



Interset

© Copyright 2020 Micro Focus or one of its affiliates.

The only warranties for products and services of Micro Focus and its affiliates and licensors ("Micro Focus") are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Except as specifically indicated otherwise, this document contains confidential information and a valid license is required for possession, use or copying. If this work is provided to the U.S. Government, consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Contents

Introduction	5
Supported Environments	5
Supported Data Sources	6
Intended Audience	7
How to Use This Guide	7
Getting Further Assistance	8
Prepare to Upgrade Interset 5.7.x to Interset 5.9.2	9
Offline Upgrade: Download and Extract the Interset 5.9.2 Packages	9
Update Your Local Interset Mirror Repository	10
Prepare to Upgrade	12
Upgrade Interset	13
Stop the Flume Service	13
Upgrade the interset User	14
Install and Configure the PostgreSQL External Database	14
Upgrade Elasticsearch to Version 6.8	17
Remove the Elasticsearch Plug-ins	17
Upgrade Elasticsearch to Version 6.8.1	17
Upgrade Kibana to Version 6.8.1	19
Upgrade Interset Reporting to Version 5.9.2	19
Upgrade the Interset Exports Service	21
Upgrade the Interset Analytics Component to Version 5.9.2	21
Upgrade the Interset EDR Server to Version 5.9.2	23
Upgrade Interset Data Ingest to Version 5.9.2	25
Upgrade Stream Nodes	25
Upgrade NiFi to Version 1.7.1	26
Update Flume Data Ingest Configuration Information	28
Upgrade Interset Elasticsearch Index Data, Mappings, and Templates	28
Upgrade Workflow to Version 5.9.2	29

Suspend the Active Workflows	29
Upgrade Workflow	29
Restart Ambari Services	31
Appendix A: Install and Configure the Schema Registry	32
Install the Schema Registry	32
Configure the ConfluentSchemaRegistry Controller Service	32
Configure Readers to Use the ConfluentSchemaRegistry Controller Service	33
Configure a New Data Source	34
Configure an SSL Context Service	35
Create a Process Group	37
Configure the Schema Registry	37
Enable the Controller Services	38
Configure the Data Flow	38
Configure the GetLocalCSVRecord Process Group	39
Configure the FormatCSVData Process Group	40
Configure the AsyncActiveDirectoryHbaseProcessor	40
Configure the LoadElasticsearch Processor	42
Configure the PublishKafkaRecord_0_10 Processor	43
Start the Data Flow	44
Configure Violations Loading	45
Update Workflow for Schema Registry	45
Update the EDR Server for Schema Registry	46
Appendix B: Flume Configuration Changes	48
Extract Configurations	48
Transform Configurations	51
Loading Configurations	57
Index	61

Introduction

This guide describes how to upgrade from Interset 5.7.x to version 5.9.2, and also provides information about supported environments and data sources.

Supported Environments

Interset 5.9.2 is supported in the following (x86_64) environments:

- CentOS 7.6
- Red Hat Enterprise 7.6

Interset 5.9.2 is supported with the following third-party components:

- Oracle OpenJDK 8u201/211
- Elasticsearch 6.8.1

Interset 5.9.2 is supported with HDP 3.1.0, including the following components:

Ambari	2.7.3.0
AsyncHBase	1.8.2
Avro	1.8.2
Hadoop	3.1.1
HBase	2.0.2
Hortonworks Schema Registry	0.5.3
Kafka	2.0.0
NiFi	1.10.0
Phoenix	5.0.0
Scala	2.11.8
Spark	2.3.2
Storm	1.2.1
ZooKeeper	3.4.6

Interset 5.9.2 is supported with CDH 6.1.1, including the following components:

AsyncHBase	1.8.2
Avro	1.8.2
Hadoop	3.0.0
HBase	2.1.1
Hortonworks Schema Registry	0.5.3
Kafka	2.0
NiFi	1.10.0
Phoenix	5.0.0-HBase-2.0-cdh6.1.1
Scala	2.11.8
Spark	2.4

Storm	1.2.1
ZooKeeper	3.4.5

Interset 5.9.2 supports the following Web browsers:

- Google Chrome 74 and above
- Mozilla Firefox 67 and above

Supported Data Sources

Interset 5.9.2 supports the following data sources. For .csv data sources, the delimiter can be customized.

