search for and view all the routes that a message can take within a production. See Viewing Interface Maps in *Monitoring Productions* for details.
- *Search for Interface References* — Users can search to find where production components are referenced by other production components. See Finding Interface References in *Monitoring Productions* for details.
- *Data Transformation Testing Enhancements* — When testing data transformations, users can unit test record maps in the Data Transformation Editor by allowing raw text input in the Test Transform dialog and can enter values for aux, context, and process system objects as if the data transformation was invoked with these objects instantiated. See Using the Transformation Testing Page in *Developing DTL Transformations* for details.
- *DTL Editor Enhancements* — The usability of the Data Transformation Editor has been enhanced with the addition of switch/case actions, the ability to group actions together, the ability to collapse/expand groups, and the ability to add comments to the data transformation. See Adding a Switch Action, Working with Groups of Actions, and Adding a Comment Action in *Developing DTL Transformations* for details.

- *Unit Testing of Routing Rules* — This enhancement introduces a unit testing capability to the Rule Editor, whereby a user can feed a message through a business rule and view rule execution results without having to run the message through the entire production. See *Testing Routing Rules* in *Developing Business Rules* for details.
- *Download Multiple Messages to Local Computer* — Users can select multiple messages in the Message Viewer and download them to their local computer. See *Exporting Messages* in *Monitoring Productions* for details.
- *Download Event Logs to Local Computer* — Users can download event logs to their local computer. Previously event logs could only be downloaded to the server. See *Introduction to the Event Log Page* in *Monitoring Productions* for details.
- *Rule Editor Enhancements* — The usability of the Rule Editor has been enhanced with the ability to add comments to a business rule and the ability to view and edit a Data Transformation (DTL) directly from the Rule Editor when the given DTL is used in a business rule. See *Selecting the Transformation and Target of a Send Action* in *Developing Business Rules* for more details about opening the DTL from the Rule Editor.
- *Queue Wait Alert Modification* — The Queue Wait Alert setting now specifies the length of time that a message can wait in the business host's queue or be the active message before an alert is triggered. Previously, the setting only applied to messages in the queue, not the active message. See the Queue Wait Alert setting in *Configuring Productions* for details.
- *Restrict Access to System Default Settings* — Administrators can control whether users can create, edit, or delete system default settings. See *Security for System Default Settings* in *Managing Productions* for details.
- *Export Productions to Local Computer* — Users can export productions to their local computer. Previously productions could only be exported to the server. See *Exporting a Production* in *Configuring Productions* for details.
- *Deploy Productions from Local Computer* — Users can deploy productions from their local computer. Previously productions could only be deployed from the server. See *Deploying a Production on a Target System* in *Developing Productions* for details.
- *Enhanced Navigation from Production Configuration Page* — Links have been added to the tabs of the Production Configuration window to quickly open related items in a separate window. On the **Queue** tab, clicking the message ID opens a window to display the visual trace for the message. On the **Messages** tab, clicking the Session ID opens a window to display the visual trace of the message. On the **Jobs** tab, clicking the message ID opens a window to display the visual trace for the message, and clicking the Job ID opens a window to display the Process Details for the job.
- *Business Host Wizard Enhancements* — To enhance user productivity, additional options have been added to the wizards used to create business hosts. Users can use the business host wizards to automatically assign system default values when fields are left blank and to define a package prefix for auto-generated routing rules. See *Wizard Options* in *Configuring Productions* for details on new options in the business host wizards.

10.8 System Performance and Capabilities

This release contains the following system security, performance, and efficiency enhancements:

- Substantial scalability and performance improvements, particularly for large-scale Non-Uniform Memory Access (NUMA) systems. This includes changes to improve scalability in statistics tracking and global buffer management, performance improvements in use of subscript level mapping, and more effective optimizations to avoid traversing global pointer blocks. To enable these improvements there are minor changes to memory utilization and system statistics described in the [Incompatibility History](#) document.

These enhancements increase the amount of memory allocated for global buffer metadata by 64 bytes per buffer on Intel systems and by 128 bytes per buffer on IBM Power systems. For example, with 8K buffer sizes, the shared

memory allocated for a global buffer increases by 0.75% on Intel systems and by 1.5% on IBM Power systems. These enhancements also cause minor changes in statistics displayed by utilities and the Management Portal.

