



# HealthShare Health Connect Release Notes

Version 2024.1  
2024-05-02

*HealthShare Health Connect Release Notes*  
HealthShare Version 2024.1 2024-05-02  
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](mailto:support@InterSystems.com)

# Table of Contents

<b>1 New in HealthShare Health Connect 2024.1</b> .....	<b>1</b>
1.1 Release Information for 2024.1 .....	1
1.2 Private Web Server (PWS) .....	1
1.3 Healthcare Interoperability .....	2
1.3.1 SMART Scopes .....	2
1.3.2 JsonAdvSQL .....	2
1.3.3 New FHIR Object Model .....	2
1.3.4 Interoperability Usage Metrics .....	3
1.3.5 HL7 Production Validator .....	3
1.4 Enhancing Developer Experience .....	3
1.4.1 Using vectors in ObjectScript .....	3
1.4.2 Support for JSON_TABLE .....	4
1.4.3 SQL Development .....	4
1.5 Enhancing Analytics and AI .....	5
1.5.1 Vector Search .....	5
1.5.2 Business Intelligence Improvements .....	5
1.6 Enhancing Speed, Scale and Security .....	6
1.6.1 Multi-Volume Databases .....	6
1.6.2 Fast Online Backup .....	6
1.7 Platform Updates .....	6
1.7.1 Minimum Supported CPU Models .....	6
1.7.2 Operating System Updates .....	7
<b>2 New in HealthShare Health Connect 2023.1</b> .....	<b>9</b>
2.1 Release Information for 2023.1 .....	9
2.2 Enhancing Analytics and AI .....	9
2.2.1 Columnar Storage .....	9
2.3 Enhancing Speed, Scale and Security .....	10
2.3.1 Memory Settings .....	10
2.3.2 Platform Scalability .....	10
2.4 Platform Updates .....	10
<b>3 New in HealthShare Health Connect 2022.3</b> .....	<b>11</b>
3.1 Enhancing Developer Experience .....	11
3.1.1 SQL Development .....	11
3.2 Enhancing Analytics and AI .....	11
3.3 Enhancing Speed, Scale and Security .....	12
3.4 Platform Updates .....	12
<b>4 New in HealthShare Health Connect 2022.2</b> .....	<b>13</b>
4.1 Enhancing Analytics and AI .....	13
4.1.1 Columnar Storage (Experimental Feature) .....	13
4.2 Enhancing Speed, Scale, and Security .....	13
4.2.1 SQL Enhancements .....	13
4.3 Platform Updates .....	14
<b>5 New in Health Connect 2022.1</b> .....	<b>17</b>
5.1 Release Information for 2022.1 .....	17
5.2 Enhancing Interoperability .....	17

5.2.1 Kafka Messaging Support .....	17
5.2.2 Embedded Python .....	17
5.2.3 Interoperability Productions in Python .....	18
5.2.4 Visual Studio Code ObjectScript Extension Pack Updates .....	18
5.2.5 Make SQL Queries With Minimal Code in Interoperability Productions .....	19
5.3 Enhancing Cloud and Operations .....	19
5.3.1 Cloud Connectors .....	19
5.3.2 IKO Enhancements .....	19
5.4 Enhancing Speed, Scale, and Security .....	19
5.4.1 Adaptive SQL Optimizer .....	19
5.4.2 Saving on Storage .....	20
5.4.3 TLS 1.3 Support (OpenSSL 1.1.1) .....	20
5.4.4 New ^TRACE Utility .....	20
5.5 Other Enhancements and Efficiency Improvements .....	21
<b>6 New in Health Connect 2021.2 .....</b>	<b>23</b>
6.1 Enhancing Interoperability .....	23
6.1.1 Embedded Python .....	23
6.1.2 Interoperability Productions in Python .....	24
6.1.3 Visual Studio Code ObjectScript Extension Pack Updates .....	24
6.1.4 Make SQL Queries With Minimal Code in Interoperability Productions .....	24
6.2 Enhancing Cloud and Operations .....	25
6.2.1 Cloud Connectors .....	25
6.2.2 IKO Enhancements .....	25
6.3 Enhancing Speed, Scale, and Security .....	25
6.3.1 Adaptive SQL Optimizer .....	25
6.3.2 Saving on Storage .....	25
6.3.3 TLS 1.3 Support (OpenSSL 1.1.1) .....	26
6.3.4 New ^TRACE Utility .....	26
6.4 First Continuous Delivery Releases of Health Connect .....	26
6.5 Other Enhancements and Efficiency Improvements .....	27
<b>7 New in Health Connect 2021.1 .....</b>	<b>29</b>
7.1 Release Information for 2021.1 .....	29
7.2 Enhancing Health Interoperability .....	29
7.2.1 Enhancing FHIR® and IHE Profiles .....	29
7.2.2 Enhancing FHIR R4 Data Transformations .....	30
7.2.3 FHIR Repository and FHIR Interoperability Adapter .....	30
7.2.4 APIs for Client-Side FHIR Operations .....	30
7.2.5 New Configuration UI for FHIR Server .....	30
7.2.6 Support for IHE RMU Profile .....	30
7.2.7 IHE Connectathon Updates .....	30
7.2.8 Enhancing HL7 Productivity Tools .....	30
7.3 Enhancements that Improve Interoperability .....	31
7.4 Enhancing Operations .....	31
<b>8 New in Health Connect 2020.1 .....</b>	<b>33</b>
8.1 Healthcare Interoperability .....	34
8.1.1 FHIR R4 Support .....	34
8.1.2 HL7 Productivity Tools .....	34
8.2 API Management .....	34
8.2.1 InterSystems API Manager .....	35

8.2.2 Open API/Swagger Specification-First REST Development .....	35
8.3 In-Place Conversion from Caché, Ensemble, and Health Connect .....	35
8.4 New Look in the Management Portal .....	36
8.5 SQL Enhancements .....	36
8.5.1 Universal Query Cache .....	36
8.6 Interoperability Production Enhancements .....	37
8.6.1 New PEX Framework for Coding Production Components in Java and .NET .....	37
8.6.2 Port Authority for Monitoring Port Usage in Interoperability Productions .....	37
8.6.3 X12 Validation Enhancements .....	37
8.6.4 Enhanced DTL Support for X12 .....	38
8.6.5 Import X12 Schemas from XSD Files .....	38
8.6.6 MQTT Adapters .....	38
8.7 Infrastructure and Cloud Deployment Improvements .....	38
8.8 New Automatic Configuration Customization .....	39
8.9 Analytics Enhancements .....	39
8.9.1 Selective Cube Build .....	39
8.9.2 PowerBI Connector .....	39
8.9.3 Pivot Table Preview .....	39
8.10 Improved Performance and Scalability of the Database .....	40
8.11 Other Enhancements and Efficiency Improvements .....	40
<b>9 New in Health Connect 2019.1 .....</b>	<b>41</b>
9.1 New Features in Health Connect 2019.1 .....	41
9.1.1 FHIR STU3 Support .....	41
9.1.2 Java Business Hosts .....	41
9.1.3 Managed File Transfer (MFT) .....	41
9.1.4 Containerization of Health Connect .....	42
9.1.5 New Features in 2019.1.1 Maintenance Release .....	42
9.2 Other Changes in Health Connect 2019.1 .....	43
9.2.1 Enhancements .....	43
9.2.2 Class Removal and Deprecation .....	45
9.2.3 Method Replacement .....	47
9.2.4 Other Deprecations .....	47
9.3 Corrections in Health Connect 2019.1 .....	47
9.3.1 FHIR DSTU2 and STU3 Sorting .....	47
9.3.2 HTTP Response for Unknown Resource Type .....	48
9.3.3 Correct URL in MHD Find Document Manifest Request .....	48
9.3.4 New Maximum ID Length in XDS.b Message .....	48
9.3.5 Excluding Identifier Types during AA Code Resolution .....	48
9.3.6 Time Zone Offsets in DSUB Transactions .....	48
9.3.7 C-CDA v2.1 Export Correction for BirthTime .....	49
9.3.8 C-CDA v2.1 Import Converts the Name Qualifier BR to Birth .....	49
9.4 Upgrading from Health Connect 15.03 to Health Connect 2019.1 .....	49
<b>10 Upgrading to This Release .....</b>	<b>51</b>
10.1 New “Configure Secure Communication” Option in the Installer Wizard .....	51
10.2 Upgrading Containers .....	51
<b>11 Known Issues and Notes .....</b>	<b>53</b>
11.1 Unable to Open Documentation from the Launcher .....	53
11.2 Regression in NULL-Handling SQL Functions .....	53
<b>12 Deprecated and Discontinued Features .....</b>	<b>55</b>

