



Informatica® Cloud Integration Hub
December 2022

Cloud Integration Hub

© Copyright Informatica LLC 2016, 2022

This software and documentation contain proprietary information of Informatica LLC and are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright law. Reverse engineering of the software is prohibited. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC. This Software may be protected by U.S. and/or international Patents and other Patents Pending.

Use, duplication, or disclosure of the Software by the U.S. Government is subject to the restrictions set forth in the applicable software license agreement and as provided in DFARS 227.7202-1(a) and 227.7702-3(a) (1995), DFARS 252.227-7013⁽¹⁾(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

The information in this product or documentation is subject to change without notice. If you find any problems in this product or documentation, please report them to us in writing.

Informatica, Informatica Platform, Informatica Data Services, PowerCenter, PowerCenterRT, PowerCenter Connect, PowerCenter Data Analyzer, PowerExchange, PowerMart, Metadata Manager, Informatica Data Quality, Informatica Data Explorer, Informatica B2B Data Transformation, Informatica B2B Data Exchange Informatica On Demand, Informatica Identity Resolution, Informatica Application Information Lifecycle Management, Informatica Complex Event Processing, Ultra Messaging, Informatica Master Data Management, and Live Data Map are trademarks or registered trademarks of Informatica LLC in the United States and in jurisdictions throughout the world. All other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties, including without limitation: Copyright DataDirect Technologies. All rights reserved. Copyright © Sun Microsystems. All rights reserved. Copyright © RSA Security Inc. All Rights Reserved. Copyright © Ordinal Technology Corp. All rights reserved. Copyright © Aandacht c.v. All rights reserved. Copyright Genivia, Inc. All rights reserved. Copyright Isomorphic Software. All rights reserved. Copyright © Meta Integration Technology, Inc. All rights reserved. Copyright © Intalio. All rights reserved. Copyright © Oracle. All rights reserved. Copyright © Adobe Systems Incorporated. All rights reserved. Copyright © DataArt, Inc. All rights reserved. Copyright © ComponentSource. All rights reserved. Copyright © Microsoft Corporation. All rights reserved. Copyright © Rogue Wave Software, Inc. All rights reserved. Copyright © Teradata Corporation. All rights reserved. Copyright © Yahoo! Inc. All rights reserved. Copyright © Glyph & Cog, LLC. All rights reserved. Copyright © Thinkmap, Inc. All rights reserved. Copyright © Clearpace Software Limited. All rights reserved. Copyright © Information Builders, Inc. All rights reserved. Copyright © OSS Nokalva, Inc. All rights reserved. Copyright Edifecs, Inc. All rights reserved. Copyright Cleo Communications, Inc. All rights reserved. Copyright © International Organization for Standardization 1986. All rights reserved. Copyright © ej-technologies GmbH. All rights reserved. Copyright © Jaspersoft Corporation. All rights reserved. Copyright © International Business Machines Corporation. All rights reserved. Copyright © yWorks GmbH. All rights reserved. Copyright © Lucent Technologies. All rights reserved. Copyright © University of Toronto. All rights reserved. Copyright © Daniel Veillard. All rights reserved. Copyright © Unicode, Inc. Copyright IBM Corp. All rights reserved. Copyright © MicroQuill Software Publishing, Inc. All rights reserved. Copyright © PassMark Software Pty Ltd. All rights reserved. Copyright © LogiXML, Inc. All rights reserved. Copyright © 2003-2010 Lorenzi Davide, All rights reserved. Copyright © Red Hat, Inc. All rights reserved. Copyright © The Board of Trustees of the Leland Stanford Junior University. All rights reserved. Copyright © EMC Corporation. All rights reserved. Copyright © Flexera Software. All rights reserved. Copyright © Jinfonet Software. All rights reserved. Copyright © Apple Inc. All rights reserved. Copyright © Telerik Inc. All rights reserved. Copyright © BEA Systems. All rights reserved. Copyright © PDFlib GmbH. All rights reserved. Copyright © Orientation in Objects GmbH. All rights reserved. Copyright © Tanuki Software, Ltd. All rights reserved. Copyright © Ricebridge. All rights reserved. Copyright © Sencha, Inc. All rights reserved. Copyright © Scalable Systems, Inc. All rights reserved. Copyright © jQWidgets. All rights reserved. Copyright © Tableau Software, Inc. All rights reserved. Copyright © MaxMind, Inc. All Rights Reserved. Copyright © TMate Software s.r.o. All rights reserved. Copyright © MapR Technologies Inc. All rights reserved. Copyright © Amazon Corporate LLC. All rights reserved. Copyright © Highsoft. All rights reserved. Copyright © Python Software Foundation. All rights reserved. Copyright © BeOpen.com. All rights reserved. Copyright © CNRI. All rights reserved.

This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>), and/or other software which is licensed under various versions of the Apache License (the "License"). You may obtain a copy of these Licenses at <http://www.apache.org/licenses/>. Unless required by applicable law or agreed to in writing, software distributed under these Licenses is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the Licenses for the specific language governing permissions and limitations under the Licenses.

This product includes software which was developed by Mozilla (<http://www.mozilla.org/>), software copyright The JBoss Group, LLC, all rights reserved; software copyright © 1999-2006 by Bruno Lowagie and Paulo Soares and other software which is licensed under various versions of the GNU Lesser General Public License Agreement, which may be found at <http://www.gnu.org/licenses/lgpl.html>. The materials are provided free of charge by Informatica, "as-is", without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose.

The product includes ACE(TM) and TAO(TM) software copyrighted by Douglas C. Schmidt and his research group at Washington University, University of California, Irvine, and Vanderbilt University, Copyright (©) 1993-2006, all rights reserved.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (copyright The OpenSSL Project. All Rights Reserved) and redistribution of this software is subject to terms available at <http://www.openssl.org> and <http://www.openssl.org/source/license.html>.

This product includes Curl software which is Copyright 1996-2013, Daniel Stenberg, <daniel@haxx.se>. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://curl.haxx.se/docs/copyright.html>. Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

The product includes software copyright 2001-2005 (©) MetaStuff, Ltd. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://www.dom4j.org/license.html>.

The product includes software copyright © 2004-2007, The Dojo Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://dojotoolkit.org/license>.

This product includes ICU software which is copyright International Business Machines Corporation and others. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://source.icu-project.org/repos/icu/icu/trunk/license.html>.

This product includes software copyright © 1996-2006 Per Bothner. All rights reserved. Your right to use such materials is set forth in the license which may be found at <http://www.gnu.org/software/kawa/Software-License.html>.

This product includes OSSP UUID software which is Copyright © 2002 Ralf S. Engelschall, Copyright © 2002 The OSSP Project Copyright © 2002 Cable & Wireless Deutschland. Permissions and limitations regarding this software are subject to terms available at <http://www.opensource.org/licenses/mit-license.php>.

This product includes software developed by Boost (<http://www.boost.org/>) or under the Boost software license. Permissions and limitations regarding this software are subject to terms available at http://www.boost.org/LICENSE_1_0.txt.

This product includes software copyright © 1997-2007 University of Cambridge. Permissions and limitations regarding this software are subject to terms available at <http://www.pcre.org/license.txt>.

This product includes software copyright © 2007 The Eclipse Foundation. All Rights Reserved. Permissions and limitations regarding this software are subject to terms available at <http://www.eclipse.org/org/documents/epl-v10.php> and at <http://www.eclipse.org/org/documents/edl-v10.php>.

This product includes software licensed under the terms at <http://www.tcl.tk/software/tcltk/license.html>, <http://www.bosrup.com/web/overlib/?License>, <http://www.stlport.org/doc/license.html>, <http://asm.ow2.org/license.html>, <http://www.cryptix.org/LICENSE.TXT>, <http://hsqldb.org/web/hsqldbLicense.html>, <http://httpunit.sourceforge.net/doc/license.html>, <http://jung.sourceforge.net/license.txt>, http://www.gzip.org/zlib/zlib_license.html, <http://www.openldap.org/software/release/license.html>, <http://www.libssh2.org>, <http://slf4j.org/license.html>, <http://www.sente.ch/software/OpenSourceLicense.html>, <http://fusesource.com/downloads/license-agreements/fuse-message-broker-v-5-3-license-agreement>, <http://antlr.org/license.html>, <http://aopalliance.sourceforge.net/>, <http://www.bouncycastle.org/licence.html>, <http://www.jgraph.com/jgraphdownload.html>, <http://www.jcraft.com/jsch/LICENSE.txt>, http://jotm.objectweb.org/bsd_license.html, <http://www.w3.org/Consortium/Legal/2002/copyright-software-20021231>, <http://www.slf4j.org/license.html>, <http://nanoxml.sourceforge.net/orig/copyright.html>, <http://www.json.org/license.html>, <http://forge.ow2.org/projects/javaservice/>, <http://www.postgresql.org/about/licence.html>, <http://www.sqlite.org/copyright.html>, <http://www.tcl.tk/software/tcltk/license.html>, <http://www.jaxen.org/faq.html>, <http://www.jdom.org/docs/faq.html>, <http://www.slf4j.org/license.html>, <http://www.iodbc.org/dataspace/iodbc/wiki/IODBC/License>, <http://www.keplerproject.org/md5/license.html>, <http://www.toedter.com/en/jcalendar/license.html>, <http://www.edankert.com/bounce/index.html>, <http://www.net-snmp.org/about/license.html>, <http://www.openmdx.org/#FAQ>, http://www.php.net/license/3_01.txt, <http://srp.stanford.edu/license.txt>, <http://www.schneider.com/blowfish.html>, <http://www.jmock.org/license.html>, <http://xsom.java.net>, <http://benalman.com/about/license/>, <https://github.com/CreateJS/EaselJS/blob/master/src/easeljs/display/Bitmap.js>, <http://www.h2database.com/html/license.html#summary>, <http://jsoncpp.sourceforge.net/LICENSE>, <http://jdbc.postgresql.org/license.html>, <http://protobuf.googlecode.com/svn/trunk/src/google/protobuf/descriptor.proto>, <https://github.com/rantav/hector/blob/master/LICENSE>, <http://web.mit.edu/Kerberos/krb5-current/doc/mitK5license.html>, <http://jibx.sourceforge.net/jibx-license.html>, <https://github.com/lyokato/libgeohash/blob/master/LICENSE>, <https://github.com/hjiang/jsonxx/blob/master/LICENSE>, <https://code.google.com/p/lz4/>, <https://github.com/jedisct1/libsodium/blob/master/LICENSE>, <http://one-jar.sourceforge.net/index.php?page=documents&file=license>, <https://github.com/EsotericSoftware/kryo/blob/master/license.txt>, <http://www.scala-lang.org/license.html>, <https://github.com/tinkerpop/blueprints/blob/master/LICENSE.txt>, <http://gee.cs.oswego.edu/dl/classes/EDU/oswego/cs/dl/util/concurrent/intro.html>, <https://aws.amazon.com/ssl/>, <https://github.com/twbs/bootstrap/blob/master/LICENSE>, <https://sourceforge.net/p/xmlunit/code/HEAD/tree/trunk/LICENSE.txt>, <https://github.com/documentcloud/underscore-contrib/blob/master/LICENSE>, and <https://github.com/apache/hbase/blob/master/LICENSE.txt>.

This product includes software licensed under the Academic Free License (<http://www.opensource.org/licenses/afl-3.0.php>), the Common Development and Distribution License (<http://www.opensource.org/licenses/cddl1.php>), the Common Public License (<http://www.opensource.org/licenses/cpl1.0.php>), the Sun Binary Code License Agreement Supplemental License Terms, the BSD License (<http://www.opensource.org/licenses/bsd-license.php>), the new BSD License (<http://opensource.org/licenses/BSD-3-Clause>), the MIT License (<http://www.opensource.org/licenses/mit-license.php>), the Artistic License (<http://www.opensource.org/licenses/artistic-license-1.0>) and the Initial Developer's Public License Version 1.0 (<http://www.firebirdsql.org/en/initial-developer-s-public-license-version-1-0/>).

This product includes software copyright © 2003-2006 Joe Walnes, 2006-2007 XStream Committers. All rights reserved. Permissions and limitations regarding this software are subject to terms available at <http://xstream.codehaus.org/license.html>. This product includes software developed by the Indiana University Extreme! Lab. For further information please visit <http://www.extreme.indiana.edu/>.

This product includes software Copyright (c) 2013 Frank Balluffi and Markus Moeller. All rights reserved. Permissions and limitations regarding this software are subject to terms of the MIT license.

See patents at <https://www.informatica.com/legal/patents.html>.

DISCLAIMER: Informatica LLC provides this documentation "as is" without warranty of any kind, either express or implied, including, but not limited to, the implied warranties of noninfringement, merchantability, or use for a particular purpose. Informatica LLC does not warrant that this software or documentation is error free. The information provided in this software or documentation may include technical inaccuracies or typographical errors. The information in this software and documentation is subject to change at any time without notice.

NOTICES

This Informatica product (the "Software") includes certain drivers (the "DataDirect Drivers") from DataDirect Technologies, an operating company of Progress Software Corporation ("DataDirect") which are subject to the following terms and conditions:

1. THE DATADIRECT DRIVERS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT.
2. IN NO EVENT WILL DATADIRECT OR ITS THIRD PARTY SUPPLIERS BE LIABLE TO THE END-USER CUSTOMER FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR OTHER DAMAGES ARISING OUT OF THE USE OF THE ODBC DRIVERS, WHETHER OR NOT INFORMED OF THE POSSIBILITIES OF DAMAGES IN ADVANCE. THESE LIMITATIONS APPLY TO ALL CAUSES OF ACTION, INCLUDING, WITHOUT LIMITATION, BREACH OF CONTRACT, BREACH OF WARRANTY, NEGLIGENCE, STRICT LIABILITY, MISREPRESENTATION AND OTHER TORTS.

Publication Date: 2022-12-07

Table of Contents

Chapter 1: Introduction to Cloud Integration Hub.....	9
Cloud Integration Hub architecture.	11
System Requirements.	13
Cloud Integration Hub user interface	14
Hub Overview diagram.	15
Navigator.	16
Explore page.	16
Accessing Cloud Integration Hub.	17
Cloud Integration Hub topics.	17
Publication repositories.	18
Hosted and private publication repositories.	18
Hosted publication repository size specifications.	19
Relational and file store publication repositories.	20
Data Integration tasks	20
File ingestion tasks.	21
Cloud Integration Hub publications and subscriptions.	21
Publication and Subscription process.	21
Informatica Resources.	22
Informatica Network.	22
Informatica Knowledge Base.	22
Informatica Documentation.	23
Informatica Product Availability Matrices.	23
Informatica Velocity.	23
Informatica Marketplace.	23
Informatica Global Customer Support.	23
 Chapter 2: Hub administration.....	 24
Organization management.	24
Before you begin.	25
Setting up the organization in Informatica Cloud Integration Hub.	25
Editing organization settings in Cloud Integration Hub.	27
Salesforce Accelerator package deployment.	27
Before you begin.	28
Salesforce Accelerator package components.	28
Deploying the Salesforce Accelerator package.	29
Set up a private publication repository.	29
Bypass the Publication Repository Service in a private publication repository.	30
Changing the port number of the publication repository service.	31
Using a customized Java KeyStore with a private publication repository.	31
Enable zero downtime for a private publication repository.	32

Intermediate staging policy for subscriptions.	32
Setting up Cloud Integration Hub to show Data Integration Hub events.	32
Configure load balancer URL.	33
System Properties.	33

Chapter 3: Project and Asset Management..... 35

Viewing an asset.	35
Editing an asset.	36
Editing a topic.	36
Moving folders and assets.	36
Deleting projects, folders, and assets.	36
User roles.	37
Privileges.	38
Permissions.	39
Rules and guidelines for permissions.	40
Configuring permissions.	40
Asset migration.	42
Dependent assets.	42
Exporting assets.	42
Importing assets.	43
Migrating assets between organizations.	44
Migration error handling.	44
Asset dependencies.	45
Tags.	45
Creating tags.	46
Assigning tags.	46
Editing and deleting tags.	47

Chapter 4: Applications..... 49

Application management.	49
Creating an Application.	49
Adding a publication or a subscription to an existing application.	49
Application properties.	50

Chapter 5: Topics..... 51

Topic structure in a relational database.	51
Create topic tables in a relational database.	52
Using metadata files to create topic tables in a relational database.	52
Topic structure updates in a relational database.	54
Topic data retention.	56
Topic management.	56
Creating a topic.	56
Subscribing to a topic.	57

Topic properties.	58
Topic Diagram.	58
General Details properties.	59
Topic structure properties.	60
Publications properties.	63
Subscriptions properties.	63
Chapter 6: Data Integration tasks.	65
Data Integration Task Types.	65
Data Integration tasks rules and guidelines.	66
Synchronization tasks with Cloud Integration Hub.	67
Creating a synchronization task for a publication.	67
Creating a synchronization task for a subscription.	69
Mapping configuration tasks with Cloud Integration Hub.	71
Mapping task configuration process.	71
Creating a mapping and task for a publication.	71
Creating a mapping and task for a subscription.	72
Chapter 7: File ingestion tasks.	75
File ingestion tasks rules and guidelines.	75
Creating a file ingestion task for a publication.	76
Step 1. Define task details.	76
Step 2. Configure publication source.	76
Step 3. Configure publication target.	77
Creating a file ingestion task for a subscription.	77
Step 1. Define task details.	77
Step 2. Configure subscription source.	78
Step 3. Configure subscription target.	78
Step 4. Configure file processing actions (optional).	78
Chapter 8: Publications.	79
Publication methods.	79
Publication types.	80
Publication processes.	80
Publication process for publications that trigger Data Integration tasks.	80
Publication process for publications that trigger file ingestion tasks.	80
Publication process for publications that publish data with an API.	81
Publication mapping.	81
Publication sources.	81
Publication schedules.	81
Publication management.	82
Creating a publication that triggers a Data Integration task.	82
Creating a publication that triggers a file ingestion task.	83

Creating a publication that publishes data with an API.	84
Running a publication manually.	85
Disabling and enabling a publication.	85
Publication properties.	86

Chapter 9: Subscriptions..... 88

Subscription methods.	88
Subscription types.	89
Subscription processes.	89
Subscription process for subscriptions that trigger Data Integration tasks.	89
Subscription process for subscriptions that trigger file ingestion tasks.	90
Subscription process for subscriptions that consume data with an API.	90
Subscription mapping.	90
Subscription targets.	91
Subscription schedules.	91
Subscription retry policy.	91
Subscription management.	92
Creating a subscription that triggers a Data Integration task.	92
Creating a subscription that triggers a file ingestion task.	94
Creating a subscription that consumes data with an API.	95
Running a subscription manually.	96
Getting previous publications for a subscription.	97
Disabling and enabling a subscription.	97
Subscription properties.	97

Chapter 10: Tracking and monitoring..... 101

Publication and Subscription Events.	101
Event Types.	102
Event statuses.	102
Event Consumption Statuses.	103
Event History.	103
Event Session Log.	104
Event Processing Information.	104
System Event Maintenance Report.	104
Event Filters.	104
Managing Events.	104
Event Properties.	106
Open Data Integration Hub assets from Cloud Integration Hub.	107
Event Monitors.	107
Monitoring Rules.	107
Managing Monitoring Rules.	107
Monitoring Rule Properties.	109

Chapter 11: Cloud Integration Hub REST APIs.....	111
Authorization Header.	112
Run Publication Subscription REST API	112
Run Publication Subscription REST API Request.	112
Run Publication Subscription REST API Action Response.	114
Publish Data REST API.	114
Publish Data REST API Action Response.	115
Topic Swagger Structure for Publish Data REST API.	115
Consume Data REST API.	116
Consume Data REST API request.	116
Consume Data REST API action response.	117
Topic Swagger Structure for Consume Data REST API.	118
Change Publication Subscription Mode REST API.	118
Change Publication Subscription Mode REST API Action Response.	119
Reprocess Event REST API.	120
Reprocess Event REST API Action Response.	120
Event Status REST API.	121
Event Status API Response.	121
Events REST API.	123
Events REST API Request.	123
Events REST API Response.	126
Cloud Integration Hub Catalog REST API.	130
Cloud Integration Hub Catalog API Response.	131
 Chapter 12: Glossary.....	 136
 Index.	 138

CHAPTER 1

Introduction to Cloud Integration Hub

Cloud Integration Hub is a cloud-based application integration solution that you can use to share and synchronize data between different applications in an organization.

To publish data to Cloud Integration Hub, first define the data set that you want to manage, for example, sales, customers, or orders. You define a data set by defining a topic. A topic defines the structure of the data that Cloud Integration Hub stores in the publication repository and the type of publication repository where data is stored. You can manage multiple topics that represent different data sets in Cloud Integration Hub. Applications publish data to topics and subscribe to data sets that are represented by topics.

Multiple applications can publish to the same topic, for example, different ordering applications can publish their orders to the same Orders topic. Multiple subscribers can consume the data from a topic. Different subscribing applications can consume the data in different formats and in different latencies based on a defined schedule.

Cloud Integration Hub stores the data that applications publish to topics in the Cloud Integration Hub publication repository in the following ways:

- For each publication instance, the retention period for consumed data starts if all the subscribers have either successfully consumed or discarded the data. That is, after all the events that are associated with the publication instance are either in a Complete or in a Discarded event status. If all the subscribers consume or discard the data, Cloud Integration Hub stores the consumed data in the publication repository until the retention period for consumed data expires, and then deletes the consumed data from the publication repository.
- Cloud Integration Hub stores unconsumed data in the publication repository until the retention period for unconsumed data expires, and then deletes the unconsumed data from the publication repository.

Applications can use PowerExchange® adapters and Informatica Intelligent Cloud Services™ connectors to share data from different sources, such as database tables, files, or any sources that Informatica supports. Each application can be a publisher and a subscriber to different topics.

Publications publish to a specific topic. A publication defines the data source type and the location from where Cloud Integration Hub retrieves the data that the application publishes. Subscriptions subscribe to one or more topics. A subscription defines the data target type and the location in the subscribing application to where Cloud Integration Hub sends the published data.

Examples

Your organization uses multiple applications. Some of the applications are located on-premises and some are located on the cloud. Your applications require the following data:

Marketing application

Requires data about campaigns, accounts, contracts, and employees for operational purposes.

Data warehouse

Requires data about campaigns and contracts for analytical purposes.

Business Intelligence (BI) application

Requires data about campaigns and orders for analytical purposes.

Customer relationship management (CRM) application

Requires data about sales department employees, including sales representatives, for operational purposes.

With Cloud Integration Hub, you can address the following use-cases:

Share daily accounts data.

You can share the daily account updates from the CRM application with the marketing application, as follows:

1. Create an Accounts topic.
2. Define a publication that publishes account details from the CRM application to the Accounts topic and set the schedule to publish the data daily.
3. Define a subscription from the marketing application to the Accounts topic and set the subscription to consume the published data when it is available in Cloud Integration Hub.

Share campaign details as required.

You can share the campaign details from the CRM application with the marketing, data warehouse, and CRM applications at varying schedules, as follows:

1. Create a Campaigns topic.
2. Define a publication that publishes campaign details from the CRM application to the Campaigns topic and set the schedule to publish the data daily.
3. Define a subscription from the marketing application to the Campaigns topic, and set the schedule to consume the data when it is published.
4. Define a subscription from the data warehouse application to the Campaigns topic, and set the schedule to consume the data twice a week.
5. Define a subscription from the BI application to the Campaigns topic, and set the schedule to consume the data once a week.

Share weekly contract details.

You can share the weekly contract details from the CRM application with the marketing and data warehouse applications, as follows:

1. Create a Contracts topic.
2. Define a publication that publishes contract details from the CRM application to the Contracts topic and set the schedule to publish the data weekly.
3. Define a subscription from the marketing application to the Contracts topic, and set the schedule to consume the data when it is published.
4. Define a subscription from the data warehouse application to the Contracts topic, and set the schedule to consume the data when it is published.

Share bi weekly orders data.

You can share the daily order updates from the CRM application with the marketing application, as follows:

1. Create an Orders topic.

2. Define a publication that publishes order details from the CRM application to the Orders topic and set the schedule to publish the data every two weeks on the last day of the week.
3. Define a subscription from the BI application to the Orders topic and set the subscription to consume the published data when it is available in Cloud Integration Hub.

Share monthly employee details.

You can share the monthly employee details from the HR application with the CRM application, as follows:

1. Create an Employees topic.
2. Define a publication that publishes employee details from the HR application to the Employees topic and set the schedule to publish monthly, on the first day of the month.
3. Define a subscription from the CRM application to the Employees topic, and filter the subscription to consume data pertaining to sales department employees only. Set the subscription schedule to consume the data when it is published.

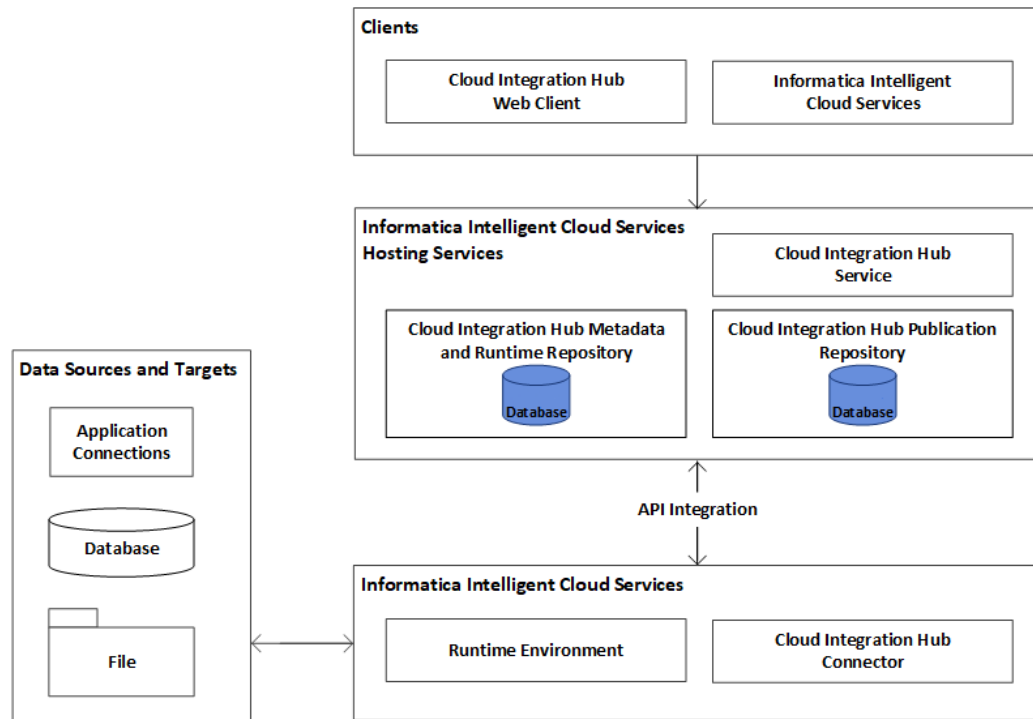
Cloud Integration Hub architecture

The Cloud Integration Hub environment consists of user interface clients, Cloud Integration Hub service and repositories that are hosted on Informatica Intelligent Cloud Services Hosting Services, and the Informatica Intelligent Cloud Services Secure Agent and Cloud Integration Hub connector that are located on Informatica Intelligent Cloud Services.

You can select to host the publication repository for the organization on-premises or on a private cloud. In that case, the repository is not hosted on Informatica Intelligent Cloud Services Hosting Services but is installed and managed by the organization.

The following image shows the Cloud Integration Hub components:

Figure 1.



Cloud Integration Hub contains the following components:

Cloud Integration Hub Web client

User interface to manage applications, topics, publications, and subscriptions, and to monitor publications, subscriptions, and events. Administrators use the Web client to create the organization in Cloud Integration Hub.

Informatica Intelligent Cloud Services user interfaces

User interfaces to define sources and targets and to create connections, mappings, and tasks.

Informatica Intelligent Cloud Services Hosting Services

Services that host the Cloud Integration Hub service and repositories. The services stores all task and organization information.

Cloud Integration Hub service

A service that manages publication and subscription processing in Cloud Integration Hub.

Cloud Integration Hub metadata and runtime repository

Database that stores metadata and runtime data for Cloud Integration Hub applications, topics, publications, subscriptions, and events.

Publication repository

Where Cloud Integration Hub stores published data or files until the retention period for the data or files expires. You can use the following types of publication repositories:

- **Relational.** Stores relational data. You can use a relational publication repository on a hosted or a private repository.

- File store. Stores files as-is. You can use a file store publication repository with an on-premises Secure Agent.

Data sources and targets

Sources and targets that you use to publish and consume data. You can use the following types of sources and targets:

- Database. Tables and columns.
- File. Binary, text, or unstructured files.
- Applications connections. Connection objects for applications.

System Requirements

The following table describes the minimum system requirements for Cloud Integration Hub.

Verify that the system meets the requirements that are applicable for the setup of the organization.

Component/Use Case	Minimum Requirement
Informatica Intelligent Cloud Services Secure Agent	<ul style="list-style-type: none"> - 8 GB memory - Two CPU cores
Network between the Secure Agent and the private publication repository database	Ping latency of less than 10 ms

Component/Use Case	Minimum Requirement
Access via a proxy gateway	<p>The following URL is accessible from the machine where the Secure Agent is installed:</p> <pre>https://<pod>.<baseUrl>/</pre> <p>Where:</p> <ul style="list-style-type: none"> - <pod> is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: <code>cih-pod1</code>, or <code>emw11-cih</code>. - <baseUrl> is the Informatica Intelligent Cloud Services URL. For example: <code>dm-us.informaticacloud.com/</code>. <p>For example:</p> <pre>https://cih-pod1.dm-us.informaticacloud.com/</pre> <p>Tip: You can copy the values of <pod> and <baseUrl> from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.</p> <p>Informatica recommends that you add the URL to the allowlist of the proxy server.</p>
Cloud Integration Hub private publication repository	<p>You can use one of the following database systems:</p> <ul style="list-style-type: none"> - Oracle - Microsoft SQL Server - MySQL <p>Note: For more information about supported editions and versions, see the Product Availability Matrix (PAM).</p> <p>The Cloud Integration Hub private publication repository requires at least 512 MB of disk space for the publication repository database based on the number of publications and publication instances that you need to retain.</p> <p>Note: Unicode data requires twice as much storage than single-byte character sets.</p> <p>Multiple database connections for the private publication repository must always be available. The number of required connections depends on the number of publications and subscriptions that run concurrently. Use the following formula to calculate the number of required database connections:</p> $\text{NumberOf ConcurrentPublicationsOrSubscriptions} \times 3 + 2$ <p>If you do not have enough database connections available, Cloud Integration Hub might fail or encounter database deadlocks.</p>

For more information about system requirements, see the Product Availability Matrix (PAM) for Informatica Intelligent Cloud Services. PAMs indicate the versions of operating systems, databases, and other types of data sources and targets that a product release supports. You can access the PAMs on the Informatica Network at <https://network.informatica.com/community/informatica-network/product-availability-matrices/>.