- Active Directory
 - Active Directory event logs stored in McAfee® Enterprise Security Manager (ESM)
 - Active Directory event logs stored in Splunk®
 - Active Directory event logs stored in Micro Focus ArcSight Logger
 - Active Directory event logs stored in IBM QRadar
 - Windows Security event logs (.csv)
 - Interset-extracted Windows event logs (.csv)
 - Universal Windows event logs (.csv)
 - Windows Event Viewer-extracted event logs (.csv)
 - Active Directory authentication logs
- Universal Alerts stored in third-party DLP systems (.csv)
- NetFlow
 - Version 5
 - Version 9
 - Version 10 (IPFIX)
- Repository
 - Perforce
 - P4AUDIT logs
 - Perforce Structured Server audit logs
 - GitHub Enterprise audit logs
 - Universal repository logs (.csv)
- Pluggable Authentication Module (PAM) AuditD logs (.csv)
- Printer logs
 - Windows printer events stored in Splunk
 - Windows event logs (.csv)
 - Universal logs (.csv)
- Universal Web Proxy (.csv)

- Violations
- Expense Data
- Email Data

Interset 5.9.2 data ingest uses NiFi for data extraction, transformation, and loading. It supports the processing of data set files in the following compression formats:

- tar
- gzip
- tar gzip

To ingest packaged data from other containers such as Arcsight, IBM QRadar, McAfee ESM, and Splunk, please contact Micro Focus Interset Support at interset.support@microfocus.com.

Intended Audience

This Guide assumes that you are an experienced system administrator with sound Linux skills and are familiar with your organization's server environment, security infrastructure, and data sources.

You should also be familiar with the business needs of your organization.

How to Use This Guide

The instructions in this Upgrade Guide can be used when your Interset cluster is representative of a standard, out-of-the-box deployment as described in the Interset 5.7.x *Installation and Configuration Guide*. If your cluster is not a standard, out-of-the-box deployment, please contact Micro Focus Interset Support at interset.support@microfocus.com for assistance with your Interset 5.9.2 upgrade.

The instructions in this Guide refer to a **<mirror>** location, a repository with the upgrade packages and files required for this Interset upgrade:

- In an offline upgrade scenario, this **<mirror>** location must be created on a local server accessible to your Interset cluster.
- In an online upgrade scenario, you configure this **<mirror>** location to be the online Interset repository location. If you require access credentials, please contact Micro Focus Interset Support at interset.support@microfocus.com.

As you use this Guide to prepare for and execute your Interset upgrade, we recommend that you work through the sections in the order in which they are presented.



Important: The scripts and commands provided throughout the upgrade instructions are designed to be copied from this Guide to the command prompt in your console. However, if you are viewing this Guide in a program other than Adobe Acrobat Reader, the scripts may not copy correctly. ***As a result, Interset strongly recommends that you use Adobe Acrobat Reader to view this Guide.***

Getting Further Assistance

Should you experience unexpected results or identify issues that are not addressed in this document, please contact Micro Focus Intersect Support at intersect.support@microfocus.com.

Prepare to Upgrade Interset 5.7.x to Interset 5.9.2

You can use the instructions in this Upgrade Guide when your Interset cluster is representative of a standard, out-of-the-box deployment as described in the Interset 5.7.x *Installation and Configuration Guide*. If your cluster is not a standard, out-of-the-box deployment, please contact Micro Focus Interset Support at interset.support@microfocus.com for assistance with your Interset 5.9.2 upgrade.

Before you can upgrade to Interset 5.9.2, you must:

- [download and extract the Interset 5.9.2 package](#)
- [update your local Interset mirror repository](#)
- [prepare to upgrade](#)



Important: Unless instructed otherwise, all commands are run as the **interset** user.

Offline Upgrade: Download and Extract the Interset 5.9.2 Packages



The instructions in this Guide refer to a `<mirror_fqdn>` location, a repository with the upgrade packages and files required for this Interset upgrade:

- In an offline upgrade scenario, this `<mirror_fqdn>` location is created on a local Web server machine accessible to your Interset cluster.
- In an online upgrade scenario, you configure this `<mirror_fqdn>` location to be the online Interset repository location provided to you by Micro Focus Inter-set Support.