12.1 About Deprecated and Discontinued Technologies and Features .....	55
12.2 Private Web Server (PWS) .....	55
12.3 Studio .....	56
12.4 Unstructured Information Management Architecture (UIMA) Integration .....	56
12.5 Spark Connector .....	56
12.6 Atelier .....	56
12.7 Zen .....	56
12.8 Zen Reports .....	57

# 1

## New in HealthShare Health Connect 2024.1

This page describes the new and enhanced features in the 2024.1 release of HealthShare® Health Connect, which is an extended maintenance (EM) release. Some of these features were also available in the continuous delivery (CD) releases since 2023.1, the previous EM release.

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

### 1.1 Release Information for 2024.1

The current release is 2024.1. The posting for 2024.1 is build 2024.1.0.263.0.

### 1.2 Private Web Server (PWS)

Starting with this release:

- Local HealthShare Health Connect 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.
- HealthShare Health Connect 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.

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 HealthShare Health Connect 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 HealthShare Health Connect 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 deprecated; at that point, upgrades of existing HealthShare Health Connect instances will remove the private web server.

## 1.3 Healthcare Interoperability

### 1.3.1 SMART Scopes

InterSystems FHIR Server now supports SMART on FHIR OAuth2 access scopes using “CRUDS” (create/read/update/delete/search) style permissions. This enhancement aligns the application with the latest standards in healthcare interoperability and security, ensuring seamless integration with electronic health record systems and better data protection.

To learn more, refer to the HL7 documentation on scopes for using SMART to delegate access: <https://hl7.org/fhir/smart-app-launch/scopes-and-launch-context.html#scopes-for-requesting-fhir-resources>

For backward compatibility, InterSystems continues to support the 1.0 FHIR scopes.

### 1.3.2 JsonAdvSQL

InterSystems is excited to announce the release of JsonAdvSQL, a new FHIR search strategy designed to enhance performance and compliance. This release addresses critical limitations in the previous JSON strategy and introduces significant improvements.

JsonAdvSQL represents a significant step forward in FHIR search capabilities, offering faster, more compliant, and scalable solutions. This upgrade is poised to greatly enhance user experience and operational efficiency in managing healthcare data.

Some of the features of the enhanced conformance JsonAdvSQL provides are as follows:

- Multi-resource compartment searches with wildcards and `_type`
- Full support for `_include` / `_revinclude` and wildcards and `:iterate`
- Canonical searches now include version support
- New modifiers/prefixes and improved handling of implicit ranges
- Enhanced capability in text, identifier, and URI searches
- Better handling of Date/Quantity/Number search parameters, including new prefixes like ‘sa’, ‘eb’, and ‘ap’
- Improved expression of SearchInclude/SearchRevInclude wildcards and system vs. type level operations in CapabilityStatement

### 1.3.3 New FHIR Object Model

HealthShare Health Connect 2024.1 introduces a new FHIR Object Model to enhance the developer experience when working with FHIR data in ObjectScript. Key features of the new FHIR Object Model include:

- **Static Class Representation:** A package of `%RegisteredObject` classes that closely parallel the structure and naming conventions of FHIR resources and data elements. This provides better tooling support like IntelliSense, auto-completion, and documentation pop-ups when working with FHIR data structures.
- **Convenience Methods:** The model classes include convenient methods to allocate nested structures and collections on demand, following the naming conventions of the FHIR specification. This allows developers to synthesize FHIR payloads more easily, driven by IDE suggestions.

- **Abstract Set API Support:** The FHIR model classes support a subset of the Abstract Set API, providing a uniform coding pattern across dynamic (%DynamicAbstractObject) and static (%RegisteredObject) data models when working with FHIR data.
- **Data Conversion:** The classes include methods to convert between the static class model and the dynamic %DynamicAbstractObject representation, facilitating interoperability between the two models.
- **Query Engine:** The Abstract Set API integration allows developers to query FHIR data structures using the JSON Path Language (JPL), without the need for indexing or persisting the data.

The new FHIR Object Model aims to reduce the cognitive burden of working with the extensive FHIR data structures, providing better tooling support, consistent coding patterns, and efficient querying capabilities, ultimately enhancing the developer experience when building FHIR-based applications in HealthShare Health Connect. It is fully compatible with VS Code.

### 1.3.4 Interoperability Usage Metrics

InterSystems introduces Interoperability Usage Metrics in this release. With the substantial volume of data flowing through the interoperability engine, gaining insights into how the platform is utilized has become imperative for efficient operation and resource planning. Interoperability now provides a centralized monitor for counting the number of used interfaces for a given deployment. Users can access a granular breakdown of interface usage by date and namespace, this enables a more detailed understanding of how the system is utilized over specific periods and within specific contexts. The Interoperability Usage Metrics are conveniently accessible via API, facilitating the creation of comprehensive reports. API access empowers users to integrate usage data seamlessly into their reporting systems.

### 1.3.5 HL7 Production Validator

It can be expensive to bring a new interface live and hard to trust machine-generated parts of your logic (for example, logic produced by the HL7 Productivity Toolkit), especially when you lack a good testing facility. Interface verification is one of the major blockers for adopting new releases due to the cost and effort it requires. To ease the burden of upgrading systems, there is a need to have verification and trusted transformation logic to ensure correct system behavior.

To solve this common issue, InterSystems introduces the HL7 Production Validator to the 2024.1 release. The HL7 Production Validator is a terminal based utility that extracts HL7 Service headers, messages, and Operation messages to a temporary database. This production data is copied to your upgraded instance (the “target system”) and reconstituted for comparison against the original instance (the “source system”). Differences are evaluated based on recorded changes to the messages; these differences are represented to the user such that they can easily identify if the behavior is the same and identify where adjustments may be needed.