Cloud Integration Hub user interface

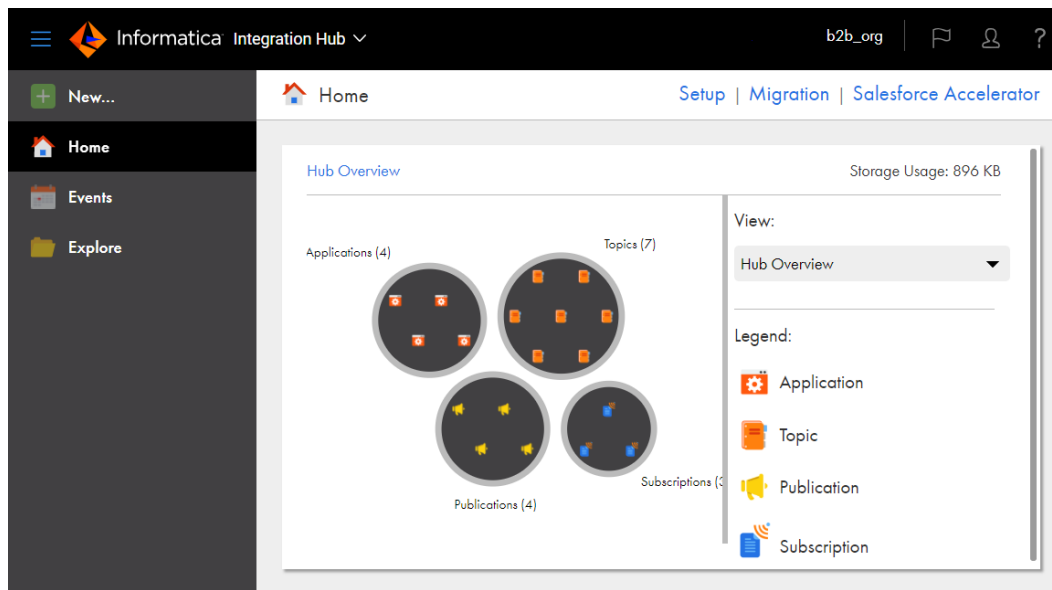
The Cloud Integration Hub Home page includes a navigator at the left of the page, the Hub Overview diagram, and filters at the right of the page. The Cloud Integration Hub Home page appears when you log in to Cloud Integration Hub.

Use the navigator to create assets, track events, and explore and perform actions on existing assets.

The Hub Overview diagram provides a visual overview of the existing assets. Use the **View** filter to filter the assets that the Hub Overview diagram shows.

If you use a hosted publication repository, the repository storage usage shows at the top right of the Hub Overview diagram.

The following image shows a sample Hub Overview diagram:



Hub Overview diagram

The Cloud Integration Hub Overview page shows the Hub Overview diagram when Cloud Integration Hub contains assets such as applications, topics, publications, or subscriptions.

The Hub Overview diagram provides a visual overview of the existing assets, grouped into categories.

When you rest on an asset in the diagram, all related assets are highlighted. For example, when you rest on a topic, the applications and the publications that publish to the topic and the subscriptions that subscribe to the topic are highlighted. When you click an asset, a drill down view of the asset and its relations to other assets appears. For example, when you click a publication, the drill down view shows the publishing application, the topic to which the publication publishes data, and the subscriptions that subscribe to the topic.

When you right-click an asset in the drill down view, an action menu opens. You can perform the following actions from the menu, based on the asset type:

- View. Applies to all assets. Opens the asset in view mode.
- Run. Applies to publications and subscriptions. Runs the publication or the subscription.

Filters

You can filter the Hub Overview diagram to the following views:





- Process errors. Entities with current error events.
- Non-valid entities. Entities that are not valid.
- Topics with no publications. Topics with no associated publications.
- Topics with no subscriptions. Topics with no associated subscriptions.
- Most used topics. Three most used topics, based on the number of publications and subscriptions that use the topic.

When you filter the diagram, entities that are not relevant to the selected filter appear in view only mode.

Navigator

Use the navigator to create assets, track events, and explore and perform actions on existing assets.

The following table lists the navigator icons and describes the functions that they perform:

Icon	Name	Function
	New	Create a new asset: application, publication, subscription, topic, or monitoring rule.
	Home	Go to the Overview page.
	Events	Go to the Events page.
	Explore	Explore existing assets and perform actions on existing assets.

Explore page

Use the **Explore** page to work with your Informatica Intelligent Cloud Services projects and assets.

Finding projects and assets on the Explore page

Use any of the following methods to find your projects and assets on the **Explore** page:

- Explore by projects and folders. View all projects or select a particular project.
- Explore by asset types. View all assets or view assets of a particular type.
- Explore by tags. View assets associated with a particular tag.
- Search for projects or assets. To search all projects, folders, and assets in the organization, view the **Explore** page by **All Projects**, and then enter a name or description in the Find box. Or, to narrow your search, filter the **Explore** page by **All Assets**, select a specific asset type, project, or folder, and then enter a name or description in the Find box.
- Sort the search results. Sort the **Explore** page by name, last update date, description, or type. When you sort by type, the **Explore** page groups assets by asset type. It does not list the asset types in alphabetical order.

You can see projects, folders, and assets for all of the services that you use. If you select an asset to open it or perform an action and the asset is created in a different service than the one you have open, the service opens in a new browser tab.

You can't use the following characters:

? ' | { } " ^ & [] / \

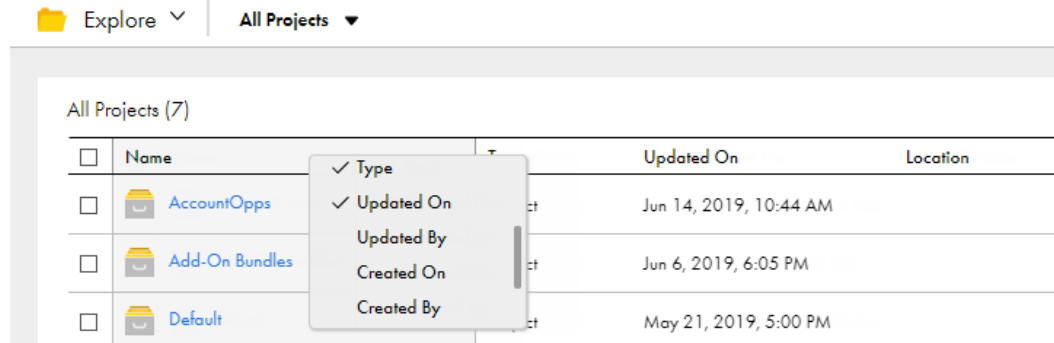
Do not use these characters in project, folder, asset, or tag names.

Working with projects and assets on the Explore page

Perform actions on projects, folders, and assets on the **Explore** page. To see what actions you can perform on an object, in the row that contains the object, click the **Actions** icon. The Actions menu lists the actions you can perform based on your user role privileges.

Customizing the Explore page

You can display, hide, or rearrange object properties on the **Explore** page. To display or hide properties, right-click the column heading area and check or uncheck the properties. The following image shows the properties menu on the **Explore** page column heading area:



To rearrange columns, click a column heading and drag it to a different location.

Accessing Cloud Integration Hub

Access Cloud Integration Hub through the Informatica Intelligent Cloud Services, from the **My Services** page.

Before you access Cloud Integration Hub for the first time, the administrator sets up the organization in Informatica Intelligent Cloud Services and then sets up the organization in Cloud Integration Hub. If, when you access Cloud Integration Hub for the first time, the **Organization Cloud Setup** dialog box shows, it is an indication that your administrator did not perform the process of provisioning the organization to the hub. Contact your administrator or follow the instructions on the screen. For details, see [“Organization management” on page 24](#).

1. On the Informatica Intelligent Cloud Services login page, enter your Informatica Intelligent Cloud Services user name and password.
2. Click **Log In**.

The Informatica Intelligent Cloud Services **My Services** page appears.

3. Select **Integration Hub**.

The Cloud Integration Hub application appears.

Note: The **Integration Hub** link appears on the **My Services** page if your organization has the required licences. If the link doesn't appear on the **My Services** page, contact your administrator.

Cloud Integration Hub topics

A Cloud Integration Hub topic is an entity that represents a data domain that is published and consumed in Cloud Integration Hub.

You can create topics that store processed data in a relational database or topics that store files as-is in a file store. In a relational database, topics define the canonical data structure and additional data definitions such as the data retention period. In a file store, topics define the data retention period.

Publication repositories

Cloud Integration Hub stores topic data and files in a publication repository, according to the type of publication and subscription tasks that you need to run.

The publication repository stores the data or files for a short intermediate period after the data or files are consumed by all subscribers.

Cloud Integration Hub stores the data and files in the publication repository in the following ways:

- For each publication instance, the retention period for consumed data or files starts if all the subscribers have either successfully consumed or discarded the data or files. That is, after all the events that are associated with the publication instance are either in a Complete or in a Discarded event status. If all the subscribers consume or discard the data or files, Cloud Integration Hub stores the consumed data or files in the publication repository until the retention period for consumed data or files expires, and then deletes the consumed data or files from the publication repository.
- Cloud Integration Hub stores unconsumed data or files in the publication repository until the retention period for unconsumed data or files expires, and then deletes the unconsumed data or files from the publication repository.

Hosted and private publication repositories

Choose to use a hosted or a private publication repository based on the functionality that your organization requires.

Hosted publication repository

Cloud Integration Hub hosts and manages the publication repository on Informatica Intelligent Cloud Services Hosting Services. Storage usage of the repository shows on the Cloud Integration Hub home page. Use a hosted publication repository to host a relational database. If you use an on-premises Secure Agent, you can use the hosted publication repository to host a file store publication repository.

Private publication repository

Use your own, private repository. A private publication repository can reside on-premises or on the organization's private cloud. You can use a private publication repository to host a relational database. If the repository resides on-premises, you can use it to host a file store publication repository. For more information about setting up a private publication repository, see [“Set up a private publication repository” on page 29](#).

The following table describes some of the differences in functionality between hosted and private publication repositories:

Functionality	Description	Hosted	Private
Cloud Integration Hub processor license	License required to access Cloud Integration Hub processor	Not required	Required
Data encryption	Process of encrypting data	Enabled	Not applicable
Staging database type	Type of database used for data processing operations	Not applicable	User choice

Functionality	Description	Hosted	Private
Zero downtime	Environment that ensures uninterrupted access to data	Enabled	Enabled for publications and subscriptions that trigger a Data Integration task
Informatica Intelligent Cloud Services Hosted Secure Agent	Secure Agent hosted by Informatica Intelligent Cloud Services	Accessible	Not applicable

Hosted publication repository size specifications

If you use a hosted publication repository when you create topic tables, verify that they don't exceed the maximum allowed row and storage sizes. If topic tables exceed the maximum allowed row or storage size on a hosted publication repository, topic table creation fails.

Maximum row size

Row size is the total length of all the columns in a topic table and is dependent on the character set.

The maximum allowed row size for a topic table is 8,000 bytes.

Maximum storage size

The storage requirements for individual columns determine how many columns fit within the maximum storage size for a topic table.

The maximum allowed storage size for a topic table row is 64,000 bytes, and the total allowed storage size is 25 GB.

Refer to the following table for the row and storage sizes that the columns occupy, based on their data type:

Data Type	Row Size in Bytes	Storage Size in Bytes
String, precision lower than or equal to 85	For data that contains ASCII characters: Precision + 1 For data that contains Unicode characters: (Precision × 3) + 2	(Precision × 3) + 2
String, precision higher than 85	22	(Precision × 3) + 2
Decimal	20	20
Double	8	8
INT32	4	4
INT64	8	8
Text	22	12
Date time	8	8

An encrypted column in a topic table occupies a row size of 22 bytes and a storage size of 12 bytes.

For example, if a topic table contains 100 columns of data type string with a precision of 80, then it exceeds the maximum allowed row size of 8,000 bytes. If a topic table contains 84 columns of data type string with a precision of 255, then it exceeds the allowed storage size of 64,000 bytes.

Refer to the following table for the row sizes and storage sizes that the mandatory topic table fields occupy:

Field	Row Size in Bytes	Storage Size in Bytes
PUBLICATION_DATE_FIELD_TYPE	8	8
PUBLICATION_INSTANCE_ID	20	20

The maximum allowed response size for API-based subscriptions is 5 MB.

Relational and file store publication repositories

You can use a relational publication repository to store relational data. If you use an on-premises publication repository, you can also use a file store publication repository to store files as-is.

Relational publication repository

Cloud Integration Hub stores relational data for the subscribers to consume in a relational publication repository. Use a relational publication repository if you use publications and subscriptions that run Data Integration tasks or publications and subscriptions that publish data with an API. You can use a relational publication repository on a hosted repository.

File store publication repository

Cloud Integration Hub stores files as-is for the subscribers to consume in a file store publication repository. Use a file store publication repository if you use publications and subscriptions that run file ingestion tasks. You can use a file store publication repository with an on-premises Secure Agent. To use a file store publication repository, your organization must have a valid Mass Ingestion Files license. For more information, contact Informatica Global Customer Support.

Data Integration tasks

Cloud Integration Hub uses Data Integration tasks to publish data from source applications to the Cloud Integration Hub publication repository and to consume data from the publication repository by target applications.

You develop Data Integration tasks for Cloud Integration Hub in the same way that you develop other Data Integration tasks. You use the Cloud Integration Hub connection as the target in publication tasks and as the source in subscription tasks.

For more information, see [“Data Integration tasks rules and guidelines” on page 66](#).

File ingestion tasks

Cloud Integration Hub uses file ingestion tasks to publish files as-is to the file store publication repository and to consume files as-is from the file store publication repository.

You develop file ingestion tasks for Cloud Integration Hub in a similar way that you develop other file ingestion tasks in Mass Ingestion.

For more information, see [“File ingestion tasks rules and guidelines” on page 75](#).

Cloud Integration Hub publications and subscriptions

Publications and subscriptions are entities that define how applications publish data and files to Cloud Integration Hub and how applications consume data and files from Cloud Integration Hub. Publications publish data and files to a defined topic and subscriptions subscribe to topics.

Publications and subscriptions control the data and file flow and the schedule of data and file publication or data and file consumption. An application can be a publisher and a subscriber. Multiple applications can publish to the same topic. Multiple applications can consume data or files from the same topic.

Publications and subscriptions can publish from and subscribe to any type of source and target that Informatica Intelligent Cloud Services supports. You can publish from and subscribe to different sources of data and files. Because the publishing process and the consuming process are completely decoupled, the publishing source and the consuming target do not have to be of the same data or file type. For example, you can publish data from a file and consume it into a database.

Publications and subscriptions can publish and consume data or files by triggering a Data Integration task, by triggering a file ingestion task, or with an API. For publications and subscriptions that trigger a Data Integration task, you create the tasks in Data Integration. You then select a Data Integration task when you create the publication or subscription in Cloud Integration Hub. For publications and subscriptions that trigger a file ingestion task, you create the tasks in Mass Ingestion. You then select a file ingestion task when you create the publication or subscription in Cloud Integration Hub. For publications and subscriptions that are triggered by an API, you run the API manually.

Publication and Subscription process

The publication process starts on the schedule that you define in the publication, when an external process triggers the publication, or when you manually run the publication.

When data transfer is complete, the topic data set is ready for subscribers. The subscription process starts when one of the following conditions exist, based on the configuration of data consumption in the subscriptions:

- When data is published to the topic.
- When all publishers that publish to the topic finish publishing.

If the topic to which the data is published has subscribers, Cloud Integration Hub triggers a Cloud Integration Hub subscription workflow for each subscriber, to consume the data.

Cloud Integration Hub generates events to track the progress of the publication and subscription process. When an application publishes data, Cloud Integration Hub creates a parent publication event. When the publication process ends and the published data is ready to consume, Cloud Integration Hub generates a child event for each subscription.

The events change status as the publication and subscription process progresses, and reach a completed status after the process ends successfully. You also use events to monitor and troubleshoot issues that might occur during the process.

During the publication or the subscription process, Cloud Integration Hub communicates with Informatica Intelligent Cloud Services, going through the following stages:

- When a cloud application publishes a data set, the Cloud Integration Hub server triggers the Data Integration task that is defined for the publication through an Informatica Intelligent Cloud Services REST API.
- For cloud publications, the target is defined using a Cloud Integration Hub cloud connector. The publication process uses the connector to write the data to Cloud Integration Hub.
- If the topic to which the data is published has subscribers, Cloud Integration Hub triggers the subscription workflows to consume the data.
- For cloud subscriptions, the source is defined using a Cloud Integration Hub cloud connector. The subscription process uses the connector to read data from Cloud Integration Hub.
- Cloud Integration Hub monitors the task for processing status.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Network

The Informatica Network is the gateway to many resources, including the Informatica Knowledge Base and Informatica Global Customer Support. To enter the Informatica Network, visit <https://network.informatica.com>.

As an Informatica Network member, you have the following options:

- Search the Knowledge Base for product resources.
- View product availability information.
- Create and review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit <https://search.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit <https://docs.informatica.com>.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Product Availability Matrices

Product Availability Matrices (PAMs) indicate the versions of the operating systems, databases, and types of data sources and targets that a product release supports. You can browse the Informatica PAMs at <https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services and based on real-world experiences from hundreds of data management projects. Informatica Velocity represents the collective knowledge of Informatica consultants who work with organizations around the world to plan, develop, deploy, and maintain successful data management solutions.

You can find Informatica Velocity resources at <http://velocity.informatica.com>. If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that extend and enhance your Informatica implementations. Leverage any of the hundreds of solutions from Informatica developers and partners on the Marketplace to improve your productivity and speed up time to implementation on your projects. You can find the Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

You can contact a Global Support Center by telephone or through the Informatica Network.

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:

<https://www.informatica.com/services-and-training/customer-success-services/contact-us.html>.

To find online support resources on the Informatica Network, visit <https://network.informatica.com> and select the eSupport option.

CHAPTER 2

Hub administration

Before the organization can use Cloud Integration Hub, you must set up an organization in Informatica Intelligent Cloud Services and then set up the organization in Cloud Integration Hub.

After you set up the organization in Informatica Intelligent Cloud Services, you can perform one or more of the following tasks:

- Deploy the Cloud Integration Hub Salesforce Accelerator package for rapid synchronization of data from Salesforce to other applications through Cloud Integration Hub. Deploying the package creates the components that are required to connect the Salesforce application to Cloud Integration Hub. Some of the components are created in Cloud Integration Hub and some are created in Informatica Intelligent Cloud Services.
- Set up a private publication repository to store topic data.
- Enable zero downtime for a private publication repository to ensure uninterrupted access to data.
- Modify the policy for writing data to intermediate staging in subscription flows.
- To view Data Integration Hub publication and subscription events in Cloud Integration Hub, configure Cloud Integration Hub system properties.
- Configure an external load balancer URL as the base API URL of publications and subscriptions that publish and consume data with an API to a private publication repository.

Organization management

Before the organization can use Cloud Integration Hub, you must set up an organization in Informatica Intelligent Cloud Services and then set up the organization in Cloud Integration Hub.

When you set up the organization in Cloud Integration Hub, Cloud Integration Hub creates the connection **Cloud Integration Hub** in Informatica Intelligent Cloud Services.

Warning: Do not rename the connection. The only connection property that you can change is the option **Do not use intermediate staging for subscription flows**. For more details, see [“Intermediate staging policy for subscriptions” on page 32](#).

Editing other connection properties or renaming the connection might result in errors at run time.

If you select to use a hosted publication repository, Cloud Integration Hub creates the Cloud Integration Hub publication repository on Informatica Intelligent Cloud Services Hosting Services.

Before you begin

Before you set up the organization in Cloud Integration Hub verify that the following conditions exist in Informatica Intelligent Cloud Services.

Before you can set up an organization in Cloud Integration Hub you must set up the organization in Informatica Intelligent Cloud Services. For details about setting up an organization, see the *Informatica Intelligent Cloud Services Administrator* help.

Configuration

From the **Configure** menu, under **Runtime Environments**, verify that the Secure Agent is running.

Administration

From the **Administer** menu, under **Licenses**, verify that following conditions exist:

License Category	Required Condition
REST API license	Maximum Concurrent Sessions is set to a high value, for example, 100 sessions.
Connector license	A valid Cloud Integration Hub connector.

Proxy Settings

If your organization uses an outgoing proxy server to connect to the internet, set the following JVM options on the Secure Agent:

Name	Value
JVMOption1	-Dhttp.proxyHost=<proxy host>
JVMOption2	-Dhttp.proxyPort=<proxy port>
JVMOption3	-Dhttp.useProxy=true
JVMOption4	-Dhttp.proxyUser=<proxy user name>
JVMOption5	-Dhttp.proxyPassword=<proxy password>

After the Secure Agent restarts, check the agent core log file to verify that the correct proxy server is used. The agent core log file is the following file:

```
<Secure Agent installation directory>\apps\agentcore\agentcore.log
```

To find the proxy information, search for "proxy" in the log file.

Setting up the organization in Informatica Cloud Integration Hub

Set up the organization in Cloud Integration Hub to create and manage Cloud Integration Hub assets. You must define the settings required by the organization.

1. Access Cloud Integration Hub and accept the license agreement.

The **Organization Cloud Setup** dialog box appears.

- Define the following settings:

Property	Description
Organization Name	Name of the organization in Informatica Intelligent Cloud Services. Appears in view only mode.
Organization ID	ID of the organization in Informatica Intelligent Cloud Services. Appears in view only mode.
Informatica Cloud User	Name of the Informatica Intelligent Cloud Services user to use at run time. The user must have an Admin user role in Informatica Intelligent Cloud Services. The user must also have Update and Execute permissions for the Data Integration assets associated with publications and subscriptions.
Informatica Cloud Password	Password for the Informatica Intelligent Cloud Services user to use at run time.
Runtime Environment	Informatica Intelligent Cloud Services Secure Agent runtime environment to use at run time.
Rotate Key	Click Rotate Key to rotate the encryption key used for data encryption.

- To use a private publication repository, define the following settings:

Property	Description
Repository Type	Select Private .
Repository Database Type	Choose Oracle, Microsoft SQL Server, or a MySQL database. If you use an SQL Server database, the collation is case-insensitive.
Repository URL	JDBC URL of the repository, based on the database type: <ul style="list-style-type: none"> - Oracle: jdbc:informatica:oracle://<ip>:<port>;sid=<sid>; - Microsoft SQL Server: jdbc:informatica:sqlserver://<ip>:<port>; DatabaseName=<DatabaseName>; - MySQL: jdbc:mysql://<ip>:<port>/<DatabaseName>; To prevent failure while establishing a connection to the proxy service, add the following string to the repository URL: useSystemProxyOptions=false;
User	Name of the user to access the repository.
User Role	Role granted to the user to access the repository, based on the database type: <ul style="list-style-type: none"> - On an Oracle database, the user must be granted CONNECT and RESOURCE roles. - On a Microsoft SQL Server database, the user must be granted db_datareader, db_datawriter, and db_ddladmin roles.
Password	Password of the user.
Repository Schema	Schema used with the repository. For more information, see “Bypass the Publication Repository Service in a private publication repository” on page 30 .

4. Optionally, define the following file store publication repository settings:

Property	Description
Enable File Store	If you use an on-premises Secure Agent, select to enable the hub to store published files as-is in a file store publication repository.
File Store Repository Location	File directory path of the file store publication repository. The Informatica Cloud Secure Agent must be able to access the location.

5. Click **Save**.

Editing organization settings in Cloud Integration Hub

Edit the organization settings in Cloud Integration Hub.

1. Access Cloud Integration Hub.
2. Click the **Setup** link in the upper right corner of the page.
The **Setup** page appears.
3. Edit the required settings and then click **Save**.

Note:

- When you change the publication repository hosting option, for example, from a hosted repository to a private repository, Cloud Integration Hub deletes all data from the current publication repository. Subscribers can no longer consume the data that existed in the publication repository before the change.
- If you change the Organization's user credentials (password), you must update it in the CIH Settings page.

Salesforce Accelerator package deployment

For rapid synchronization of data from Salesforce to other applications through Cloud Integration Hub deploy the Cloud Integration Hub Salesforce Accelerator package.

The package includes components required to connect the Salesforce application to Cloud Integration Hub, including the following components:

- Publishing and subscribing applications.
- The topic to which to publish and from which to subscribe.
- Publication and subscription, including Informatica Intelligent Cloud Services mappings and tasks.

Some of the components are created in Cloud Integration Hub and some are created in Informatica Intelligent Cloud Services.

After you deploy the package, you can use the Salesforce Accelerator components to publish the Contacts, Accounts, and Opportunities tables from Salesforce to the topic in the hub and use the sample subscribing application to consume the data and write it to a file.

Before you begin

Before you deploy the Cloud Integration Hub Salesforce Accelerator package create the following connections in Informatica Intelligent Cloud Services:

- **CIH_Salesforce.** A connection to the organization's Salesforce cloud application.
- **CIH_FF_target.** A connection that the sample subscription mapping uses as the target where Cloud Integration Hub places the consumed data in flat file format.

Note: You must name the connections **CIH_Salesforce** and **CIH_FF_target**.

In addition, verify that the organization's Salesforce cloud application includes the tables Accounts, Contacts, and Opportunities, and that the Cloud Integration Hub user has privileges to read the tables.

Salesforce Accelerator package components

When you deploy the Salesforce Accelerator package, Cloud Integration Hub creates the following components in Cloud Integration Hub and in Informatica Intelligent Cloud Services:

Component	Description	Deployed To
CIH_Salesforce	An Cloud Integration Hub application that represents the organization's Salesforce cloud application. This is the publishing application.	Cloud Integration Hub.
Salesforce	An Cloud Integration Hub topic that includes the following Salesforce tables: Accounts, Contacts, Opportunities. Note: If you use a hosted publication repository, by default, all topic fields are encrypted. After you deploy the package, you can edit the topic and turn off the encryption for specific columns. For example, for columns you plan to use as filters in your mappings.	Cloud Integration Hub.
Pub_sfcdc_Acct_Contact_Opp	Publication from the Salesforce application, from the Accounts, Contacts, and Opportunities tables to the Salesforce topic. The publication schedule is set to the option Manually or by an external trigger . If required, you can change the publication scheduling option in Cloud Integration Hub.	Cloud Integration Hub.
Sub_app	A sample Cloud Integration Hub application that subscribes to the Salesforce topic.	Cloud Integration Hub
Sub_sfcdc_Acct_Contact_Opp	Subscription to the Salesforce topic that reads the Accounts, Contacts, and Opportunities tables and writes them into a flat file, based the definition of the target connection. The subscription schedule is set to the option When published data is ready . If required, you can change the subscription scheduling option in Cloud Integration Hub.	Cloud Integration Hub

Component	Description	Deployed To
cih_pub_Account_Contact_Opportunity	An Informatica Intelligent Cloud Services mapping that publishes data from Salesforce to Cloud Integration Hub.	Informatica Intelligent Cloud Services
mct_CIH_pub_Account_Contact_Opportunity	A mapping task that publishes data from the Salesforce application to Cloud Integration Hub.	Informatica Intelligent Cloud Services
cih_sub_Account_Contact_Opportunity	An Informatica Intelligent Cloud Services mapping that consumes data from the Cloud Integration Hub Salesforce topic to the flat file target.	Informatica Intelligent Cloud Services
mct_CIH_sub_Account_Contact_Opportunity	A mapping task that consumes data from the Cloud Integration Hub Salesforce topic to the flat file connection.	Informatica Intelligent Cloud Services

Note: If any of the Salesforce Accelerator package components exist in Informatica Intelligent Cloud Services or in Cloud Integration Hub, the deploy operation fails.

Deploying the Salesforce Accelerator package

Deploy the Salesforce Accelerator package to Cloud Integration Hub and Informatica Intelligent Cloud Services.

1. Click the **Salesforce Accelerator** link in the upper right corner of the screen.
2. Click **Yes** in the confirmation message.

Set up a private publication repository

You can set up a private publication repository to store topic data on-premises or on your organization's private cloud. If you use a private publication repository, verify the following requirements:

Database

A private publication repository must reside on an Oracle, Microsoft SQL Server, or MySQL database. The repository must be accessible through the Informatica Intelligent Cloud Services Secure Agent. To optimize performance, set up the Secure Agent and the private repository on the same machine.

When you use a private publication repository, Cloud Integration Hub doesn't encrypt data. For more information about publication repository functionality, see ["Hosted and private publication repositories" on page 18](#).

Note: If the private repository resides on a MySQL database, and the MySQL server and the Secure Agent use different time zones, add the property `connectionTimeZone` to the connection string to define a common time zone.

For example, configure the connection string as follows:

```
jdbc:mysql://<HOST_NAME>:<port>/<DB_NAME>?connectionTimeZone=<AGENT_TIMEZONE>
```

where `<AGENT_TIMEZONE>` is the time zone of the node where the Secure Agent is installed.

The following table lists the databases that you can use to set up a private publication repository:

Database	Versions	Edition	System
Oracle	19c, 18c, 12cR2, 12c	Enterprise	64-bit
Microsoft SQL Server	2019, 2017, Azure SQL Database	Enterprise	-
MySQL	5.7	-	InnoDB

Database user account

Verify that you have the user name and password for the required database user account that you create. The database user account must have privileges to perform the following actions:

- Select data from tables and views.
- Insert data into tables, delete data from tables, and update data in tables.
- Create, change, and delete the following elements:
 - Tables
 - Views
 - Synonyms
 - Indexes
 - Custom data types
 - Triggers
- Create, change, delete, and run stored procedures and functions.

If you use a Microsoft SQL Server database, consider granting database owner privileges to the database user account.