Tip: Within this *Upgrade Guide*, the mirror server is referred to as `<mirror_fqdn>`.

This section provides instructions for performing an offline upgrade of your Interset cluster, and for updating your existing offline 5.7.x repository for version 5.9.2.

1. In a Web browser, navigate to <https://repo.interset.com/5.9.2/interset>.

If you require access credentials, please contact Micro Focus Interset Support at inter-set.support@microfocus.com.

2. Download and extract the **repo570.tar** file to your offline repository server.



Important: You must download and extract the Interaset 5.9.2 files as the **interaset** user. If you download and extract these files as a different user within your environment, you will be unable to run the Interaset scripts necessary to perform the upgrade.

3. Extract the **repo570.tar** file to the root level of your repository.
4. In a Web browser, ensure that you can access the extracted contents of the **repo570.tar** file at **http://<mirror_fqdn>/5.7.0/**.
5. On the Monitoring node in your Interaset cluster, download the **interaset_installer_5.7.0.tar** file from **http://<mirror_fqdn>/5.7.0/interaset**, and extract it to the **/tmp** directory.



Important:

- The umask on your system must be set to **027**, or a more permissive value to ensure a successful installation.
- The Interaset upgrade requires access to the CentOS / RHEL base repository to retrieve multiple dependencies that are not bundled within the Interaset mirror repository. Please ensure that this access is available either by an off-line, locally-mirrored repository for the relevant operating system, or through the use of an RHEL satellite server.
- Each node in the Interaset cluster must be able to access the local repositories.

Update Your Local Interaset Mirror Repository

When you establish the location for your new Interaset 5.9.2 mirror repository, you must update your local offline Interaset repository with the URLs for the new Interaset 5.9.2 mirror location.

For reference, the following table lists the Interaset components and the mirror URLs for both the 5.7.x and 5.9.2 versions.

**Tips:**

- In the table below, **\$EL_VER** represents the major version of Enterprise Linux. Wherever it appears in a path, substitute this value with the appropriate version number.
- InterSet 5.9.2 supports only the CentOS 7.6 and Red Hat Enterprise 7.6 operating systems.
- When you run the **update_repos.sh** script to update your repository mirror location, the MySQL repository will be disabled. If your organization requires that the MySQL repository remain enabled, do the following:

- Run the following command to identify the status of your MySQL repository and sub-repositories:

```
yum repolist all | grep mysql
```

- For each disabled MySQL repository or sub-repository, run the following command:

```
sudo yum-config-manager --enable <repository_name>
```

Component	5.7.x Path	5.9.2 Path
deploy.sh	<mirror>/5.7.x/interaset/	<mirror>/5.9.2/interaset/
ambari	<mirror>/5.7.x/AMBAR/AMBAR-2.6.1.3/centos\$EL_VER/2.6.1.3-3/	<mirror>/5.9.2/AMBAR/AMBAR-2.7.3.0/centos7/2.7.3.0-139//
datastax	<mirror>/5.7.x/datastax/	<mirror>/5.9.2/datastax/
elasticsearch	<mirror>/5.7.x/elasticsearch/6.3.1/	<mirror>/5.9.2/elasticsearch/6.8.1
hdp	<mirror>/5.7.x/HDP/centos\$EL_VER/2.x/updates/2.6.4.0-91/	<mirror>/5.9.2/HDP/centos7/2.x/updates/2.6.4.0-91/
hdp-utils	<mirror>/5.7.x/HDP-UTILS/HDP-UTILS-1.1.0.22/repos/centos\$EL_VER/	<mirror>/5.9.2/HDP-UTILS/HDP-UTILS-1.1.0.22/repos/centos7/
hdp-gpl	<mirror>/5.7.x/HDP-GPL/centos\$EL_VER/2.6.4.0-91/	Not required
interasetepel	<mirror>/5.7.x/EPEL/el\$EL_VER/	<mirror>/5.9.2/EPEL/el7/
nginx	<mirror>/5.7.x/nginx/\$EL_VER/	<mirror>/5.9.2/nginx/7/