The HL7 Production Validator is particularly beneficial to customers who are migrating from HealthShare HealthConnect to InterSystems IRIS for Health or upgrading their current HealthShare Health Connect instance.

## 1.4 Enhancing Developer Experience

### 1.4.1 Using vectors in ObjectScript

As part of the new Vector Search capability (see related note in the Enhancing Analytics and AI section), InterSystems is making the \$vector ObjectScript language feature available to all developers. InterSystems introduced \$vector as an internal-only feature in 2022.2 to address the need for a new datatype and corresponding set of operators to work with large arrays of elements of a single datatype, and its highly efficient encoding format and use of SIMD functions for bulk operations contributed to the order-of-magnitude performance improvements brought by Columnar Storage. The Vector Search

capability introduced with this release offers another use case for working with this kind of datatype, notably to capture embeddings and to perform fast bulk operations on the entire vector at once. Given the feature addresses two distinct use cases in modern computing challenges well, InterSystems has decided to make the feature available to all ObjectScript developers.

You can find full documentation of the new syntax and the current set of operators in the ObjectScript reference.

At first sight, you may think \$vector has a lot of overlap with \$list and wonder which one to use. They are, however, distinct features that server different use cases. In general, if you are dealing with a single datatype and always access and operate on data elements in bulk, you may fit the \$vector use case.

## 1.4.2 Support for JSON\_TABLE

HealthShare Health Connect 2024.1 introduces support for the **JSON\_TABLE** function according to the SQL standard. This function, intended for use in the FROM clause of a query, maps JSON values to columns in a relational table, enabling you to query JSON data in the context of SQL and providing very powerful capabilities in combination with lateral joins.

## 1.4.3 SQL Development

The SQL standard includes a **WITH** clause that helps in breaking complex nested queries down into small blocks, similar to how you'd use SQL views, but restricted to a single statement. In 2024.1, you can now use such Common Table Expressions (CTEs) in **SELECT** statements to add more structure and transparency to your queries. Many tools that generate SQL statement can also use CTEs.

As with each release, HealthShare Health Connect 2024.1 includes a number of enhancements that improve query performance:

- When executing queries, any intermediate results that can be shared between different worker threads are now cached. This saves significantly on I/O when running complex queries on large datasets.
- Space efficiency for tables that use columnar storage where all table columns meet the requirements for pure columnar storage has been optimized.
- The Adaptive Parallel Execution mechanism introduced in prior releases now applies to a broader set of queries, including those that employ complex conditions, and virtually eliminates the cases in which traditional parallel subqueries would be used in the background.
- The set of cases where filter and other query predicates for statements involving foreign tables can be pushed down to a remote database has been expanded. This makes foreign tables more performant by limiting the amount of data that gets passed back to your instance of HealthShare Health Connect.

### 1.4.3.1 Flexible Embedded Python Runtime

Flexible Python Runtime allows an administrator to choose the Python runtime that Embedded Python uses. This is perfect for upgrading to a specific version of Python or using a distribution like Anaconda. The Python runtime library can be configured through the Configuration Parameter File. This feature is available only for Linux-based builds of HealthShare Health Connect in 2024.1

### 1.4.3.2 Embedded Python for AIX

Embedded Python is now available on AIX systems.

### 1.4.3.3 Python BPL Editor

The Business Process Language editor now gives developers the option to write their business processes using Python. While most BPL scripts require no coding, there are a few blocks that enable the developer to implement their own logic in ObjectScript, JavaScript, or (new in the 2024.1 release) Python.

### 1.4.3.4 WSGI Web Apps

HealthShare Health Connect 2024.1 includes, on an experimental basis, the ability to create, secure, and host Web Applications that conform to the Python WSGI standard. With this enabled, you can run WSGI-compliant frameworks, such as flask or Django, inside your instance and take advantage of all the power of Embedded Python for modern Python web development. Even better, authentication is handled by the Web Gateway just like with all web applications. This feature is experimental in the 2024.1 release.

InterSystems welcomes feedback on how well this feature meets your needs.

## 1.5 Enhancing Analytics and AI

### 1.5.1 Vector Search

Vector Search is a new experimental feature in HealthShare Health Connect 2024.1 that provides the core functionality for semantic and vector-powered search of structured data, the cornerstone of generative AI applications, such as those implementing the Retrieval Augmented Generation (RAG) pattern. Vector search in 2024.1 comprises a new VECTOR SQL datatype and a set of related functions, including TO\_VECTOR (which converts strings to vectors) and both VECTOR\_DOT\_PRODUCT and VECTOR\_COSINE (which compute similarity between vectors).

In a typical application, specialized statistical language models called “embedding models” (smaller versions of LLMs like ChatGPT) convert sequences of text (or images, if the embedding model has been trained to process images) into embedding vectors, dense numeric arrays that specify a point in high-dimensional “latent embedding” space. Vectors are then stored in a SQL table column and the similarity functions are used to compare them to an input or query vector, the result of applying the same embedding model to the query string. The results of that search are the most similar vectors, and this can be a higher quality search capability than traditional “keyword” search because the embedding models have been trained to capture the “semantics” of the text. Future releases of InterSystems IRIS will increase the speed of search and introduce SQL syntax that abstracts the process of converting stored text as well as queries into vectors thereby providing a convenient way to implement powerful “semantic search” over data in InterSystems IRIS without having to deal with the underlying vectors directly.

### 1.5.2 Business Intelligence Improvements

InterSystems continues to improve the performance of InterSystems IRIS BI. In this release, the Analyzer user interface’s filter list management has been enhanced. In addition, this release includes:

- Handling of deeply nested AND/OR conditions in the Advanced filters in Analyzer. This change introduces a means for computing structures consisting of deeply nested CROSSJOIN and %OR functions in MDX slicer clauses.
- Enhanced APIs to allow accessing pivot table and dashboard metadata to allow for other visualization user interfaces to access that stored data.

## 1.6 Enhancing Speed, Scale and Security

### 1.6.1 Multi-Volume Databases

As customer databases grow, so do database files. To avoid those files becoming unmanageably large, or hit hard filesystem limits, HealthShare Health Connect now supports splitting your database across multiple physical “volumes” transparently. This new capability is easily configured: for any database, you can now configure a threshold size. When your initial .DAT file is about to hit this size and new global data needs to be written, HealthShare Health Connect will transparently create a new “database volume” file and start writing new data to that file. When that volume hits the threshold, another file is created, and so on.

There is no impact on applications and code accessing the data, as they continue to see the full database's content, no matter how many volumes it might be spread across. Furthermore, you can configure which directory the next volume should be created in, and when needed, you can rearrange database volumes across directories as a maintenance operation. Combined with planned work to increase the overall maximum database size, this will ensure your data remains easy to manage, well into the petabyte range.

### 1.6.2 Fast Online Backup

InterSystems for a long time has offered two main backup options: External Backup and Online Backup. With External Backup, users can freeze the write daemon for a brief period of time such that external solutions can create a consistent snapshot of all database, journal and other files. Those external solutions may track what changed versus the previous snapshot to produce an incremental snapshot that is much smaller than a full copy. In Online Backup, HealthShare Health Connect tracks which blocks changed, and can produce incremental backups while the system remains fully online (with only a very brief pause in writing).