Language support

You can enable UTF-8 character encoding on Oracle Database. Configure the database to use the character set `<AMERICAN_AMERICA.AL32UTF8>`.

If you use a Linux operating system on the Secure Agent machine, configure it to use the character set `<EN_US.UTF8>`.

Bypass the Publication Repository Service in a private publication repository

If you use a private publication repository, you can configure Cloud Integration Hub to bypass the Publication Repository Service (PRS) to improve system performance.

When you use a private publication repository, by default, Cloud Integration Hub writes published data to the publication repository and reads data from the publication repository through the PRS.

To configure Cloud Integration Hub to bypass the PRS in publication and subscription flows, select **Use JDBC for Private Publication Repository** in the Cloud Integration Hub connection. In the **Repository Schema** field of the **Organization Cloud Setup** dialog box, you can enter either a custom repository schema name or the default repository schema name to connect to the repository.

Warning:

- Do not edit any of the other connection properties unless you are instructed to do so when performing other tasks.
- Do not rename the connection.

Editing connection properties unnecessarily or renaming the connection might result in errors at run time.

Changing the port number of the publication repository service

When you select to use a private publication repository, Cloud Integration Hub communicates with the repository via the publication repository service on the Secure Agent.

By default, the port number of the publication repository service is 19443. You can change the port number.

1. In Administrator, select **Runtime Environments**, and then, on the **Runtime Environments** page, click the name of the Secure Agent that Cloud Integration Hub uses at run time.

Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.

2. On the **Details** tab, in the upper right corner, click **Edit**.
3. In the **System Configuration Details** area, select **CIH Processor**.
4. Click the **Edit Agent Configuration** icon next to **api-port** and enter the port number.
5. Click **Save**.

Using a customized Java KeyStore with a private publication repository

When you select to use a private publication repository, Cloud Integration Hub assigns a default Java KeyStore (JKS) as the repository of security certificates.

You can assign a different keystore to use with the publication repository.

1. Place the customized keystore in the following location:
`<Secure Agent installation directory>\apps\CIHProcessor\conf\`
2. In Administrator, select **Runtime Environments**, and then, on the **Runtime Environments** page, click the name of the Secure Agent that Cloud Integration Hub uses at run time.

Note: You might have to expand the Secure Agent group to see the list of Secure Agents within the group.

3. On the **Details** tab, in the upper right corner, click **Edit**.
4. In the **System Configuration Details** area, select **CIH Processor**.
5. Click the **Edit Agent Configuration** icon next to **keystore-filename** and enter the name of the keystore.
6. Click the **Edit Agent Configuration** icon next to **keystore-password** and enter the password to the keystore.
7. Click **Save**.

Enable zero downtime for a private publication repository

If you use a private publication repository, you can configure a zero downtime environment to ensure uninterrupted access to data.

You can enable zero downtime for publications and subscriptions that trigger a Data Integration task. You can't enable zero downtime for publications and subscriptions that publish data with an API.

To configure zero downtime for a private publication repository, in Administrator, select the option **Use JDBC for Private Publication Repository** in the Cloud Integration Hub connection properties.

On a hosted publication repository, Cloud Integration Hub applies zero downtime by default for all publication and subscription types.

Intermediate staging policy for subscriptions

During the subscription process, the Data Integration task reads the data from Cloud Integration Hub and then writes the data to the target application.

For performance tuning purposes, when the application consumes the data from the publication repository, Cloud Integration Hub writes the data to a local folder and then writes the data to the target location.

Cloud Integration Hub deletes the data from the local server at the end of the subscription process.

You can disable writing to intermediate staging on the local server in the Cloud Integration Hub connection. When intermediate staging is not used, the Data Integration task reads the data from Cloud Integration Hub and then writes the data directly to the target location. Disabling writing to intermediate staging might affect system performance.

To disable writing to intermediate staging, select the option **Do not use intermediate staging for subscription flows** in the Cloud Integration Hub connection.

Warning:

- Do not edit any of the other connection properties unless you are instructed to do so when performing other tasks.
- Do not rename the connection.

Editing connection properties unnecessarily or renaming the connection might result in errors at run time.

Setting up Cloud Integration Hub to show Data Integration Hub events

Set up Cloud Integration Hub to show Data Integration Hub events on the Cloud Integration Hub **Events** page.

1. On the Cloud Integration Hub **Home** page, click **System Properties**.
The **System Properties** page appears.

2. Configure the following properties:

System Property	Description
dih.console.accessmode	Enter cihprocessor or direct .
dih.console.url	Enter the URL of the Data Integration Hub console.
dih.console.username	Enter the user name of the user account of the Data Integration Hub console.
dih.console.password	Enter the password of the user account of the Data Integration Hub console.

Data Integration Hub events show on the Cloud Integration Hub **Events** page.

Configure load balancer URL

You can configure an external load balancer URL as the base API URL for publications and subscriptions that publish and consume data with an API to a private publication repository.

If the load balancer system property is not configured, publications and subscriptions that publish and consume data with an API use the first agent URL as the base API URL.

To configure the load balancer URL, add the system property **cih.api.loadbalancer.base.url** on the Cloud Integration Hub **System Properties** page. Enter the value as the URL of the load balancer. The URL of the load balancer is used as the base API URL of all existing or new publications and subscriptions that publish and consume data with an API.

System Properties

System properties determine Cloud Integration Hub behavior, such as showing events and identifying load balancer. You can access the **System Properties** page from the **System Properties** link on the top right of the

Cloud Integration Hub **Home** page. To configure and edit the system properties in Cloud Integration Hub, you must be assigned the Admin role.

The following table describes the system properties:

System Property	Description
dih.console.accessmode	<p>Access mode for the Data Integration Hub console to show Data Integration Hub events in Cloud Integration Hub.</p> <p>If the Cloud Integration Hub server can access Data Integration Hub REST APIs, set the value to direct.</p> <p>If the Cloud Integration Hub server can't access Data Integration Hub REST APIs, set the value to cihprocessor. Your organization must have a valid <code>CIHProcessor</code> license in Informatica Intelligent Cloud Services, and <code>CIHProcessor</code> must be able to access the Data Integration Hub REST APIs.</p>
dih.console.url	<p>URL of the Data Integration Hub console.</p> <p>The host can contain either the IP address or the host name, for example:</p> <p><code>https://dihhost:18443/dih-console</code></p>
dih.console.username	User name of the user account of the Data Integration Hub console.
dih.console.password	Password of the user account of the Data Integration Hub console.
cih.api.loadbalancer.base.url	<p>URL of the load balancer.</p> <p>Configure this property to use the URL of an external load balancer as the base API URL of publications and subscriptions that publish and consume data with an API to a private publication repository.</p>
cih.api.swagger.base.url	<p>Base URL of the Swagger structure.</p> <p>Configure this property to add a base URL to Swagger structures for publications and subscriptions that publish and consume data with an API.</p>
cih.pubrepo.recreate.disable	<p>Determines whether or not Cloud Integration Hub re-creates assets when you switch publication repositories.</p> <p>By default, Cloud Integration Hub doesn't re-create the assets. To enable asset re-creation, set the value to false.</p>

CHAPTER 3

Project and Asset Management

Manage projects, and the assets and folders within them, on the **Explore** page. The **Explore** page is a Informatica Intelligent Cloud Services feature that is available for most services. If you use multiple services, you might see projects, folders, and assets for all of your services on the **Explore** page.

You can manage your Informatica Intelligent Cloud Services projects and assets in the following ways:

- View assets.
- Edit assets.
- Move folders or assets to other locations on the **Explore** page.
- Delete projects, folders, or assets.
- Export assets, import assets, and migrate assets from one organization to another organization. Assets include applications, topics, publications, subscriptions, and monitoring rules.
- Apply tags so you can filter for related assets on the **Explore** page.

For more information about additional actions that you can perform on assets and for asset properties, see the chapters relevant to the asset type.

Viewing an asset

Use the **Explore** page to view assets, such as applications, topics, publications, and subscriptions. When you view a topic, the topic diagram appears by default. The topic diagram displays a graphical representation of the topic and the applications, publications, and subscriptions that are associated with the topic.

1. On the **Explore** page, navigate to the object that you want to view.
2. In the row that contains the object, click **Actions** and select **View**.

Tip: You can also view a publication or a subscription from the topic that the asset is associated with by right-clicking the asset on the **Publications** or **Subscriptions** area of the topic page and selecting **View**.

The asset appears.

Editing an asset

Use the **Explore** page to edit assets.

1. On the **Explore** page, navigate to the object that you want to edit.
2. In the row that contains the object, click **Actions** and select **Edit**.
The asset appears.
3. Edit the asset and then click **Save**.

Editing a topic

You can edit a topic to change the topic structure.

1. On the **Explore** page, in the row that contains the object, click **Actions** and select **Edit**.
The topic page appears. You can expand or collapse the areas of the page.
2. Perform one or more of the following tasks:
 - To edit the general details of the topic, scroll to the **General Details** area.
 - To edit the topic structure, scroll to the **Topic Structure** area.
 - To create, edit, disable, or delete publications that publish to the topic, scroll to the **Publications** area.
 - To create, edit, disable, or delete subscriptions that subscribe to the topic, scroll to the **Subscriptions** area.
3. Click **Save**.

Moving folders and assets

You can move folders and assets on the **Explore** page.

1. On the **Explore** page, navigate to the folder or assets that you want to move.
2. To move a folder or a single asset, in the row that contains the folder or asset, click **Actions** and select **Move To**, and then browse to the new location and click **Select**.
3. To move multiple assets, select the assets, click **Selected** and select **Move To**, and then browse to the new location and click **Select**.

Deleting projects, folders, and assets

You can delete a project, folder, or asset if you no longer need it. However, before you delete it, verify that no users in the organization plan to use it. You cannot retrieve projects, folders, or assets after you delete them.

You cannot delete an asset in the following situations:

- The asset is a task that is currently running.

- The asset is a mapping that is currently running.
- The asset is used by another asset. You must first delete the dependencies of the asset before you can delete the asset.
For information about viewing asset dependencies, see [“Asset dependencies” on page 45](#).
- The asset has associated publications or subscriptions.

Delete a project, folder, or asset from the **Explore** page, as shown in the following image:

1. On the **Explore** page, navigate to the object that you want to delete.
2. In the row that contains the project, folder, or asset, click **Actions** and select **Delete**.

Tip: You can also delete a publication or a subscription from the topic that the asset is associated with by right-clicking the asset on the **Publications** or **Subscriptions** area of the topic page and selecting **Delete**.

3. To delete multiple assets, select the assets, click **Selected** and select **Delete**.

User roles

A role is a collection of privileges that you can assign to users and groups. To ensure that every user can access assets and perform tasks in your organization, assign at least one role to each user or user group.

Administrators assign roles for the organization in Administrator. For more information, see *User roles* in the Administrator help.

To perform actions on Cloud Integration Hub assets, including applications, monitoring rules, publications, subscriptions, and topics, Cloud Integration Hub users need privileges for the assets that they will use. For example, to run publications, users need run privileges for the Hub Publication asset. The Informatica Intelligent Cloud Services system-defined roles Designer, Admin, and Monitor define access privileges for Cloud Integration Hub assets.

Designer and Admin roles

The Designer and Admin roles grant the following privileges for Cloud Integration Hub assets:

Asset type	Create	Read	Update	Delete	Run	Set Privilege
Hub Application	Yes	Yes	Yes	Yes	Not applicable	Yes
Hub Monitoring Rule	Yes	Yes	Yes	Yes	Not applicable	Yes
Hub Publication	Yes	Yes	Yes	Yes	Yes	Yes
Hub Subscription	Yes	Yes	Yes	Yes	Yes	Yes
Hub Topic	Yes	Yes	Yes	Yes	Not applicable	Yes

To configure and edit the system properties, users must be assigned the Admin role.

Monitor role

The Monitor role grants read privileges for all Cloud Integration Hub assets.

Privileges

Privileges determine the access a user has at the object level. You can configure privileges for object types at the user group-level or configure privileges for specific objects in object-level privileges. Privileges add additional or custom security for an object. Privileges define which users and groups can read, update, delete, execute, and change privilege on the object.

Administrators assign privileges for the organization in Administrator. For more information, see the Administrator help.

Required privileges for Cloud Integration Hub users

To perform actions in Cloud Integration Hub, Cloud Integration Hub users need the following privileges:

Administrator service

Read privileges for Organization, Secure Agent, Secure Agent Group, and User assets.

Data Integration service

Read privileges for Connection, Mapping Task, and Synchronization Task assets.

Integration Hub service

- Integration Hub feature is enabled.
- Read privileges for Hub Application, Hub Monitoring Rule, Hub Publication, Hub Subscription, and Hub Topic.
- Create, update, and delete privileges for Hub Application, Hub Monitoring Rule, and Hub Topic, based on the tasks that users need to perform on each asset type.
- Create, update, delete, and run privileges for Hub Publication and Hub Subscription, based on the tasks that users need to perform on each asset type.

You can assign privileges for Cloud Integration Hub assets by assigning user roles to users and user groups. You can either use the Informatica Intelligent Cloud Services system-defined roles Designer, Admin, or Monitor, or define custom roles. For more information about user roles in Cloud Integration Hub, see [“User roles” on page 37](#).

Required privileges for Data Integration users

To perform actions in Informatica Intelligent Cloud Services for Cloud Integration Hub operations, for example, to develop mappings and to create tasks, Informatica Intelligent Cloud Services users need the following privileges:

Administrator service

Read privileges for Secure Agent and Runtime Environment assets.

Data Integration service

- Read, create, update, and delete privileges for Connection asset.
- Read, create, update, delete, and run privileges for mapping task, synchronization task, and Mapping assets.

- The following features are enabled:
 - Data -preview
 - Debug logs - view
 - Job Results - view

Permissions

Permissions determine the access rights that a user has for a Secure Agent, Secure Agent group, connection, schedule, or asset. Permissions add additional or custom security for an object. Permissions define which users and groups can read, update, delete, execute, and change permissions on the object.

To configure permissions on an object, you need the following licenses and privileges:

- To configure permissions at the project level for all assets in a project, your organization must have the Set/Unset Security Permissions at Project Level license.
- To configure permissions at the folder level for all assets in a folder, your organization must have the Set/Unset Security Permissions at Folder Level license.
- To configure permissions on individual assets, your organization must have the Fine Grained Security license.
- The role assigned to your user account or to a group in which you are a member must have the Set Permission privilege for the object type. For example, to configure permissions on a Secure Agent, you must be assigned a role that has the Set Permission privilege for Secure Agents.

To configure permissions on an object, navigate to the object and set the appropriate permissions. For example, you want only users in the Development Team user group to have access to assets in the Development Data folder. Navigate to the folder, edit the permissions, and grant the Development Team user group permissions on the folder.

Permissions apply to the objects for which you configure them but not to copies of the object. Therefore, when you copy or export an asset, the permissions are not copied or exported with the asset. For example, you export a mapping task in which only user rjones has execute permission. When you import the mapping task, the imported mapping has no permissions assigned to it. Therefore, any user with privileges to run mapping tasks can run the imported task.

You can configure the following permissions on an object:

Permission	Description
Read	Open and view the object. If the object is source controlled, this permission allows the user or group to pull or check out the object from the source control repository. If you select a task, this permission also allows the user or group to use a connection or schedule in the task.
Update	Edit the object. If the object is source controlled, this permission allows the user or group to check in, check out, pull, unlink, or roll back the object. Requires read permission, which is automatically granted.
Delete	Delete the object.

Permission	Description
Execute	Run the object. Applies to mappings, tasks, taskflows, and Cloud Integration Hub assets. Monitor, stop, and restart instances of the mapping, task, or taskflow.
Change permissions	Change the permissions that are assigned to the object.

Note: These permissions control permissions within Informatica Intelligent Cloud Services. They do not control operating system permissions, such as the ability to start, stop, or configure the Secure Agent on Windows or Linux.

Rules and guidelines for permissions

Use the following rules and guidelines for permissions:

- When you configure permissions on an object, verify that the user or group to which you grant permissions is assigned a role with the appropriate privileges for the object type. For example, if you grant a user with the Service Consumer role Update privilege on a particular folder, the user cannot update the folder because the Service Consumer role does not have update privileges for folders.
- To edit an asset, the user must have read permission on all assets used within the asset. For example, when you assign a user Read and Update permissions on a synchronization task, verify that the user also has Read permission on the connections, mapplets, schedules, and saved queries that are used in the task.
- When a user edits a task, assets without Read permission are not displayed. To avoid unexpected results, the user should cancel all changes and avoid editing the task until the user is granted the appropriate Read permissions.
- When configuring a taskflow, a user needs Execute permission on all tasks to be added to the taskflow.
- To edit a taskflow, a user needs Execute permission on all tasks in the taskflow. Without Execute permission on all tasks, the user cannot save changes to the taskflow.
- To run a taskflow, a user needs Read and Execute permissions on taskflows.
- To monitor jobs or to stop a running job, a user needs Execute permission on the mapping, task, or taskflow.
- If you assign custom permissions to a Data Integration task and invoke the Data Integration task through an Application Integration process or a guide, you must complete either of the following tasks:
 - Give the Application Integration anonymous user permission to run the associated Data Integration asset.
 - Add the Application Integration anonymous user to a user group that has permission to run the associated Data Integration asset.

Configuring permissions

You can configure permissions on an object if you are assigned a role with the Set Permission privilege for the object type. For example, to configure permissions on a folder, you must be assigned a role that has the Set Permission privilege for folders.

1. Navigate to the object for which you want to configure permissions.

For example:

- To configure permissions on a Secure Agent or Secure Agent group, in Administrator, select **Runtime Environments**.
 - To configure permissions on a connection, in Administrator, select **Connections**.
 - To configure permissions on a mapping, in Data Integration, open the project and folder that contain the mapping.
 - To configure permissions on a Cloud Integration Hub asset, open the project and folder that contain the asset. For example, to configure permissions on a topic, open the project and folder that contain the topic.
2. In the row that contains the object, either click **Actions** and select **Permissions**, or click the **Change Permission** icon.

The **Permissions** dialog box lists the users and groups that have permissions on the object.

If the **Permissions** dialog box lists no users or groups, then no permissions are configured for the object. Any user with appropriate privileges for the object type can access the object.

The following image shows the **Permissions** dialog box for a mapping:

m_BostonCustomers Permissions

Click the tabs below to see what users and groups have access to this asset, or to modify the list and any associated permissions. Everyone except the following would have no access to the asset.

Users Groups

<input type="checkbox"/> User Name	First Name	Last Name	Read	Update	Delete	Execute	Change Permissions
<input type="checkbox"/> mclark	Melissa	Clark	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input type="checkbox"/> ajones	Adam	Jones	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
<input type="checkbox"/> dsmith	David	Smith	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>

Add Remove

Save Cancel

3. To configure user permissions on the object:
- a. Select **Users**.
 - b. If the user does not appear in the **Users** list, click **Add**, and select a user.
 - c. Enable or disable the appropriate permissions on the user.

Note: When you grant any user permissions on the object, Informatica Intelligent Cloud Services also adds you as a user with permissions on the object. This prevents you from losing access to the object when you configure permissions.

4. To configure user group permissions on the object:
- a. Select **Groups**.
 - b. If the group does not appear in the **Groups** list, click **Add**, and select a group.
 - c. Enable or disable the appropriate permissions on the group.

Note: When you grant any group permissions on the object, Informatica Intelligent Cloud Services also adds you as a user with permissions on the object. This prevents you from losing access to the object when you configure permissions.

5. To remove all permissions restrictions for the object, remove all users and groups from the **Permissions** dialog box.

When you remove all users and groups, any user with appropriate privileges for the object type can access the object.

6. Click **Save**.

Asset migration

You can migrate Cloud Integration Hub assets from one organization to another organization. Assets include applications, topics, publications, subscriptions, and monitoring rules.

The process to migrate assets depends on whether or not the source and target organizations reside on the same PoD (Point of Delivery):

- To migrate assets between organizations that reside on different PoDs, you export the assets from the source organization and then import the assets into the target organization. For more information, see [“Exporting assets” on page 42](#) and [“Importing assets” on page 43](#).
- To migrate assets between organizations that reside on the same PoD, you run the org to org migration process. For more information, see [“Migrating assets between organizations” on page 44](#).

Before you start the migration process, note the following considerations:

- When you migrate publications and subscriptions that publish and consume data with an API, Cloud Integration Hub changes the API URL based on the URL of the target organization. Be sure to inform API users of the new URL. After the migration is complete, you can copy the new URL from the publication or subscription page.
- You cannot migrate a publication or subscription with the same name as a publication or subscription that you used previously and later renamed or deleted.

Dependent assets

When you export, import, or migrate applications, publications, or subscriptions, Cloud Integration Hub also exports, imports, or migrates dependent Cloud Integration Hub assets.

Dependent assets can include applications, topics, publications, and subscriptions.

Cloud Integration Hub does not export, import, or migrate assets that users created in other Informatica Intelligent Cloud Services services and that the users later associated with Cloud Integration Hub assets. For example, Cloud Integration Hub does not export, import, or migrate Data Integration mappings or tasks. For more information about asset dependencies, see [“Asset dependencies” on page 45](#).

Exporting assets

Export Cloud Integration Hub assets from the organization to an export file. You can select a single asset or multiple assets to export, or you can export all assets in the organization. You can then import the assets to another organization.

1. Click the **Migration** link in the upper-right corner of the Home page.

2. In the **Export** tab, click **Select Entities**.
The **Select Entities** page appears.
3. From the **Entity Type** list, select the types of assets to export. You can select **All** to export all asset types.
Assets of the type or types that you select show in the **Available Entities** list.
4. In the **Available Entities** list select the assets to export and then click **Add**. To select all assets, click **Add All**.
The assets to export show in the **Selected Entities** list.
5. In the **Select Entities** page, click **OK**.
The assets to export show in the **Export** tab. If there are conflicts, a conflict resolution shows next to the relevant asset.
Note: You cannot remove dependent Cloud Integration Hub assets from the export list without removing the parent asset.
6. Click **Export**.
7. In the **Save As** dialog box, define the location and name of the file to export the assets to, and then click **Save**.
Cloud Integration Hub exports the assets and their dependent Cloud Integration Hub assets to the export file.

Importing assets

Import Cloud Integration Hub assets to the organization from a Cloud Integration Hub export file.

1. Click the **Import** link in the upper-right corner of the Home page and then select the **Import** tab.
2. In the **Conflict Resolution Rules** area, choose the actions to take when assets that you select to import exist in the organization. Select one of the following resolutions for each asset type:
 - **Overwrite**. Overwrite the asset with the imported asset. Overwritten assets cannot be recovered.
 - **Reuse**. Do not import the asset and keep the existing asset.
 - **Cancel**. Cancel the import operation.
3. In the **Select Entities** page, click **OK**.
The assets to import show in the **Import** tab. If there are conflicts, a conflict resolution shows next to the relevant asset.
4. Click **Import**, select the export file in the **Open** dialog box, and then click **Open**.
Cloud Integration Hub imports the selected assets and their dependent Cloud Integration Hub assets to the organization. If a selected asset exists in the organization, the action that Cloud Integration Hub takes depends on the conflict resolution that you defined for the asset type. Import results and conflicts appear in the **Import** tab.

Migrating assets between organizations

Migrate assets from one organization to another organization that resides on the same PoD (Point of Delivery). You can select a single asset or multiple assets to migrate, or select to migrate all the assets in the organization.

Before you begin the migration process, verify that the following conditions exist:

- You have Informatica Intelligent Cloud Services login credentials for the source organization.
 - The source organization is provisioned to Cloud Integration Hub.
1. Click the **Migration** link in the upper right corner of the Home page and then select the **Org to Org Asset Migration** tab.
 2. In the **Source Organization** area, click **Log in**, and then log in to the Informatica Intelligent Cloud Services organization that contains the assets to migrate.
 3. In the **Conflict Resolution Rules** area, choose the actions to take when assets that you select to migrate exist in the target organization. Select one of the following resolutions for each asset type:
 - **Overwrite**. Overwrite the target asset with the source asset. Overwritten assets cannot be recovered.
 - **Reuse**. Don't migrate the source asset and keep the existing target object.
 - **Cancel**. Cancel the entire migration operation.
 4. In the **Entities to Migrate** area, click **Select**.
The **Select Entities** page appears.
 5. From the **Entity Type** list select the types of assets to migrate, or select **All** to migrate all asset types.
Assets of the selected types show in the **Available Entities** list.
 6. In the **Available Entities** list select the assets to migrate and then click **Add**. To select all assets, click **Add All**.
The assets to migrate show in the **Selected Entities** list.
 7. In the **Select Entities** page, click **OK**.
The assets to migrate show in the **Review Locations** area.
 8. In the **Review Locations** area, select the target folder from each **Target Project** list to migrate the assets to. If the target project contains a folder that corresponds to the source folder, Cloud Integration Hub migrates the assets to the target folder. If the target project doesn't contain a folder that corresponds to the source folder, Cloud Integration Hub creates the target folder and migrates the assets to it.
 9. Click **Migrate**.
Cloud Integration Hub migrates the selected assets and their dependant Cloud Integration Hub assets to the target organization. If a selected asset exists in the target organization, the action that Cloud Integration Hub takes depends on the conflict resolution that you defined for the asset type. Migration conflicts and results appear in the **Org to Org Asset Migration** tab.

Migration error handling

When you import or migrate a topic, Cloud Integration Hub creates the topic structure in the publication repository.

If Cloud Integration Hub encounters a problem in creating or updating the structure, the state of the topic might change to not valid. To make the topic valid, perform one of the following actions:

- Re-run the import or migration process.
- Edit and save the topic in the topic wizard.

Asset dependencies

You can view object dependencies for an asset. You might want to view object dependencies before performing certain operations on an asset.

For example, you cannot delete an asset if another object depends on the asset. You must first delete the dependent objects and then delete the asset. You can find the dependent objects by viewing the asset dependencies.

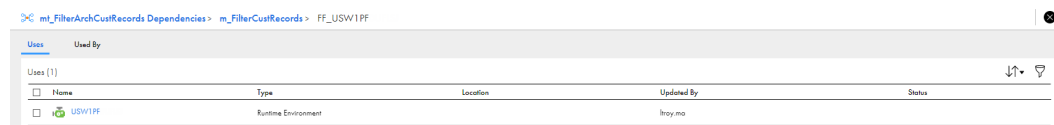
You can view object dependencies for Cloud Integration Hub assets from the topic or application pages and from the relationship diagram on the Hub Overview page. To view object dependencies, click an asset. The topic page, application page, or relationship diagram opens, showing the object dependencies.

The **Uses** tab lists the objects that the selected asset uses.

The **Used By** tab lists the objects that use the selected asset.

To drill down to the lowest level dependency, you can continue to show dependencies for each asset that appears on the **Dependencies** page. At the top of the **Dependencies** page, a breadcrumb shows the chain of dependencies.

The following image shows that the asset mt_FilterArchCustRecords is dependent on m_FilterCustRecords, which is dependent on FF_USW1PF:



Name	Type	Location	Updated By	Status
FF_USW1PF	Runtime Environment	lroy.ma		

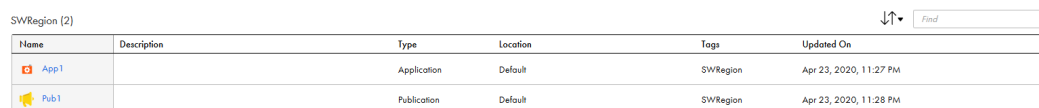
To view or delete an asset, in the row that contains the asset, click **Actions** and select the action.

Tags

A tag is an asset property that you can use to group assets. Create tags to filter for assets that share a common attribute on the **Explore** page.

For example, each of your organization's assets includes a tag that identifies the regional office that manages the asset. You want to view all of the assets that the Southwest regional office manages. On the **Explore** page, you explore by tag and then click the SW Region tag.

The following image shows this configuration:



Name	Description	Type	Location	Tags	Updated On
App1		Application	Default	SWRegion	Apr 23, 2020, 11:27 PM
Pub1		Publication	Default	SWRegion	Apr 23, 2020, 11:28 PM

You can assign tags to all asset types. An asset can have up to 64 tags.

You can find all of the assets that have a particular tag using one of the following methods:

- Click the name of the tag in the **Tags** column, in any row.
- Explore by tag, and then in the list of tags that shows on the page, click the name of the tag.

The following image shows an **Explore** page that lists all the tags created for the organization:

Explore ▾ All Tags ▾

All Tags (3)			
<input type="checkbox"/> Name	Asset Count	Description	Updated On
<input type="checkbox"/> NE Region	1		Mar 29, 2018, 6:44 PM
<input type="checkbox"/> NW Region	2		Mar 29, 2018, 7:29 PM
<input type="checkbox"/> SW Region	2		Mar 29, 2018, 6:48 PM

Creating tags

Use an asset's **Properties** dialog box to create and assign tags for that asset or to create tags to be available for future use.

Perform the following steps to create multiple tags without assigning them to an asset:

1. On the **Explore** page, browse by asset type.
2. In a row that contains an asset, click **Actions** and select **Properties**.
3. In the **Tags** field, enter the name of a tag that you want to create, and then press Enter.

A tag can have a maximum of 255 characters.

You can't use following characters in project, folder, asset, or tag names:

? ' | { } " ^ & [] / \

4. Continue to enter the desired tags. Press Enter after each tag name to add it to the tag list.

Properties ✕

Name:

Description:

Tags:

Save Cancel

5. After you have entered the tags, delete the tags from the **Tags** field so that the asset does not become associated with the tags. The tags will still appear in the list of available tags.
6. Click **Save**.


Assigning tags

You can assign a tag to one asset at a time or assign a tag to multiple assets at the same time. You can also assign multiple tags to one asset.