To update your local, offline InterSet repositories with the new 5.9.2 mirror URLs, do the following:

1. On the Monitoring node in your InterSet cluster, navigate to the **/tmp/interaset_installer/scripts/update_repos/** directory.
2. Edit the **repo_config** file to populate the **ALL_HOSTS** and **INTERSET_REPO** variables, following the examples in the file comments.
3. If your InterSet cluster uses TLS or a port other than the standard **:8080** for Ambari, do the following:

- As the **interaset** user, open the **update_repos.sh** file for editing.
- On line 28 and line 44, change the following line:

```
result=$(curl -sH "X-Requested-By: ambari" -X GET -u $ambari_login  
"http://$ambari:8080/api/v1/clusters/interaset/stack_versions/" | tr -d ' \n\t')
```

to

```
result=$(curl -skH "X-Requested-By: ambari" -X GET -u $ambari_login "https://$ambari:<configured_  
port>/api/v1/clusters/interaset/stack_versions/" | tr -d ' \n\t')
```

substituting your Ambari port number for **<configured_port>**.



Important: Ensure that you make the above change on both lines 28 and 44.

- Save and close the **update_repos.sh** file.
4. As the **interaset** user, run the following script to update the repositories:

```
./update_repos.sh
```

This script validates your inputs, and then invokes subsequent scripts that check for the presence of InterSet repositories, validate the operating system, and update the repositories with any new mirror URLs. If repositories do not exist, they are created.

When the repositories have all been updated, you will see the following response:

```
COMPLETED ON ALL NODES
```

```
INTERSET CLUSTER REPO DOT FILES ALL UP TO DATE
```

Prepare to Upgrade

Before you proceed to upgrade InterSet to version 5.9.2, ensure that all InterSet services are running.

Upgrade Interset

This section provides information to upgrade your Interset cluster from version 5.7.x to version 5.9.2. Upgrading Interset to version 5.9.2 includes:

- [stopping the Flume service](#)
- [upgrading the **interset** user](#)
- [installing and configuring the PostgreSQL external database](#)
- [upgrading Elasticsearch to version 6.3.1](#)
- [installing the New Interset Exports Service](#)
- [upgrading Reporting](#)
- [upgrading Interset Analytics](#)
- [upgrading the EDR Server](#)
- [upgrading Data Ingest](#)
- [upgrading Elasticsearch index data, mappings, and templates](#)
- [upgrading Workflow](#)
- [restarting the Ambari services](#)

After the Interset 5.9.2 upgrade is complete and working as expected, you can then decide whether to [install and configure the 5.9.2 schema registry](#).

Stop the Flume Service

Before you begin to upgrade your Interset cluster, you must stop the active data ingest and back up the Flume configs.

1. In a Web browser, navigate to **http://<ambari_fqdn>:8080**, and log in to Ambari as an administrator.
2. In the Ambari header, select **Services** and then, from the dropdown list, select **Flume**.
3. Click **Service Actions** and then, in the dropdown list, click **Stop** to stop the Flume service.
4. On the **Configs** tab, click the **Group** dropdown list, and then select a Interset Flume config group.
5. In the **flume.conf** text box, select the entire text and copy it to a text editor to create a backup of the config.
6. After you save the config backup text file, select the entire text in the **flume.conf** text box, and press **Delete** to remove the config.
7. Repeat Steps 4 through 6 for each of your Interset Flume config groups.
8. In the upper right of the Flume page, click **Service Actions** and then, in the dropdown list, click **Start** to restart the Flume service.

Upgrade the interset User

1. On each node in the InterSet cluster, run the following command to upgrade the **interset** user:

```
sudo yum upgrade -y interset-user
```

Install and Configure the PostgreSQL External Database

InterSet 5.9.2 requires the use of an external PostgreSQL database. This section provides the instructions to upgrade your existing Ambari-embedded PostgreSQL 9.2 database to an external PostgreSQL 10.4 database.