This release introduces a new mechanism for incremental and full backups to the Online Backup feature, writing separate backup files per database using parallel processes. For large systems with many databases, this provides very significant improvements in the overall time it takes to perform a backup. The new `Backup.Online` class is the entry point for the new capability.

**Important:** This change is the first part of a broader major project improving the performance and interface of the Online Backup feature, and covers the parallelism described above. As we intend to make further changes to the API and output format, the new API is currently labeled as experimental.

InterSystems welcomes feedback from customers through the [Early Adopter Program](#) and look forward to incorporating such feedback in an upcoming release.

During this phase, the pre-existing API in `Backup.General` and other utilities continue to invoke Online Backup's pre-existing implementation and can continue to be used without any changes.

## 1.7 Platform Updates

### 1.7.1 Minimum Supported CPU Models

HealthShare Health Connect 2024.1 now has set a minimum CPU instruction set policy for Intel and AMD (amd64/x86\_64) processors. InterSystems now requires all CPUs to have the AVX and BMI instructions, which are generally available on the following CPU architectures:

- For Intel processors: Haswell and up
- For AMD processors: Steamroller and up

InterSystems is taking advantage of newer instructions to improve product performance. HealthShare Health Connect 2024.1 takes advantage of the AVX instruction to reliably speed up vector operations.

## 1.7.2 Operating System Updates

HealthShare Health Connect 2024.1 will support Ubuntu 24.04 shortly after the OS is generally available.

MacOS Sonoma is added as a supported version of MacOS.



# 2

## New in HealthShare Health Connect 2023.1

This page describes the new and enhanced features in the 2023.1 release of HealthShare® Health Connect, 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](#).

### 2.1 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.

### 2.2 Enhancing Analytics and AI

#### 2.2.1 Columnar Storage

*Columnar Storage* is a new storage option for InterSystems 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 HealthShare Health Connect 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](#).

## 2.3 Enhancing Speed, Scale and Security

### 2.3.1 Memory Settings

New installations of Health Connect 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.

### 2.3.2 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.

## 2.4 Platform Updates

This release adds support for the following new server platforms:

- macOS 13 (Ventura)

# 3

## New in HealthShare Health Connect 2022.3

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

### 3.1 Enhancing Developer Experience

#### 3.1.1 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.

### 3.2 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)

## 3.3 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.

## 3.4 Platform Updates

This release adds support for the following new server platforms:

- Oracle Linux 9
- SUSE 15 SP4

# 4

## New in HealthShare Health Connect 2022.2

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

### 4.1 Enhancing Analytics and AI

#### 4.1.1 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 are running 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 [SQL Table Storage Layout](#).

### 4.2 Enhancing Speed, Scale, and Security

#### 4.2.1 SQL Enhancements

The 2022.2 release includes many enhancements to SQL.

### 4.2.1.1 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.

### 4.2.1.2 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.

### 4.2.1.3 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.

### 4.2.1.4 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 style 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).

## 4.3 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.



# 5

## New in Health Connect 2022.1

This page describes the new and enhanced features in the 2022.1 release of HealthShare® Health Connect. This 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](#).

### 5.1 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.

### 5.2 Enhancing Interoperability

#### 5.2.1 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.

#### 5.2.2 Embedded Python

Health Connect 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 Health Connect 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 Health Connect persistence to store objects in the database, making them available in future sessions until the objects are explicitly deleted.

Embedded Python augments the Health Connect 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 Health Connect 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.

## 5.2.3 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)

## 5.2.4 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-objectscript/](https://intersystems-community.github.io/vscode-objectscript/). (first in 2021.2)

## 5.2.5 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)

# 5.3 Enhancing Cloud and Operations

## 5.3.1 Cloud Connectors

This release contains cloud connectors that make it easier for you to manage Health Connect 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)
- Outbound adapter for Healthlake (Amazon health data service) (first in 2021.2)

## 5.3.2 IKO Enhancements

This release makes it easier to deploy and manage Health Connect 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 Health Connect cluster. This makes it easier to administer and scale your system. (first in 2021.2)
- IKO can deploy locked down Health Connect 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)

# 5.4 Enhancing Speed, Scale, and Security

## 5.4.1 Adaptive SQL Optimizer

The Health Connect 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 Health Connect 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, Health Connect 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, Health Connect SQL will now gather table statistics on-the-fly when a table has none to ensure appropriate query plans. (first in 2021.2)

## 5.4.2 Saving on Storage

In this release, stream and journal compression can significantly reduce storage needed for your Health Connect 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 Health Connect data integrity strategy. See [Journaling Best Practices](#) for more information. (first in 2021.2)

## 5.4.3 TLS 1.3 Support (OpenSSL 1.1.1)

With this version, Health Connect 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. (first in 2021.2)

Beginning with Health Connect 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 Health Connect 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 Health Connect. For example, on MacOS, you can install OpenSSL using Homebrew, see <https://formulae.brew.sh/formula/openssl@1.1>.

## 5.4.4 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)

## 5.5 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.
- 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)



# 6

## New in Health Connect 2021.2

This topic describes the new and enhanced features in the 2021.2 release of HealthShare® Health Connect. It addresses the new features in 2021.2 that were not present in the 2021.1 version. This is the first [Continuous Delivery \(CD\) release](#) of Health Connect.

### 6.1 Enhancing Interoperability

#### 6.1.1 Embedded Python

Health Connect 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 Health Connect 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 Health Connect persistence to store objects in the database, making them available in future sessions until the objects are explicitly deleted.

Embedded Python augments the Health Connect 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.

## 6.1.2 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](#).

## 6.1.3 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](https://intersystems-community.github.io/vscode-objectscript/) <https://intersystems-community.github.io/vscode-objectscript/>.

## 6.1.4 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.

## 6.2 Enhancing Cloud and Operations

### 6.2.1 Cloud Connectors

This release contains cloud connectors that make it easier for you to manage Health Connect 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
- Outbound adapter for Healthlake (Amazon health data service)

### 6.2.2 IKO Enhancements

This release makes it easier to deploy and manage Health Connect 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 Health Connect cluster. This makes it easier to administer and scale your system.
- IKO can deploy locked down Health Connect 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.

## 6.3 Enhancing Speed, Scale, and Security

### 6.3.1 Adaptive SQL Optimizer

The Health Connect 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 Health Connect 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, Health Connect 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, Health Connect SQL will now gather table statistics on-the-fly when a table has none to ensure appropriate query plans.

### 6.3.2 Saving on Storage

In this release, stream and journal compression can significantly reduce storage needed for your Health Connect 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 Health Connect data integrity strategy. See [Journaling Best Practices](#) for more information.

### 6.3.3 TLS 1.3 Support (OpenSSL 1.1.1)

With this version, Health Connect 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 Health Connect 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 Health Connect 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 Health Connect. For example, on MacOS, you can install OpenSSL using Homebrew, see <https://formulae.brew.sh/formula/openssl@1.1>.

### 6.3.4 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.

## 6.4 First Continuous Delivery Releases of Health Connect

Health Connect 2021.2 is the first continuous delivery release of Health Connect. There are two streams of Health Connect releases:

- Continuous delivery releases — these releases provide access to new features and are ideal for developing and deploying applications that are continuously being updated and can take immediate advantage of new features.
- 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.