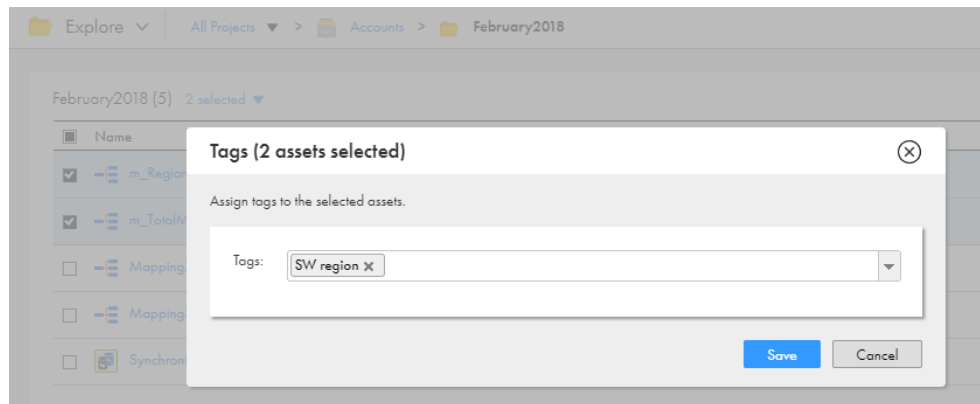
When you assign tags to an asset, you can choose an existing tag or create a new one.

1. On the **Explore** page, navigate to the asset or assets.

2. Perform one of the following tasks depending on whether you want to assign tags to one asset or assign tags to multiple assets at the same time.
 - To assign tags to one asset, in the row that contains the asset, click **Actions** and select **Properties**.



- To assign tags to multiple assets at the same time, in the row for each asset, select the check box. After you have selected all of the assets, from the Selection menu, select **Tags**.



3. Select an existing tag or enter the name of a new tag.
Continue adding tags or creating new tags until you have assigned all of the desired tags.
4. Click **Save**.

Editing and deleting tags

You can edit or delete a tag on the **Explore** page.

Edit a tag name or description in the tag properties. When you edit a tag, the properties for associated assets update as well. For example, if your `m_sales` asset has the NorthWest tag and you change the name of the tag to NW, the name of the tag changes to NW in the `m_sales` asset properties.

If you delete a tag, the tag no longer appears in the asset properties.

1. On the **Explore** page, browse by tags.

2. In the row that contains the tag, perform one of the following tasks:
 - To edit a tag, click **Actions** and select **Edit**. After you make your changes, click **Save**.
 - To delete a tag, click **Actions** and select **Delete**.

CHAPTER 4

Applications

An application represents an entity in your organization that needs to share data with other applications in your organization, such as sales applications or customer service applications. In Cloud Integration Hub, an application is a container for publications and subscriptions.

An application can publish data to a defined topic and can subscribe to data from a topic. For example, a sales application can publish sales reports and subscribe to inventory updates from an operations application. When you add a publication to an application, you define the schedule according to which topic data will be published from the application. You also define the schedule according to which topic data will be retrieved from the application and published to the Cloud Integration Hub publication repository. When you add a subscription to an application, you define the topic to which the application subscribes and the schedule and scope of data that the application consumes from the topic. The topic defines the structure of the data that the associated publications and subscriptions publish and consume.

Application management

Create applications and add a publication or a subscription to an application.

Creating an Application

Use the Navigator to create applications.

1. In the Navigator, click **New > Application**.
The **New Application** page appears.
2. Enter the application name, optionally, enter a description for the application, and then click **Save**.
3. To add a publication to the application, click **New Publication** and then define and save the publication.
4. To add a subscription to the application, click **New Subscription** and then define and save the subscription.

Adding a publication or a subscription to an existing application

Use the **Explore** page to add publications and subscriptions to existing applications.

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Applications**.
The **Explore** page shows all existing applications. You can sort the display by name, description, or last modified.

2. Rest on the application, click the Actions menu at the right end of the line, and then, from the menu, select **Add Publication** or **Add Subscription**.

The **New Publication** or **New Subscription** page shows. Define and save the publication or subscription.

Application properties

Application properties include general information about the application, a list of the publications that are associated with the application, and a list of the subscriptions that are associated with the application.

The application page includes the following properties:

Application Name

Name of the application. The name can contain up to 60 characters and can contain special characters.

Description

Description of the application. The description can contain up to 255 characters.

CHAPTER 5

Topics

A topic is an entity that represents a domain that is published and consumed in Cloud Integration Hub.

You can use Cloud Integration Hub to store relational data in a relational database and store files as-is in a file store. In a relational database, a topic defines the data structure and additional data definitions, such as the data retention period.

For example: Create an Accounts topic into which two CRM applications, a current application and a legacy application, publish accounts data. The marketing application and the data warehouse subscribe to the data in the Accounts topic.

You can use Cloud Integration Hub to store files as-is in a file store. In a file store, a topic defines the data retention period.

Multiple applications can publish to the same topic and an application can subscribe to a specific topic in a relational database or a file store. An application can subscribe to multiple topics in a relational database.

Topic structure in a relational database

When you create a topic in a relational database, you define the data structure on the publication repository where publications publish data to and subscriptions consume data from. The topic structure can consist of multiple tables.

Cloud Integration Hub generates the tables in the publication repository where it retains the data that is published for the topic. Cloud Integration Hub uses the data structure for the publications and subscriptions that are associated with the topic.

For example: Create a Sales topic that represents sales data. Applications from all the stores in the organization publish sales data to the Sales topic. The accounting application subscribes to the Sales topic and consumes published sales data from all stores, or, if a filter is applied, from specific stores.

Before you define publications and subscriptions for the data that is published and consumed in Cloud Integration Hub, you need to define the canonical structure that will hold the data that is published to Cloud Integration Hub in the Cloud Integration Hub publication repository. You define the canonical structure when you define the topic. You can define multiple topics that represent different source data sets.

Create topic tables in a relational database

You can use the following methods to create topic tables in a relational database:

- Create a table from a connection. Use this method when the structure of a table in the data domain that the topic represents exists in a connection object. You can use relational, flat file, and Salesforce connections to create topic tables. The **Add Table From Connection** dialog box shows up to 200 tables.
- Create a table from a flat file. Use this method when the structure of a table in the data domain that the topic represents exists in a flat file.
- Create a table from a metadata file. Use this method when the structure of a table in the data domain that the topic represents exists in a JSON, XML, XLS, or XLSX file. For more information, see [“Using metadata files to create topic tables in a relational database” on page 52](#).
- Create a new table. Use this method to define the structure manually if the structure of the table does not exist in a compatible file.

You can use more than one method to create tables in a single topic. For example, create two tables from a flat file, create three tables from a metadata file, and create a new table.

Note: If you add a table or table column to a topic with associated publications or subscriptions, to publish and to consume the additional data, edit the mapping to include the additional table or column. If you do not update the mapping, Cloud Integration Hub won't publish the additional data to the publication repository and subscribers won't receive it.

Using metadata files to create topic tables in a relational database

You can load a metadata file to Cloud Integration Hub and create a topic table in a relational database that is based on the structure of the file.

When you use a metadata file to create a topic table, you can define table attributes in the file before you load it to Cloud Integration Hub. For example, define column data type and precision, or define a column as a filter accelerator that is not encrypted.

You can use JSON, XML, XLS, and XLSX metadata files to create topic tables.

The metadata file must contain the following fields, and must not contain any other fields:

columnName

Mandatory. Name of the table column. The name must begin with an alphabetic character or underscore and can contain only alphanumeric characters or underscores.

filterAccelerator

Optional. Indicates that the column will be used in subscription queries and requires performance-related handling by Cloud Integration Hub. Use this indicator with topics that you plan to use for unbound subscriptions. By default, false.

When you use filter accelerators, consider the following guidelines:

- Filter accelerators slow down the writing of publication data to the Cloud Integration Hub publication repository.
- Filter accelerators have no impact on subscriptions that do not use filters.
- On a hosted Cloud Integration Hub publication repository, by default, Cloud Integration Hub encrypts the topic data. To use a column as a filter accelerator, you must define the value of the column's `encryption` field to `false`.

datatype

Optional. Data type of the field. By default, string.

The file can contain fields of the following data types:

- String
- Decimal
- Double
- Int32
- Int64
- Date_time
- Text

precision

Optional. Applies to data types that support precision. The default precision value depends on the data type of the field:

- String: 255
- Decimal: 15
- Text: 10000

scale

Optional. Applies to data types that support data scaling. The default scale value depends on the data type of the field:

- Decimal: 0. If the decimal field isn't encrypted and the published value is less than the topic decimal scale value, Cloud Integration Hub adds zeros to the decimal value. For example, if you define a scale of 4 for a decimal of 15.4, Cloud Integration Hub publishes and consumes the data with a decimal of 15.4000.
- Date_time: 0. For example, if you define a scale of 2 for a date_time of 2022-02-04 00:40:23.946, Cloud Integration Hub rounds the date_time to 2022-02-04 00:40:23.950. Applies to a hosted publication repository.
- All other data types: empty

encryption

Optional. Applies to a hosted publication repository.

If a file doesn't contain all the required fields, or contains non-required fields, loading the file to Cloud Integration Hub fails.

If a file contains identical rows, Cloud Integration Hub adds only the first row to the topic table.

Note: In a topic table, encrypting fields increases the storage space, thus impacting the publication and subscription run time.

Example table in a JSON file

```
[
  {
    "columnName": "id", "filterAccelerator": "false", "dataType": "int32",
    "encryption": false},
  {
    "columnName": "name", "filterAccelerator": "false", "dataType": "string",
    "precision" : 100, "encryption": false},
  {
    "columnName": "age", "filterAccelerator": "true", "dataType": "decimal",
    "precision" : 3, "encryption": "true"},
  {
    "columnName": "city", "filterAccelerator": "True", "dataType": "string",
    "precision" : 50, "encryption": "FALSE"},
  {
    "columnName": "salary", "filterAccelerator": false, "dataType": "decimal",
```

```
"precision" : 15, "scale":2, "encryption":true}
]
```

Example table in an XML file

```
<table>
<column>
<columnName>id</columnName>
<dataType>int32</dataType>
<encryption>>false</encryption>
<filterAccelerator>true</filterAccelerator>
</column>
<column>
<columnName>name</columnName>
<dataType>String</dataType>
<encryption>true</encryption>
<precision>100</precision>
<filterAccelerator>>false</filterAccelerator>
</column>
</table>
```

Example table in an XLS or XLSX file

columnName	filterAccelerator	dataType	precision	scale	encryption
id	TRUE	Int32			FALSE
name	FALSE	String	255		FALSE

Topic structure updates in a relational database

When you edit the structure of a topic with associated publications or subscriptions in a relational database, it might affect the associated publications and subscriptions. Topic structure changes might also impact the data in the publication repository and sometimes cause data loss.

Based on the nature of the update, you might have to edit the associated publications and subscriptions to align with the updated topic structure. Topic updates require that no publication or subscription events be in Processing or Delayed status. Search for any publication or subscription events that are in Processing status on the **Events** page and change their status to Discarded before you update them. For more information, see [“Discarding all events in Processing status” on page 105](#).

The following table describes the effects of topic structure updates on data in the publication repository and the resulting optional or required changes to the associated publications and subscriptions.

Topic Structure Update	Effect on Data in Publication Repository ¹	Optional/Required Changes to Associated Publications and Subscriptions
Add table	Table added	Optional: To publish and to consume the additional table, edit the mapping to include the additional table. If you do not update the mapping, data in the table will not be published to the publication repository and subscribers will not receive it.
Delete table	Table deleted, including data that was published to the table	Remove references to the table from the mapping of publications and from the mapping and the filter of subscriptions.

Topic Structure Update	Effect on Data in Publication Repository ¹	Optional/Required Changes to Associated Publications and Subscriptions
Add column	Column added	Optional: To publish and to consume the additional column, edit the mapping to include the additional column. If you do not update the mapping, data in the column will not be published to the publication repository and subscribers will not receive it.
Delete column	Column deleted, including data that was published to the column	Remove references to the column from the mapping of publications and from the mapping and the filter of subscriptions.
Rename column	Column deleted, including data that was published to the column, and another column created with new name	Remove references to the changed column from the mapping of publications and from the mapping and the filter of subscriptions. Optional: To publish or to consume the column that is created with a new name, edit the mapping to include the new column. If you do not update the mapping, data in the new column will not be published to the publication repository and subscribers will not receive it.
Change column data type	Column deleted, including data that was published to the column, and another column created with new data type	Remove references to the changed column from the mapping of publications and from the mapping and the filter of subscriptions. Optional: To publish or to consume the column that is created with the new data type, edit the mapping to include the new column. If you do not update the mapping, data in the new column will not be published to the publication repository and subscribers will not receive it.
Increase column precision, scale unchanged	Column updated	Open the publication or the subscription page for all associated publications and subscriptions. You do not need to edit any of the publication or subscription settings.
Increase column precision, increase scale by a lower value than the precision increase or by the same value as the precision increase	Column updated	Open the publication or the subscription page for all associated publications and subscriptions. You do not need to edit any of the publication or subscription settings.
Any other precision or scale updates	Column deleted, including data that was published to the column, and another column created with updated precision or scale	Remove references to the changed column from the mapping of publications and from the mapping and the filter of subscriptions. Optional: To publish or to consume the column that is created with the new precision or the new scale, edit the mapping to include the new column. If you do not update the mapping, data in the new column will not be published to the publication repository and subscribers will not receive it.

¹. Deleting columns in the publication repository might take a long time, based on the number of rows in the table.

Topic data retention

The data retention of a topic defines how long Cloud Integration Hub retains data that applications publish to the topic in the Cloud Integration Hub publication repository.

The retention period for consumed data defines how long Cloud Integration Hub retains consumed data in the publication repository if all the subscribers consume it. For each publication instance, the retention period for consumed data starts if all the subscribers have either successfully consumed or discarded the data, and is between 1 and 90 days. That is, after all the events that are associated with the publication instance are either in a Complete or in a Discarded event status. If all the subscribers consume or discard the data, Cloud Integration Hub stores the consumed data in the publication repository until the retention period for consumed data expires, and then deletes the consumed data from the publication repository.

The retention period for unconsumed data defines how long Cloud Integration Hub retains unconsumed data in the publication repository before it deletes the data. The retention period for unconsumed data is between the retention period for consumed data and 90 days.

Topic management

Create topics, add publications and subscriptions to topics, and subscribe to topics.

Creating a topic

Use the Navigator to create topics.

1. In the Navigator, click **New > Topic**.
The **New Topic** page appears.
2. Enter the topic name. The name must begin with an alphabetic character or underscore and can contain only alphanumeric characters or underscores. Optionally, enter a description for the topic.
3. Click **Browse** to select a location.
4. Select the publication repository type.
 - **Relational**. Stores published data in a relational database. Select this repository type for topics that you use for publications and subscriptions that publish and consume relational data.
 - **File Store**. Stores published files in a file store. Select this repository type for topics that you use for publications and subscriptions that publish and consume files as-is.
5. If you use a relational publication repository, select the topic type.
 - **Incremental Load**. The topic instance contains only the latest data changes. If you choose this topic type, verify that the data source includes delta indicators.
 - **Full Load**. The topic instance contains all of the data changes that occurred after the last publication.
6. Choose whether to prevent new publications and new subscriptions to the topic. If you choose this option you cannot create publications and subscriptions that publish to and subscribe from the topic.
7. Enter the number of days to retain consumed data in the publication repository in the **Retention period for consumed data** field. Enter a value between 1 and 90 days. For each publication instance, the retention period for consumed data starts if all the subscribers have either successfully consumed or discarded the data. That is, after all the events that are associated with the publication instance are either in a Complete or in a Discarded event status.

8. Enter the number of days to retain unconsumed data in the publication repository in the **Retention period for unconsumed data** field. Enter a value between the retention period for consumed data and 90 days.
9. Click **Create Table From** and select one of the following methods:
 - Create a table from a connection. Use this method when the structure of a table in the data domain that the topic represents exists in a connection object. You can use relational, flat file, and Salesforce connections to create topic tables. The **Add Table From Connection** dialog box shows up to 200 tables.
 - Create a table from a flat file. Use this method when the structure of a table in the data domain that the topic represents exists in a flat file.
 - Create a table from a metadata file. Use this method when the structure of a table in the data domain that the topic represents exists in a JSON, XML, XLS, or XLSX file. For more information, see [“Using metadata files to create topic tables in a relational database” on page 52](#).
 - Create a new table. Use this method to define the structure manually if the structure of the table does not exist in a compatible file.
10. Define the table in the create table dialog box and then click **OK**.
The structure of the table shows in the **Topic Structure** area.
11. Add the number of tables that you require to the topic. You must add at least one table to the topic. You can use multiple methods to add tables to the topic.

To edit or to delete a topic table, rest on a row in the table and click the Action menu at the right end of the line. From the menu select the required action: add row, rename table, delete row, or delete table.
12. Click **Save**.

The topic page shows the Topic Diagram.
13. Optionally, add publications and subscriptions to the topic. Perform one or both of the following actions:
 - To add a publication to the topic, expand the **Publications** area and click **New Publication**. For more information about creating publications, see [Creating a publication Use the Navigator to create publications](#).
 - To add a subscription to the topic, expand the **Subscriptions** area and click **New Subscription**. For more information about creating subscriptions, see [Creating a subscription Use the Navigator to create subscriptions](#).

Subscribing to a topic

Use the **Explore** page to subscribe to a topic.

1. On the **Explore** page, navigate to the object that you want to subscribe to a topic.
2. In the row that contains the object, click **Actions** . Select **Subscribe** and then configure the subscription.

Topic properties

Topic properties include general information about the topic, the topic structure, and the publications and subscriptions that are associated with the topic.

The topic page includes the following areas:

- Topic Diagram. Provides a visual overview of the topic and its relations to other assets. You can perform actions on assets in the diagram. For more information, see [“Topic Diagram” on page 58](#).
- General Details. General information about the topic. For more information, see [“General Details properties” on page 59](#).
- Topic Structure. List of topic tables, including details about each table. You add topic tables to the topic in this area. For more information, see [“Topic structure properties” on page 60](#).
- Publications. List of publications that publish data to the topic, including information about each publication. You can perform actions on existing publications and create new publications in this area. For more information, see [“Publications properties” on page 63](#).
- Subscriptions. List of subscriptions that subscribe to data from the topic, including information about each subscription. You can perform actions on existing subscriptions and create new subscriptions in this area. For more information, see [“Subscriptions properties” on page 63](#).

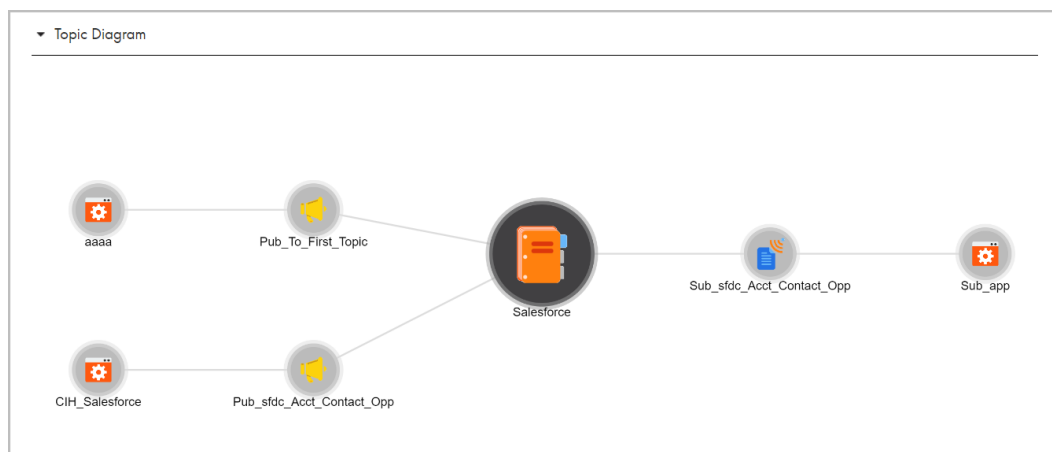
You can collapse and expand each area on the topic page.

Topic Diagram

The topic page shows the Topic Diagram. The diagram provides a visual overview of the topic and its relations to other assets, including the following assets:

- Applications that publish data to the topic
- Publications that publish the data from the applications to the topic
- Subscriptions that subscribe to data from the topic
- Applications that consume the data from the topic through the subscriptions

The following image shows a sample Topic Diagram:



When you click an asset, the properties page for the asset appears. For example, when you click a publication, the publication page appears.

When you right-click an asset, you can open it for viewing and editing. You can also run publications and subscriptions that trigger Data Integration tasks.

General Details properties

The **General Details** area of topic page includes the following properties:

Topic Name

Name of the topic. The name must begin with an alphabetic character or underscore and can contain only alphanumeric characters or underscores.

Description

Optional description of the topic.

Publication Repository Type

Type of publication repository where the hub stores published data or files to consume.

Choose one of the following options:

- Relational. Stores published data in a relational database. Use the relational publication repository for publications and subscriptions that run Data Integration tasks or that publish data with an API.
- File store. Stores published files as-is in a file store. Use the file store publication repository for publications and subscriptions that run file ingestion tasks.

Topic Type

Type of the topic. Topic type depends on the type of data that applications publish to the topic and has an impact on the delivery options to the subscribers to the topic.

Choose one of the following options:

- Incremental Load. The topic instance contains only the latest data changes. If you choose this topic type, verify that the data sources include delta indicators.
- Full Load. The topic contains all of the data changes that occurred after the last publication.

Prevent new publications and new subscriptions to this topic

Prevent new publications from publishing to the topic and prevent new subscriptions from subscribing to the topic. For example, when you plan to delete the topic. The topic is not available for selection when creating publications and subscriptions.

Existing publications can publish data to the topic and existing subscriptions can consume data from the topic.

Retention period for consumed data

Determines how long Cloud Integration Hub retains consumed data in the publication repository before it deletes the data. The retention period for consumed data must be between 1 and 90 days.

For each publication instance, the retention period for consumed data starts if all the subscribers have either successfully consumed or discarded the data. That is, after all the events that are associated with the publication instance are either in a Complete or in a Discarded event status.

Retention period for unconsumed data

Determines how long Cloud Integration Hub retains unconsumed data in the publication repository before it deletes the data. The retention period for unconsumed data must be between the retention period for consumed data and 90 days.

For each publication instance, the retention period for unconsumed data starts after the data is published.

Topic structure properties

The **Topic Structure** area of topic page includes the following properties:

Create Table From

Add tables to the topic. The topic must contain at least one table.

You can use one or more of the following methods to add tables to the topic:

- Create a table from a connection. Use this method when the structure of a table in the data domain that the topic represents exists in a connection object. You can use relational, flat file, and Salesforce connections to create topic tables. The **Add Table From Connection** dialog box shows up to 200 tables.
- Create a table from a flat file. Use this method when the structure of a table in the data domain that the topic represents exists in a flat file.
- Create a table from a metadata file. Use this method when the structure of a table in the data domain that the topic represents exists in a JSON, XML, XLS, or XLSX file. For more information, see [“Using metadata files to create topic tables in a relational database” on page 52](#).
- Create a new table. Use this method to define the structure manually if the structure of the table does not exist in a compatible file.

Show

Lists the tables in the topic. You can select to show a specific table.

The list of topic tables shows the following properties for each table:

Table

Name of the topic table. A topic table name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

Column

Name of the table column. The name must begin with an alphabetic character or underscore and can contain only alphanumeric characters or underscores.

Filter Accelerator

Indicates that the column will be used in subscription queries and requires performance-related handling by Cloud Integration Hub. Use this indicator with topics that you plan to use for unbound subscriptions.

When you use filter accelerators, consider the following guidelines:

- Filter accelerators slow down the writing of publication data to the Cloud Integration Hub publication repository.
- Filter accelerators have no impact on subscriptions that do not use filters.
- In a hosted publication repository, by default, Cloud Integration Hub encrypts the topic data. To use a column as a filter accelerator you must change the value of **Encrypted** to **No** for the column.

Data Type

Select from the list of available data types. By default, Cloud Integration Hub reads the data as string.

Precision

Available for string, decimal, and text data types.

For a string data type, the maximum precision is 4,000 characters.

For a decimal data type, the maximum precision is 38 characters.

For a text data type, the maximum precision is 4,000,000 characters if the **Encrypted** option is not selected, and 100,000 characters if the **Encrypted** option is selected.

Scale

Enabled only for data types that support data scaling.

For a decimal data type, the maximum scale is 30 characters if you use a hosted publication repository, and 38 characters if you use a private publication repository.

Encrypted

Determines whether Cloud Integration Hub encrypts the column. Applies to a hosted publication repository.

By default, Cloud Integration Hub encrypts all columns except for columns with a DATE_TIME data type. To use a column as a filter in your mappings, change the value of **Encrypted** to **No** for the column.

Add Table from a Connection object

Add a topic table from a connection object that contains the structure of a table in the data domain that the topic represents. You can add tables from objects in relational, flat file, and Salesforce connections.

The **Add Table from a Connection** page includes the following properties:

Connection

Connection that contains the object to create the topic table from.

Source Object

Object to create the topic table from.

Formatting Options

Applies to flat file connections. Defines the delimiter, text qualifier, and escape character that are used in the file.

Table Name

Name of the topic table. The name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

Add Table from Flat File properties

Add a topic table from a flat file that contains the structure of a table in the data domain that the topic represents.

The **Add Table from Flat File** page includes the following properties:

File

Name of the file that contains the structure of the data domain that the topic represents.

Drop a file into the **File** field or click **Choose File** to browse to and choose the sample file on which to base the table structure.

Table Name

Name of the topic table. The name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

Import Column Names

Optional. Select this option to use the column names in the file as the default column headers in the table. Enter the number of the lines that serves as the file's header line in the **From Line** field.

Code page

Character encoding used in the file.

Default text length

Optional. Length of the text fields in the table.

Delimiter

Delimiter used in the file to separate between columns. Select a predefined delimiter or select **Custom** to define a custom delimiter.

Text qualifier

Optional. Symbols used in the file to enclose a string.

Load File

Loads the selected file and shows a preview of the file.

Preview

Shows the columns that will be added to the table after you load the file.

Add Table from Metadata File properties

Add a topic table from a metadata file that contains the structure of a table in the data domain that the topic represents.

The **Add Table from Metadata File** page includes the following properties:

File

Name of the file that contains the structure of the data domain that the topic represents.

Drop a file into the **File** field or click **Choose File** to browse to and choose the sample file on which to base the table structure.

Table Name

Name of the topic table. The name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

Load File

Loads the selected file and shows the status of the file, valid or invalid. If a file is valid and Cloud Integration Hub converts source values to Cloud Integration Hub default values, the changes are listed in the **Create Table from Metadata File** page. For more information, see ["Using metadata files to create topic tables in a relational database" on page 52](#).

Create New Table properties

Add a topic table and define the structure of a topic table manually.

The **Create New Table** page includes the following properties:

Table Name

Name of the table. The name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

Number of columns

Number of columns in the table.

Publications properties

The **Publications** area of topic page includes the following properties:

New Publication

Create a publication that publishes data to the topic. For more information about creating publications, see [Creating a publication Use the Navigator to create publications.](#)

Publication list

List of publications that publish data to the topic. When you right-click a publication, an actions menu opens. From the menu you can run, view, disable or enable, and delete the publication.

The publication list shows the following properties for each publication:

Name

Name of the publication.

Description

Description of the publication.

Mode

Publication mode, enabled or disabled. A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.

Last Modified

Date and time when the publication was last modified.

Subscriptions properties

The **Subscriptions** area of topic page includes the following properties:

New Subscription

Create a subscription that consumes data from the topic. For more information about creating subscriptions, see [Creating a subscription Use the Navigator to create subscriptions.](#)

Subscription list

List of subscriptions that consume data from the topic. When you right-click a subscription, an actions menu opens. From the menu you can run, view, disable or enable, and delete the subscription. You can also get data that was published before the subscription subscribed to the topic and therefore was not consumed by the subscriber.

The subscription list shows the following properties for each subscription:

Name

Name of the subscription.

Description

Description of the subscription.

Mode

Subscription mode, enabled or disabled. A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

Last Modified

Date and time when the subscription was last modified.

CHAPTER 6

Data Integration tasks

Cloud Integration Hub uses Data Integration synchronization and mapping tasks to publish data from source applications to the Cloud Integration Hub relational publication repository and to consume data from the relational publication repository into target applications.

To use a Data Integration task in a publication, you create a synchronization task or a mapping task in Data Integration before you create the publication. You select the task when you create the publication.

To use a Data Integration task in a subscription, you can use one of the following methods:

- Create a synchronization task or a mapping task in Data Integration before you create the subscription, and select the task when you create the subscription.
- Create a synchronization task when you create the subscription. Cloud Integration Hub saves the task in Data Integration.

Note: Publications and subscriptions that publish and consume data with an API use the Cloud Integration Hub REST APIs. For more information, see [Chapter 11, “Cloud Integration Hub REST APIs” on page 111](#).

Data Integration Task Types

You can use synchronization tasks and mapping tasks in Cloud Integration Hub publications and subscriptions to and from cloud-based applications.

Use a synchronization task for a publication or a subscription where the publication or subscription process requires mappings and filters that synchronization tasks support. For example, to read data from a CRM application and publish the data as is.

Use a mapping task for a publication or a subscription if you want to use an advanced ETL (Extract, Transform, and Load) process for the Cloud Integration Hub publication or subscription process. For example, you can use a mapping task to perform the following actions on a publication or subscription:

- Run data quality rules on the data.
- Add data from an additional source to the data that a publication publishes or to the data that a subscription consumes.

When a publication uses a mapping task, you can create a compound publication that publishes data to multiple topics.

Data Integration tasks rules and guidelines

When you develop Data Integration mappings and tasks to use in Cloud Integration Hub publications and subscriptions, consider these rules and guidelines.

General rules and guidelines

Consider the following rules and guidelines when you create Data Integration mappings and tasks:

- Don't run tasks that you create for Cloud Integration Hub from within Informatica Intelligent Cloud Services. You must run the tasks from Cloud Integration Hub by running the publication or the subscription to which the task is associated.
- Don't check the Data Integration task into a source control repository if you use it in a Cloud Integration Hub publication or subscription.
- When you use the Cloud Integration Hub connection, the target object in a publication mapping or task and the source object in a subscription mapping or task presents the list of topics defined in Cloud Integration Hub. The format of the list is `TopicName/tableName`.

Warning: When you set up the organization in Cloud Integration Hub, Cloud Integration Hub creates the connection **Cloud Integration Hub** in the organization in Informatica Intelligent Cloud Services. Do not rename or edit this connection. Editing the connection or changing the connection name might result in errors at run time.

