Informatica

Informatica® Cloud Data Integration
July 2024

Troubleshooting

Informatica Cloud Data Integration Troubleshooting
July 2024
April 2024

© Copyright Informatica LLC 2006, 2024

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. 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.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Informatica, Informatica Cloud, Informatica Intelligent Cloud Services, PowerCenter, PowerExchange, and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at https://www.informatica.com/trademarks.html. 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. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2024-07-12

Table of Contents

Preface 4

Informatica Resources 4

- Informatica Documentation 4
- Informatica Intelligent Cloud Services web site 4
- Informatica Intelligent Cloud Services Communities 4
- Informatica Intelligent Cloud Services Marketplace 4
- Data Integration connector documentation 5
- Informatica Knowledge Base 5
- Informatica Intelligent Cloud Services Trust Center 5
- Informatica Global Customer Support 5

Chapter 1: Troubleshooting 6

Troubleshooting a Secure Agent 7

Troubleshooting an advanced cluster 9

Troubleshooting an advanced cluster on AWS 10

Troubleshooting an advanced cluster on Microsoft Azure 13

Troubleshooting a connection 14

- Connection errors 14

Troubleshooting assets 14

Troubleshooting a task 14

- Task errors 15

Troubleshooting a mapping task 16

Troubleshooting an advanced cluster subtask 16

Troubleshooting a synchronization task 18

- Synchronization task errors 18

Troubleshooting a replication task 19

- Replication task errors 20

Troubleshooting masking task errors 21

Troubleshooting a PowerCenter task 23

- PowerCenter task errors 23

Troubleshooting a mapping 24

Troubleshooting a mapping in advanced mode 25

Troubleshooting scheduled tasks 25

Troubleshooting security 25

Index 27

Preface

Refer to *Troubleshooting* for information about troubleshooting errors in Informatica Cloud® Data Integration.

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 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 Intelligent Cloud Services web site

You can access the Informatica Intelligent Cloud Services web site at http://www.informatica.com/cloud. This site contains information about Informatica Cloud integration services.

Informatica Intelligent Cloud Services Communities

Use the Informatica Intelligent Cloud Services Community to discuss and resolve technical issues. You can also find technical tips, documentation updates, and answers to frequently asked questions.

Access the Informatica Intelligent Cloud Services Community at:

https://network.informatica.com/community/informatica-network/products/cloud-integration

Developers can learn more and share tips at the Cloud Developer community:

https://network.informatica.com/community/informatica-network/products/cloud-integration/cloud-developers

Informatica Intelligent Cloud Services Marketplace

Visit the Informatica Marketplace to try and buy Data Integration Connectors, templates, and mapplets:

https://marketplace.informatica.com/

Data Integration connector documentation

You can access documentation for Data Integration Connectors at the Documentation Portal. To explore the Documentation Portal, visit https://docs.informatica.com.

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 Intelligent Cloud Services Trust Center

The Informatica Intelligent Cloud Services Trust Center provides information about Informatica security policies and real-time system availability.

You can access the trust center at https://www.informatica.com/trust-center.html.

Subscribe to the Informatica Intelligent Cloud Services Trust Center to receive upgrade, maintenance, and incident notifications. The Informatica Intelligent Cloud Services Status page displays the production status of all the Informatica cloud products. All maintenance updates are posted to this page, and during an outage, it will have the most current information. To ensure you are notified of updates and outages, you can subscribe to receive updates for a single component or all Informatica Intelligent Cloud Services components. Subscribing to all components is the best way to be certain you never miss an update.

To subscribe, on the Informatica Intelligent Cloud Services Status page, click **SUBSCRIBE TO UPDATES**. You can choose to receive notifications sent as emails, SMS text messages, webhooks, RSS feeds, or any combination of the four.

Informatica Global Customer Support

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

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at https://www.informatica.com/services-and-training/support-services/contact-us.html.

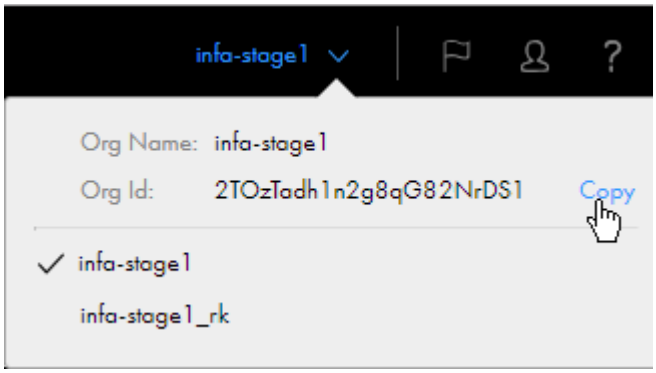
CHAPTER 1

Troubleshooting

Use the following sections to troubleshoot errors in Data Integration.

Note: To get support for Data Integration, you might need to give your organization ID to Informatica Global Customer Support. You can find your organization ID through the **Organization** menu in the upper right corner.

The following image shows the **Organization** menu:



To copy the organization ID, click the **Copy** option that appears when you hover the cursor to the right of the **Org ID** field.

You can also find your organization ID on the **Organization** page in Administrator.

Troubleshooting a Secure Agent

The Secure Agent did not install or start correctly.

If the Secure Agent does not install or start correctly, complete the following tasks:

1. Review the following logs:

File	Description
<Secure Agent installation directory>\apps\agentcore \infaagent.log	Contains startup and shutdown information about the Secure Agent.
<Secure Agent installation directory>\apps\agentcore \agentcore.log	Contains information about the activities related to the Secure Agent, including details of all the services enabled for the agent.
<Secure Agent installation directory>\apps \Data_Integration_Server\logs\tomcat \tomcat<version>.log	Contains details related to the Tomcat process for the Data Integration Server. It records the task execution details, including the start time, end time, and statistics for the tasks. Also includes design time and metadata-related activities for the Data Integration assets.
<Secure Agent installation directory>\apps \Data_Integration_Server\<version>\tomcat.out	Contains basic information about the Tomcat process for the Data Integration Server, including when the Tomcat process started and stopped. Also includes Secure Agent connection details, such as certificate and SSL-related information.
<Secure Agent installation directory>\apps \Data_Integration_Server\<version>\scripts.log	Contains information about the Data Integration Server, including details about how the script used by the Data Integration Server was run. Use this log if there is an issue with the Data Integration Server.