Both CD and EM releases are provided on all supported platforms (see Supported Platforms Guide) including UNIX, Windows, the cloud platforms, and the OCCI-compliant container format. Previous CD releases were only available on cloud platforms and the OCCI-compliant container format.

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.

## 6.5 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.



# 7

## New in Health Connect 2021.1

This topic describes the new and enhanced features in the 2021.1 release of HealthShare® Health Connect. It addresses the new features in 2021.1 that were not present in the 2020.1.0 version.

### 7.1 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.

### 7.2 Enhancing Health Interoperability

#### 7.2.1 Enhancing FHIR® and IHE Profiles

InterSystems HealthShare Health Connect 2021.1 provides significant enhancements for InterSystems HealthShare Health Connect's FHIR® server framework, which provides a standard, user-friendly mechanism for incorporating and using FHIR profiles, including the ability to load a FHIR package and use it to configure a FHIR server endpoint. InterSystems HealthShare Health Connect supports externally published FHIR profiles as well as custom profiles, with out-of-the-box support for US Core Profiles (US Core Implementation Guide v3.1.0). For more details, see [FHIR Profiles and Adaptations](#).

HealthShare Health Connect 2021.1 introduces support for FHIRPath, a navigation and extraction language for FHIR that is similar to XPath for XML. FHIRPath is used in the FHIR Specification to express schema-level conditions & search parameter paths, etc. InterSystems HealthShare Health Connect includes APIs for parsing and evaluating FHIRPath expressions against FHIR data, supporting a subset of the various functions and operations defined in the official FHIRPath Specification.

This release enhances InterSystems HealthShare Health Connect's FHIR® server framework to provide a standard, user-friendly mechanism for incorporating and using FHIR profiles, including the ability to load a FHIR package and use it to configure a FHIR server endpoint. InterSystems HealthShare Health Connect supports externally published FHIR profiles as well as custom profiles, with out-of-the-box support for US Core Profiles (US Core Implementation Guide v3.1.0). For more details, see [FHIR Profiles and Adaptations](#).

RMD (Remove Metadata and Documents) is a new IHE profile that defines a process for removing metadata from an XDS.b Document Registry and documents from an XDS.b Document Repository that are no longer required to be discoverable within a patient's care record. InterSystems HealthShare Health Connect now supports Document Registry and Document Repository actors of this IHE profile.

## 7.2.2 Enhancing FHIR R4 Data Transformations

This release of InterSystems HealthShare Health Connect provides bi-directional data transformations between FHIR R4 and SDA. This includes transformation API methods as well as base Business Process components for performing the data transformation. For details, see [SDA-FHIR Transformations](#) .

## 7.2.3 FHIR Repository and FHIR Interoperability Adapter

Some solutions require a FHIR server that routes requests to an internal repository, but sometimes the requirements are only to receive a FHIR request and forward it to an external FHIR server without ever storing its payload. This release provides two important capabilities that handle these two different requirements:

- FHIR Repository — the [FHIR Resource Repository](#) is the default storage strategy for a FHIR server, allowing you to install a fully functioning FHIR® server without further development tasks. The FHIR Resource Repository is an add-on component that is available with Health Connect in this release.
- FHIR Interoperability Adapter — the [FHIR Interoperability Adapter](#) creates a new interoperability REST endpoint that uses special business hosts to process FHIR requests in a production. If your Health Connect license does not include the FHIR Resource Repository, the FHIR Interoperability Adapter provides an easy way to receive FHIR requests in your productions.

Although previous releases could handle FHIR requests, these new capabilities make it easier to do this and reduce the need to develop custom code in productions.

## 7.2.4 APIs for Client-Side FHIR Operations

New APIs for sending and receiving FHIR request/response messages, allowing your production to perform client-side FHIR operations. For more details, see [FHIR Clients](#).

## 7.2.5 New Configuration UI for FHIR Server

A new UI page has been added to enable the creation and configuration of a FHIR server directly from the Management Portal. To access the page, navigate to **Health > FHIR Configuration > Server Configuration** .

## 7.2.6 Support for IHE RMU Profile

RMU (Restricted Metadata Update) is a new IHE profile that provides a mechanism for modifying Document Sharing metadata both within and across community boundaries in a controlled manner. InterSystems HealthShare Health Connect now supports both Update Initiator and Update Responder actors of this profile .

## 7.2.7 IHE Connectathon Updates

This release of InterSystems HealthShare Health Connect includes all software updates and testing results from IHE North American Connectathon 2020. This includes integration of various fixes related to official IHE Change Proposals approved prior to January 2020 .

## 7.2.8 Enhancing HL7 Productivity Tools

eGate Support in the HL7 Migration Tooling. Migrates transformation logic from the eGate interface engine to InterSystems HealthShare Health Connect. For more details, see [HL7 Migration Tool](#).

## 7.3 Enhancements that Improve Interoperability

With InterSystems HealthShare Health Connect 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 .
- Improved X12 error handling .
- 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](#) .
- 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.

## 7.4 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 Health Connect-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 and SAM deployments.
- This release include asynchronous mirroring support for sharded clusters. 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.
- 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](#)”.
- You can now manage Work Queues from the System Management Portal.
- 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, Health Connect 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 Health Connect. As part of a closer integration, InterSystems Reports now uses the same user accounts as Health Connect for managing, building and executing reports. In addition, all configuration and management data for InterSystems Reports uses Health Connect if the setup scripting is used. A script to complete the initial configuration of Health Connect 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.
- 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. See [Memory Usage Changes for Global and Routine Buffers](#) for details.
- 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.

# 8

## New in Health Connect 2020.1

This page describes the new and enhanced features in the 2020.1 release of HealthShare® Health Connect. It addresses the new features in 2020.1 that were not present in the 2019.1.0 version. 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:

- [Healthcare Interoperability](#)
  - [FHIR R4 Base Standard Support](#)
  - [HL7 Productivity Tools](#)
- [API Management](#)
  - [InterSystems API Manager](#)
  - [Open API/Swagger Specification-First REST Development](#)
- [In-Place Conversion from Caché and Ensemble](#)
- [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](#)
- [Infrastructure and Cloud Deployment Improvements](#)
- [New Automatic Configuration Customization](#)
- [Analytics Enhancements](#)
  - [Selective Cube Build](#)

- [PowerBI Connector](#)
- [Pivot Table Preview](#)
- [Improved Performance and Scalability of the Database](#)
- [Other Enhancements and Efficiency Improvements](#)

## 8.1 Healthcare Interoperability

This release includes two new healthcare interoperability features:

- [FHIR R4 Support](#)
- [HL7 Productivity Tools](#)

### 8.1.1 FHIR R4 Support

A new FHIR server architecture provides enhanced performance and FHIR RESTful API support for all FHIR R4 base resources. Transformations between FHIR R4 and SDA will be supported in a subsequent release. For more information, see [FHIR Support in InterSystems Products](#).

### 8.1.2 HL7 Productivity Tools

This release includes three new tools designed to accelerate the process of building HL7 interfaces or migrating them from another interface engine.