- Cloud Integration Hub determines the scheduling of the publication or the subscription based on the settings that the operator defined for the publication or the subscription. When you create the Data Integration task, in the **Schedule** page of the task wizard, verify that the option **Do not run this task on a schedule** is selected.
- To distinguish between publication tasks and subscription tasks, indicate the type of the task in the task name. When you select a task for a publication or for a subscription, you can easily select an appropriate task.
For example, name a publication task `Pub_relational_<TaskName>`, and name a subscription task `Sub_relational_<TaskName>`.

Synchronization task rules and guidelines

Consider the following rules and guidelines when you create synchronization tasks and mappings:

- The task operation for publication tasks is an insert operation.
- When you create a publication task, select the Cloud Integration Hub connection as the target. When you create a subscription task, select the Cloud Integration Hub connection as the source.
- Synchronization tasks do not support multiple sources. Therefore, when you create a synchronization task for a publication or a subscription with multiple sources, create a relationship between the sources for the following use cases:
 - Publications: when you publish from multiple tables.
 - Subscriptions: when you subscribe to multiple tables, or when the subscription is a compound subscription.
- Cloud Integration Hub supports the following connection types in synchronization tasks that you create for subscriptions in Cloud Integration Hub:
 - Relational database
 - Salesforce
 - Flat file

Mapping task rules and guidelines

Consider the following rules and guidelines when you create mapping tasks and mappings:

- The mapping operation is an insert operation for both publication and subscription mappings.
- When you create a publication mapping, select the Cloud Integration Hub connection when you configure the target properties. When you create a subscription mapping, select the Cloud Integration Hub connection when you configure the source properties.

Synchronization tasks with Cloud Integration Hub

Use synchronization tasks for publications and subscriptions where the publication or subscription process requires only mapping and filtering. For example, to read data from a CRM application and publish the data as is.

For publications and subscriptions that require additional data processing and for compound publications, use mapping tasks.

Creating a synchronization task for a publication

To create a synchronization task for a publication, perform the following tasks:

- Define task details.
- Select the publication source. The source is the cloud application from which you want to publish data.
- Select the publication target. The target is the topic table in the Cloud Integration Hub publication repository, into which the cloud application publishes the data. The topic must exist in Cloud Integration Hub before you create the task.
- Optionally, define data filters. Cloud Integration Hub does not support the use of advanced filters in synchronization tasks.
- Configure field mapping. Map source fields to topic fields.
- Save and close the task.

Step 1. Define task details

Define task details in the **Definition** page of the **Synchronization Task Wizard**.

1. Click **Task Wizards > Data Synchronization**.

The **Synchronization Task Wizard** appears.

2. Specify the following details:

Task Name

Enter a name for the task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub publication workflow, you select a publication task. For example, name the task Pub_<TaskName>.

Description

Optionally, enter a description for the task. The description can contain up to 255 characters.

Task Operation

Choose **Insert**.

3. Click **Next**

The **Source** page appears.

Step 2. Select publication source

Select the publication source in the **Source** page of the **Synchronization Task Wizard**.

- Specify the following details and then click **Next**:

Connection

Select a source connection that connects to the source from which you want to publish data.

Source Type

The source type depends on the number of tables that you want to publish:

- To publish a single table, select **Single**.
- To publish multiple tables, select **Multiple** and then create a relationship between the tables.

Source Object

Select the source from which you want to publish data.

The **Target** page appears.

Step 3. Select publication target

Select the publication target in the **Target** page of the **Synchronization Task Wizard**. The publication target is the topic table in the Cloud Integration Hub publication repository to which you want to publish data.

1. Specify the following details:

Connection

Select the **Cloud Integration Hub** connection.

Target Object

Select the topic table to which you want to publish data. The format of the target object is `TopicName/tableName`.

2. Click **Next**.

The **Data Filters** page appears.

3. Optionally, configure data filters. You configure data filters for Cloud Integration Hub publications in the same way that you configure data filters for other Data Integration tasks.

4. Click **Next**.

The **Field Mapping** page appears.

Step 4. Configure field mapping

Map source fields to topic fields in the **Field Mapping** page of the **Synchronization Task Wizard**.

1. Map fields in the **Source** column to fields in the **Target** column and then click **Next**.
The **Schedule** page appears.
2. Verify that the option **Do not run this task on a schedule** is selected. The task runs according to the schedule of the publication that uses the task.
3. Select **Save > Save and Close** to save the task.

Creating a synchronization task for a subscription

To create a synchronization task for a subscription in Data Integration, perform the following tasks:

- Define task details.
- Select the subscription source. The source is the topic table in the Cloud Integration Hub publication repository, from which you want to consume data. The topic must exist in Cloud Integration Hub before you create the task.
- Select the subscription target. The target is the cloud application that you want to consume the data.
- Optionally, define data filters. Cloud Integration Hub does not support the use of advanced filters in synchronization tasks.
- Configure field mapping. Map topic fields to target fields.
- Save and close the task.

Tip: You can also create a synchronization task for a subscription in Cloud Integration Hub. For more information, see [“Creating a subscription that triggers a Data Integration task” on page 92](#).

Step 1. Define task details

Define task properties in the **Definition** page of the **Synchronization Task Wizard**.

1. Click **Task Wizards > Data Synchronization**.
The **Synchronization Task Wizard** appears.
2. Specify the following details and then click **Next**:

Task Name

Enter a name for the synchronization task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub subscription workflow, you select a subscription task. For example, name the task Sub_<TaskName>.

Description

Optionally, enter a description for the task. The description can contain up to 255 characters.

Task Operation

Choose **Insert**.

The **Source** page appears.

Step 2. Select subscription source

Select the subscription source in the **Source** page of the **Synchronization Task Wizard**. The subscription source is the topic table in the Cloud Integration Hub publication repository from which you want to consume data.

- Specify the following details and then click **Next**:

Connection

Select the **Cloud Integration Hub** connection.

Source Type

The source type depends on the number of tables that you want to consume and on the subscription type:

- To consume a single table, select **Single**.
- To consume multiple tables, or when the subscription is a compound subscription, select **Multiple** and then create a relationship between the tables.

Source Object

Select the topic table from which you want to consume data. The format of the object is `TopicName/tableName`.

The **Target** page appears.

Step 3. Select subscription target

Select the subscription target in the **Target** page of the **Synchronization Task Wizard**.

1. Specify the following details:

Connection

Select a target connection that connects to the target into which you want to consume data.

Target Object

Select the target into which you want to consume the data.

2. Click **Next**.
The **Data Filters** page appears.
3. Optionally, configure data filters. You configure data filters for Cloud Integration Hub subscriptions in the same way that you configure data filters for other Data Integration tasks.
4. Click **Next**.
The **Field Mapping** page appears.

Step 4. Configure Field Mapping

Map topic fields to target fields in the **Field Mapping** page of the **Synchronization Task Wizard**.

1. Map fields in the **Source** column to fields in the **Target** column and then click **Next**.
The **Schedule** page appears.
2. Verify that the option **Do not run this task on a schedule** is selected. The task runs according to the schedule of the publication that uses the task.
3. Select **Save > Save and Close** to save the task.

Mapping configuration tasks with Cloud Integration Hub

Use mapping tasks for publications and subscriptions if you want to add an ETL (Extract, Transform, and Load) process to the Cloud Integration Hub publication or subscription process.

For publications and subscriptions that require mapping and filtering only, use synchronization tasks.

Mapping task configuration process

To use mapping tasks with Cloud Integration Hub Connector, you create the mapping in Informatica Intelligent Cloud Services Mapping Designer and then create a mapping task that uses the mapping.

To use mapping configuration with Cloud Integration Hub Connector, perform the following tasks:

1. Create a mapping in Mapping Designer.
When you create a mapping for a publication, the source is the publishing cloud application and the target is the topic table in the Cloud Integration Hub publication repository into which to publish the data.
When you create a mapping for a subscription, the source is the topic table in the Cloud Integration Hub publication repository to consume data from and the target is the subscribing cloud application.
The topic must exist in Cloud Integration Hub before you create the mapping.
2. Create a mapping task and select the appropriate mapping.

Creating a mapping and task for a publication

The mappings and tasks that you create to use in Cloud Integration Hub publications include a source and a target.

In publication mappings and tasks the source is the cloud application from which to publish data and the target is the topic table in the Cloud Integration Hub publication repository to which the publication publishes data.

Creating a mapping for a publication

Create the mapping to use in the mapping task for the publication.

The topic must exist in Cloud Integration Hub before you create the mapping.

1. Click **Design > Mappings**, and then click **New Mapping**.
2. In the **New Mapping** dialog box, enter the mapping name and description, and click **OK**.
You can use alphanumeric characters and underscores (_) in the mapping name.
3. Add a source to the mapping canvas and configure source properties.
4. Add a target to the mapping canvas and configure target properties.
 - a. In the **Properties** panel, on the **General** tab, you can enter a name and description.
 - b. Click the **Target** tab. From the **Connection** list, select the **Cloud Integration Hub** connection.
 - c. Click **Select** next to the **Object** field, select a topic table in the **Select Object Target** dialog box, and then click **OK**.

- d. Optionally, open the **Connection** list and select a parameter from the **Parameters** list. If no parameters exist in the list click **New Parameter** and then name the parameter and click **OK**.
- e. Click the **Field Mapping** tab and map fields from the source to the connection.
5. On the mapping canvas, connect the source to the target.
6. Click **Save > Save and Close**.

Creating a mapping task for a publication

When you create a publication task you select mapping that you created for the publication and select the **Cloud Integration Hub** connection as the target of the task.

The mapping must exist in Mapping Designer before you create the task.

1. In Data Integration, click **New > Task > Mapping Task > Create**.
2. Specify the following task details:

Task Name

Enter a name for the task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub publication workflow, you select a publication task. For example, name the task Pub_<TaskName>.

Description

Optionally, enter a description for the task. The description can contain up to 255 characters.

Runtime Environment

Runtime environment that contains the Secure Agent to run the task.

Mapping

Mapping associated with the task. Select the publication mapping.

To select a mapping, click **Select**. The **Select a Mapping** dialog box displays up to 200 mappings. If the mapping you want to use does not display, enter a search string to reduce the number of mappings that display.

Select a mapping and click **OK**.

An image of the mapping displays below the mapping name.

3. Select the **Targets** step and then, from the **Connection** list, select the **Cloud Integration Hub** connection.
4. From the **Object** list, select the topic table into which to publish data.
5. Click **Finish**.

Creating a mapping and task for a subscription

The mappings and tasks that you create to use in Cloud Integration Hub subscriptions include a source and a target.

In subscription mappings and tasks the source is the topic table in the Cloud Integration Hub publication repository from where to consume data and the target is the cloud application that consumes the data.

The topic must exist in the Cloud Integration Hub before you create the mapping and task.

Creating a mapping for a subscription

Create the mapping to use in the mapping task for the subscription.

The topic must exist in Cloud Integration Hub before you create the mapping.

1. Click **Design > Mappings**, and then click **New Mapping**.
2. In the **New Mapping** dialog box, enter the mapping name and description, and click **OK**.
You can use alphanumeric characters and underscores (_) in the mapping name.
3. Add a source to the mapping canvas and configure source properties.
 - a. In the **Properties** panel, on the **General** tab, you can enter a name and description.
 - b. Click the **Source** tab. From the **Connection** list, select the **Cloud Integration Hub** connection.
 - c. Click **Select**. To consume multiple topic tables select **Multiple Objects** and then, in the actions menu, click **Add Source Object**.
The **Select Source Object** dialog box shows.
 - a. Select the database table or tables to consume and then click **OK**.
 - b. Click **Partitions** and enter the number of partitions to process data in parallel.
4. Add a target to the mapping canvas and configure target properties.
5. On the mapping canvas, connect the source to the target.
6. Click **Save > Save and Close**.

Creating a mapping task for a subscription

When you create a subscription task you select mapping that you created for the subscription and select the **Cloud Integration Hub** connection as the source of the task.

The mapping must exist in Mapping Designer before you configure the task.

1. In Data Integration, click **New > Mapping Task > Create**.
2. Specify the following task details:

Task Name

Enter a name for the task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub publication workflow, you select a publication task. For example, name the task Pub_<TaskName>.

Description

Optionally, enter a description for the task. The description can contain up to 255 characters.

Runtime Environment

Runtime environment that contains the Secure Agent to run the task.

Mapping

Mapping associated with the task. Select the subscription mapping.

To select a mapping, click **Select**. The **Select a Mapping** dialog box displays up to 200 mappings. If the mapping you want to use does not display, enter a search string to reduce the number of mappings that display.

Select a mapping and click **OK**.

An image of the mapping displays below the mapping name.

3. Select the **Sources** step and then, from the **Connection** list, select the **Cloud Integration Hub** connection.
4. From the **Object** list, select the topic table from which to consume data.
5. Optionally, you can run multiple instances of a mapping task simultaneously. To enable simultaneous task runs, on the **Schedule** page, select **Allow the mapping task to be executed simultaneously**.
6. Click **Finish**.

CHAPTER 7

File ingestion tasks

Cloud Integration Hub uses file ingestion tasks to publish files as-is from source connections to the Cloud Integration Hub file store publication repository and to consume files from the file store publication repository into target connections.

To use a file ingestion task in a publication, you create a file ingestion task in Mass Ingestion Files before you create the publication. You select the task when you create the publication.

To use a file ingestion task in a subscription, you create a file ingestion task in Mass Ingestion Files before you create the subscription. You select the task when you create the subscription.

File ingestion tasks rules and guidelines

When you develop file ingestion tasks to use in Cloud Integration Hub publications and subscriptions, consider the following rules and guidelines:

- Verify that your organization has a valid Mass Ingestion Files license.
- Don't run tasks that you create for Cloud Integration Hub from within Mass Ingestion. You must run the tasks from Cloud Integration Hub by running the publication or the subscription to which the task is associated.
- The publication source must be a Mass Ingestion Files connection.
- The publication target directory must be a local folder with the `${PMTargetFileDir}` parameter added.
- The subscription source directory must be a local folder with a by file list pickup method and the `${PMSourceFileDir}` parameter added.
- The subscription target must be a Mass Ingestion Files connection.
- Cloud Integration Hub determines the scheduling of the publication or the subscription based on the settings that the operator defined for the publication or the subscription. When you create the file ingestion task, on the **Schedule** page of the task wizard, verify that the option **Do not run this task on a schedule** is selected.
- Don't select **Check File Stability** when you create the file ingestion task, on the **Source** page of the task wizard.

Creating a file ingestion task for a publication

To create a file ingestion task for a publication in Mass Ingestion, perform the following tasks:

- Define task details.
- Configure the publication source. The source is the Mass Ingestion Files connection to publish files from.
- Configure the publication target. The target directory is a local folder to publish files to.
- Optionally, set the runtime options.
- Save and close the task.

Step 1. Define task details

Define task details on the **Definition** page of the **File Ingestion Task Wizard**.

1. Log in to Mass Ingestion. Click **New > File Ingestion Task**.
The **Definition** page of the **File Ingestion Task Wizard** appears.
2. Specify the following details:

Task Name

Enter a name for the task.

The task name can contain alphanumeric characters, spaces, and underscores. Names must begin with an alphabetic character or underscore.

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub publication workflow, you select a publication task. For example, name the task Pub_<TaskName>.

Location

Select a local folder where the task will reside.

Description

Optionally, enter a description of the task. The description can contain up to 1,024 characters.

Runtime Environment

Select the same runtime environment that contains the Secure Agent that runs the task.

3. Click **Next**.
The **Source** page appears.

Step 2. Configure publication source

Configure the publication source on the **Source** page of the **File Ingestion Task Wizard**. The options that appear on the page vary based on the type of source connection that you select.

1. On the **Source** page, from **Source Type**, select **Source Connection**.
2. From the **Connection Type** list, select a connection type.
Don't select the connection type **File Listener**.
3. From the **Connection** list, select a source connection. Based on the connection that you select, enter the source options.
Don't select **Check File Stability**. Selecting it can cause performance issues.

4. Click **Next**.
The **Target** page appears.

Step 3. Configure publication target

Configure the publication target on the **Target** page of the **File Ingestion Task Wizard**. The publication target is the local folder in the Cloud Integration Hub publication repository to publish files to.

1. On the **Target** page, from **Connection Type**, select **Local Folder**.
2. From the **Target Directory** list, click **Add Parameter** and select **TargetFileDir**. Click **Ok**.
3. Click **Save**.

Creating a file ingestion task for a subscription

To create a file ingestion task for a subscription in Mass Ingestion, perform the following tasks:

- Define task details.
- Configure the subscription source. The source directory is a local folder to consume files from.
- Enter a dummy value for the file list.
- Configure the subscription target. The target is a Mass Ingestion Files connection that consumes the files.
- Optionally, configure one or more file processing actions.
- Save and close the task.

Step 1. Define task details

Define task properties on the **Definition** page of the **File Ingestion Task Wizard**.

1. Log in to Mass Ingestion. Click **New > File Ingestion Task**.
The **Definition** page of the **File Ingestion Task Wizard** appears.
2. Specify the following details:

Task Name

Enter a name for the task.

The task name can contain alphanumeric characters, spaces, and underscores. Names must begin with an alphabetic character or underscore.

Tip: Indicate the type of the task in the task name. This will ensure that when you select a task to use in a Cloud Integration Hub subscription workflow, you select a subscription task. For example, name the task Sub_<TaskName>.

Location

Select a local folder where the task will reside.

Description

Optionally, enter a description for the task. The description can contain up to 1,024 characters.

Runtime Environment

Select the same runtime environment that contains the Secure Agent that runs the task.

3. Click **Next**

The **Source** page appears.

Step 2. Configure subscription source

Configure the subscription source on the **Source** page of the **File Ingestion Task Wizard**.

1. On the **Source** page, from **Source Type**, select **Source Connection**.
Don't select the source type **File Listener**.
2. From the **Connection Type** list, select **Local Folder** as the connection type.
3. In the **Source Options** area, select **By File List** as the file pickup method.
4. From the **Source Directory** list, click **Add Parameter** and select **SourceFileDir**. Click **Ok**.
5. Select **File List**. In the input field, enter a dummy value. Cloud Integration Hub can't run a file ingestion task without a dummy value.
Don't select **Check File Stability**. Selecting it can cause performance issues.
6. Click **Next**.
The **Target** page appears.

Step 3. Configure subscription target

Configure the subscription target on the **Target** page of the **File Ingestion Task Wizard**. Options that appear on the page vary based on the type of target connection that you select.

1. On the **Target** page, from **Connection Type**, select a connection type.
2. Select a connection.
3. Based on the target connection that you select, enter the target options.
4. Click **Save**.

Step 4. Configure file processing actions (optional)

You can configure one or more file processing actions on the **Actions** page of the **File Ingestion Task Wizard**. For example, you can choose to flatten the file structure of the files you consume. For more information about file processing actions, see the Mass Ingestion help.

1. On the **Actions** page, click **Add**.
The **Action Details** window appears.
2. Select **FileStructure** from the **Actions** list.
3. Select **Flatten** from the **Action Type** list and click **Save**.

CHAPTER 8

Publications

Publications are entities that define how applications publish data to Cloud Integration Hub, including the type, format, and schedule of data publication. Publications publish data to topics. Multiple publications can publish to the same topic. The topic defines the structure to which the data is published.

Publications can publish from any type of source that Informatica Intelligent Cloud Services can access.

Publication methods

You can use the following methods to publish data and files with Cloud Integration Hub:

Publications that trigger a Data Integration task

When the publication runs, the Cloud Integration Hub server triggers the Data Integration task that is defined for the publication and instructs the Informatica Intelligent Cloud Services data engine to retrieve the data from the publishing application. The data engine runs the Data Integration task, and transfers the source data to the topic on the Cloud Integration Hub relational publication repository.

To publish the data to multiple topics with a single task, create a compound publication. To create a compound publication, use a Data Integration mapping task.

Publications that trigger a file ingestion task

When the publication runs, the Cloud Integration Hub server triggers the file ingestion task that is defined for the publication and instructs Mass Ingestion to retrieve the files from the publishing source. Mass Ingestion runs the file ingestion task, and transfers the source files as-is to the topic on the Cloud Integration Hub file store publication repository.

Publications that publish data with an API

The Cloud Integration Hub Publish Data API publishes to a specific topic on the Cloud Integration Hub relational publication repository.

Use this method to publish small transactions from within a workflow, for example, from within Application Integration.

Publication types

You can publish data and files using the following publication types in Cloud Integration Hub:

Simple Publication

Use this type of publication to publish single or multiple data sets and files to a single topic with a Data Integration or a Mass Ingestion task or with an API.

Compound Publication

Use this type of publication to publish single or multiple data sets to multiple topics with a single Data Integration mapping task.

Publication processes

The publication process depends on the publication type.

Publication process for publications that trigger Data Integration tasks

For publications that trigger Data Integration tasks, the publication process includes retrieving the data from the publisher, running the publication task, and writing the data to the relevant topic in the publication repository. After the publication process ends, each subscriber consumes the published data according to the schedule and the filter that you define when you create the subscription.

The publication process includes the following stages:

1. When the publication is triggered, either according to schedule or by an external API, the Cloud Integration Hub server triggers the Data Integration task that is defined for the publication through an Informatica Intelligent Cloud Services REST API.
2. The publication process uses the Cloud Integration Hub cloud connector to write the data to Cloud Integration Hub.
3. The Cloud Integration Hub server changes the status of the publication event to complete and triggers subscription processing.

Publication process for publications that trigger file ingestion tasks

For publications that trigger file ingestion tasks, the publication process includes running the publication task and writing the files to the relevant file store topic in the publication repository. After the publication process ends, each subscriber consumes the published files according to the schedule that you define when you create the subscription.

The publication process includes the following stages:

1. When the publication is triggered, either according to schedule or by an external API, the Cloud Integration Hub server triggers the file ingestion task that is defined for the publication.
2. The file ingestion task reads the files from the source and then writes the files to the file store publication repository.

3. The Cloud Integration Hub server changes the status of the publication event to complete and triggers subscription processing.

Publication process for publications that publish data with an API

For publications that publish data with an API, you run the Publish Data API. The API retrieves the data from the publisher and writes the data to the relevant topic in the relational publication repository. After the publication process ends, each subscriber consumes the published data according to the schedule and the filter that you define when you create the subscription.

The publication process includes the following stages:

1. The user triggers the Publish Data API.
2. The Publish Data API runs the publication, retrieves the data from the publishing applications, and writes the data to the topic that is defined in the publication.
3. The Cloud Integration Hub server changes the status of the publication event to complete and triggers subscription processing.

Publication mapping

For publications that trigger a Data Integration task, mapping is the data mapping between the publishing source and the Cloud Integration Hub publication repository.

A publication runs a Data Integration task that reads from the source and publishes to the topic tables. Task targets must include at least one of the topic tables, and must not include any target table that is not defined in the topic.

You create the task in Data Integration and then select it when you create the publication in Cloud Integration Hub. Cloud Integration Hub uses an Informatica Intelligent Cloud Services REST API to trigger the task, and the Cloud Integration Hub cloud connector writes the published data to Cloud Integration Hub.

Publication sources

Publications can publish from any type of source that Informatica Intelligent Cloud Services supports.

Publication schedules

For publications that trigger a Data Integration task, the publication schedule defines the frequency of the publication. You can publish the data manually or by an external trigger, or publish the data at defined intervals.

For file publications that are published manually, by an external trigger, or at defined intervals, and that publish multiple files, all the files must be present in the source location when the publication starts.

The publication starts when one of the following conditions is true:

- The scheduled start time arrives.
- You run the publication manually.
- You start the publication from a REST API.

Publication management

Create, disable, and enable publications, and run publications manually, including disabled publications.

Creating a publication that triggers a Data Integration task

Use the Navigator to create publications that trigger a Data Integration task to retrieve the data from the publishing application and write the data to the topic on the Cloud Integration Hub relational publication repository.

The following conditions must exist before you create a publication:

- An application to publish the data from must exist. You can either use an existing application, or create and save a new application.
- A relational topic to publish data to must exist. You can either use an existing topic, or create and save a new topic.
- A publication Data Integration task must exist. For more information, see [“Creating a mapping and task for a publication” on page 71](#).

1. In the Navigator, click **New > Publication**.

The **New Publication** page appears.

2. Enter the publication name. Optionally, enter a description for the publication.
3. Select the publication mode, enabled or disabled.

A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.

4. Select **Publish data with a Data Integration task**.
5. Select the application that publishes the data.
6. Select a topic that stores published data in a relational database. If you use a Data Integration mapping task for the publication, you can select multiple topics to publish the data to.
7. Select the task that defines the publication mapping.
8. If the publication publishes large amounts of data, increase the write batch size to optimize the performance of the publication.

Note: Increasing the batch size increases the memory consumption of the Secure Agent and might impact the performance of the Secure Agent machine.

9. Select the method and the frequency of data publishing.

Manually or by an external trigger

No schedule. You can use the following methods to run the publication:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the publication.

For file publications that use this scheduling option and that publish multiple files, all the files must be present in the source location when the publication starts.

By schedule

Runs the publication according to the defined schedule. Select one of the following options:

- Every n minutes. Runs the publication in intervals of up to 60 minutes. You select the number of minutes from the list.
- Hourly. Runs the publication in intervals of up to 24 hours. You select the number of hours from the list. The publication runs at the beginning of the hour. For example, if you enter 2, the publication runs at 00:00, 02:00, and at consecutive two-hour intervals.
- Daily. Runs the publication at the same hour every day.
- Weekly. Runs the publication every week on one or more days at the same hour. Select the check boxes of the days of the week that the publication runs on. For example, select Saturday and Sunday to schedule the publication to run on weekends.
- Monthly. Runs the publication every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the publication runs. For example, schedule the publication to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the publication in the selected time zone. You can change the default time zone in the user profile.

Define the publication intervals in the **Repeat running** area.

10. Click **Save**.

Creating a publication that triggers a file ingestion task

Use the Navigator to create publications that trigger a file ingestion task to retrieve the files from the publishing application and write the files to a specific topic on the Cloud Integration Hub file store publication repository.

The following conditions must exist before you create a publication:

- An application to publish the files from must exist. You can either use an existing application, or create and save a new application.
- A file store topic to publish files to must exist. You can either use an existing topic, or create and save a new topic
- A publication file ingestion task must exist.

1. In the Navigator, click **New > Publication**.

The **New Publication** page appears.

2. Enter the publication name. Optionally, enter a description for the publication.
3. Select the publication mode, enabled or disabled.

A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.

4. Select **Publish files with a file ingestion task**.
5. Select the application that publishes the files.
6. Select a topic that stores files in a file store to publish the files to.
7. Select the file ingestion task that defines the publication mapping.
8. Select the method and the frequency of files publishing.

Manually or by an external trigger

No schedule. You can use the following methods to run the publication:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the publication.

For publications that use this scheduling option and that publish multiple files, all the files must be present in the source location when the publication starts.

By schedule

Runs the publication according to the defined schedule. Select one of the following options:

- Every n minutes. Runs the publication in intervals of up to 60 minutes. You select the number of minutes from the list.
- Hourly. Runs the publication in intervals of up to 24 hours. You select the number of hours from the list. The publication runs at the beginning of the hour. For example, if you enter 2, the publication runs at 00:00, 02:00, and at consecutive two-hour intervals.
- Daily. Runs the publication at the same hour every day.
- Weekly. Runs the publication every week on one or more days at the same hour. Select the check boxes of the days of the week that the publication runs on. For example, select Saturday and Sunday to schedule the publication to run on weekends.
- Monthly. Runs the publication every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the publication runs. For example, schedule the publication to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the publication in the selected time zone. You can change the default time zone in the user profile.

Define the publication intervals in the **Repeat running** area.

9. Click **Save**.

Creating a publication that publishes data with an API

Use the Navigator to create publications that use the Publish Data REST API to publish the data to a specific topic in the Cloud Integration Hub relational publication repository.

The following conditions must exist before you create a publication:

- An application to publish the data from must exist. You can either use an existing application, or create and save a new application.
 - A relational topic to publish data to must exist. You can either use an existing topic, or create and save a new topic.
1. In the Navigator, click **New > Publication**.
The **New Publication** page appears.
 2. Enter the publication name. Optionally, enter a description for the publication.

3. Select the publication mode, enabled or disabled. A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.
4. Select **Publish data with an API**.
5. Select the application that publishes the data.
6. Select a topic that stores published data in a relational database. If you use a Data Integration mapping task for the publication, you can select multiple topics to publish the data to.
7. Click **Save**.

You can copy the following URLs and use them in the request that runs the publication:

- URL of the REST API. Use this URL to publish the data.
- URL of the Swagger structure for the topic that the publication publishes data to. Use the structure in the publication request. If a Swagger structure base URL is configured for the organization, Cloud Integration Hub appends the base URL to the topic Swagger structure URL. For more information, see [“System Properties” on page 33](#).

Running a publication manually

Use the **Explore** page to manually run publications that trigger Data Integration tasks.

Tip: You can also run publications manually on the topic page. For more information, see [“Topic properties” on page 58](#).

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Publications**.
The **Explore** page shows all existing publications. You can sort the display by name, description, mode, topic, or last modified.
2. Rest on the publication and click the Actions menu at the right end of the line. From the menu select **Run**.

Disabling and enabling a publication

Use the **Explore** page to disable and enable publications. A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.

Tip: You can also disable and enable publications on the topic page. For more information, see [“Publications properties” on page 63](#).

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Publications**.
The **Explore** page shows all existing publications. You can sort the display by name, description, mode, topic, or last modified.
2. Rest on the publication click the Actions menu at the right end of the line. From the menu select **Disable** or **Enable**, as required.

Publication properties

Publication properties include general information about the publication, the application and topic to use for the publication, and, for publications that trigger a Data Integration task or a file ingestion task, the task to run and the publication scheduling.

The publication page can include the following properties:

Publication Name

Name of the publication. The name can contain up to 60 characters and can contain special characters.

Description

Description of the publication. The description can contain up to 255 characters.

Mode

Publication mode, enabled or disabled. A disabled publication does not run according to schedule or by an external API. You can only run a disabled publication from the **Explore** page or from the topic page of the topic that the publication publishes to.