2. For Secure Agents that run on Windows, view the application logs in the Windows Event Viewer.

I started the Secure Agent, but the status is inactive.

The Secure Agent might take a few minutes to start. The status refreshes every 5 seconds. If the Secure Agent does not become active, complete the following tasks:

- If your organization uses a proxy server to access the internet, verify that the proxy settings are set correctly.
- View the details in infaagent.log in the directory where you installed the Secure Agent.

I installed the Secure Agent, but I want to install another on a different machine. How do I do that?

On the new machine, use your login to connect to Informatica Intelligent Cloud Services. Then, download and install the Secure Agent.

One of my services shows an error status after I restarted the service successfully.

If a service fails with an error status, the error status for the service might continue to display in the Agent Service Details after the service starts up successfully. The error stays on the page until an internal job that cleans up obsolete messages runs. You can ignore the error.

I am trying to uninstall the Secure Agent, but the Secure Agent status still shows "Up and Running."

When you uninstall the Secure Agent without first stopping the Secure Agent, the Agent Core and other services might continue to run for several minutes. To avoid this issue, stop the Secure Agent before you uninstall it.

Why does my Secure Agent always display "Agent Core Upgrading" in Administrator?

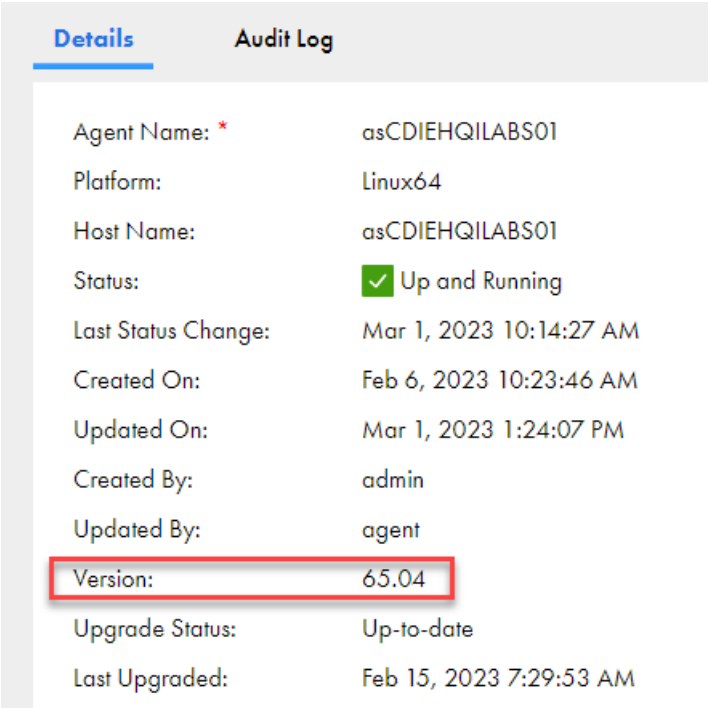
On the **Runtime Environments** page in Administrator, the status of an agent always displays "Agent Core Upgrading". You see the following message in the `agentcore.log` file:

```
2022-10-11 17:02:57,560 GMT tid="21" tn="Agent Core State Machine Thread" ERROR
[com.informatica.saas.infaagent.agentcore.AgentCoreStateMachine] - Authentication failed
due to IO error: [cannot decrypt null or empty string].
```

This issue occurs when the agent missed one or more previous major upgrades. For example, you stopped an agent that was on version 62.x and the current version is 65.x when you restart it. The automatic upgrade only supports upgrading from the previous major version, 64.x. Since your version is older than version 64.x, the automatic upgrade fails.

To resolve the issue, either reregister or reinstall the Secure Agent.

You can see the agent version in the Details tab of a Secure Agent:



Troubleshooting an advanced cluster

What should I do if the status of the advanced cluster is Unknown?

When the cluster status is Unknown, first verify that the Secure Agent is running. If the agent is not running, enable the agent and check whether the cluster starts running.

If the cluster does not start running, an administrator can run the command to list clusters. If the command output returns the cluster state as partial or in-use, the administrator can run the command to delete the cluster.

For more information about the commands, see the Administrator help.

I looked at the ccs-operation.log file to troubleshoot the advanced cluster, but there wasn't enough information. Where else can I look?

You can look at the `cluster-operation` logs that are dedicated to the instance of the advanced cluster. When an external command set begins running, the `ccs-operation` log displays the path to the `cluster-operation` logs.

For example:

```
2020-06-15 21:22:36.094 [reqid:] [T:000057] INFO   :
c.i.c.s.c.ClusterComputingService        [CCS_10400] Starting to run command set
[<command set>] which contains the following commands: [
    <commands> ;
]. The execution log can be found in the following location: [/data2/home/cldagnt/
SystemAgent/apps/At_Scale_Server/35.0.1.1/ccs_home/3xukm9iqp5zeahyrb7rqoz.k8s.local/infa/
cluster-operation.log].
```

The specified folder contains all `cluster-operation` logs that belong to the instance of the cluster. You can use the logs to view the full `stdout` and `stderr` output streams of the command set.

The number in the log name indicates the log's generation and each `cluster-operation` log is at most 10 MB. For example, if the cluster instance generated 38 MB of log messages while running external commands, the folder contains four `cluster-operation` logs. The latest log has 0 in the file name and the oldest log has 3 in the file name. You can view the messages in the `cluster-operation0.log` file to view the latest errors.

If you set the log level for the Elastic Server to DEBUG, the `ccs-operation` log shows the same level of detail as the `cluster-operation` logs.

How do I find the initialization script logs for the nodes where the init script failed?