- **Production Generator** — Generates the infrastructure of new interfaces (business services, routers, rules, transformations, and business operations) based on information in CSV files. Can also be used for bulk updates. For more details, see [Production Generator](#).
- **Message Analyzer** — Scans HL7 messages to determine if they conform with an HL7 schema. If the messages deviate, the tool can make changes to a custom schema at the field, data structure, and code table level. For more details, see [Message Analyzer](#).
- **Migration Tool** — Migrates transformation logic from the Cloverleaf interface engine to Health Connect. Future releases will support additional interface engines. For more details, see [Migration Tool](#).

## 8.2 API Management

This release includes two new API Management features:

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

## 8.2.1 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](#).

## 8.2.2 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)

## 8.3 In-Place Conversion from Caché, Ensemble, and Health Connect

This release of Health Connect allows you to convert an existing instance of Caché, Ensemble, or Ensemble-based Health Connect to Health Connect. 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.

The following conversion paths are supported:

- Health Connect (HSAP) 15.03x to Health Connect
- Ensemble 2016.2 (or above) to Health Connect
- Caché 2016.2 (or above) to Health Connect

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

## 8.4 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)

## 8.5 SQL Enhancements

As with every release, Health Connect 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. For details, see “Software Version Upgrade Automatically Freezes Plans” in the *InterSystems SQL Optimization Guide*.

- [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)

### 8.5.1 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.

## 8.6 Interoperability Production Enhancements

### 8.6.1 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:

- EnsLib.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](#) .

### 8.6.2 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)

### 8.6.3 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.

## 8.6.4 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](#).

## 8.6.5 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](#).

## 8.6.6 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](#).

# 8.7 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 Health Connect 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)

## 8.8 New Automatic Configuration Customization

A new Health Connect configuration feature enables customization of the configuration parameter file (CPF) of an Health Connect 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 Health Connect., 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)

## 8.9 Analytics Enhancements

This release contains the following analytics enhancements:

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

### 8.9.1 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.

### 8.9.2 PowerBI Connector

InterSystems customers can now use Microsoft Power BI to access tabular and cube data stored on Health Connect. This allows combining the data visualization capabilities offered by Power BI with the high-performance data management and querying capabilities offered by Health Connect. While the connector leverages ODBC, it will also allow customers to access Health Connect 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 .

### 8.9.3 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)

## 8.10 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.

## 8.11 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. See [Class Compiler Validates Global Name Length Limit](#) for a related compatibility issue. (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.

# 9

## New in Health Connect 2019.1

Health Connect 2019.1 is the first release of Health Connect that is powered by InterSystems IRIS. This page describes its features, corrections, and upgrade checklist items.

### 9.1 New Features in Health Connect 2019.1

The following major features are new in Health Connect 2019.1:

- [FHIR STU3 Support](#)
- [Java Business Hosts](#)
- [Managed File Transfer \(MFT\)](#)
- [Containerization of Health Connect](#)
- [New Features in 2019.1.1 Maintenance Release](#)

#### 9.1.1 FHIR STU3 Support

This release of HealthShare Health Connect adds support for FHIR STU3. With this new feature, Health Connect can now support both FHIR DSTU2 and STU3, providing a greater coverage of customer use cases involving FHIR. The support for FHIR STU3 applies to all FHIR-related components and functionalities presently covered by FHIR DSTU2, including FHIR server and client components, FHIR message and object models, data transforms to and from SDA, and FHIR-based IHE profiles. A standard mechanism for customizing the DTL classes that perform SDA to FHIR STU3 and FHIR STU3 to SDA transformations is provided. These transformations can be invoked with built-in business processes or by calling transformation APIs. Please see [FHIR Support in HealthShare Health Connect](#) for more information.

#### 9.1.2 Java Business Hosts

This feature enables users to create business services and operations entirely in Java without the need for any ObjectScript coding. Users can now add new protocols and perform complex operations using existing Java libraries.

#### 9.1.3 Managed File Transfer (MFT)

Managed File Transfer (MFT) is a brand new integration option that allows a third-party file transfer service to be used directly from within a production. This feature provides business hosts that can support Box, DropBox and Accellion

kiteworks, and users can configure these business hosts in a production to perform various file transfer operations, such as retrieving files from an end-user account.

## 9.1.4 Containerization of Health Connect

HealthShare Health Connect now supports deployment in a Docker container. This allows Health Connect to be deployed as a platform-independent, fully portable runtime solution, suitable for deployment on public cloud platforms such as Google, Amazon and Azure. Please see [First Look: InterSystems Products in Docker Containers](#) for more information.

## 9.1.5 New Features in 2019.1.1 Maintenance Release

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

### 9.1.5.1 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.

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.

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

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

### 9.1.5.2 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*.

## 9.2 Other Changes in Health Connect 2019.1

Additional changes in Health Connect 2019.1 include:

- [Enhancements](#)
- [Class Removal and Deprecation](#)
- [Method Replacement](#)
- [Other Deprecations](#)

### 9.2.1 Enhancements

Enhancements in Health Connect 2019.1 include:

- [Healthcare Interoperability Enhancements](#)
- [System Interoperability Enhancements](#)

#### 9.2.1.1 Healthcare Interoperability Enhancements

- **HL7 Schema Editor Enhancements** — The HL7 Schema Editor now provides a more user-friendly graphical user interface, allowing a user to edit HL7 schemas simply by dragging and dropping components. See “Defining a New Message Type and Structure Type” in *Routing HL7 Version 2 Messages in Productions* for details.
- **Source Control for HL7 Schemas** — The Management Portal can be integrated with a third-party source control system to place HL7 schemas under version control. See “Overview of HL7 Schemas and Messages” in *Routing HL7 Version 2 Messages in Productions* for details.
- **Verification of Access Token Audience by FHIR REST Handler** — The FHIR REST Handler classes for FHIR DSTU2 and STU3 (HS.FHIR.vDSTU2.REST.Handler and HS.FHIR.vSTU3.REST.Handler) now include support for verifying the audience of a given access token against the current CSP application endpoint. This update conforms to the OAuth 2.0 specification, which strongly recommends verifying the audience of an access token.
- **IHE Connectathon Updates** — This version of Health Connect has been updated with all changes and testing results from the 2018 North American and European IHE Connectathons. This includes integration of all fixes related to IHE Change Proposals approved prior to April 2018. For a full list of supported IHE profiles, please refer to <http://www.intersystems.com/ihe>.
- **XDS.b Registry Validation Setting for Repository OID** — A new Business Operation setting called **ValidateRepositoryOID** has been added to the IHE XDS.b Registry Operation component (HS.HC.IHE.XDSb.Registry.Operations). This check box indicates whether or not the repository OID needs to be validated; if unspecified, the business operation defaults to the previous behavior of obtaining this information from the “\IHE\XDSb\Registry\ValidateRepositoryOID” entry in the Config Registry.