- *Key Management Interoperability Protocol (KMIP)* — In this release, InterSystems IRIS® can be a client to an enterprise key management server and use the Key Management Interoperability Protocol (KMIP) to store and retrieve keys on the server. KMIP, an OASIS standard, gives you the power of centralized key management. You can use keys from a KMIP server to encrypt data at rest — for both database encryption and data-element encryption. They are available for all the same activities as keys from key files, such as journal file encryption. InterSystems IRIS also allows you to copy keys from the KMIP server to local files, so that there can be local backup copies. For more information, see [Managing Keys with the Key Management Interoperability Protocol \(KMIP\)](#) in the *Encryption Guide*.

Note: InterSystems IRIS does not support KMIP on the macOS platform.

- *DataMove* — Enables you to move data from one database to another, revise the mappings to access the data, and delete the data from its old location.
- *Support for large JSON strings.*
- *IRIS Studio support for other InterSystems products.*
- *Support for Microsoft Integrated Windows Authentication for HTTP Connections (SPNEGO)* — This new enhancement allows %Net.HttpRequest to use windows based authentication over HTTP 1.1 to establish a connection to a secure server. Users can provide credentials or, if no credentials are provided, the system will try to authenticate using the current logged in context. Client may initiate a connection to the server with an "Authorization" header or try to establish a connection without that header and process the 401 status code with its associated WWW-Authenticate header and then respond with the appropriate authentication mechanism. The supported authentication schemes are Negotiate (Kerberos & NTLM), NTLM, and Basic. For more information see [Providing Authentication](#) in *Using Internet Utilities*.
- *Journaling efficiency improvements.*
- *Async I/O efficiency improvements.*

10.9 New Features in 2019.1.1 Release

This section describes new features that are only available in the InterSystems IRIS 2019.1.1 maintenance release and later maintenance releases. If you are running release 2019.1.0, you do not have these features.

10.9.1 In-Place Conversion from Caché and Ensemble

This release of InterSystems IRIS allows you to convert an existing instance of Caché or Ensemble to InterSystems IRIS. The conversion process may require some changes to application code, configuration scripts, and other procedures, but will be relatively easy for the majority of cases. As with any major upgrade, you should thoroughly test your custom code, including any production business services, processes, and operations, in a test environment before deploying to a live production environment.

Before performing an in-place conversion, it is important that you read the *IRIS In-Place Conversion Guide* and the *InterSystems IRIS Adoption Guide* for background information on the differences between Caché or Ensemble and InterSystems IRIS. You can download these documents from the InterSystems Worldwide Response Center [documents distribution page](#).

Important: InterSystems IRIS does not support the HL7 and DICOM features and the X12 health schemas that are available in Ensemble. This support is included in the HealthShare Health Connect and InterSystems IRIS for Health products. Consequently, if your Ensemble productions use HL7, DICOM, or the X12 health schemas, you should not perform an in-place conversion to InterSystems IRIS. If you are using Ensemble as an integration engine, you should first upgrade to HealthShare Health Connect 15.03x built on the Caché/Ensemble platform and then perform an in-place conversion to Health Connect 2019.1 on the InterSystems IRIS platform. If you are using Ensemble as a general-purpose data platform for health care, you should wait for the release of InterSystems IRIS for Health that supports in-place conversion. In order to ensure that the in-place conversion is relatively easy, InterSystems performs significant testing, including test conversions at customer sites. We are deferring support of in-place conversion to InterSystems IRIS for Health until after we have completed this testing.

10.9.2 InterSystems API Manager

This release includes the InterSystems API Manager (IAM) enabling you to monitor and control traffic to and from your web-based APIs.

If you are building service-oriented application layers, you are very likely to find the number of APIs you are using quickly rise. The more distributed your environment the more critical it becomes to properly govern and monitor your API traffic. The API Manager enables you easily route all your traffic through a centralized gateway and forward API request to appropriate target nodes. This enables you to:

- Monitor all your API traffic in a central spot.
- Plan, document, and update the list of APIs you are using and the servers that provide them.
- Identify issues before they become critical.
- Control API traffic by throttling throughput, configuring allowed payload sizes, whitelist and blacklist IP addresses and domains, and quickly taking an endpoint into maintenance mode.
- Onboard internal and external developers by providing interactive API documentation through a dedicated and customizable developer portal.
- Secure your API's in a central place.