To find the init script logs, complete the following tasks:

1. Locate the `ccs-operation.log` file in the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
   ```

2. In the `ccs-operation.log` file, find a message that is similar to the following message:

   ```
   Failed to run the init script for cluster [<cluster instance ID>] on the following
   nodes: [<cluster node IDs]. Review the log in the following S3 file path: [<cloud
   platform location>].
   ```

3. Navigate to the cloud platform location that is provided in the message.
4. Match the cluster node IDs to the init script log file names for the nodes where the init script failed.

How are the resource requirements calculated in the following error message for an advanced cluster?

```
2019-04-26T19:04:11.762+00:00 <Thread-16> SEVERE: java.lang.RuntimeException:
[java.lang.RuntimeException: The Cluster Computing System rejected the Spark task
[InfaSpark0] due to the following error: [[CCS_10252] Cluster
[6bjwune8v4bkt3vneokii9.k8s.local] doesn't have enough resources to run the application
[spark--infaspark0e6674748-b038-4e39-a2a9-3fd49e63f289infaspark0-driver] which requires
a minimum resource of [(KB memory, mCPU)]. The cluster must have enough nodes, and each
node must have at least [(KB memory, mCPU)] to run this job.].]
```

The first resource requirement is the total number of resources that are required by the Spark driver and the Spark executor.

The second resource requirement is calculated based on the minimum resource requirements on each worker node to run a minimum of one Spark process.

The resources are calculated using the following formulas:

```
Memory: MAX(driver_memory, executor_memory)
CPU: MAX(driver_CPU, executor_CPU)
```

The Spark process can be either a Spark driver process or a Spark executor process. The cluster must have two nodes where each node fulfills the minimum requirements to run either the driver or the executor, or the cluster must have one node with enough resources to run both the driver and the executor.

Note: The resource requirements for the driver and executor depend on how you configure the following advanced session properties in the mapping task:

- spark.driver.memory
- spark.executor.memory
- spark.executor.cores

For more information about minimum resource requirements, see the Administrator help.

I shut down the Secure Agent machine on my cloud platform, but some jobs are still running.

When you shut down the agent machine, the agent starts on a new machine, but jobs do not carry over to the new machine.

In Monitor, cancel the jobs and run them again. The agent on the new machine will start processing the jobs.

To avoid this issue, see the instructions to shut down the agent machine in the Administrator help.

Troubleshooting an advanced cluster on AWS

Why did the advanced cluster fail to start?

To find out why the advanced cluster failed to start, use the `ccs-operation.log` file in the following directory on the Secure Agent machine:

```
<Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
```

The following table lists some reasons why a cluster might fail to start:

Reason	Possible Cause
The cluster operator failed to update the cluster.	The VPC limit was reached on your AWS account.
The master node failed to start.	The master instance type isn't supported in the specified region or availability zone or in your AWS account.
All worker nodes failed to start.	The worker instance type isn't supported in the specified region or availability zone or in your AWS account.
The Kubernetes API server failed to start.	The user-defined master role encountered an error.

When a cluster fails to start due to at least one of these reasons, the `ccs-operation.log` file displays a BadClusterConfigException.

For example, you might see the following error:

```
2019-06-27 00:50:02.012 [T:000060] SEVERE : [CCS_10500] [Operation of <cluster instance
ID>: start_cluster-<cluster instance ID>]:
com.informatica.cloud.service.ccs.exception.BadClusterConfigException: [[CCS_10207] The
cluster configuration for cluster [<cluster instance ID>] is incorrect due to the
following error: [No [Master] node has been created on the cluster. Verify that the
instance type is supported.]. The Cluster Computing System will stop the cluster soon.]
```

If the cluster encounters a BadClusterConfigException, the agent immediately stops the cluster to avoid incurring additional resource costs and to avoid potential resource leaks. The agent does not attempt to recover the cluster until the configuration error is resolved.

I ran a job to start the advanced cluster, but the VPC limit was reached.

When you do not specify a VPC in the advanced configuration for a cluster, the Secure Agent creates a new VPC on your AWS account. Because the number of VPCs on your AWS account is limited for each region, you might reach the VPC limit.

If you reach the VPC limit, edit the advanced configuration and perform one of the following tasks:

- Provide a different region.
- Remove the availability zones. Then, provide an existing VPC and specific subnets within the VPC for the cluster to use.

Any cloud resources that were provisioned for the cluster will be reused when the cluster starts in the new region or the existing VPC. For example, the Secure Agent might have provisioned Amazon EBS volumes before it received an error for the VPC limit. The EBS volumes are not deleted, but they are reused during the next startup attempt.

I ran a job to start the advanced cluster, but the cluster failed to be created with the following error:

```
Failed to create cluster [<cluster instance ID>] due to the following error:
[[CCS_10302] Failed to invoke AWS SDK API due to the following error: [Access Denied
(Service: Amazon S3; Status Code: 403; Error Code: AccessDenied; Request ID: <request
ID>; S3 Extended Request ID: <S3 extended request ID>)].].]
```

The Secure Agent failed to create the advanced cluster because Amazon S3 rejected the agent's request.

Make sure that the S3 bucket policies do not require clients to send requests that contain an encryption header.

How do I troubleshoot a Kubernetes API Server that failed to start?

If the Kubernetes API Server fails to start, the advanced cluster fails to start. To troubleshoot the failure, use the Kubernetes API Server logs instead.

To find the Kubernetes API Server logs, complete the following tasks:

1. Connect to the master node from the Secure Agent machine.
2. On the master node, locate the Kubernetes API Server log files in the directory `/var/log/`.

I updated the staging location for the advanced cluster. Now mappings fail with the following error:

```
Error while executing mapping. ExecutionId '<execution ID>'. Cause: [Failed to start
cluster for [01000D25000000000005]. Error reported while starting cluster [Cannot apply
cluster operation START because the cluster is in an error state.].].
```

Mappings fail with this error when you change the permissions to the staging location before you change the S3 staging location in the advanced configuration.

If you plan to update the staging location, you must first change the S3 staging location in the advanced configuration and then change the permissions to the staging location on AWS. If you used role-based security, you must also change the permissions to the staging location on the Secure Agent machine.

To fix the error, perform the following tasks:

1. Revert the changes to the permissions for the staging location.
2. Edit the advanced configuration to revert the S3 staging location.
3. Stop the cluster when you save the configuration.
4. Update the S3 staging location in the configuration, and then change the permissions to the staging location on AWS.

I updated the staging location for the advanced cluster. Now the following error message appears in the agent job log:

```
Could not find or load main class com.informatica.compiler.InfaSparkMain
```

The error message appears when cluster nodes fail to download Spark binaries from the staging location due to access permissions.

Verify access permissions for the staging location based on the type of connectors that the job uses:

Connectors with direct access to Amazon data sources

If you use credential-based security for advanced jobs, make sure that the credentials in the Amazon S3 V2 and Amazon Redshift V2 connections can be used to access the staging location.

If you use role-based security for advanced jobs, make sure that the advanced cluster and the staging location exist under the same AWS account.

Connectors without direct access to Amazon data sources

If you use a user-defined worker role, make sure that the worker role can access both the staging location and the data sources in the advanced job.

If you use the default worker role, make sure that the Secure Agent role can access both the staging location and the data sources in the advanced job.

I restarted the Secure Agent machine and now the status of the advanced cluster is Error.

Make sure that the Secure Agent machine and the Secure Agent are running. Then, stop the advanced cluster in Monitor. In an AWS environment, the cluster might take 3 to 4 minutes to stop. After the cluster stops, you can run an advanced job to start the cluster again.

Is there anything I should do before I use a custom AMI to create cluster nodes?

If you use a custom AMI (Amazon machine image) to create cluster nodes, make sure that the AMI contains an installation of the AWS CLI.

The Secure Agent uses the AWS CLI to propagate tags to Amazon resources and to aggregate logs. The cluster nodes also use the AWS CLI to run initialization scripts.

For information about how to use a custom AMI, contact Informatica Global Customer Support.

Troubleshooting an advanced cluster on Microsoft Azure

I restarted the Secure Agent machine and now the status of the advanced cluster is Error.

Make sure that the Secure Agent machine and the Secure Agent are running. Then, stop the advanced cluster in Monitor. In an Azure environment, the cluster might take 10 minutes to stop. After the cluster stops, you can run a job to start the cluster again.

The init script failed with the following standard error on some nodes in the advanced cluster:

```
Created symlink from /etc/systemd/system/apt-daily.service to /dev/null.
Created symlink from /etc/systemd/system/apt-daily-upgrade.service to /dev/null.
Removed symlink /etc/systemd/system/timers.target.wants/apt-daily.timer.
Removed symlink /etc/systemd/system/timers.target.wants/apt-daily-upgrade.timer.
E: Could not get lock /var/lib/dpkg/lock-frontend - open (11: Resource temporarily
unavailable)
E: Unable to acquire the dpkg frontend lock (/var/lib/dpkg/lock-frontend), is another
process using it?
```

The init script failed because the node was running an internal process at the same time as the init script. If you continue to see the error, wait for the internal process to complete by placing a sleep command for the required duration in your init script.

For example, you might use a sleep command as follows:

```
#!/bin/sh

while(sudo lsof /var/lib/dpkg/lock-frontend)
do
echo "Sleeping 10s"
sleep 10
done

sudo apt-get -y update
sudo apt-get install -y expect
```

Troubleshooting a connection

Review the rules and guidelines for the connections. To find general information about connections, see *Connections*. To find connection-specific information, see the help for the appropriate connector.

Connection errors

The connection I want to use does not display in the task wizard.

You do not have read permission on the connection. You need read permission on all assets that you want to use in a task or taskflow. Ask your organization administrator for the appropriate permissions.

Troubleshooting assets

I want to create a new asset but New does not display in the navigation bar.

When a license expires, you cannot access the features associated with the license. Contact your organization administrator.

I want to copy multiple assets, but the operation fails.

If you copy an asset and its duplicate with the suffix "Copy 1" into a folder where assets with the same name already exist and choose to keep both assets, Informatica Intelligent Cloud Services might not copy the assets.

For example, you want to copy "Mapping_1" and "Mapping_1-Copy 1" into Folder_2, but "Mapping_1" already exists in Folder_2. "Mapping_1-Copy 1" copies successfully, but "Mapping_1" might fail because both "Mapping_1" and "Mapping_1-Copy 1" now also exist in Folder_2.

Change the name of the conflicting assets and try again, or copy the assets one at a time.

Troubleshooting a task

You can view the number of rows that failed to load into the target when you view the details for a job. For more information about the reason for each failure, view the error rows file.

To view job details, in Data Integration, select **My Jobs** and then click the job name. You can also view job details in Monitor on the **All Jobs** or **Running Jobs** page.

You can find the error rows files in the following directory:

```
<Secure Agent installation directory>\apps\Data_Integration_Server\data\error
```

If a job does not complete, you can stop the job on the **All Jobs**, **Running Jobs**, or **My Jobs** page.

Task errors

My task failed to complete and the following message appears in the job details:

```
DTM process killed for workflow [wf_dss_<workflow id>] session [s_dss_<session id>]
because agent is shutdown.
```

The Secure Agent that was running the task was stopped or restarted while the task was in progress.

My tasks are failing after I installed a 64-bit Secure Agent on Windows. The following message appears in the job details:

```
Failed to run job because of missing VC++ libraries on the machine that runs the Secure
Agent.
```

The Secure Agent machine that runs the task is missing the `MSVCR.DLL` files. Install the Microsoft Visual C++ 2008 Redistributable Package (x86) on the Secure Agent machine.

My task failed to complete and the following message appears in the job details:

```
The DTM process failed to start due to [CreateProcess error=14001, The application has
failed to start because its side-by-side configuration is incorrect. Please see the
application event log or use the command-line sxstrace.exe tool for more detail].
'Failed to run job because of missing VC++ libraries on the machine that runs the Secure
Agent'
```

The Secure Agent machine that runs the task is installed on a Windows 2008 machine that is not correctly configured. Install the Microsoft Visual C++ 2008 Redistributable Package (x86) on the Windows 2008 machine.

I need to review details for a task, but I can't find the task anywhere!

The role to which you are assigned might not have privileges to read this type of task, or you might not have Read permission for the task. To view task details, you must be assigned a role that has Read privilege for the task type. If permissions are set for the task, you also need Read permission for the task. Ask your organization administrator for the appropriate privileges and permissions.

I can view task details, but I cannot edit a task.

The role to which you are assigned might not have privileges to update this type of task, or you might not have Update permission for the task. To edit a task, you need to be assigned a role that has Update privilege for the task type and Read privilege for all asset types used in the task. If permissions are set for the task, you need Update and Read permissions for the task. If permissions are set for any of the assets in the task, you also need Read permission for the assets. Ask your organization administrator for the appropriate privileges and permissions.

When I edit a task, one of the connections displays a blank field instead of the connection name.

The role to which you are assigned might not have the Read privilege for connections, or you might not have Read permission for the connection. When you edit a task, you should have Read privileges for all asset types used in the task. If permissions are set for assets in the task, you should also have read permission for the assets. To avoid unexpected results, cancel all changes and ask your organization administrator for the appropriate privileges and permissions.

After you cancel out of the wizard, you can view the task details to determine the names of the objects for which you need permissions.

When creating a task, the connection that I want to use does not appear in the list of available connections.

The role to which you are assigned might not have the Read privilege for connections, or you might not have Read permission for the connection. When you create a task, you need Read privileges for all asset types that you want to use in the task. If permissions are set for the assets you want to use, you also need Read permission for the assets. Ask your organization administrator for the appropriate privileges and permissions.

The following error appears:

```
Unknown SQL datatype for field [<Object name.Field name>]. [<Number that represents the datatype>]
```

Data Integration does not recognize the data type of a source or target field.

Troubleshooting a mapping task

When I run a mapping task, the task fails with the following message:

```
TE_7073: Expecting Keys to be ascending.
```

This occurs when you configure an Aggregator transformation for sorted data, but do not sort the data earlier in the data flow. For some connection types, you can configure a Source transformation to sort data. You can also use a Sorter transformation.

When I run a mapping task, the task completes successfully with the following row error message:

```
FR_3016 Row [<row>]: Record length [<row length in bytes>] is longer than line sequential buffer length [<maximum buffer length in bytes>] for [<file name>]. Record will be rejected.
```

This message appears when the number of bytes the task reads for the row exceeds the calculated maximum line sequential buffer length.

For example, if the code page of the source flat file and the code page defined for the connection don't match, the row length can exceed the calculated maximum line sequential buffer length, and Data Integration rejects the row.

To resolve the error, configure the Line Sequential Buffer Length advanced session property and set the value to the record length in the error message.

Troubleshooting an advanced cluster subtask

The job failed but there are many logs I can view. Where do I start?

Troubleshoot the job by examining the logs in the following order:

1. Execution plan. Debug the Scala code for the job.

2. Session log. Debug the logic that compiles the job and generates the Spark execution workflow.
3. Agent job log. Debug how the Secure Agent pushes the Spark execution workflow to the advanced cluster for processing.
4. Spark driver and executor logs. Debug how the advanced cluster runs the job.

You can download the execution plan, session log, agent job log, and Spark driver log in Monitor.

To find the Spark executor log, copy the advanced log location for a specific Spark task that failed. Then, navigate to the log location on your cloud platform and download the log.

I can't find all of the log files for the job that failed. I've tried to download the logs from both Monitor and the log location on my cloud platform.

The logs that are available for the job depend on the step where the job failed during processing.

For example, if the job fails before the job is pushed to the advanced cluster, the Spark driver and executor logs are not generated in the log location, and Monitor cannot query the logs from the cloud platform either.

You can recover some of the log files, but you might have to use other types of logs to troubleshoot the job.

I can't find the Spark driver and Spark executor logs. Can I recover them?

If you can't download the Spark driver log from the user interface, you can recover the log using the Spark driver Pod. You cannot recover Spark executor logs.

When the Secure Agent pushes a job to an advanced cluster, the Secure Agent creates one Spark driver Pod and multiple Spark executor Pods to run the Spark tasks. You can use the Spark driver Pod to recover the Spark driver log, but you cannot recover the Spark executor logs. The Spark driver Pod deletes the Spark executor Pods immediately after a job succeeds or fails.

Note: When a job succeeds or fails, the Spark driver Pod is deleted after 5 minutes by default. If you need to increase the limit to assist troubleshooting, contact Informatica Global Customer Support.

To recover the Spark driver log, perform the following tasks:

1. Find the name of the Spark driver Pod in the agent job log. For example, see the name of the Spark driver Pod in the following message:

   ```
   2019/04/09 11:10:15.511          : INFO :Spark driver pod [spark-
   passthroughparquetmapping-veryvery-longlongname-1234567789-
   infaspark02843891945120475434-driver] was successfully submitted to the cluster.
   ```

   If you cannot download the agent job log in Monitor, the log is available in the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/logs/job-logs/
   ```

   The file name of the agent job log uses the format *AgentLog-<Spark job ID>.log*. You can find the Spark job ID in the session log. For example, the Spark job ID is *0c2c5f47-5f0b-43af-a867-da011452c19dInfaSpark0* in the following message of the session log:

   ```
   2019-05-09T03:07:52.129+00:00 <LdtmWorkflowTask-pool-1-thread-9> INFO: Registered
   job to status checker with Id 0c2c5f47-5f0b-43af-a867-da011452c19dInfaSpark0
   ```