- **Improved XDS.b Registry Query Performance** — Several updates have been made to the existing implementation of XDS.b Registry to improve Registry Query performance on large datasets.
- **Reference ID Option for XDS.b Registry** — This release of Health Connect adds the Reference ID option to the existing XDS.b Registry implementation, enabling the Document Registry to respond to FindDocumentsByReferenceId queries from Document Consumers.
- **DSUB Notification Broker Enhancements** — DSUB Notification Broker has been enhanced to support subscription for SubmissionSets, which allows the use of "ihe:SubmissionSetMetadata" as a TopicExpression. Additionally, the Notification Broker's existing support for subscription for DocumentEntry has been extended to include the \$XDSDocumentEntryReferenceIdList filter parameter.
- **Shared/National Patient Identifier Query and Feed Option in XCPD** — This release of Health Connect enhances the existing IHE XCPD (Cross-Community Patient Discovery) implementation by adding a new mode of operation called Shared/National Patient Identifier Query and Feed. This mode allows a national patient identifier to be used in place of patient demographics in a Cross Gateway Patient Discovery (ITI-55) query, applicable to both XCPD Initiating Gateway and Responding Gateway. For ease of configuration, a new host-level setting (check box) called "National-PatientIdentifier" has been added to HS.IHE.XCPD.InitiatingGateway.Process and HS.IHE.XCPD.RespondingGateway.Process. Additionally, a new Config Registry key named "\IHE\NationalPatientIdentifierAA" has been added for specifying the Assigning Authority to use when handling national patient identifiers.
- **New Data Fields in PIXv3 Patient Identity Feed Message** — This enhancement extends the existing PIXv3 implementation by allowing additional demographics data to be included in a Patient Identity Feed HL7 V3 (ITI-44) message, whose corresponding internal message is HS.Message.AddUpdateHubRequest. This change is applicable to both PIX Manager and Patient Identity Source. More specifically, the following additional data fields are now supported: Birth Name (Family, Given, Prefix), Aliases, MothersName, FathersName and SpousesName.
- **File-Based WSDLs for IHE Services** — This update introduces file-based WSDLs for various IHE services supported in Health Connect. The new WSDL files, such as HS.IHE.PIXv3.Manager.Services.CLS.wsdl and HS.IHE.XCA.InitiatingGateway.Services.CLS.wsdl, are found under <Install Dir>/dev/wsdl, with all the applicable XSD schemas that are referenced from those WSDLs now locally stored under /wsdl/schema. Previously WSDLs for IHE services were generated and served up from individual web services, with references going out to XSDs found externally on intersystems.com. Recent product-level security updates have made it harder to access WSDLs via the old mechanism without changing certain security settings.
- **Enhancements to XUA** — This enhancement allows the customer to have greater control and configurability when setting up XUA in Health Connect. More specifically, new HS.HC.IHE.XUA.Creator and HS.HC.IHE.XUA.Processor classes have been introduced, and these classes are different from their predecessors in that they do not require the user/clinician registry. The previous implementation of XUA was closely tied to HealthShare's user/clinician registry, which made it difficult to support alternate ways of validating users.

### 9.2.1.2 System Interoperability Enhancements

Health Connect 2019.1 includes the following new system interoperability capabilities that speed configuring and troubleshooting of productions:

- **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 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” 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” for details on new options in the business host wizards.
- **Support for Custom Database Locations** — This enhancement allows a user to specify a custom database location when installing a Foundation production. Previously there was no officially supported way to indicate a database location that was different from the default location found under <Install Dir>/mgr. A new field called **Alternate Database Location** has been added to the Installer Wizard in the Management Portal, and this optional field is visible to the user only when a new Foundation production is being created. Additionally, corresponding changes have been made in the install script to support the alternate location, defining a new variable called "DBInstallDirectory". Please note that this enhancement does not support moving the database location of an existing production.

## 9.2.2 Class Removal and Deprecation

The following classes have been removed or deprecated in Health Connect 2019.1:

Health Connect 15.03 Class	Status in Health Connect 2019.1
HS.Audit.ConsolidationServices	Removed; use HS.HC.Audit.ConsolidationServices instead.
HS.FHIR.MHD.*	Deprecated; use HS.FHIR.vDSTU2.MHD.* or HS.FHIR.vSTU3.MHD.* instead.
HS.FHIR.PDQm.*	Deprecated; use HS.FHIR.vDSTU2.PDQm.* or HS.FHIR.vSTU3.PDQm.* instead.
HS.FHIR.PIXm.*	Deprecated; use HS.FHIR.vDSTU2.PIXm.* or HS.FHIR.vSTU3.PIXm.* instead.
HS.FHIR.Model.*	Deprecated; use HS.FHIR.vDSTU2.Model.* or HS.FHIR.vSTU3.Model.* instead.
HS.FHIR.Utils.DateConversion	Deprecated.
HS.FHIR.Utils.SearchTableBuilder	Removed; use HS.FHIR.vDSTU2.SearchTableBuilder or HS.FHIR.vSTU3.SearchTableBuilder instead.
HS.FHIR.REST.Handler	Deprecated; use HS.FHIR.vDSTU2.REST.Handler or HS.FHIR.vSTU3.REST.Handler instead.
HS.FHIR.Operation.Process	Removed. If you have customized this class, move the customizations to HS.FHIR.vDSTU2.Repository.OperationProcessor or HS.FHIR.vSTU3.Repository.OperationProcessor.
HS.FHIR.Gateway.*	Removed.
HS.FHIR.Repository.*	All classes except for HS.FHIR.Repository.Operations have been removed. HS.FHIR.Repository.Operations is an extension of HS.FHIR.vDSTU2.Repository.Operations.
HS.IHE.DSUB.Publisher.Process	Removed; use HS.HC.IHE.DSUB.Publisher.Process instead.
HS.IHE.XDR.Recipient.CommonProcess	Removed; use HS.HC.IHE.XDR.Recipient.CommonProcess instead.
HS.IHE.XDSb.Consumer.Operations	Removed; use HS.HC.IHE.XDSb.Consumer.Operations instead.
HS.IHE.XDSb.Registry.Operations	Removed; use HS.HC.IHE.XDSb.Registry.Operations instead.
HS.Message.ECRUpdateRequest	Removed. If HS.Message.ECRUpdateRequest was previously used as an input message type for HS.FHIR.FromSDA.DTL.Transaction.Process, then modify the input to use HS.Message.XMLMessage or Ens.StreamContainer instead.
HS.MPI.Manager	Removed; use HS.HC.MPI.Manager instead.
HS.Test.UI.FHIR.LoggedOut	Removed.
HS.Test.UI.FHIR.ServerSelect	Removed.
HS.Test.UI.FHIR.ServerSelectOAuth2	Removed.

Health Connect 15.03 Class	Status in Health Connect 2019.1
HS.UI.AssigningAuthorities	Removed; use HS.HC.UI.AssigningAuthorities instead.
HS.UI.Home	Removed; use HS.HC.UI.Home instead.
HS.Util.Installer.Kit.FHIR.*	Removed; use HS.HC.Util.Installer.Kit.FHIR.* instead.
HS.UI.Installer.Welcome	Removed; use HS.HC.UI.Installer.Welcome instead.
HS.Util.Trace.Helper	Removed; use HS.HC.Util.Trace.Helper

## 9.2.3 Method Replacement

The following FHIR-related install methods from previous releases have been removed and replaced in this release of Health Connect:

Removed Method	New Replacement Method
HS.Util.Installer.InstallFHIRServer()	HS.HC.Util.Installer.FHIR.Install()
HS.Util.Installer.InstallOAuth2()	HS.HC.Util.Installer.OAuth2.ConfigureSampleOAuth2()

## 9.2.4 Other Deprecations

Support for QRDA (Quality Reporting Document Architecture), which was first introduced in HealthShare Core 13, is now deprecated as of this release, and all code and XSLTs associated with this feature will be removed in the next major release of Health Connect.