Publication Method

Method by which to publish the data or files:

- **Publish data with a Data Integration task.** The publication process triggers a Data Integration task to retrieve the data from the publishing application and write the data to the topic on the Cloud Integration Hub publication repository.
- **Publish files with a file ingestion task.** The publication process triggers a file ingestion task to retrieve the files from the publishing application and write the files to a specific topic on the Cloud Integration Hub publication repository.
- **Publish data with an API.** Use the Publish Data REST API to publish the data to a specific topic in the Cloud Integration Hub publication repository. Select this option for high-frequency publications of small transactions.

After you configure the publication properties, you can copy the following URLs from the publication page:

- URL of the REST API. Use this URL to publish the data.
- URL of the Swagger structure for the topic that the publication publishes data to. Use the structure in the publication request.

You use the URLs when you create the request that runs the publication.

Application

Application that publishes the data or files.

Topic

Topic to which the application publishes the data or files.

Task

Task that defines the publication mapping. Applies to publications that trigger a Data Integration task or a file ingestion task.

Write Batch Size

Number of records that the Cloud Integration Hub connector writes to the publication repository in a single batch. Applies to publications that trigger a Data Integration task.

Note: If you configure the Cloud Integration Hub connection to use JDBC for private publication repository, batch size doesn't apply.

Scheduling

Method and frequency of data publishing. Applies to publications that trigger a Data Integration task or a file ingestion task.

Manually or by an external trigger

No schedule. You can use the following methods to run the publication:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the publication.

For file publications that use this scheduling option and that publish multiple files, all the files must be present in the source location when the publication starts.

By schedule

Runs the publication according to the defined schedule. Select one of the following options:

- Every n minutes. Runs the publication in intervals of up to 60 minutes. You select the number of minutes from the list.
- Hourly. Runs the publication in intervals of up to 24 hours. You select the number of hours from the list. The publication runs at the beginning of the hour. For example, if you enter 2, the publication runs at 00:00, 02:00, and at consecutive two-hour intervals.
- Daily. Runs the publication at the same hour every day.
- Weekly. Runs the publication every week on one or more days at the same hour. Select the check boxes of the days of the week that the publication runs on. For example, select Saturday and Sunday to schedule the publication to run on weekends.
- Monthly. Runs the publication every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the publication runs. For example, schedule the publication to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the publication in the selected time zone. You can change the default time zone in the user profile.

Define the publication intervals in the **Repeat running** area.

CHAPTER 9

Subscriptions

Subscriptions are entities that define how applications consume data from Cloud Integration Hub. Subscriptions subscribe to topics. Multiple subscriptions can consume data from the same topic.

Subscriptions can consume data into any type of target that Informatica Intelligent Cloud Services can access.

Subscription methods

You can use the following methods to consume data and files with Cloud Integration Hub:

Subscriptions that trigger a Data Integration task

When the subscription runs, the Cloud Integration Hub server triggers the Data Integration task that is defined for the subscription and instructs the Informatica Intelligent Cloud Services data engine to retrieve the data from the topic or topics on the Cloud Integration Hub publication repository. The data engine runs the Data Integration task, and transfers the data to the subscribing application. To consume data from multiple topics with a single task, create a compound subscription.

For subscriptions that trigger a Data Integration task you can define the delivery behavior for the published data, for example, to aggregate all data sets to a single data set, or to consume the latest published data set. You can also configure a retry policy that defines the number of times Cloud Integration Hub retries to run the subscription in case of failure and the retry interval.

Use this method to consume batch data from the relational publication repository into files, applications, and repositories.

Subscriptions that trigger a file ingestion task

When the subscription runs, the Cloud Integration Hub server triggers the file ingestion task that is defined for the subscription and instructs Mass Ingestion to retrieve the files from a specific topic on the Cloud Integration Hub publication repository. Mass Ingestion runs the file ingestion task, and transfers the files to the subscriber target.

Use this method to consume files as-is from the file store publication repository.

Subscriptions that consume data with an API

The Cloud Integration Hub Consume Data API consumes data from a specific topic on the Cloud Integration Hub publication repository.

Use this method for high frequency, event-driven subscriptions. For example, to consume data that is published with the Publish Data API.

Subscription types

You can consume data and files using the following subscription types in Cloud Integration Hub:

Simple Subscription

Use this type of subscription to consume single or multiple data sets or files from a single topic with a Data Integration or a Mass Ingestion task or with an API.

Compound Subscription

Use this type of subscription to consume relational data sets from multiple topics with a single Data Integration mapping task.

Unbound Subscription

An unbound subscription isn't restricted to specific publication instances.

Use this type of subscription to subscribe to all the relational data that a publication publishes and to consume the data based on the subscription filter, regardless of when or in what batch the data was published.

Aggregated Subscription

An aggregated subscription consumes multiple relational data sets from the same topic with a single batch workflow.

Use this type of subscription with an automatic mapping or a custom mapping to process relational data. When you use an automatic mapping, the subscription sorts the relational data according to the publication date and time of the publication instances.

Subscription processes

The subscription process depends on the subscription type.

Subscription process for subscriptions that trigger Data Integration tasks

When a subscription triggers an Data Integration task, the subscription process includes retrieving the required data from the Cloud Integration Hub relational publication repository, running the subscription task, and writing the data to one or more subscriber targets. Cloud Integration Hub keeps the data in the publication repository until the retention period of the topic expires.

The subscription process includes the following stages:

1. When the publication is ready for subscribers, the Cloud Integration Hub server triggers the Data Integration task that is defined for the subscription through an Informatica Intelligent Cloud Services REST API.
2. The subscription process uses the Cloud Integration Hub cloud connector to read data from Cloud Integration Hub.
3. The Data Integration task reads the data from Cloud Integration Hub and then writes the data to the subscribing application.
4. The Cloud Integration Hub server changes the status of the subscription event to complete.

Note: For performance tuning purposes, Cloud Integration Hub writes the data to a folder on the local server for intermediate staging, and then writes the data to the target location. Cloud Integration Hub deletes the data from the local server at the end of the subscription process.

Subscription process for subscriptions that trigger file ingestion tasks

When a subscription triggers a file ingestion task, the subscription process includes running the subscription task and writing the files to a subscriber target. Cloud Integration Hub keeps the files in the file store publication repository until the retention period of the topic expires.

The subscription process includes the following stages:

1. When the publication is ready for subscribers, the Cloud Integration Hub server triggers the file ingestion task that is defined for the subscription.
2. The file ingestion task reads the files from the file store publication repository and then writes the files as-is to the target.
3. The Cloud Integration Hub server changes the status of the subscription event to complete.

Subscription process for subscriptions that consume data with an API

When a subscription consumes data with an API, you run the Consume Data API in order to consume the data. The API retrieves the data from the topic in the relational publication repository and writes the data to the subscribing application.

When you create or edit a subscription that consumes data with an API, you can define a notification URL. Cloud Integration Hub sends notifications to this URL when data is ready to consume. Cloud Integration Hub must be able to access the notification URL.

The subscription process includes the following stages:

1. The user triggers the Consume Data API.
2. The Consume Data API runs the subscription, retrieves the data from the topic that is defined in the subscription, and writes the data to the subscribing application.

If the subscription process fails, you can attempt to consume the published data by reprocessing the subscription Error event with the Consume Data REST API.

You can reconsume data that had previously been processed by triggering the subscription Completed event with the Consume Data REST API.

Subscription mapping

For subscriptions that trigger Data Integration tasks, mapping is the data mapping between the Cloud Integration Hub publication repository and the target that consumes the data.

A subscription runs a Data Integration task that includes information about the target data structure and the database connection. The task reads from the topic tables and consumes the data into the target application.

You can create the task in Data Integration and then select it when you create the subscription in Cloud Integration Hub, or create the task when you create the subscription. Cloud Integration Hub triggers the task

when the publication is ready for subscribers and uses the Cloud Integration Hub cloud connector to read the data from Cloud Integration Hub.

You can create a compound subscription, where the subscription consumes data sets from multiple topics. The subscription process starts after all publications from all topics publish data. You can specify the maximum time to wait for all publications to finish publishing, from the time the first publication is ready to consume.

Subscription targets

Subscriptions can consume data into any type of target that Informatica Intelligent Cloud Services supports.

Subscription schedules

For subscriptions that trigger Data Integration tasks, the subscription schedule defines the frequency of the subscription. You can consume published data when it is published, manually, by an external trigger, or at defined intervals. If you create a compound subscription, you can only choose to consume data when it is published, manually, or by an external trigger.

Consumption of data by the subscription starts when one of the following conditions exist:

- The subscription schedule is set to consumes data immediately after the publisher publishes the data to Cloud Integration Hub.
- The scheduled start time arrives.
- You start the subscription from a REST API.
- You manually run a subscription.
- You manually get previous publications.

Subscription retry policy

To improve business continuity, you can configure a retry policy for subscriptions that trigger Data Integration tasks. The policy defines the number of times Cloud Integration Hub retries to run the subscription in case of failure and the retry interval.

You can define a policy of up to nine retry attempts with a retry interval that is between five minutes and 23 hours. Cloud Integration Hub attempts to reprocess subscription events in an Error status based on the policy you define. Cloud Integration Hub doesn't attempt to reprocess Error events in the following scenarios:

- You manually change the status of an Error event to Complete.
- You manually change the status of an event to Error.
- You manually reprocess an Error event and the subscription runs successfully.

When Cloud Integration Hub attempts to run a subscription according to the policy, the details of the subscription event on the **Events** page indicate that the attempt was based on a retry policy.

When you define a retry policy for a subscription, make sure that the policy doesn't conflict with the subscription schedule. If a conflict occurs, one of the Processing events is delayed and the subscription consumes the data when it next runs according to its schedule.

Subscription management

Create, disable, and enable subscriptions, get previous publications for a subscription, and run a subscription manually, including disabled subscriptions.

Creating a subscription that triggers a Data Integration task

Use the Navigator to create subscriptions that trigger a Data Integration task to retrieve the data from the topic or topics in the Cloud Integration Hub relational publication repository.

The following conditions must exist before you create a subscription:

- An application or applications that consume data must exist. You can either use existing applications, or create and save new applications.
- A relational topic to consume data from must exist. You can either use an existing topic, or create and save a new topic.
- If the subscription triggers a mapping task, a subscription task must exist in Data Integration. If the subscription triggers a synchronization task, you can either select a subscription task that exists in Data Integration or create the task.

1. On the Navigator, click **New > Subscription**.

The **New Subscription** page appears.

2. Enter the subscription name. Optionally, enter a description for the subscription.
3. Select the subscription mode, enabled or disabled.

A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

4. Select **Consume data with a Data Integration task**.
5. Optionally, select **Unbound Subscription**.

A subscription that is not restricted to specific publication instances. It consumes all the publication events data in the publication repository for the topics that the subscription subscribes to. Applies to subscriptions that consume data with a Data Integration task.

6. Select the application that subscribes to the data.
7. Select the topic that stores data in a relational database from which the application consumes the data and then click **Add Topic**. Add as many topics as required.
8. If you added more than one topic to the subscription, specify the maximum number of hours to wait for all associated publications to finish publishing the data, after the first publication is ready for consumption.
 - If all the publications finish publishing the data during the time interval, the subscription process starts after the last publication is ready for consumption.
 - If one or more of the publications do not finish publishing the data during the time interval, the subscription process is cancelled and no data is delivered.

9. If the task that defines the subscription mapping exists in Data Integration, choose the task. If not, click **Create New Task** to create a synchronization task.
10. To create a synchronization task, enter the following properties in the **Create New Task** window and click **Create**:

Task Name

Enter a name for the task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Source

Select the topic table to consume data from. The format of the object is `TopicName/tableName`.

Connection

Select the connection that connects to the target to consume data to.

Target

Select the target table to consume the data to. The **Create New Task** window shows the first 200 tables in the list.

Cloud Integration Hub creates the task in the default folder and assigns the task to the subscription.

11. If the subscription subscribes to large amounts of data, increase the read batch size to optimize the performance of the subscription.

Note: Increasing the batch size increases the memory consumption of the Secure Agent and might impact the performance of the Secure Agent machine.

12. Select the method and the frequency of data consumption.

When published data is ready

Runs the subscription immediately after the published data is ready.

Manually or by an external trigger

No schedule. You can use the following methods to run the subscription:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the subscription.

If a file subscription uses this scheduling option and publishes multiple files, all the files must be present in the source location when the subscription starts.

By schedule

Runs the subscription according to the defined schedule. Select one of the following options:

- Every n minutes. Runs the subscription in intervals of up to 60 minutes. You select the number of minutes from the list.
- Hourly. Runs the subscription in intervals of up to 24 hours. You select the number of hours from the list.
- Daily. Runs the subscription at the same hour every day.
- Weekly. Runs the subscription every week on one or more days at the same hour. Select the check boxes of the days of the week that the subscription runs on. For example, select Saturday and Sunday to schedule the subscription to run on weekends.

- **Monthly.** Runs the subscription every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the subscription runs. For example, schedule the subscription to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the subscription in the selected time zone. You can change the default time zone in the user profile.

Define the delivery intervals in the **Repeat running** area.

Data to Consume

If a subscription runs manually or by an external trigger, or runs by a schedule, multiple files from the same publication might be available for consumption at the scheduled time, or at the time the subscription is triggered. Select one of the following options:

- **All available publications.** Consumes data or files from all available publications.
- **Latest publication only.** Consumes data or files from only the latest publication.

13. Optionally, in the **Retry Policy** area, select **Reprocess Events in Error Status** and then select the number of times Cloud Integration Hub retries to run the subscription in case of failure and the retry interval. You can define a policy of up to nine retry attempts with a retry interval that is between five minutes and 23 hours.
14. Click **Save**.

Creating a subscription that triggers a file ingestion task

Use the Navigator to create subscriptions that trigger a file ingestion task to retrieve the files from a specific topic in the Cloud Integration Hub file store publication repository.

The following conditions must exist before you create a subscription:

- An application or applications that consume files must exist. You can either use existing applications, or create and save new applications.
- A file store topic from which to consume files must exist. You can either use an existing topic, or create and save a new topic.
- A subscription file ingestion task must exist.

1. On the Navigator, click **New > Subscription**.

The **New Subscription** page appears.

2. Enter the subscription name. Optionally, enter a description for the subscription.
3. Select the subscription mode, enabled or disabled.

A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

4. Select **Consume files with a file ingestion task**.
5. Select the application that subscribes to the files.
6. Select the topic that stores files in a file store from which the application consumes the files and then click **Add Topic**.
7. Select the file ingestion task that defines the subscription mapping.
8. Select the method and the frequency of file consumption.

When published data is ready

Runs the subscription immediately after the published files are ready

Manually or by an external trigger

No schedule. You can use the following methods to run the subscription:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the subscription.

If a file subscription uses this scheduling option and publishes multiple files, all the files must be present in the source location when the subscription starts.

By schedule

Runs the subscription according to the defined schedule. Select one of the following options:

- Every n minutes. Runs the subscription in intervals of up to 60 minutes. You select the number of minutes from the list.
- Hourly. Runs the subscription in intervals of up to 24 hours. You select the number of hours from the list.
- Daily. Runs the subscription at the same hour every day.
- Weekly. Runs the subscription every week on one or more days at the same hour. Select the check boxes of the days of the week that the subscription runs on. For example, select Saturday and Sunday to schedule the subscription to run on weekends.
- Monthly. Runs the subscription every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the subscription runs. For example, schedule the subscription to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the subscription in the selected time zone. You can change the default time zone in the user profile.

Define the delivery intervals in the **Repeat running** area.

Data to Consume

If a subscription runs manually or by an external trigger, or runs by a schedule, multiple files from the same publication might be available for consumption at the scheduled time, or at the time the subscription is triggered. Select one of the following options:

- **All available publications.** Consumes data or files from all available publications.
- **Latest publication only.** Consumes data or files from only the latest publication.

9. Click **Save**.

Creating a subscription that consumes data with an API

Use the Navigator to create subscriptions that use the Consume Data REST API to consume the data from a specific topic in the Cloud Integration Hub relational publication repository.

The following conditions must exist before you create a subscription:

- An application or applications that consume data must exist. You can either use existing applications, or create and save new applications.

- A relational topic to consume data from must exist. You can either use an existing topic, or create and save a new topic.

1. On the Navigator, click **New > Subscription**.

The **New Subscription** page appears.

2. Enter the subscription name. Optionally, enter a description for the subscription.
3. Select the subscription mode, enabled or disabled.

A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

4. Select **Consume data with an API**.

5. Optionally, enter a notification URL. Cloud Integration Hub sends notifications to this URL when data is ready to consume.

For more information, see [“Subscription properties” on page 97](#).

6. Select the application that subscribes to the data.
7. Select the topic that stores data in a relational database from which the application consumes the data and then click **Add Topic**.
8. Click **Save**.

You can copy the following URLs and use them in the request that runs the subscription:

- URL of the REST API. Use this URL to consume the data.
- URL of the Swagger structure for the topic that the subscription consumes data from. If a Swagger structure base URL is configured for the organization, Cloud Integration Hub appends the base URL to the topic Swagger structure URL. For more information, see [“System Properties” on page 33](#).

Running a subscription manually

Use the **Explore** page to manually run subscriptions that trigger Data Integration tasks.

Tip: You can also run subscriptions manually on the topic page. For more information, see [“Topic properties” on page 58](#).

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Subscriptions**.

The **Explore** page shows all existing subscriptions. You can sort the display by name, description, mode, topic, or last modified.

2. Rest on the subscription and click the Actions menu at the right end of the line. From the menu select **Run**.

Getting previous publications for a subscription

Use the **Explore** page to get data that was published before the subscription subscribed to the topic and therefore was not consumed by the subscriber. The generated subscription events will run according to the subscription schedule.

Tip: You can also get previous publications for a subscription on the topic page. For more information, see [“Subscriptions properties” on page 63](#).

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Subscriptions**.
The **Explore** page shows all existing subscriptions. You can sort the display by name, description, mode, topic, or last modified.
2. Rest on the subscription for which to get previous publications and click the Action menu at the right end of the line. From the menu select **Get Previous Publications**, define the date range for which to get the publications, and then click **Run**.

Disabling and enabling a subscription

Use the **Explore** page to disable and enable subscriptions. A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

Tip: You can also disable and enable subscriptions on the topic page. For more information, see [“Subscriptions properties” on page 63](#).

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Hub Management > Subscriptions**.
The **Explore** page shows all existing subscriptions. You can sort the display by name, description, mode, topic, or last modified.
2. Rest on the subscription to disable or to enable and click the Action menu at the right end of the line. From the menu select **Disable** or **Enable**, as required.

Subscription properties

Subscription properties include general information about the subscription, the applications, topic, and task to use for the subscription, and subscription scheduling.

The subscription page can include the following properties:

Subscription Name

Name of the subscription. The name can contain up to 60 characters and can contain special characters.

Description

Description of the subscription. The description can contain up to 255 characters.

Mode

Subscription mode, enabled or disabled. A disabled subscription does not run according to schedule or by an external API. You can only run a disabled subscription from the **Explore** page or from the topic page of the topic that the subscription subscribes to.

Consumption Method

Method by which the subscription consumes data or files:

- **Consume data with a Data Integration task.** The subscription process triggers a Data Integration task to retrieve the data from the topic or topics in the Cloud Integration Hub relational publication repository and write the data to the subscribing application. Select this method to consume batch data into files, applications, and repositories.
- **Consume files with a file ingestion task.** The subscription process triggers a file ingestion task to retrieve the files from a specific topic in the Cloud Integration Hub file store publication repository and write the files to the subscribing application. Select this method to consume files into a target location.
- **Consume data with an API.** Use the Consume Data REST API to consume the data from a specific topic in the Cloud Integration Hub relational publication repository. Select this method for high frequency, event-driven subscriptions.

After you configure the subscription properties, you can copy the following URLs from the subscription page:

- URL of the REST API. Use this URL to consume the data.
- URL of the Swagger structure for the topic from which the subscription consumes data. Use the structure in the subscription request.

You use the URLs when you create the request that runs the subscription.

Unbound Subscription

A subscription that is not restricted to specific publication instances. It consumes all the publication events data in the publication repository for the topics that the subscription subscribes to. Applies to subscriptions that consume data with a Data Integration task.

Application

Application that consumes the data or files.

Topics

Topic or topics from which the application consumes the data or files.

Notification URL

URL to where Cloud Integration Hub sends notifications when data is ready to consume. Applies to subscriptions that consume data with an API.

The notification URL cannot be authenticated and the HTTP request method must be POST. The payload of the POST request must include the following parameters:

Parameter	Description
publicationEventId	ID of the event of the publication that published the data to consume.
subscriptionEventId	ID of the event of the subscription to consume the data.
subscriptionName	Name of of the subscription to consume the data.

For example:

```
{"publicationEventId":123, "subscriptionEventId" : 234, "subscriptionName" :  
"payrollSubscription"}
```

You must immediately return the notification URL with an HTTP 200 SUCCESS response.

Wait for all topics to be available for consumption for ... hours

Maximum time to wait until all published data is available from the time that the first topic is ready to consume. Applies to compound subscriptions that consume data from multiple topics.

If all of the publications in all topics finish publishing the data before the maximum time, the subscription process runs immediately after the last publication is ready to consume. If some publications are not ready to consume within the maximum time, the subscription process does not run. An error event is created, and no data is delivered.

Task

Task that defines the subscription mapping. Applies to subscriptions that trigger a Data Integration task or a file ingestion task.

Create New Task

Create a synchronization task that defines the subscription mapping. Applies to subscriptions that trigger a Data Integration task.

The Create New Task window includes the following properties:

Task Name

Enter a name for the task.

The name of the task must be unique within the organization. The task name is not case sensitive.

The task name can contain alphanumeric characters, spaces, and the following special characters:

_ . + -

Source

Select the topic table to consume data from. The format of the object is `TopicName/tableName`.

Connection

Select the connection that connects to the target to consume data to.

Target

Select the target table to consume the data to. The **Create New Task** window shows the first 200 tables in the list.

Read Batch Size

Number of records that the Cloud Integration Hub connector reads from the publication repository in a single batch. Applies to subscriptions that trigger a Data Integration task.

Scheduling

Method and frequency of data consumption. Applies to subscriptions that trigger a Data Integration task or a file ingestion task.

When published data is ready

Runs the subscription immediately after the published data is ready.

Manually or by an external trigger

No schedule. You can use the following methods to run the subscription:

- Run manually from the Cloud Integration Hub explorer.
- Run by an API. Call a REST API that starts the subscription.

If a file subscription uses this scheduling option and publishes multiple files, all the files must be present in the source location when the subscription starts.

By schedule

Runs the subscription according to the defined schedule. Select one of the following options:

- **Every n minutes.** Runs the subscription in intervals of up to 60 minutes. You select the number of minutes from the list.
- **Hourly.** Runs the subscription in intervals of up to 24 hours. You select the number of hours from the list.
- **Daily.** Runs the subscription at the same hour every day.
- **Weekly.** Runs the subscription every week on one or more days at the same hour. Select the check boxes of the days of the week that the subscription runs on. For example, select Saturday and Sunday to schedule the subscription to run on weekends.
- **Monthly.** Runs the subscription every month on a specific date or a specific day at the same hour.

For minutes and hourly intervals, you can define a period of the day when the subscription runs. For example, schedule the subscription to run during nighttime hours.

Select a time zone for the schedule. The schedule runs the subscription in the selected time zone. You can change the default time zone in the user profile.

Define the delivery intervals in the **Repeat running** area.

Data to Consume

If a subscription runs manually or by an external trigger, or runs by a schedule, multiple files from the same publication might be available for consumption at the scheduled time, or at the time the subscription is triggered. Select one of the following options:

- **All available publications.** Consumes data or files from all available publications.
- **Latest publication only.** Consumes data or files from only the latest publication.

Retry Policy

Defines the number of times Cloud Integration Hub retries to run the subscription in case of failure and the retry interval. Applies to subscriptions that trigger a Data Integration task. Configure the following parameters:

- **Reprocess Events in Error Status.** Enables the retry policy.
- **Retry ... times at ... intervals.** Select the number of retry attempts and the time interval. You can define a policy of up to nine retry attempts with a retry interval that is between five minutes and 23 hours.

CHAPTER 10

Tracking and monitoring

Cloud Integration Hub generates events as it processes publications and subscriptions to help you track and monitor the publication and subscription processes. The event list provides full visibility into the processes and alerts you to errors that might occur.

Cloud Integration Hub generates file events for files that it receives and sends.

Cloud Integration Hub generates events as it processes publications and subscriptions, and it changes the status of the events as they go through the process. You can view all events on the **Events** page. From the **Events** page you can access the event history, session log, and processing information, and reprocess events or change the event status. You can use filters to search for specific events. You can view up to 2,000 of the most recent events in search results on the **Events** page.

If your organization uses both Data Integration Hub and Cloud Integration Hub, you can view Data Integration Hub publication and subscription events on the **Events** page in Cloud Integration Hub. To set up Cloud Integration Hub to show Data Integration Hub events, see [“Setting up Cloud Integration Hub to show Data Integration Hub events” on page 32](#).

You can create rules that monitor publication and subscription events, and perform actions on events that are in a defined status. For example, you can create rules to perform the following tasks:

- Disable publications that have events with an Error status.
- Send an email to the Cloud Integration Hub administrator when a subscription event is in an Error status.

Publication and Subscription Events

The **Events** page provides detailed event processing information for every publication and subscription that Cloud Integration Hub processed in the past three months.

The Publication event is the root event and the parent event for all of the subscription events that Cloud Integration Hub generates during processing. After the published data is ready for subscribers, Cloud Integration Hub generates a Subscription child event for each subscriber that needs to consume the published data. The Publication event contains aggregated status information for all Subscription child events.

By default, the **Events** page displays root events: Publication, File, Aggregated Subscription, and Compound Subscription. After a publication is ready for subscribers, you can drill down to the associated Subscription child events of the publication.

Event Types

Cloud Integration Hub assigns the following event types to publication and subscription events:

- **Publication.** Assigned to a publication process. Acts as the parent event for all Subscription events and for File events of publications that publish multiple files.
- **Compound Publication.** Assigned to a publication process that publishes single or multiple data sets to multiple topics with a single Data Integration mapping task.
- **Subscription.** Assigned to a subscription process. Acts as a child event for a publication event.
- **Compound Subscription.** Assigned to a subscription process that consumes data sets from multiple topics with a single subscription mapping. The event contains references to all Subscription events that Cloud Integration Hub creates when each topic publication finished publishing the data set.
- **Unbound Subscription.** Assigned to a subscription process that is not restricted to specific publication instances but subscribes to all the data that a publication publishes regardless of when or in what batch the data was published.
- **Aggregated Subscription.** Assigned to a subscription process that consumes multiple data sets from the same topic with a single subscription mapping. The event contains references to all Subscription events that were created when the associated topic finished publishing each data set. The Subscription events inherit their status from the Aggregated Subscription event.
- **System.** Event generated for system notifications. For example, Cloud Integration Hub generates a system event when a compound subscription cannot consume published data from all required publications.

Event statuses

For publications, Cloud Integration Hub assigns the following event statuses:

- **Processing.** Indicates that the publication instance is running.
- **Completed.** Indicates that the publication instance finished running and that the data is ready for subscribers.
- **Error.** Indicates that the publication instance encountered errors and did not finish running.

Note: When you publish data through the Publish Data REST API to a private publication repository and the publication fails because the publication repository service is not accessible, Cloud Integration Hub returns an error to the calling application and does not create an error event.

Each Publication event also shows the consumption status of the child Subscription events. The status reflects the overall consumption and changes after all Subscription events changed status. For example, the consumption status changes to complete after all subscribers finished consuming the published data.

For subscriptions, Cloud Integration Hub assigns the following event statuses:

- **Delayed.** Indicates that the published data is ready but that the subscribing application did not start consuming the data.
- **Processing.** Indicates that the subscription instance is running.
- **Completed.** Indicates that the subscription instance finished running and that the subscribing application consumed all published data.
- **Error.** Indicates that the subscription instance encountered errors and did not finish running.
- **Discarded.** Indicates that the subscription instance was changed to Discarded.

Note: If you don't enable simultaneous task runs when you run parallel instances of a compound subscription or an aggregated subscription, Cloud Integration Hub processes the first subscription event and assigns all

the remaining subscription events the Discarded status. For more information, see ["Creating a mapping task for a subscription" on page 73](#).

When you hover over the Event Status icon on the **Events** page, event details appear. For example, the time when the event processing completed, the time when the event changed status, or the cause of the error in Error events.

Event Consumption Statuses

Cloud Integration Hub assigns the following consumption statuses to publication and subscription events:

- Processing. Cloud Integration Hub is processing the publication or the subscription.
- Final. For publications, all data is published. For subscriptions, all data is consumed.
- Delayed. Applicable for subscriptions only. Data is ready but that the subscribing application did not start consuming the data.
- Error. An error occurred during data publication or consumption.

Event History

You can view the event status history for each publication or subscription that the Cloud Integration Hub processes.

The event history shows the processing stages that the publication or subscription passed through, when each stage started, and the cumulative processing status.

The following table describes the processing stages that can show in the Event History for publications:

Stage	Description
Processing	The publication instance is running.
Complete	The publication instance finished running and data is ready for subscribers.
Error	The publication instance encountered errors and did not finish running.
Discarded	The status of the publication instance was changed to Discarded.

The following table describes the processing stages that can show in the Event History for subscriptions:

Stage	Description
Delayed	The subscription instance is delayed. Published data is ready but the subscribing application did not start consuming the data.
Processing	The subscription instance is running.
Complete	The subscription instance finished running and the subscribing application consumed all published data.
Error	The subscription instance encountered errors and did not finish running.

Stage	Description
Reprocessed	The subscription instance was reprocessed.
Discarded	The status of the subscription instance was changed to Discarded.

Event Session Log

Each time that a publication or a subscription that triggers a Data Integration task runs, Cloud Integration Hub generates a task in Informatica Intelligent Cloud Services.

You can access the task session log from the specific event.

If an error occurs during file processing, you can use the related session log to view further information about the error.

Event Processing Information

Each time that a publication or a subscription that triggers a Data Integration task runs, Cloud Integration Hub generates a task in Informatica Intelligent Cloud Services.

You can access the task processing information from the specific event.

System Event Maintenance Report

For system events, Cloud Integration Hub generates a maintenance report.

You can access the report from the Actions menu of the event.

Event Filters

You can use filters to narrow the view of the **Events** page to show events for event ID, type or status, show events for a selected application, topic, publication, or subscription, or show events for a selected time frame.