2. Confirm that the Spark driver Pod exists. If the driver Pod was deleted, you cannot retrieve the Spark driver log.

   To confirm that the driver Pod exists, navigate to the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/mercury/
   services/shared/kubernetes/kubernetes_<version>/bin
   ```

   In the directory, run the following command:

   ```
   ./kubectl get pods
   ```

3. Find the cluster instance ID in one of the following ways:
   - Locate the cluster instance ID in the session log. For example, you might see the following message:

     ```
     2019/05/07 16:22:00.20    : INFO :[SPARK_2005] Uploading the local file in the
     path [/export/home/builds/ws/yxiao_hadoopvm_ML/Mercury/platformdiscale/main/
     components/cluster/hadoop-tests/cats/edtm/spark/./target/
     hadoop3a0b1db6-76ea-4317-8272-5b3a8dfd2171_InfaSpark0/
     log4j_infa_spark.properties] to the following shared storage location: [s3a://
     soki-k8s-local-state-store/k8s-infa/testcluster2.k8s.local/staging/
     sess4280021555102778947/log4j_infa_spark.properties].
     ```

     Note the following cloud storage location that you see in the message:

     ```
     s3a://soki-k8s-local-state-store/k8s-infa/testcluster2.k8s.local/staging/
     ```

     The cluster instance ID is the entry that follows "k8s-infa." In this case, the ID is testcluster2.k8s.local.
   - Locate the cluster instance ID in the `ccs-operation.log` file. The file is located in the following directory on the Secure Agent machine:

     ```
     <Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
     ```
4. Log in to the Secure Agent machine as the sudo user that started the agent.
5. Set the environment variable KUBECONFIG on the Secure Agent machine to the following value:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/ccs_home/
   <cluster ID>/.kube/kubeconfig.yaml
   ```
6. To retrieve the Spark driver log, navigate to the following directory on the Secure Agent machine:

   ```
   <Secure Agent installation directory>/apps/At_Scale_Server/<version>/mercury/
   services/shared/kubernetes/kubernetes_<version>/bin
   ```

   In the directory, run the following command:

   ```
   ./kubectl logs <Spark driver pod name>
   ```

Troubleshooting a synchronization task

Review the rules and guidelines for the components of a synchronization task.

Synchronization task errors

The synchronization task fails.

A synchronization task might fail for the following reasons:

- You tried to run a synchronization task that was already running. You can run one instance of a synchronization task at a time.
- You tried to run a synchronization task, but the corresponding Secure Agent is not running. Verify that the Secure Agent is running.
- You tried to run a synchronization task that reads data from or writes data to a database, but the corresponding database server is not running. Verify that the database server is running.
- You tried to run a synchronization task that contains an invalid field mapping.

  If the source or target changed after you created the field mapping for a synchronization task and you try to run the synchronization task, the following error appears:

  ```
  TE_7002 Transformation Parse Fatal Error; transformation stopped...
  ```

The synchronization task fails in the following cases:

- A required field is not mapped in the field mapping. You may have changed the default field mappings or a required field may have been added to the source or target after you created the synchronization task. You must edit the synchronization task and include all required fields before you can successfully run the synchronization task. For more information about required fields in field mappings of synchronization tasks, see *Tasks*.
- A field in a field mapping no longer exists in the source or target. The structure of the source or target can change after you create the field mapping in the synchronization task. Resolve the field mappings in the synchronization task that point to nonexistent fields.

- You tried to run a synchronization task that contains an invalid field name. When a field name is invalid, the task fails with the following error:

  ```
  [DTM_10004] The DTM process terminated unexpectedly.
  ```

- You tried to run a synchronization task and the field in the data filter changed in one of the following ways:
  - The field no longer exists in the source.
  - The datatype of the field changed.

  Edit the data filters in the synchronization task definition based on the latest version of the source.

My flat file target contains data from the latest job, but not from previous jobs.

The synchronization task truncates a flat file target before writing target data to the file.

To avoid overwriting existing target data, you might use a post-session command to merge target data with a master target file in a different location.

The data preview does not show the latest changes to the dataset for the synchronization task.

The synchronization task wizard caches the dataset for connections when you configure the task. The synchronization task wizard deletes the cache after you save the task and refreshes the cache each time you configure the task. To refresh the cache, save or cancel the task, and then edit the task.

For FTP/SFTP connections, data preview shows the latest changes for the dataset when you update the local file to match the remote file.

I do not see the target when creating the task.

The list of targets are based on the task operation. For more information, see *Tasks*.

The synchronization task inserted all data into the target, instead of performing the update/upsert/delete task operation configured for the task.

When you use an active mapplet with a synchronization task that includes a saved query, the synchronization task ignores the configured target option for the task and tries to insert data to the target.

Troubleshooting a replication task

Review the rules and guidelines for the components of a replication task.

Replication task errors

I chose a source connection while creating a replication task, but I cannot see any available objects to replicate.

When you choose a source connection in the replication task, Data Integration connects to the source to retrieve objects that it can replicate. If the connection is not valid or the connection does not have privilege to view the object in the source system, Data Integration cannot connect to the source and shows no objects.

An error message appears while I configure the target for the replication task.

You cannot configure multiple replication tasks to replicate the same source object to the same target object. For example, you cannot configure two replication tasks to write Salesforce Account data to the SF_ACCOUNT Oracle database table. When you select the same target connection and target object that is used by another replication task, an error appears.

If you replicate all objects for the source connection, the following error appears:

```
Duplicate data replication task. The data replication task <name of the other data
replication task> is also replicating the same objects from the same source to the same
target with the same target prefix as this data replication task. Change the target or
enter a different target prefix.
```

If you replicate specific objects for the source connection, the following error appears:

```
Duplicate data replication task for <object name> object with a target prefix of <target
prefix>. The data replication task <name of the other data replication task> is already
replicating an object whose target prefixed name is same as this one. Change the target
connection or enter a different target prefix or remove the object from the task.
```

To write data from sources with the same name to different target objects in the same target connection, configure different target prefixes. For more information, see *Tasks*.

I ran a replication task, but I did not get results in the expected time period.

A replication task may take longer than expected for the following reasons:

- The replication task is in queue because too many replication tasks are running simultaneously. Contact Informatica Global Customer Support to check the status of the replication task.
- You have a large amount of data in the source.
- Database servers have an outage or are experiencing heavy traffic.

Why did my replication task fail?

A replication task might fail for the following reasons:

- The source or target connection might have timed out or it might be down.
- The Secure Agent is not running.
- The target database server is not running. If you restart the database server, the test connection to the database succeeds, but the replication task fails, restart the Secure Agent.
- Your organization's license to use the replication application expired.
- You tried to run a replication task that was already running.