# 9.3 Corrections in Health Connect 2019.1

Significant corrections in this version include:

- [FHIR DSTU2 and STU3 Sorting](#)
- [HTTP Response for Unknown Resource Type](#)
- [Correct URL in MHD Find Document Manifest Request](#)
- [New Maximum ID Length in XDS.b Message](#)
- [Excluding Identifier Types during AA Code Resolution](#)
- [Time Zone Offsets in DSUB Transactions](#)
- [C-CDA v2.1 Export Correction for BirthTime](#)
- [C-CDA v2.1 Import Converts the Name Qualifier BR to Birth](#)

Also see [Health Connect Maintenance Release Changes \(2019.1.4\)](#).

## 9.3.1 FHIR DSTU2 and STU3 Sorting

This release of Health Connect fixes resource sorting issues found in the previous implementation of the FHIR server and correctly implements the "\_sort" search parameter for both FHIR DSTU2 and STU3. FHIR STU3 makes a few changes to

how the `_sort` parameter on a search URL works. In DSTU2, multiple `_sort` values could be submitted by repeating the `_sort` parameter. In STU3, there should only ever be one `_sort` parameter per request, and the value is a comma-delimited list of search parameters. Also, in DSTU2, each `_sort` parameter could have the modifier "asc" or "desc" to specify the sort order, with no modifier being the same as specifying "asc". In STU3, there are no modifiers to the `_sort` parameter. If a value in the `_sort` list is prefixed by a "-", then it indicates decreasing order, otherwise increasing order is assumed.

### 9.3.2 HTTP Response for Unknown Resource Type

The FHIR server API behavior has been corrected to better conform to the FHIR DSTU2 and STU3 specification on how to respond to requests for unknown resource types. The following response behavior has been implemented for different FHIR interactions:

- read: HTTP status 404 Not Found and OperationOutcome
- vread: HTTP status 404 Not Found and OperationOutcome
- update: HTTP status 404 Not Found only
- delete: HTTP status 204 No Content
- history: HTTP status 404 Not Found only
- create: HTTP status 404 Not Found only
- search: HTTP status 404 Not Found and OperationOutcome

### 9.3.3 Correct URL in MHD Find Document Manifest Request

The previous implementation of MHD Document Responder was populating the `fullURL` element with a relative URL value in its response to a Find Document Manifest request. This issue, which was seen only with Find Document Manifest, is now fixed.

### 9.3.4 New Maximum ID Length in XDS.b Message

Previously Health Connect internally restricted ID values in an XDS.b message to 73 characters, and this can result in an error if, for example, an XDS.b Provide and Register request is received with a very long ID. This limit on the ID length is now increased to 256 characters, based on the eBRIM specification, as IHE does not directly specify a limit on the length.

### 9.3.5 Excluding Identifier Types during AA Code Resolution

A new host-level setting (text box) called **ExcludeIdentifierTypes** has been added to PIX and PIXv3 Consumer and Manager processes (HS.IHE.PIX.Consumer.Process, HS.IHE.PIX.Manager.Process, HS.IHE.PIXv3.Consumer.Process and HS.IHE.PIXv3.Manager.Process). This new setting specifies which Identifier Types to exclude when a PIX Consumer or Manager needs to query the HS\_AssignAuth.Config table to obtain the Assigning Authority (AA) Identifier Type for a given AA code. Previously it was assumed that there would only be one AA Identifier Type associated with an AA code. This incorrect assumption has been resolved through this new setting by controlling exactly which Identifier Types to consider during AA code resolution. DL (Driver License) and DN (Doctor Number) are excluded by default.

### 9.3.6 Time Zone Offsets in DSUB Transactions

Subscription times in a DSUB transaction are expected to be in UTC or with an offset, but the previous DSUB implementation incorrectly assumed the times to be always local and ignored any time zone offsets specified. This has been corrected. Subscriptions can now properly handle both UTC and UTC +/-, with UTC always being used for internal processing.

### 9.3.7 C-CDA v2.1 Export Correction for BirthTime

The SDA to C-CDA v2.1 export has been corrected to remove the unexpected `stdc:birthTime` element appearing under the `playingEntity` node. Previously this extraneous element was causing schema validation issues.

### 9.3.8 C-CDA v2.1 Import Converts the Name Qualifier BR to Birth

The base C-CDA v2.1 import transform was previously missing the logic for converting "BR" to "Birth" when importing a C-CDA v2.1 document into SDA. "BR" is a name qualifier code used in CDA for representing birth names and this is internally represented in HealthShare as "Birth". This issue has now been corrected.

## 9.4 Upgrading from Health Connect 15.03 to Health Connect 2019.1

Health Connect 2019.1 is the first release of Health Connect that is powered by InterSystems IRIS. Because of this change in underlying technology, a special upgrade procedure is required when upgrading from a previous version of Health Connect. If you are considering upgrading to Health Connect 2019.1, see the *InterSystems IRIS In-Place Conversion Guide*, which is available from [InterSystems WRC Documents](#).



# 10

## 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.

**Note:** A special upgrade procedure is required when upgrading from Health Connect 15.03. For details, see the *InterSystems IRIS In-Place Conversion Guide*, which is available from the [InterSystems WRC Documents](#) page.

### 10.1 New “Configure Secure Communication” Option in the Installer Wizard

The **Configure SSL Access** dialog in the Installer Wizard has been renamed to **Configure Secure Communication**. In the new dialog you must now specify an SSL/TLS Configuration in order to make the secure communication settings Active. A default value of `HS.Secure.Demo` is entered for the SSL/TLS Configuration setting upon upgrade.

If you had an active configuration in the **Configure SSL Access** dialog before you upgraded, [identify the correct SSL/TLS configuration](#) for your instance in the new **Configure Secure Communication** dialog after you complete your upgrade but before you reactivate your namespaces.

### 10.2 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 Health Connect is run in a container, allowing a single instance to run sequentially in multiple containers over time. For example, if you run an Health Connect 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, you need to change some file ownerships in the host's durable directory before running Health Connect 2022.1. 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 Health Connect.

# 11

## Known Issues and Notes

This chapter describes known issues in HealthShare® Health Connect 2023.3.

### 11.1 Unable to Open Documentation from the Launcher

If you are using a new installation of 2024.1, are upgrading from 2023.2 or later, 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.

### 11.2 Regression in NULL-Handling SQL Functions

InterSystems has identified a regression in vectorized, bulk execution of null-handling SQL functions, like IFNULL and NULLIF, which can lead to wrong results. This change only affects 2024.1 and will be fixed in the next release.

Contact the [Worldwide Resource Center](#) (WRC) for assistance and further questions.



# 12

## Deprecated and Discontinued Features

### 12.1 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.

### 12.2 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.

## 12.3 Studio

As of version 2023.2, Studio has been deprecated. It is still included with this release, but it will be removed from a future release.

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](#).

## 12.4 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](#).

## 12.5 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.

## 12.6 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](#).

## 12.7 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.

## 12.8 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.