You can click the Filter icon to expand the filter pane and define the filtering criteria. After you click **Apply Filter**, the event list updates to show the relevant events.

By default the event list shows all events from last 24 hours. After you filter the view of the list, to restore the default view, click **Restore Defaults**.

Managing Events

Reprocess an event and change the status of an event.

Note: You can only perform these operations on Cloud Integration Hub events.

Reprocessing an Event

Use the **Events** page to reprocess events. You can reprocess only subscription events, to re-consume data that was already consumed.

1. In the Navigator, click **Events**.
The **Events** page appears. By default the page shows all events from the last 24 hours.
2. Use the filter pane to filter the view of the page.
3. Rest on the event to reprocess and click the Action menu at the right end of the line. From the menu select **Reprocess** and then confirm the action.

Changing Event Status

Use the **Events** page to change the status of events.

1. In the Navigator, click **Events**.
The **Events** page appears. By default the page shows all events from the last 24 hours.
2. Rest on the event to reprocess and click the Action menu at the right end of the line. From the menu select **Change Event Status**.
3. In the **Change Event Status** dialog box select the new event status then click **OK**.

Discarding all events in Processing status

Use the **Events** page to discard all events in Processing status, for example before you update a topic.

1. In the Navigator, click **Events**.
The **Events** page appears. By default the page shows all events from the last 24 hours.
2. Use the filter pane to filter the view of the page.
3. In the **Status** field, select **Processing**.
4. In the **Topic** field, enter the topic name.
5. In the **Time Frame** field, select **All**. Click **Apply Filter**.
6. Select all of the filtered events.
7. Click **All Selected** and select **Change Event Status**.
8. In the **Change Event Status** dialog box, select **Discarded** and click **OK**.
The event status changes to Discarded for all of the selected events.

Event Properties

Event properties include general information about the event, the applications, topic, and task to use for the event, and event scheduling.

The following image shows a sample **Events** page:

The screenshot shows the 'Events' page with a table of 1103 events. The table has columns for Event ID, Asset Source, Application, Publication/Subscription, Topic, Start Time, Event Status, and Consumption Status. The first few rows show events with IDs 167661, 167655, 167648, 9081081, 9081080, 9081079, 9081078, 9081076, 9081075, 167641, 167632, 167626, and 167619. The Event Status column shows various icons: a green checkmark for successful events, a red circle with a slash for failed events, and a grey circle for processing events. The Consumption Status column shows a green checkmark for consumed events and a red circle with a slash for failed consumption. The page includes filters for Event ID, Type, Status, Asset Source, Application, Topic, Pub/Sub, and Time Frame. The Time Frame is set to 'Last 30 days'. The page also shows a summary of 25 Error events, 185 Delayed events, and 4 Processing events.

Event ID	Asset Source	Application	Publication/Subscription	Topic	Start Time	Event Status	Consumption Status
167661	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	03:20:00.350 pm 29/03/2021	✓	✓
167655	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	03:15:00.392 pm 29/03/2021	✓	✓
167648	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	03:10:00.274 pm 29/03/2021	✓	✓
9081081	DIH	Schedulers	FFP20	schedTopic2	03:06:20.731 pm 29/03/2021	○	
9081080	DIH	Schedulers	FFI12	schedTopic2	03:06:10.878 pm 29/03/2021	○	
9081079	DIH	Schedulers	FFP20	schedTopic2	03:05:50.488 pm 29/03/2021	○	
9081078	DIH	Schedulers	FFI12	schedTopic2	03:05:39.703 pm 29/03/2021	○	
9081076	DIH	NewApp	clpa_Copy	abcd	03:05:25.838 pm 29/03/2021	✓	✓
9081075	DIH	NewApp	clpa	abcd	03:05:23.936 pm 29/03/2021	✗	
167641	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	03:05:00.319 pm 29/03/2021	✓	✓
167632	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	03:00:00.404 pm 29/03/2021	✓	✓
167626	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	02:55:00.301 pm 29/03/2021	✓	✓
167619	CIH	MCT_APP	MCT_PUB1	DL_MCT_TOP	02:50:00.283 pm 29/03/2021	✓	✓

The **Events** page includes the following properties:

Event ID

ID of the event.

By default, the Events page shows only parent events. To show the list of subscription events for a publication event, expand the publication event.

Asset Source

Source of the asset that generated the event.

This filter appears when there are Data Integration Hub events on the **Events** page.

Application

For publication events, the application that publishes the data. For subscription events, the application that consumes the data.

Publication/Subscription

Name of the publication of subscription for which Cloud Integration Hub generates the event.

Topic

For publication events, the topic that the application publishes the data to. For subscription events, the topic or topics from which the application consumes the data.

Start Time

Time when the event started.

Event Status

Status of the event.

Consumption Status

Applicable for publication events. Data consumption status for the event.

Open Data Integration Hub assets from Cloud Integration Hub

You can open Data Integration Hub publication and subscription assets from the **Events** page in Cloud Integration Hub.

Click the asset in a Data Integration Hub event on the Cloud Integration Hub **Events** page to open Data Integration Hub in a new tab and view the Data Integration Hub asset that generated the event.

Event Monitors

You can create event monitors that track publications and subscriptions based on their event status, and instigate actions when an event is in a defined status.

You create monitoring rules that define which entities to monitor, what are the event statuses for which to take action, and what actions Cloud Integration Hub takes when an event reaches the defined status.

You can create rules that monitor publication and subscription events, and perform actions on events that are in a defined status. For example, you can create rules to perform the following tasks:

- Disable publications that have events with an Error status.
- Send an email to the Cloud Integration Hub administrator when a subscription event is in an Error status.

Monitoring Rules

A monitoring rule defines which assets to monitor, the event statuses that trigger actions, and the actions to take when an event is in a defined status.

When you create a monitoring rule, you define the following elements:

- Asset or assets to which the rule applies. A rule can apply to a single publication, to multiple publications, or to all current and future publications, or to a single subscription, to multiple subscriptions, or to all current and future subscriptions
- Event status or statuses to which the rule applies. Cloud Integration Hub applies the rule only to events that are in a final state.
- Rule action or actions. You can select one or more of the following actions:
 - Send email notification. You define the user or users to which Cloud Integration Hub sends an email notification when the rule conditions are true.
 - Pause subscriptions or disable publications and subscriptions that are in the status or statuses to which the rule applies.

Managing Monitoring Rules

Create, edit, view, disable, enable, and delete monitoring rules.

Creating a monitoring rule

Use the Navigator to create monitoring rules.

1. In the Navigator, click **New > Monitoring Rule**. Then click **Create**.
The **New Monitoring Rule** page appears.
2. Enter the rule name. Optionally, enter a description for the rule.
3. Select the location to save the rule.
4. Choose the rule mode, enabled or disabled. A disabled rule does not perform the defined actions.
5. Select the type of asset that the rule affects, publication or subscription, and then select the asset or assets to which the apply the rule. You must apply the rule to at least one publication or one subscription.
 - To apply the rule to all publications or to all subscriptions, including current publications or subscriptions and publications or subscriptions that are added to Cloud Integration Hub after you create the rule, select **Apply to all**.
 - To select a single publication or a single subscription to which to apply the rule, select the check box to the left of the publication name or the subscription name.
 - To select multiple publications or multiple subscriptions to which to apply the rule, select multiple check boxes to the left of the publication names or the subscription names.
6. Select the event statuses to monitor. You must select at least one status.
7. Select one or both of the following rule actions:

Send email notification

Send email notifications when a publication or a subscription is in one of the affected statuses. You can send notifications to existing Cloud Integration Hub users or to email addresses that you specify. You can define up to 30 email recipients.

Perform the following steps for each user:

1. Click **Add** to the right of **Send email notification**.
2. Select the name of an existing user or select a non-existing user from the **User Name** list and then enter the email address in the **Email** field.

Cloud Integration Hub sends email notifications to the recipients that you define here when events of any of the affected publications or subscriptions are in any of the affected statuses.

Disable publications and subscriptions that are in Rejected statuses

Select **Disable publications and subscriptions**.

Cloud Integration Hub disables the affected publications or subscriptions when their events are in any of the affected statuses. A disabled publication or subscription does not run according to schedule or by an external API. You can only run a disabled publication or subscription from the **Explore** page.

8. Click **Save**.

Editing a Monitoring Rule

Use the **Explore** page to edit monitoring rules.

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Monitors > Monitoring Rules**.
The **Explore** page shows all existing monitoring rules.
2. Click the name of the monitoring rule to edit.

The monitoring rule page shows.

3. Edit the monitoring rule and then click **Save**.

Disabling and enabling a monitoring rule

Use the **Explore** page to disable and enable a monitoring rule.

1. In the Navigator, click **Explore**. Click the **All Assets** list and then select **Monitors > Monitoring Rules**. The **Explore** page shows all existing monitoring rules.
2. In the row that contains the rule, click **Actions** and select one of the following actions:
 - To disable a rule select **Disable**. A disabled rule does not perform the defined actions.
 - To enable a disabled rule select **Enable**.

Monitoring Rule Properties

Monitoring rule properties include general information about the monitoring rule, the asset or assets to which the rule applies, the event statuses that the rule monitors, and the rule action or actions.

The following image shows a sample monitoring rule page:

The monitoring rule page includes the following properties:

Rule Name

Name of the monitoring rule. The name can contain up to 60 characters and can contain special characters.

Description

Description of the monitoring rule. The description can contain up to 255 characters.

Mode

Monitoring rule mode, enabled or disabled. A disabled rule does not perform the defined actions.

Content

The conditions of the monitoring rule.

Affected Assets

The publications or subscriptions that the rule applies to.

Affected Statuses

The statuses of the affected assets that the rule applies to.

Actions

The actions that the rule performs when any of the affected assets are in any of the affected statuses.

CHAPTER 11

Cloud Integration Hub REST APIs

Use the Cloud Integration Hub REST APIs to run publications and subscriptions, to publish data directly to a specific topic and to consume data directly from a specific topic, to enable and disable publications and subscriptions, to reprocess publication and subscription events, to query the status of publication and subscription events, and to extract data from the Cloud Integration Hub catalog.

You can use the following REST APIs:

Run Publication Subscription

Starts a publication or a subscription, including disabled publications and subscriptions, and returns the event ID of the publication or the subscription event that Cloud Integration Hub generates.

You can use the Run Publication Subscription REST API to publish data and subscribe to data with publications and subscriptions that trigger a Data Integration task and with publications and subscriptions that trigger a file ingestion task. You can't use the API to publish data with publications that publish data directly to a topic or to consume data with subscriptions that consume data directly from a topic.

Publish Data

Publishes data directly to a topic on the Cloud Integration Hub publication repository. Returns the status of a publication process.

You can use the Publish Data API to publish data with publications that publish data with an API. You can't use the API with publications that trigger a Data Integration task or a file ingestion task.

Consume Data

Consumes data directly from a topic on the Cloud Integration Hub publication repository. You can't use the Consume Data API to consume files directly from a topic on a file store. You can use the Consume Data API to consume data with subscriptions that consume data with an API. You can't use the API with subscriptions that trigger a Data Integration task or a file ingestion task.

If the subscription process fails, you can attempt to consume the published data by reprocessing the subscription Error event with the API.

You can reconsume data that had previously been processed by triggering the subscription Complete event with the API.

Change Publication Subscription Mode

Changes the mode of a publication or a subscription, that is, enables a disabled publication or subscription and disables an enabled publication or subscription.

Reprocess Event

Reprocesses subscription events, including events of disabled subscriptions.

You can use the Reprocess Event REST API to reprocess events of subscriptions that trigger a Data Integration task. You can't use the API to reprocess events of subscriptions that consume data with an API.

Event Status

Returns the status of a publication or subscription event.

Catalog

Extracts data from the Cloud Integration Hub catalog, including topic, publication, and subscription metadata.

Authorization Header

Each Cloud Integration Hub REST API call must contain an authorization header.

The type of the authorization header must be Basic, and the header must include an Informatica Intelligent Cloud Services user and an Informatica Intelligent Cloud Services password.

For example:

```
{
  Username: Administrator@MyCompany.com
  Password: MyPassword
}
```

Run Publication Subscription REST API

Use the Cloud Integration Hub Run Publication Subscription REST API to run a specific publication or a specific subscription. You can run the publication or the subscription regardless of its mode, that is, you can run both enabled and disabled publications and subscriptions.

Note: You can use the Run Publication Subscription REST API to publish data and subscribe to data with publications and subscriptions that trigger a Data Integration task or a file ingestion task.

The Run Publication Subscription API returns the response code of the action that you perform. If the publication or subscription runs successfully, the API returns the event ID of the publication or the subscription event that Cloud Integration Hub generates. You can run the Cloud Integration Hub Event Status API to query the status of the publication or subscription event.

Run Publication Subscription REST API Request

Cloud Integration Hub uses different REST URLs for running a publication and for running a subscription.

To run a publication, use the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/publication/start
```

Where:

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.

- `<baseUrl>` is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/publication/start
```

To run a subscription, use the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/subscription/start
```

Where:

- `<pod>` is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- `<baseUrl>` is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/subscription/start
```

Request syntax for running a publication

To run a publication, use the following request syntax:

Request headers

Include the following headers in the REST API request:

```
Content-Type: application/json
Accept: application/json
```

Request body

Include the following in the REST API request:

```
{
    "publicationName": "<publicationName>",
    "runDisabled": "<true/false>"
}
```

The following list describes the elements of the request:

- `publicationName`. Name of the publication to run.
- `runDisabled`. Whether or not to run a publication that is in a Disabled status.

For example:

```
Content-Type: application/json
Accept: application/json
{
    "publicationName": "daily_sales",
    "runDisabled": "true"
}
```

Tip: You can copy the values of `<pod>` and `<baseUrl>` from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.

Request syntax for running a subscription

To run a subscription, use the following request syntax:

Request headers

Include the following headers in the REST API request:

```
Content-Type: application/json
Accept: application/json
```

Request body

Include the following in the REST API request:

```
{
    "subscriptionName": "<subscriptionName>",
    "runDisabled": "<true/false>"
}
```

The following list describes the elements of the request:

- `subscriptionName`. Name of the subscription to run.
- `runDisabled`. Whether or not to run a subscription that is in a Disabled status.

For example:

```
Content-Type: application/json
Accept: application/json
{
    "subscriptionName": "daily_report",
    "runDisabled": "true"
}
```

Run Publication Subscription REST API Action Response

When you use the Cloud Integration Hub Run Publication Subscription REST API to start the running of a publication or of a subscription, Cloud Integration Hub returns the response code of the action that you perform in the REST API response.

Running a publication or a subscription from the REST API returns one of the following response codes:

- **SUCCESS.** Cloud Integration Hub triggered the publication or the subscription successfully. The status message includes the event ID of the publication or the subscription event that Cloud Integration Hub generates.
- **FAILED.** Cloud Integration Hub could not trigger the publication or the subscription. The response provides the reason for the failure. For example, Cloud Integration Hub did not run the subscription because no publications are ready for consumption by the subscription.

Publish Data REST API

Use the Cloud Integration Hub Publish Data REST API to publish data transactions directly to a topic on the Cloud Integration Hub publication repository.

You can use the Publish Data API to publish data with publications that publish data directly to a topic with an API. You cannot use the API with publications that trigger a Data Integration task.

To publish data through the API, copy the URL of the API from the Publication page in Cloud Integration Hub.

Note: When you use a private publication repository, if you change the Secure Agent on which the publication repository service runs or the port number of the publication repository, the URL of the API changes accordingly. In this case, be sure to notify API users and consumers of the new URL.

Request Headers

Include the following headers in the request:

```
Accept - application/json
Content-Type - application/json
```

To support UTF-8 character encoding, for example, to use Japanese characters in table and column names, include the following headers in the request:

```
Accept-Charset: charset=utf-8
Content-Type: application/json; charset=utf-8
```

Request syntax

Use the following syntax to publish data directly to a topic:

```
{
    "<table_name>":
        [
            { "<column_name>": "<data>" }
        ]
}
```

A topic table name must begin with an alphabetic character or underscore and can contain only ASCII alphanumeric characters or underscores. The name must be unique in the Cloud Integration Hub repository.

For example:

```
{
    "Sales":
        [
            { "Opportunity_Name": "string", "Opportunity_Owner_Id": "string" }
        ],
    "Orders":
        [
            { "Account_Name": "string", "Account_Id": "string", "OrderId":
"string" }
        ]
}
```

DATETIME field

If the topic to which you publish includes a DATETIME field, you must use the following format for the DATETIME value: `yyyy-MM-dd HH:mm:ss.SSS`.

Publish Data REST API Action Response

When you use the Cloud Integration Hub Publish Data REST API to publish data directly to a topic, Cloud Integration Hub returns the response code of the action that you perform in the REST API response.

Publishing data through the REST API returns one of the following response codes:

- **SUCCESS.** Cloud Integration Hub published the data successfully. The status message includes the event ID of the publication event that Cloud Integration Hub generates, the number of row accepted, and the number of rows successfully processed.
- **FAILED.** Cloud Integration Hub could not publish the data.

Note: When you publish data through the Publish Data REST API to a private publication repository and the publication fails because the publication repository service is not accessible, Cloud Integration Hub returns an error to the calling application and does not create an error event.

Topic Swagger Structure for Publish Data REST API

The Publish Data REST API returns the Swagger structure for the topic into which the publication publishes data.

To view the Swagger structure, copy the URL of the structure from the Publication page in Cloud Integration Hub.

Consume Data REST API

Use the Cloud Integration Hub Consume Data REST API to perform the following actions for API-based subscriptions:

- Consume data from a topic on the Cloud Integration Hub publication repository.
- Reconsume data that had previously been processed by triggering the subscription Complete event.
- Reprocess a subscription Error event to consume published data if a subscription process fails.

The API can consume up to 2,000 events at a time. For example, to consume 5,000 published events, the first trigger of the subscription consumes the 2,000 oldest events, the second trigger consumes the next 2,000 oldest events, and the final trigger consumes the remaining 1,000 events.

You can't use the API with subscriptions that trigger a Data Integration task.

Note: When you use a private publication repository, if you change the Secure Agent on which the publication repository service runs or the port number of the publication repository, the URL of the API changes accordingly. In this case, be sure to notify API users and consumers of the new URL.

Consume Data REST API request

To consume data through the API, copy the URL of the API from the Subscription page in Cloud Integration Hub.

Request headers

Include the following headers in the Consume Data REST API request:

```
Accept - application/json
Content-Type - application/json
```

To support UTF-8 character encoding, for example, to use Japanese characters in table and column names, include the following headers in the request:

```
Accept: application/json; charset=utf-8
Accept-Charset: charset=utf-8
Content-Type: application/json; charset=utf-8
```

Request body

The syntax of the Consume Data REST API request body varies, based on the action you perform with the API.

Consume data

To consume data from a topic, use the following request syntax:

```
{
  "aggregated": <value>
}
```

Where <value> takes one of the following values:

- true. The subscription consumes all the available publications in each API call.
- false. The subscription consumes only the oldest publication in each API call.

For example:

```
{
  "aggregated": true
}
```

When you run multiple publications, you can add the publication event ID of a specific publication to the request body to consume only the data of the specific publication event. You can add only one publication event ID to the request body. The key "publicationEventId" is case sensitive. If you don't add the publication event ID of a specific publication to the request body, the subscription consumes all the available publication events.

To add the publication event ID of a specific publication event to the request, use the following syntax:

```
{
  "publicationEventId" : "<publicationEventId>"
}
```

For example:

```
{
  "publicationEventId" : "594210"
}
```

When you add the publication event ID of a specific publication to the request body, Cloud Integration Hub validates the key "publicationEventId" before it runs the subscription instance.

Reconsume data

To reconsume data that had previously been processed, use the following request syntax:

```
{
  "requestType" : "RECONSUME",
  "eventId" : "<subscriptionEventId>"
}
```

Where the key "eventId" is the subscription event ID to reconsume.

For example:

```
{
  "requestType" : "RECONSUME",
  "eventId" : "40559"
}
```

Reprocess subscription

To reprocess a failed subscription, use the following request syntax:

```
{
  "requestType" : "REPROCESS",
  "eventId" : "<subscriptionEventId>"
}
```

Where the key "eventId" is the subscription event ID to reprocess.

For example:

```
{
  "requestType" : "REPROCESS",
  "eventId" : "40577"
}
```

Consume Data REST API action response

When you use the Cloud Integration Hub Consume Data REST API to consume data directly from a topic, Cloud Integration Hub returns the response code of the action that you perform in the REST API response.

Consuming data through the REST API returns one of the following response codes:

SUCCESS

Cloud Integration Hub consumed the data successfully.

The response includes the consumed data in the following syntax:

```
{
  "<table_name>":
    [
      { "<column_name>": "<data>" }
    ]
}
```

For example:

```
{
  "Sales":
    [
      { "Opportunity_Name": "string", "Opportunity_Owner_Id": "string" }
    ],
  "Orders":
    [
      { "Account_Name": "string", "Account_Id": "string", "OrderId":
"string" }
    ]
}
```

A SUCCESS response also includes the aggregated event ID of the subscription event that Cloud Integration Hub generates, the number of rows successfully processed, and the number of total rows processed.

FAILURE

Cloud Integration Hub could not consume the data. For example, if there is no pending data for the subscription to consume. The response includes a description of the error that caused the failure.

Topic Swagger Structure for Consume Data REST API

The Consume Data REST API returns the Swagger structure for the topic from which the subscription consumes data.

To view the Swagger structure, copy the URL of the structure from the Subscription page in Cloud Integration Hub.

Change Publication Subscription Mode REST API

Use the Cloud Integration Hub Change Publication Subscription Mode REST API to change the mode of a publication or a subscription, that is, to enable a disabled publication or subscription or to disable an enabled publication or subscription.

To change the mode of a publication, use the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/publication/changemode
```

Where:

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- **<baseUrl>** is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/publication/changemode
```

To change the mode of a subscription, use the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/subscription/changemode
```

Where :

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- **<baseUrl>** is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/subscription/changemode
```

Tip: You can copy the values of **<pod>** and **<baseUrl>** from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.

Request syntax for changing the mode of a publication

To change the mode of a publication, use the following request syntax:

```
{
    "publicationName": "<publicationName>",
    "mode": "<enable/disable>"
}
```

For example:

```
{
    "publicationName": "daily_sales",
    "mode": "enable"
}
```

Request syntax for changing the mode of a subscription

To change the mode of a subscription, use the following request syntax:

```
{
    "subscriptionName": "<subscriptionName>",
    "mode": "<enable/disable>"
}
```

For example:

```
{
    "subscriptionName": "daily_reports",
    "mode": "disable"
}
```

Change Publication Subscription Mode REST API Action Response

When you use the Cloud Integration Hub REST API to change the mode of a publication or of a subscription, Cloud Integration Hub returns the response code of the action that you perform in the REST API response.

Changing the mode of a publication or a subscription from the REST API returns one of the following response codes:

- When Cloud Integration Hub changes the mode of the publication or the subscription successfully, the API returns a **SUCCESS** response.
- When Cloud Integration Hub fails to change the mode of the publication or the subscription, the response provides the reason for the failure. For example, when you do not have sufficient privileges to perform the operation.

Reprocess Event REST API

Use the Cloud Integration Hub Reprocess Event REST API to reprocess events of subscriptions that trigger a Data Integration task and consume published data, including events of disabled subscriptions.

To reprocess a subscription event, use the following REST URL:

```
https://https://<pod>.<baseUrl>/dih-console/api/v1/event/reprocess
```

Where:

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- **<baseUrl>** is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/event/reprocess
```

Use the following syntax to reprocess an event:

```
{
    "eventId" : "<eventId>"
}
```

For example:

```
{
    "eventId" : "40558"
}
```

Tip: You can copy the values of **<pod>** and **<baseUrl>** from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.

Reprocess Event REST API Action Response

When you use the Cloud Integration Hub Reprocess Event REST API to reprocess a subscription event, Cloud Integration Hub returns the response code of the action that you perform in the REST API response.

The response includes the following information:

Property	Description
responseCode	Response of the action: <ul style="list-style-type: none">- 0: success.- Any number higher than 0: error.
reprocessEventId	New event ID that Cloud Integration Hub generates for the subscription when it reprocesses the existing event.
message	Error message. If the response code is 0 (success), the API returns the message <code>null</code> .

Event Status REST API

When you use a Cloud Integration Hub Run Publication Subscription API to start the running of a publication or of a subscription and the action succeeds, Cloud Integration Hub returns the event ID of the publication or the subscription event that it generates.

When you run the REST API, Cloud Integration Hub returns the event ID in the REST API response.

You can use the Cloud Integration Hub Event Status REST API to query the status of the publication or subscription event according to the event ID. You can see whether the publication or subscription process is still running, and after the process is complete, you can see whether it completed successfully. If the process fails, the response to the query includes the cause of the failure.

Note: For a list of event statuses, see [“Event statuses” on page 102](#).

To query the status of an event, use a GET command with the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/event/<eventId>
```

Where:

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- **<baseUrl>** is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/event/2435
```

Tip: You can copy the values of **<pod>** and **<baseUrl>** from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.

Event Status API Response

When you use the Cloud Integration Hub Event Status API to query the status of a publication or a subscription event, the API returns the event response in an `EventResponse.java` model class.

The following table describes the response properties:

Property	Description
responseCode	Response of the Run Publication Subscription API action.
eventId	ID of the event that Cloud Integration Hub generates for the publication or for the subscription.
eventType	Type of the event that Cloud Integration Hub generates for the publication or for the subscription.
topicName	Name of the topic that is associated with the publication or with the subscription.
publicationName or subscriptionName	Name of the publication or of the subscription.
applicationName	Name of the publishing or of the subscribing application.

Property	Description
eventStatus	Status of the event that Cloud Integration Hub generates for the publication or for the subscription.
eventStartTimeLong	Time when the publication or the subscription event started. System time in milliseconds as returned by Java API <code>java.lang.System.currentTimeMillis</code> .
eventEndTimeLong	Time when the publication or the subscription event ended. System time in milliseconds as returned by Java API <code>java.lang.System.currentTimeMillis</code> .
referencedEventsList	Applicable for file publication events, aggregated subscription events, and compound subscription events. List of event IDs that are related to the file publication, the aggregated subscription, or the compound subscription event. For example, the <code>referencedEventsList</code> of a file publication event includes the file events of the files that are published as part of the publication event.
isFinal	Is the event in a final state.
isError	Is the event in Error status.
sourceSuccessRows	Number of source rows that Cloud Integration Hub read successfully.
sourceFailedRows	Number of source rows that Cloud Integration Hub failed to read.
targetFailedRows	Number of target rows that Cloud Integration Hub failed to write.
targetSuccessRows	Number of target rows that Cloud Integration Hub wrote successfully.
detailedMessage	Applicable for events in an Error status. If the error is caused by Cloud Integration Hub, <code>detailedMessage</code> returns the error message from the Cloud Integration Hub event. For any other error, for example an authentication failure or an incorrect REST URL request, <code>detailedMessage</code> includes a message that describes the cause of the error.

Sample Event Status API Responses

Response to a request to query the status of publication event 4003:

```
{
  "responseCode": "SUCCESS",
  "eventId": 4003,
  "eventType": "Publication",
  "topicName": "top_120",
  "publicationName": "ng_pub_120_1",
  "applicationName": "appl",
  "eventStatus": "Complete",
  "eventStartTimeLong": 1431078308560,
  "eventEndTimeLong": 1431078313780,
  "isFinal": true,
  "isError": false,
  "sourceSuccessRows": 10,
  "sourceFailedRows": 0,
  "targetFailedRows": 0,
  "targetSuccessRows": 10}
```

Response to a request to query the status of aggregated subscription event 3009, which includes subscription events 3008 and 3007:

```
{
  "responseCode": "SUCCESS",
  "eventId": 3009,
  "eventType": "Aggregated Subscription",
  "topicName": "topic1",
  "subscriptionName": "sub1",
  "applicationName": "appl",
  "eventStatus": "Complete",
  "eventStartTimeLong": 1431065700088,
  "eventEndTimeLong": 1431065704372,
  "referencedEventsList": "3008,3007"
  "isFinal": true,
  "isError": false,
  "sourceSuccessRows": 15,
  "sourceFailedRows": 0,
  "targetFailedRows": 0,
  "targetSuccessRows": 15
}
```

Response to a request to query the status of publication event 3016, where the publication process failed:

Response:

```
{
  "responseCode": "SUCCESS",
  "eventId": 3016,
  "eventType": "Publication",
  "topicName": "top_120",
  "publicationName": "ng_pub_120_1",
  "applicationName": "appl",
  "eventStatus": "Error",
  "eventStartTimeLong": 1431066353202,
  "eventEndTimeLong": 1431066357162,
  "isFinal": true,
  "isError": true,
  "sourceSuccessRows": 2,
  "sourceFailedRows": 1,
  "targetFailedRows": 1,
  "targetSuccessRows": 2
  "detailedMessage": "Error while copying several rows :\nSrcFailedRows:
1\nTgtFailedRows: 1\nSrcSuccessRows: 2\nTgtSuccessRows: 2\nPowerCenter workflow:
s__DIH_pub_ng_pub_120_1\nPowerCenter session: s__DIH_pub_ng_pub_120_1\n\nCheck the
PowerCenter session log for more information."
}
```

Events REST API

Use the Cloud Integration Hub Events REST API to get the details of an event based on the search criteria. The API returns event details and the child events.

Events REST API Request

To get the details of an event, use the following REST API URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/events
```

Where:

- <pod> is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`
- <baseUrl> is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/
```

To get the details of an event, use the following request syntax:

Request headers

Include the following headers in the Events REST API request:

```
Content-Type: application/json
Accept: application/json
```

Request body

Include the following in the Events REST API request:

```
{
  "applicationName": "String",
  "topicName": "String",
  "pubSubName": "String",
  "timeframe": Custom - For more details, refer to the following table,
  "firstResult": int,
  "maxResults": int,
  "useExactMatch": boolean,
  "eventType": "String",
  "eventStatus": "String"
}
```

The request to get details of an event can include the following optional properties:

Property	Description
applicationName	Name of the publishing or subscribing application.
topicName	Name of the topic that is associated with the publication or subscription.
pubSubName	Name of the publication or subscription.

Property	Description
timeframe	<p>You can either enter the UTC Timezone format (yyyy-MM-ddTHH:mm:ss.SSSZ) or select the range from the list of enum values.</p> <pre>{ "range": enum values }</pre> <p>You can select the range from the following enum values:</p> <ul style="list-style-type: none"> - NO_RANGE - LAST_HOUR - LAST_TWO_HOURS - LAST_24_HOURS - LAST_7_DAYS - LAST_30_DAYS - TODAY - YESTERDAY - CUSTOM <p>Default range is LAST_24_HOURS.</p> <p>You must mention the <code>fromDate</code> and <code>toDate</code> in the request body if the timeframe range is set to CUSTOM.</p> <pre>{ "fromDate": "Date", "toDate": "Date", "range": "CUSTOM" }</pre> <p>The <code>fromDate</code> and <code>toDate</code> must be in UTC Timezone format (yyyy-MM-ddTHH:mm:ss.SSSZ).</p>
firstResult	The number of rows to skip. For example, you might want to skip the first three rows. Default is zero.
maxResults	Enter the maximum number of results to display. Default is 2000.
useExactMatch	<p>Exact match of the search expression when performing the search. Applies to <code>pubSubName</code>, <code>topicName</code>, and <code>applicationName</code> properties.</p> <p>You can enter one of the following values:</p> <ul style="list-style-type: none"> - true - false <p>Default is false.</p>

Property	Description
eventType	<p>The type of the event that Cloud Integration Hub generates for the publication or for the subscription. Select the value from the following enum values:</p> <ul style="list-style-type: none"> - ALL - PUBLICATION - SUBSCRIPTION - COMPOUND_SUBSCRIPTION - COMPOUND_PUBLICATION - UNBOUND_SUBSCRIPTION - AGGREGATED_SUBSCRIPTION <p>Default is ALL.</p>
eventStatus	<p>The status of the event that Cloud Integration Hub generates for the publication or for the subscription. Enter one of the following enum values:</p> <ul style="list-style-type: none"> - ALL - COMPLETE - DELAYED - DISCARDED - ERROR - PROCESSING - REPROCESSED <p>Default is ALL.</p>

For example:

```

Basic Auth
{
  Username: Administrator@MyCompany.com
  Password: MyPassword
}

Accept - application/json
Content-Type - application/json

{
  "eventStatus": "COMPLETE",
  "eventType": "PUBLICATION",
  "applicationName": "API_App",
  "topicName": "APIT",
  "pubSubName": "APIPUB",
  "timeframe": {
    "range": "NO_RANGE"
  },
  "eventType": "ALL",
  "useExactMatch": true
}

```

The rate limit policy controls the number of times API consumers can invoke the API during a designated time period. This API has a rate limit policy of 5 invocations per minute. Cloud Integration Hub rejects attempts to access the API after it reaches the 5 invocations per minute limit.

Events REST API Response

When you use the Cloud Integration Hub Events API to get the details of a publication or a subscription event, the API returns the following event response:

The following is the response syntax to a request to query the details of an event:

Response header

The Events REST API response includes the following header:

Retry-After - <time in seconds>

Retry-After is the time in seconds to wait before making the follow-up request. Retry-After appears if the response code is `TOO_MANY_REQUESTS`.

Response body

The Events REST API response includes the following body:

```
{
  "eventId": "Long",
  "eventType": "String",
  "applicationName": "String",
  "topicName": "String",
  "pubSubName": "String",
  "createTime": "Event StartTime :: Date",
  "completionTime": "Event Completed Time :: Date",
  "parentEventId": "Long",
  "descendantState": "enum mentioned in the request body.",
  "childEvents": "List<EventDetailResponse>",
  "eventHistory": {
    "eventId": "Long",
    "timestamp": "Date",
    "eventStatusName": "String",
    "comments": "String"
  }
}
```

The following table describes the elements of the response:

Property	Description
eventId	ID of the event that Cloud Integration Hub generates for the publication or subscription.
eventType	The type of event that Cloud Integration Hub generates for the publication or subscription.
applicationName	Name of the application that is associated with the publication or subscription.
topicName	Name of the topic that is associated with the publication or subscription.
pubSubName	Name of the publication or subscription.
createTime	The time when Cloud Integration Hub created the publication or subscription event.
completionTime	The time when the publication or subscription event ended.
eventStatus	The status of the publication or subscription event.
parentEventId	Creates a child event based on a profile ID and returns the event ID of the child event.
descendantState	The cumulative status of all the child events.
childEvents	The events available within the parent event. Applies to publication events.

Property	Description
eventHistory	The details of the complete event process. The following syntax appears in the eventHistory response: <pre> { "eventId": "Long", "timeStamp": "Date", "eventStatusName": "String", "comments": "String" } </pre>
sourceSuccessCount	The number of source rows or files that Cloud Integration Hub read successfully.
sourceFailedCount	The number of source rows or files that Cloud Integration Hub failed to read.
targetSuccessCount	The number of target rows or files that Cloud Integration Hub wrote successfully.
targetFailedCount	The number of target rows or files that Cloud Integration Hub failed to write.

Response codes

When you use the Cloud Integration Hub Events API to get the details of an existing event, Cloud Integration Hub returns the response code of the action that you perform in the API response.

The API can return the following response codes:

Code	Description
200	SUCCESS. The request succeeded.
429	TOO_MANY_REQUESTS. The number of failed requests exceeded the allowed limit for your organization.

For example:

```

[
  {
    "eventId": 2179004,
    "createTime": "2022-10-06T05:56:43.223Z",
    "completionTime": "2022-10-06T05:56:54.213Z",
    "eventType": "PUBLICATION",
    "eventStatus": "Complete",
    "descendantState": "FINAL",
    "applicationName": "API_App",
    "pubSubName": "APIPUB",
    "topicName": "APIT",
    "sourceSuccessCount": 1,
    "sourceFailedCount": 0,
    "targetSuccessCount": 1,
    "targetFailedCount": 0,
    "childEvents": [
      {
        "eventId": 2184003,
        "createTime": "2022-10-06T06:02:35.018Z",
        "completionTime": "2022-10-06T06:02:40.442Z",
        "parentEventId": 2179004,
        "eventType": "SUBSCRIPTION",
        "eventStatus": "Complete",
        "descendantState": "NONE",
        "applicationName": "API_App",
        "pubSubName": "APISUB",
        "topicName": "APIT",
      }
    ]
  }
]

```



```

        "sourceSuccessCount": 1,
        "sourceFailedCount": 0,
        "targetSuccessCount": 1,
        "targetFailedCount": 0,
        "eventHistory": [
            {
                "eventId": 2184003,
                "timeStamp": "2022-10-06T06:02:35.036Z",
                "eventStatusName": "Processing",
                "comments": ""
            },
            {
                "eventId": 2184003,
                "timeStamp": "2022-10-06T06:02:40.442Z",
                "eventStatusName": "Complete",
                "comments": "Subscription to topic completed successfully"
            }
        ]
    },
    {
        "eventId": 2183690,
        "createTime": "2022-10-06T05:56:54.243Z",
        "completionTime": "2022-10-06T05:57:06.678Z",
        "parentEventId": 2179004,
        "eventType": "SUBSCRIPTION",
        "eventStatus": "Reprocessed",
        "descendantState": "NONE",
        "applicationName": "API_App",
        "pubSubName": "APISUB",
        "topicName": "APIT",
        "sourceSuccessCount": 1,
        "sourceFailedCount": 0,
        "targetSuccessCount": 1,
        "targetFailedCount": 0,
        "eventHistory": [
            {
                "eventId": 2183690,
                "timeStamp": "2022-10-06T05:56:54.244Z",
                "eventStatusName": "Processing",
                "comments": ""
            },
            {
                "eventId": 2183690,
                "timeStamp": "2022-10-06T05:56:54.276Z",
                "eventStatusName": "Delayed",
                "comments": "Delayed"
            },
            {
                "eventId": 2183690,
                "timeStamp": "2022-10-06T05:56:55.768Z",
                "eventStatusName": "Processing",
                "comments": "This event is a part of a manually-processed event
hierarchy."
            },
            {
                "eventId": 2183690,
                "timeStamp": "2022-10-06T05:57:06.678Z",
                "eventStatusName": "Complete",
                "comments": "Subscription to topic completed successfully"
            },
            {
                "eventId": 2183690,
                "timeStamp": "2022-10-06T06:02:35.001Z",
                "eventStatusName": "Reprocessed",
                "comments": "This event is a part of a manually-processed event
hierarchy."
            }
        ]
    }
],
"eventHistory": [

```

```

        {
            "eventId": 2179004,
            "timeStamp": "2022-10-06T05:56:43.230Z",
            "eventStatusName": "Processing",
            "comments": ""
        },
        {
            "eventId": 2179004,
            "timeStamp": "2022-10-06T05:56:54.213Z",
            "eventStatusName": "Complete",
            "comments": "Publishing to topic completed successfully"
        }
    ]
},
{
    "eventId": 2179003,
    "createTime": "2022-10-06T05:55:29.583Z",
    "completionTime": "2022-10-06T05:55:40.412Z",
    "eventType": "PUBLICATION",
    "eventStatus": "Complete",
    "descendantState": "NONE",
    "applicationName": "API_App",
    "pubSubName": "APIPUB",
    "topicName": "APIT",
    "sourceSuccessCount": 1,
    "sourceFailedCount": 0,
    "targetSuccessCount": 1,
    "targetFailedCount": 0,
    "eventHistory": [
        {
            "eventId": 2179003,
            "timeStamp": "2022-10-06T05:55:29.589Z",
            "eventStatusName": "Processing",
            "comments": ""
        },
        {
            "eventId": 2179003,
            "timeStamp": "2022-10-06T05:55:40.412Z",
            "eventStatusName": "Complete",
            "comments": "Publishing to topic completed successfully"
        }
    ]
}
]
}
]

```

Cloud Integration Hub Catalog REST API

Use the Catalog REST API to extract data from the Cloud Integration Hub catalog, including topic metadata and metadata about the publications and subscriptions that are associated with each topic.

You can extract metadata pertaining to topics, publications, and subscriptions for which you have both View and Read privileges.

To extract data from the catalog, use the following REST URL:

```
https://<pod>.<baseUrl>/dih-console/api/v1/catalog/topics
```

Where:

- **<pod>** is the name of the Informatica Intelligent Cloud Services point of delivery (PoD) where you access Cloud Integration Hub. For example: `cih-pod1`, or `emw1-cih`.
- **<baseUrl>** is the Informatica Intelligent Cloud Services URL. For example: `dm-us.informaticacloud.com/`.

For example:

```
https://cih-pod1.dm-us.informaticacloud.com/dih-console/api/v1/catalog/topics
```

Tip: You can copy the values of `<pod>` and `<baseUrl>` from the Cloud Integration Hub URL after you access it from the My Services page of Informatica Intelligent Cloud Services.

Cloud Integration Hub Catalog API Response

When you use the Cloud Integration Hub Catalog API to extract data from the Cloud Integration Hub catalog, the API returns a JSON string that contains metadata about all the topics for which you have the required privileges.

The string includes the following data for each topic in the response:

topicName

Name of the topic.

topicDesc

Textual description of the topic.

topicType

Type of the topic: Delta or Full.

topicTables

For each table in the topic, an entry with the table name and detailed information about each of the table fields.

publications

For each publication that is associated with the topic, the following data is provided:

publicationName

Name of the publication.

publicationDesc

Textual description of the publication.

applicationName

Application from which the publication publishes data or files.

publicationSourceType

Type of publication source.

publicationConnectionName

For relational database publications and for HDFS publications: name of the connection from where the publication workflow reads the data or the files to be published.

publicationDBType

For relational database publications: type of database.

subscriptions

For each subscription that is associated with the topic, the following data is provided:

subscriptionName

Name of the subscription.

subscriptionDesc

Textual description of the subscription.

applicationName

Application that consumes data or files from the topic.

subscriptionTargetType

Type of subscription target.

subscriptionConnectionName

For relational database subscriptions and for HDFS subscriptions: name of the connection to where the subscription workflow writes the data or the files that the application consumes.

subscriptionDBType

For relational database subscriptions: type of database.

Sample Data Integration Hub Catalog API Response

The following example shows a response to a request to extract data from the Cloud Integration Hub catalog:

```
{
  "responseCode": "SUCCESS",
  "catalogTopics": [
    {
      "topicName": "FileTopic",
      "topicDesc": null,
      "topicType": "Delta"
      "topicTables": [
        {
          "tableName": "Orders"
          "tableFields": [
            {
              "name": "_Name_"
              "fieldType": "STRING"
              "nullable": false
              "scale": -1
              "precision": 255
              "length": 255
              "primaryKey": false
              "filterAccelerator": false
              "encrypted": true
            },
            {
              "name": "_Type_"
              "fieldType": "STRING"
              "nullable": false
              "scale": -1
              "precision": 255
              "length": 255
              "primaryKey": false
              "filterAccelerator": false
              "encrypted": true
            },
            {
              "name": "_ParentId_"
              "fieldType": "STRING"
              "nullable": false
              "scale": -1
              "precision": 255
              "length": 255
              "primaryKey": false
              "filterAccelerator": false
              "encrypted": true
            }
          ]
        },
        {
          "publicationName": "FilePub",
```

```

        "publicationDesc": null,
        "applicationName": "FileApp",
        "publicationSourceType": "CUSTOM",
        "publicationConnectionName": null,
        "publicationDBType": null
    }
},
"subscriptions": [
    {
        "subscriptionName": "FileSub",
        "subscriptionDesc": null,
        "applicationName": "FileApp",
        "subscriptionTargetType": "CUSTOM",
        "subscriptionConnectionName": null,
        "subscriptionDBType": null
    }
]
},
{
    "topicName": "OrderTopic",
    "topicDesc": null,
    "topicType": "Delta",
    "topicTables": [
        {
            "tableName": "OrderTable"
            "tableFields": [
                {
                    "name": "_Name_"
                    "fieldType": "STRING"
                    "nullable": false
                    "scale": -1
                    "precision": 255
                    "length": 255
                    "primaryKey": false
                    "filterAccelerator": false
                    "encrypted": true
                },
                {
                    "name": "_Type_"
                    "fieldType": "STRING"
                    "nullable": false
                    "scale": -1
                    "precision": 255
                    "length": 255
                    "primaryKey": false
                    "filterAccelerator": false
                    "encrypted": true
                },
                {
                    "name": "_ParentId_"
                    "fieldType": "STRING"
                    "nullable": false
                    "scale": -1
                    "precision": 255
                    "length": 255
                    "primaryKey": false
                    "filterAccelerator": false
                    "encrypted": true
                },
                {
                    "name": "_StartDate_"
                    "fieldType": "STRING"
                    "nullable": false
                    "scale": -1
                    "precision": 255
                    "length": 255
                    "primaryKey": false
                    "filterAccelerator": false
                    "encrypted": true
                }
            ]
        }
    ]
}

```

```

        "name": "_EndDate "
        "fieldType": "STRING"
        "nullable": false
        "scale": -1
        "precision": 255
        "length": 255
        "primaryKey": false
        "filterAccelerator": false
        "encrypted": true
    },
    {
        "tableName": "CustomerTable"
        "tableFields": [
            {
                "name": "_Name_"
                "fieldType": "STRING"
                "nullable": false
                "scale": -1
                "precision": 255
                "length": 255
                "primaryKey": false
                "filterAccelerator": false
                "encrypted": true
            },
            {
                "name": "_Type_"
                "fieldType": "STRING"
                "nullable": false
                "scale": -1
                "precision": 255
                "length": 255
                "primaryKey": false
                "filterAccelerator": false
                "encrypted": true
            },
            {
                "name": "_ParentId_"
                "fieldType": "STRING"
                "nullable": false
                "scale": -1
                "precision": 255
                "length": 255
                "primaryKey": false
                "filterAccelerator": false
                "encrypted": true
            },
            {
                "name": "_ExpectedRevenue_"
                "fieldType": "STRING"
                "nullable": false
                "scale": -1
                "precision": 255
                "length": 255
                "primaryKey": false
                "filterAccelerator": false
                "encrypted": true
            },
            {
                "name": "_IsActive_"
                "fieldType": "STRING"
                "nullable": false
                "scale": -1
                "precision": 255
                "length": 255
                "primaryKey": false
                "filterAccelerator": false
                "encrypted": true
            }
        ]
    }
]

```

```

    },
  ],
  "publications": [
    {
      "publicationName": "OrdersPublication",
      "publicationDesc": null,
      "applicationName": "OrderPublications",
      "publicationSourceType": "CUSTOM",
      "publicationConnectionName": " null",
      "publicationDBType": " null"
    }
  ],
  "subscriptions": [
    {
      "subscriptionName": "OrdersSubscription",
      "subscriptionDesc": null,
      "applicationName": "OrderSubscriptions",
      "subscriptionTargetType": "CUSTOM",
      "subscriptionConnectionName": null,
      "subscriptionDBType": null
    },
    {
      "subscriptionName": "OrderSubs",
      "subscriptionDesc": null,
      "applicationName": "OrderSubscriptions",
      "subscriptionTargetType": "CUSTOM",
      "subscriptionConnectionName": " null",
      "subscriptionDBType": " null"
    }
  ]
}

```

CHAPTER 12

Glossary

aggregated subscription

A subscription that consumes multiple data sets from the same topic with a single batch workflow. An aggregated subscription can use an automatic mapping or a custom mapping to process data. When you use an automatic mapping, the subscription sorts the data according to the publication date and time of the publication instances.

application

An entity that represents a system in your organization that needs to share data with other systems. An application can be a publisher and a subscriber. An application can publish multiple data sets.

child event

An event within the hierarchy of another event that acts as a parent event. The child event is a subsidiary of the parent event.

Cloud Integration Hub repository

A relational database table set that contains the metadata required to process publications and subscriptions in Cloud Integration Hub. It also contains the events that Cloud Integration Hub generates while it processes publications and subscriptions.

compound publication

A publication that publishes single or multiple data sets to multiple topics with a single Data Integration mapping task.

compound subscription

A subscription that consumes data sets from multiple topics with a single task.

Data Integration task

A Data Integration task is a process that you configure to analyze, extract, transform, and load data. In Cloud Integration Hub, a Data Integration task is a task that reads from a file, a database, or another source and writes to a target. Use Data Integration tasks to process Cloud Integration Hub publications and subscriptions with Informatica Intelligent Cloud Services.

When you use a Data Integration task to process publications, you use the Cloud Integration Hub cloud connector as the publication target. When you use a Data Integration task to process subscriptions, you use the Cloud Integration Hub cloud connector as the subscription source.

event

An occurrence of a publication or subscription at each stage of processing. The Cloud Integration Hub server generates the event and updates the event status while it processes the publication or subscription.

parent event

An event at the top level of a hierarchy of events.

publication

An entity that defines data flow from a data source to the Cloud Integration Hub publication repository and the data publishing schedule. The publication publishes the data to a topic that defines the structure of the data in the publication repository. When a publication runs, Cloud Integration Hub extracts the data set from the application, processes the data, and writes the data to the publication repository. You can then create one or more subscriptions to process and write the published data set to target applications.

publication repository

A relational database table set that stores published data sets that subscribers can consume. Cloud Integration Hub stores the data in the publication repository until the retention period for the data expires.

simple publication

A publication that publishes single or multiple data sets and files to a single topic with a Data Integration or a Mass Ingestion task or with an API.

simple subscription

A subscription that consumes single or multiple data sets or files from a single topic with a Data Integration or a Mass Ingestion task or with an API.

subscription

An entity that defines the type, format, and schedule of data flow from the Cloud Integration Hub publication repository to a data target. When a subscription runs, Cloud Integration Hub extracts the data set from the publication repository, processes the data, and writes the data to the target application. You can subscribe to one or more topics. Each topic to which you subscribe can contain data from multiple publishers.

topic

An entity that represents a data domain that applications publish and consume through Cloud Integration Hub. A topic defines the data structure and additional data definitions such as the data retention period. Multiple applications can publish to the same topic. Multiple applications can consume data from the same topic.

unbound subscription

A subscription that is not restricted to specific publication instances. It subscribes to all the data that a publication publishes and consumes the data based on the subscription filter, regardless of when or in what batch the data was published.

INDEX

A

- administration
 - Salesforce Accelerator package [27](#)
- API
 - authorization header [112](#)
 - consume data [116–118](#)
 - publication subscription mode [118](#)
 - publish data [114, 115](#)
 - reprocess event [120](#)
 - run publication subscription [112](#)
- application
 - add publication [49](#)
 - add subscription [49](#)
 - creating [49](#)
 - definition [49](#)
 - editing [36](#)
 - managing [49](#)
 - properties [50](#)
 - viewing [35](#)
- architecture
 - components [11](#)
- asset
 - export [42](#)
 - import [42](#)
 - migration [42](#)
- assets
 - assigning tags [46](#)
 - creating tags [46](#)
 - deleting [36](#)
 - dependent [42](#)
 - export [42](#)
 - import [43](#)
 - migrate [44](#)
 - moving [36](#)
 - tags [45](#)

C

- catalog API
 - description [130](#)
 - response [131](#)
- catalog API response
 - example [132](#)
- change mode API
 - REST API [119](#)
- change mode REST API
 - action status [119](#)
- change mode REST API action
 - status [119](#)
- change status
 - event [105](#)
- cloud
 - task [65](#)

- cloud mapping
 - creating [71, 73](#)
 - publication [71](#)
 - subscription [73](#)
- cloud task
 - definition [65](#)
 - types [65](#)
- connection
 - topic table [61](#)
- consume data
 - REST API [116–118](#)
- consume data API
 - response [117](#)
 - REST API [116](#)
 - Swagger [118](#)
- consume data REST API
 - request [116](#)
- create
 - application [49](#)
 - data synchronization task [67, 69](#)
 - monitoring rule [108](#)
 - publication [82, 84](#)
 - subscription [92, 95](#)
 - topic [56](#)
 - topic table [61–63](#)
- create topic table
 - from connection [61](#)
 - from flat file [61](#)
 - from metadata file [62](#)
 - manually [63](#)
- creating
 - tags [46](#)

D

- Data Integration tasks
 - description [20](#)
- data synchronization
 - publication [67](#)
- data synchronization task
 - creating [69](#)
 - creating a [67](#)
 - publication [67](#)
 - subscription [69](#)
- deleting
 - assets [36](#)
 - folders [36](#)
 - projects [36](#)
 - tags [47](#)
- dependencies
 - viewing dependencies [45](#)
- dependent
 - assets [42](#)
- disable
 - monitoring rule [109](#)

disable (*continued*)
 publication [85](#)
 subscription [97](#)

E

edit
 application [36](#)
 monitoring rule [108](#)
 publication [36](#)
 subscription [36](#)
 topic [36](#)
editing
 tags [47](#)
enable
 monitoring rule [109](#)
 publication [85](#)
 subscription [97](#)
error handling
 migration [44](#)
event
 changing status [105](#)
 monitoring [107](#)
 properties [106](#)
 reprocessing [105](#)
 tracking [107](#)
event details API
 REST API [123](#)
event monitoring
 overview [101](#)
event status API
 process status [121](#), [122](#), [126](#)
 response [121](#), [122](#), [126](#)
events
 consumption statuses [103](#)
 filter [104](#)
 history [103](#)
 managing [104](#)
 overview [101](#)
 processing information [104](#)
 publication [101](#)
 publications and subscriptions [102](#), [103](#)
 session logs [104](#)
 statuses [102](#)
 subscription [101](#)
 system [104](#)
 types [102](#)
events API
 request [123](#)
example
 catalog API response [132](#)
Explore page
 tags [45](#)
export
 assets [42](#)

F

filter
 events [104](#)
filters
 Hub Overview diagram [15](#)
flat file
 topic table [61](#)
folders
 deleting [36](#)

folders (*continued*)
 moving [36](#)

G

get previous publications [97](#)
glossary
 of terms [136](#)

H

hardware
 requirements [13](#)
Hub Overview
 diagram [15](#)
Hub Overview diagram
 description [15](#)
 filters [15](#)

I

import
 assets [42](#), [43](#)
Informatica Intelligent Cloud Services
 mappings [66](#)
 tasks [66](#)
intermediate staging
 policy [32](#)
intermediate staging policy
 subscription [32](#)

J

Java KeyStore
 private publication repository [31](#)

L

log in
 description [17](#)

M

mapping
 publication [81](#)
 subscription [90](#)
mapping configuration
 publication [71](#), [72](#)
 subscription [72](#), [73](#)
mapping task
 configuration process [71](#)
 creating [72](#), [73](#)
 usage [71](#)
mappings
 guidelines [66](#)
metadata file
 topic table [52](#), [62](#)
migrate
 assets [44](#)
migration
 assets [42](#)
 error handling [44](#)

- monitor
 - events [107](#)
- monitoring
 - overview [101](#)
 - rules [107](#)
- monitoring rule
 - creating [108](#)
 - disabling [109](#)
 - editing [108](#)
 - enabling [109](#)
 - managing [107](#)
 - properties [109](#)
- moving
 - assets and folders [36](#)

N

- Navigator
 - description [16](#)
- network
 - requirements [13](#)

O

- organization
 - management [24](#)
 - setup [25](#), [27](#)
- organization management
 - before you begin [25](#)
 - description [24](#)
 - setup [25](#), [27](#)
- organization setup
 - before you begin [25](#)
- overview
 - description [9](#)

P

- permissions
 - best practices [40](#)
 - configuring for objects [40](#)
 - for copied assets [39](#)
 - for imported assets [39](#)
 - overview [39](#)
 - permission descriptions [39](#)
 - rules and guidelines [40](#)
- policy
 - intermediate staging [32](#)
- private publication repository
 - Java KeyStore [31](#)
 - port number [31](#)
 - Publication Repository Service [30](#)
- privileges
 - Cloud Integration Hub users [38](#)
 - Data Integration users [38](#)
- Product Availability Matrix [13](#)
- project folders [16](#)
- projects
 - deleting [36](#)
- proxy
 - requirements [13](#)
- publication
 - add to application [49](#)
 - creating [82](#), [84](#)
 - definition [21](#)

- publication (*continued*)
 - disabling [85](#)
 - editing [36](#)
 - enabling [85](#)
 - events [101](#)
 - management [82](#)
 - mapping [81](#)
 - overview [79](#)
 - process [21](#), [80](#)
 - properties [86](#)
 - run manually [85](#)
 - running [85](#)
 - schedule [81](#)
 - source [81](#)
 - type [79](#)
 - viewing [35](#)
- publication event
 - history [103](#)
- publication process
 - publish with API [81](#)
 - task-triggering publications [80](#)
- publication repository
 - private [31](#)
- Publication Repository Service
 - private publication repository [30](#)
- publication subscription mode
 - REST API [118](#)
- publications
 - event consumption statuses [103](#)
 - event statuses [102](#)
 - event types [102](#)
- publish data
 - REST API [114](#), [115](#)
- publish data API
 - response [115](#)
 - Swagger [115](#)

R

- reprocess
 - event [105](#)
- reprocess event
 - REST API [120](#)
- reprocess event API
 - REST API [120](#)
- reprocess event API REST API
 - action status [120](#)
- reprocess event API REST API action
 - status [120](#)
- requirements
 - hardware [13](#)
 - network [13](#)
 - Product Availability Matrix [13](#)
 - proxy [13](#)
- REST API
 - authorization header [112](#)
 - consume data [116–118](#)
 - events [123](#)
 - publication subscription mode [118](#)
 - publish data [114](#), [115](#)
 - reprocess event [120](#)
 - run publication subscription [112](#)
- REST APIs
 - description [111](#)
- retry policy
 - subscription [91](#)

- rule
 - monitoring [107](#)
- run
 - publication [85](#)
 - subscription [96](#)
- run publication subscription API
 - events [123](#)
 - process status [121](#), [122](#), [126](#)
 - REST API [112](#), [114](#)
- run publication subscription process
 - get event details [123](#)
 - status [121](#), [122](#), [126](#)
- run publication subscription REST API
 - action status [114](#)
 - request [112](#)
- run publication subscription REST API action
 - status [114](#)
- run publication subscription status
 - event status API [121](#), [122](#), [126](#)

S

- Salesforce
 - Accelerator package [27–29](#)
- Salesforce Accelerator package
 - before you begin [28](#)
 - components [28](#)
 - deploying [29](#)
- schedule
 - publication [81](#)
 - subscription [91](#)
- source
 - publication [81](#)
- subscribe to
 - topic [57](#)
- subscription
 - add to application [49](#)
 - creating [92](#), [95](#)
 - definition [21](#)
 - disabling [97](#)
 - editing [36](#)
 - enabling [97](#)
 - events [101](#)
 - get previous publications [97](#)
 - intermediate staging policy [32](#)
 - management [92](#)
 - mapping [90](#)
 - overview [88](#)
 - process [21](#), [89](#)
 - properties [97](#)
 - retry policy [91](#)
 - run manually [96](#)
 - running [96](#)
 - schedule [91](#)
 - targets [91](#)
 - type [88](#)
 - viewing [35](#)
- subscription event
 - history [103](#)
- subscription process
 - consume with API [90](#)
 - task-triggering subscriptions [89](#)
- subscriptions
 - event consumption statuses [103](#)

- subscriptions (*continued*)
 - event statuses [102](#)
 - event types [102](#)
- Swagger
 - topic structure [115](#), [118](#)
- synchronization task
 - usage [67](#)
- system event
 - maintenance report [104](#)
- system requirements [13](#)

T

- tags
 - assigning [46](#)
 - creating [46](#)
 - deleting [47](#)
 - editing [47](#)
 - properties [47](#)
- targets
 - subscription [91](#)
- tasks
 - deleting [36](#)
 - guidelines [66](#)
- topic
 - creating [56](#)
 - definition [51](#)
 - diagram [58](#)
 - editing [36](#)
 - management [56](#)
 - properties [58](#)
 - publication repository [18](#)
 - structure [51](#), [52](#), [61–63](#)
 - subscribing to [57](#)
 - table [52](#)
 - tables [52](#), [61–63](#)
 - viewing [35](#)
- Topic Diagram
 - description [58](#)
- topic structure
 - Swagger [115](#), [118](#)
- topic table
 - create [61–63](#)
 - metadata file [52](#)
- track
 - events [107](#)

U

- user interface
 - description [14](#)
- user roles
 - Cloud Integration Hub users [37](#)

V

- view
 - application [35](#)
 - publication [35](#)
 - subscription [35](#)
 - topic [35](#)