If you verify that the replication task did not fail for any of the above reasons, contact Informatica Global Customer Support and provide the scenario and, if available, error messages.

I did not receive an email notification that stated that my replication task failed.

You can configure the task to send an email to you if the session completes, completes with errors, or fails. If you should receive email for all tasks, ask the organization administrator to configure the organization properties to send you email notification.

The target table name is truncated.

If the length of the source object name exceeds the maximum number of characters allowed for a table name in a relational target, the replication task truncates the corresponding table name. For more information, see *Tasks*.

The table in the target database does not contain the same data as the corresponding source object.

The replication task overwrites the data in a target table when two source objects from different replication tasks share the same target table. When multiple source objects write to the same target table, the replication task overwrites table data each time you replicate one of the source objects.

To avoid overwriting tables, use a different target table name prefix for each replication task. For more information, see *Tasks*.

The column name in the target database is truncated.

If the length of the source object name exceeds the maximum number of characters, the replication task truncates the corresponding column name based on the maximum number of characters allowed for the column name in the target database. For more information, see *Tasks*.

The following error appears when I try to create the target for the replication task:

```
Please use Reset Target first and then try Create Target again.
```

This message appears for the following reasons:

- You click the Create Target option twice.
- You edit the replication task, and then click Create Target.

To resolve, click Reset Target, and then click Create Target.

Troubleshooting masking task errors

When you perform a masking task, the following errors can occur:

Either the target does not contain the Account object, or there is a problem connecting to the source or target. Verify the source connection, target connection, and source and target objects. If the problem persists, contact Informatica Global Customer Support.

This error occurs when the masking task is not able to connect to the target or there are some objects are missing in the target connection.

REQUEST_LIMIT_EXCEEDED: Total Requests Limit exceeded.

This error occurs when you exceed the API usage limit for the day.

Error loading into target [Object_Name] : Error received from salesforce.com. Fields []. Status code [STORAGE_LIMIT_EXCEEDED]. Message [storage limit exceeded].

This error occurs when the target object is full or exceeds the limit.

Reattempt the Salesforce request [handleBulkApiError] due to the error [Invalid session id].

This error occurs when there is a temporary problem with Salesforce.com or the network. To check whether there are any known server issues, you can refer to http://trust.salesforce.com.

Error loading into target [RecordType] : Error received from salesforce.com. Fields []. Status code [DUPLICATE_VALUE]. Message [duplicate value found: <unknown> duplicates value on record with id: <unknown>].

This error occurs when a Record Type with the same name exists in the target.

The [Object_Name] object does not contain any fields that you can update. Remove the object and try again.

This error occurs when you cannot update an object or add a record. You must remove the object and run the task. This error can also occur when you have read-only access to the object. Verify that you have the right permissions.

Error : Failed to create the external ID field in the target [Object_Name](Error : External ID creation for this object is not currently supported. Select another object. If the problem persists, contact Informatica Global Customer Support.)

This error occurs when the application fails to create an external ID in the target. To upsert the data into target object, you need an external ID field. If the application does not allow to create the external ID fields, then you cannot upsert data into the target objects.

The selected object repeatedly calls one or more mandatory objects that are not supported. Remove the object and run the task.

This error occurs when the Salesforce objects defined in the composite objects do not exist in the source organization.

Cloud Masking mapgen failed. Contact Informatica Global Customer Support.

com.informatica.ilm.tdms.mapgen.exceptions.CyclicRelationException: Cyclic schemas are not supported on Salesforce.

This error occurs due to cyclic relationships within the object sets. Remove the objects that contain such cyclic relationships.

INSUFFICIENT_ACCESS_ON_CROSS_REFERENCE_ENTITY

This error occurs when the target does not contain all the related objects and there are no sufficient permissions to write on to the target.

Masking page is empty for inplace and instream masking.

This error occurs when you have read-only access to source for inplace masking or when you have read-only access to target for instream masking. When you create a masking task, the masking page is blank and you cannot view any fields.

SFDC_31102 [FATAL] Query failed. [MALFORMED_QUERY: SOQL statements cannot be longer than 20000 characters.]

This error occurs when the SOQL exceeds 20000 characters. You must contact Salesforce to increase the SOQL character limit.

Error loading into target [OpportunityLineItem]: Error received from salesforce.com. Fields [UnitPrice]. Status code [FIELD_INTEGRITY_EXCEPTION]. Message [field integrity exception: UnitPrice (only one of unit price or total price may be specified)].

This issue occurs because in the OpportunityLineItem object, both the UnitPrice and TotalPrice fields are visible and can be updated. The task fails if you insert records into both the fields. To resolve the issue, apply nullification masking on the UnitPrice field or the TotalPrice field while you insert a record into the OpportunityLineItem object.

Error loading into target [OpportunityContactRole] : Error received from salesforce.com. Fields []. Status code [INVALID_FIELD]. Message [Foreign key external ID: 003o000000a67jyaay not found for field DMASK_EXTERNAL_ID__c in entity Contact].

This issue occurs because you can neither create an external ID nor select a unique field for the OpportunityContactRole object when you perform an insert operation. Perform an upsert operation and select the idlookup field for the OpportunityContactRole object to avoid the error.

Troubleshooting a PowerCenter task

Review the rules and guidelines for the components of a PowerCenter task.

PowerCenter task errors

The following error appears when you import a workflow XML file:

```
The PowerCenter XML file is empty, please select a valid PowerCenter XML file.
```

The workflow XML file is empty.

The following error appears when you import a workflow XML file:

```
The PowerCenter file content type is not text/xml, please select a valid PowerCenter XML
file.
```

The workflow XML file is not valid.

The following error appears when you run a PowerCenter task with more than 64 partitions:

```
TM_6710 Error: Unable to allocate storage for statistics for transformation
<transformation name>.
```

The PowerCenter workflow cannot contain more than 10 partitions for sources and targets.