The API Manager is interoperable, reliant, intuitive, and scalable. You can perform all configuration using a simple web-based user interface, but can also configure the API Manager using API calls, which makes it easy to perform remote deployments,

The API Manager is released in its own container. You can configure the API Manager as a cluster of multiple nodes, but even a single node can handle the load of multiple tens of thousands of requests per second.

For more information, see InterSystems API Manager.

Note: The API Manager is only available in a Docker container distribution. You can use it with an InterSystems IRIS system that is installed on any of the InterSystems IRIS Supported Platforms, including UNIX, Windows, the cloud platforms, and the Docker container.

10.9.3 X12 Element Validation in Interoperability Productions

This release provides enhanced X12 validation. In previous releases, you could only validate that the required segments are in the correct order and that there are no segments present that are prohibited, but there was no mechanism to validate the contents of the segment. This enhancements enables you to validate that:

- Required fields are present and that all fields are allowed by the schema.

- Number of fields within a segment and how they are repeated are allowed by the schema.
- Datatypes for fields and components are correct.
- Field values conform to the code tables specified.
- Field and components conform to length restrictions.

For details, see Validation in *Routing X12 Documents in Productions*.

11

Upgrading to This Release

This section contains specific instructions applicable to upgrading to this major release. These instructions supplement [Upgrading from an Earlier Version](#) in the *Installation Guide*.

If you are upgrading to this release from a version prior to the last major release, you should also refer to the [upgrade instructions for the last major release](#), as well as any other major release which your upgrade skips over.

11.1 Upgrading Containers

Because a containerized application is isolated from the host environment, it does not write persistent data; whatever it writes inside the container is lost when the container is removed and replaced by a new container. Therefore, an important aspect of containerized application deployment is arranging for data to be stored outside of the container and made available to other and future containers.

The durable %SYS feature enables persistent storage of instance-specific data — such as user definitions, audit records, and the log, journal, and WIJ files — when InterSystems IRIS is run in a container, allowing a single instance to run sequentially in multiple containers over time. For example, if you run an InterSystems IRIS container using durable %SYS, you can upgrade the instance by stopping the original container and running a new one that uses the instance-specific data created by the old one. For information about upgrading, see [Upgrading InterSystems IRIS Containers](#); for detailed information on durable %SYS, see [Durable %SYS for Persistent Instance Data](#).

Important: In this release, the distribution container has a nonroot default user. This improves the security of your container. If you are using a durable %SYS from a 2019.2 or earlier instance with this release, you need to change some file ownerships in the host's durable directory before running InterSystems IRIS 2023.3. Please contact your InterSystems sales engineer or the InterSystems [Worldwide Response Center](#) for instructions on changing the file ownerships. If you do not make these changes, the container will encounter an error starting InterSystems IRIS.

12

Known Issues and Notes

This topic describes known issues in InterSystems IRIS® 2023.3.

12.1 Unable to Open Documentation from the Launcher

If you are using a fresh installation of 2023.3, are upgrading from 2023.2, or are using Microsoft IIS as your web server, you cannot open Documentation through the cube by default. To fix this issue, see the documentation for how to [manually redirect documentation links with Microsoft IIS](#).

Additionally, this issue affects Community Edition installations of 2023.3 with the Apache Private Web server. To fix it, see the documentation on how to [manually redirect documentation links on Apache](#).

12.2 AmazonCloudWatchLogs Classes are Deprecated

Two classes, `EnsLib.AmazonCloudWatchLogs.BusinessOperation` and `EnsLib.AmazonCloudWatchLogs.OutboundAdapter`, are discontinued and will be removed in future releases.

13

Deprecated and Discontinued Features

This topic lists features deprecated features in InterSystems IRIS®; it also lists any features that are discontinued or removed.

13.1 About Deprecated and Discontinued Technologies and Features

From time to time, InterSystems stops further development of a technology when newer and better options are available. *However, product support for these capabilities continues* in the same way that it does for products beyond our [Minimum Supported Version](#).

Deprecated designates a feature or technology that InterSystems no longer actively develops and for which better options exist. Deprecated items should not be used for new development. The *deprecated* designation indicates that customers should plan to eliminate use of the feature or technology. InterSystems maintains the staff expertise to support *deprecated* product capabilities.

Discontinued designates a feature or technology that is no longer viable for use, even in existing applications. InterSystems feels that continued use of such technology is a risk for our customers. The reasons for this include but are not limited to:

- Usage has declined to a small number of customers.
- The feature has become incompatible with current technologies or security practice.
- Incompatibilities between the feature or technology and our current product implementation make application maintenance prohibitive.
- The feature or technology depends on discontinued content from a third party.

13.2 InterSystems Cloud Manager

As of version 2023.3, InterSystems Cloud Manager (ICM) has been deprecated. InterSystems will continue to support existing customers using this technology, but it is no longer recommended for new deployments.

Customers in the cloud who are interested in deploying and managing an InterSystems IRIS deployment with many systems are encouraged to consider using a Kubernetes platform and the [InterSystems Kubernetes Operator](#), which greatly simplifies the deployment of an InterSystems IRIS sharded cluster, distributed cache cluster, or standalone instance on Kubernetes.

13.3 InterSystems IRIS Natural Language Processing

As of version 2023.3, InterSystems IRIS NLP, also known as iKnow, has been deprecated. All existing functionality continues to be available as part of InterSystems IRIS and InterSystems will continue to support existing customer deployments. However, InterSystems is no longer actively developing the capability and does not recommend starting new development projects with the technology. The [open-source version](#) of the core engine, packages as a Python module, can be used independently of InterSystems IRIS and will continue to be available. Customers looking for guidance on their current use of the technology are invited to reach out to their account teams.

InterSystems IRIS SQL Search, also known as iFind, is only partly affected. Only the Semantic and Analytic index types make use of the iKnow engine and therefore are deprecated. All other functionality and index types are not affected by this and continue to be the recommended choice for applications requiring a flexible and high-performance full text search capability.

13.4 Private Web Server (PWS)

As of version 2023.2, local installations of InterSystems IRIS will no longer install the private web server and InterSystems IRIS containers will no longer provide the private web server.

InterSystems strongly recommends that you install an external web server and configure it according to your needs, instead of using the private web server. Note that this is not a change in existing InterSystems recommendations and there is no change in the list of supported web servers.

Effective with the first EM release in 2026, the private web server will be discontinued; at that point, upgrades of existing InterSystems IRIS instances will remove the private web server.

For more details on this change, see the [Private Web Server \(PWS\)](#) topic in the 2023.2 new features listing.

13.5 Studio

As of version 2023.2, Studio has been deprecated. Version 2024.1 will be the last release that includes Studio as part of the Windows distribution. After that release, Studio will continue to be available as a standalone tool on the [WRC component download](#) page.

InterSystems recommends [Visual Studio \(VS\) Code](#) with the [InterSystems ObjectScript Extension Pack](#), which provides a fully-featured integrated development environment for InterSystems applications. The InterSystems ObjectScript Expansion Pack includes the following extensions: [InterSystems ObjectScript](#), [InterSystems Language Server](#), and [Server Manager](#) extensions. See the documentation for the [InterSystems VS Code extensions](#) for information on [migrating from IRIS Studio to VS Code](#).

13.6 Unstructured Information Management Architecture (UIMA) Integration

As of version 2023.1, the capability to implement the Unstructured Information Management Architecture (UIMA) has been deprecated. UIMA integration capability is still included in this release, but it will be removed in a future release.

As of version 2023.1, JAR files required to implement UIMA have been [removed from releases](#). To obtain these files, see [Build Updated Apache UIMA JAR files](#).

13.7 Spark Connector

As of version 2022.1, the Spark connector has been deprecated. As of version 2022.3, it has been discontinued. As of version 2023.1, it has been removed from releases.

Spark now has a built-in JDBC connector. If you use the Spark connector, you should modify your code to use the Spark JDBC connector.

13.8 Atelier

As of version 2021.1, Atelier has been deprecated.

InterSystems recommends [Visual Studio \(VS\) Code](#) with the [InterSystems ObjectScript Extension Pack](#), which provides a fully-featured integrated development environment for InterSystems applications. The InterSystems ObjectScript Expansion Pack includes the following extensions: [InterSystems ObjectScript](#), [InterSystems Language Server](#), and [Server Manager](#) extensions. See the documentation for the [InterSystems VS Code extensions](#) for information on [migrating from IRIS Studio to VS Code](#).

13.9 Zen

As of version 2018.1, Zen has been deprecated. It is still included with this release, but it will be removed from a future release.

13.10 Zen Reports

As of version 2018.1, Zen Reports has been deprecated. It is still included with this release, but it will be removed from a future release.