The following error appears when you run a PowerCenter task:

```
PowerCenter task <PowerCenter task name> failed to run. Another instance of the task is
currently executing.
```

You cannot run multiple instances of a PowerCenter task simultaneously. If you configured the task to run on a schedule, increase the time interval between the scheduled tasks to prevent multiple instances of the task from running simultaneously. If you run the task manually, wait for the currently running instance of the task to complete before starting it again. You can view running tasks on the **All Jobs** or **Running Jobs** page in Monitor or on the **My Jobs** page in Data Integration.

The PowerCenter task does not complete.

PowerCenter tasks run by a Secure Agent on Windows 7 (64 bit) might not complete. To resolve the issue, configure a network login for the Secure Agent service.

Troubleshooting a mapping

Review the rules and guidelines for configuring a mapping.

I made a lot of changes to source field metadata that I want to revert. I tried the Synchronize icon, but nothing's happening. What's going on?

Synchronization does not restore field metadata that was updated in a Source transformation. You can manually update the field metadata. Or, you can start over: Clear and select source object or create a new Source transformation.

A field name conflict error message displays in a transformation. In the transformation, I created a field rule that renamed the incoming fields, but the error continues.

A field name conflict can continue throughout a mapping after the initial problem. Locate the transformation with the first field name conflict error message and rename the fields before they enter the transformation.

A mapping failed because the number of fields in the compiled mapping exceeded the threshold. What should I do?

Mappings can have up to 7,000,000 input and output fields. When a mapping includes the Hierarchy Processor transformation, the mapping is more likely to reach this limit because of the complex data involved. The more joins, child fields, nested fields, and flattened arrays that the Hierarchy Processor transformation contains, the more likely the mapping will exceed the field limit.

If the mapping exceeds the limit, the following message appears in the mapping compilation log:

```
[LDTM_0502] The mapping [<mapping name>] failed because the number of fields in the
compiled mapping exceeds the threshold: [7,000,000]. Number of fields: [<actual
number>]. Create multiple mappings to process the data incrementally.
```

To resolve the error, create multiple mappings to process the complex data incrementally, and reduce the size of the mapping.

The Source transformation returns a JSON parsing error, even though my JSON file is good. What should I do?

When you select a JSON file in the source object, you receive the following error:

```
error [The file or partition directory [<path>/<JSON file>] is not valid. The parser
encountere the following error while parsing the content: [Unexpected character 1 at 1].
Select a valid [Json] file or partition director.]
```

The error occurs when there is an array of primitives without a field name in the JSON file. For example, the following array causes a parsing error: `[10,20,30,40,50,0]`

The Source transformation parses the following array successfully: `{"numbers":[10,20,30,40,50,0]}`

As a workaround, wrap the top-level array in a struct element before reading the file in the Source transformation.

Troubleshooting a mapping in advanced mode

Review the rules and guidelines for configuring a mapping in advanced mode.

The Source transformation returns a JSON parsing error, even though my JSON file is good. What should I do?

When you select a JSON file in the source object, you receive the following error:

```
error [The file or partition directory [<path>/<JSON file>] is not valid. The parser
encountere the following error while parsing the content: [Unexpected character 1 at
1]. Select a valid [Json] file or partition director.]
```

The error occurs when there is an array of primitives without a field name in the JSON file. For example, the following array causes a parsing error: `[10,20,30,40,50,0]`

The Source transformation parses the following array successfully: `{"numbers":[10,20,30,40,50,0]}`

As a workaround, wrap the top-level array in a struct element before reading the file in the Source transformation.

Troubleshooting scheduled tasks

The task does not run at the scheduled time.

A task does not run at the scheduled time if another instance of it is already running when the schedule tries to start the task. For example, you schedule a task to run every 5 minutes. The first task starts at 12 p.m., but does not complete until 12:06 p.m. The second instance of the task does not run at 12:05 p.m. because the first instance has not completed. Data Integration starts the next task at 12:10 p.m.

To resolve this issue, change the schedule to allow the task to complete before starting the next task run.

Troubleshooting security

I received the following security violation error:

```
There may have been a security violation while accessing the site. Verify that there are
no malicious scripts running in your browser. This error also appears when you submit
the form multiple times through a browser reload.
```

This error appears when you click an option on a page while the page is still loading from a previous click. Click the Here link to return to Data Integration.

When I try to view the details about an object, such as a connection or replication task, the Object Not Found page displays.

The object was recently deleted. The Object Not Found page appears when an object no longer exists. Refresh the page to display current objects.

When I try to perform a task, the Access Denied page displays.

The Access Denied page displays when you try to perform a task that is not allowed for your user account. You might not have the appropriate role or asset permissions to perform the task. If you need to perform the task, ask your organization administrator to review your user account.

INDEX

A

advanced cluster subtasks
  troubleshooting 16
advanced clusters
  AWS 10
  Microsoft Azure 13
  troubleshooting 10, 13

C

Cloud Application Integration community
  URL 4
Cloud Developer community
  URL 4
column names
  truncating 20
connections
  troubleshooting 14

D

Data Integration community
  URL 4
Data Replication tasks
  errors 20
  troubleshooting 19
Data Synchronization tasks
  troubleshooting 18

E

error rows files
  location 14

I

Informatica Global Customer Support
  contact information 5
Informatica Intelligent Cloud Services
  web site 4

L

login denied
  troubleshooting 25

M

maintenance outages 5
mapping tasks
  troubleshooting 16

O

organization ID
  finding 6
organizations
  finding your organization ID 6

P

PowerCenter tasks
  troubleshooting 23

S

Secure Agent
  troubleshooting 7
security
  troubleshooting 25
status
  Informatica Intelligent Cloud Services 5
system status 5

T

table names
  truncating 20
tasks
  troubleshooting 14
troubleshooting
  advanced cluster subtasks 16
  advanced clusters 10, 13
  connections 14
  Data Replication tasks 19
  Data Synchronization tasks 18
  error rows files 14
  PowerCenter tasks 23
  Secure Agent 7
  security 25
  tasks 14
trust site
  description 5

U

upgrade notifications 5

W

web site 4