1. On the Ambari node in your cluster, navigate to `/tmp/interset_installer/etc` and open the new 5.7.0 version of your InterSet **config** file.
 2. In the **config** file, edit the following settings:
- Verify that the **INTERSET_VERSION**, **THEME**, AND **INTERSET_REPO** settings are configured as follows:

```
INTERSET_VERSION="5.9.1"
THEME="interset"
INTERSET_REPO="http://<mirror_fqdn>"
```

For an offline installation, `<mirror_fqdn>` is the server where you installed the offline mirror repository.

For an online installation, `<mirror_fqdn>` is the URL of the InterSet repository (`http://repo.interset.com/5.9.2`), and must include your access credentials. For example,

```
INTERSET_REPO="http://<username>:<password>@repo.interset.com/5.9.2"
```

- Enter the appropriate Hadoop monitoring system for the cluster, **HDP** or **CDH**. For example,

```
HADOOP_ENV="HDP"
```

- Set the Hadoop security protocol to **KERBEROS**:

```
# HADOOP_SECURITY should be NONE or KERBEROS
HADOOP_SECURITY=KERBEROS
```

- For each role, enter the fully-qualified domain name of the server(s) that you want to allocate to that role.



Tip: For roles that have more than one physical node, enter the multiple fully-qualified domain names as a space-separated list.

```
# 1 value only
MONITORING="monitoring.interaset.com"
REPORTING="reporting.interaset.com"
POSTGRES="postgres.interaset.com"
PERFMON="perfmon.interaset.com"

# between 1 and many (space-separated) values
COMPUTE="compute.interaset.com"
MASTER="master.interaset.com"
SEARCH="search.interaset.com"
STREAM="stream.interaset.com"
NIFI="nifi.interaset.com"
```

**Notes:**

- **PERFMON** is optional and can be left blank.

3. Save and close the **config** file.
4. Navigate to **/tmp/interaset_installer** and then run the following command to launch the menu to upgrade the PostgreSQL database.

```
./upgrade.sh
```

5. In the Interaset 5.7.0 menu, enter **1** to run the PostgreSQL database upgrade.

```

      88888888      888      888      888888888      8888888888      d8888b.
      888      888      888      888      888      d88P  d88P  Y88b
      888      888      888      888      888      d88P  888  888
      888 88888b, 8888888  d88b, 8888888  d8888b, d88b, 8888888 8888888b, d88P 888 888
      888 888 "88b 888  d8P  Y8b 888P" 88K  d8P  Y8b 888      "Y88b 888888888 888 888
      888 888 888 888 888888888 888      "Y8888b, 888888888 888      888  d88P 888 888
      888 888 888 Y88b,  Y8b, 888      X88 Y8b,  Y88b,  Y88b d88P d8b d88P  d8b Y88b d88P
      88888888 888 888 "Y888 "Y8888 888 88888P' "Y8888 "Y888 "Y8888P" Y8P d88P  Y8P "Y8888P"

[ 1] Update Configs, Upgrade Database
[ 4] Schema Registry Setup
[ 6] Stream node(s) Installation
[ c] Re-transfer config file to all servers. Use this if server names have been updated in etc/config!
[ l] Copy node log files to current host
[ q] Exit the installer
```

6. When the upgrade begins and you are prompted to enter your existing password for Ambari, type the password that you normally enter to log on to Ambari as an administrator.

```
Postgres upgrade starting.
Please enter your existing password for ambari user - admin: [ ]
```

7. When prompted, enter a new password for the new PostgreSQL 10.4 database user.

```
Please enter new password for DB user - Postgres: [ ]
```

This is a new password for the PostgreSQL system user, and not the same user as the Ambari administrator user.

8. When the PostgreSQL database upgrade is complete and you are prompted to confirm that SELinux is set to 'permissive' and temporarily disabled, type **y** to continue.
9. When prompted to customize the user account for the Ambari-server daemon, [y/n] (n), type **n**.
10. When prompted to change the Oracle JDK, type **n**.
11. When prompted to enable Ambari-server to download and install GPL licensed LZ0 packages, type **n**.
12. When prompted to enter advanced database configuration, type **y**.

13. When presented with the options for configuring the external PostgreSQL database, enter **4** for **PostgreSQL**.
14. When prompted, enter the information for your existing 9.2 embedded PostgreSQL database. This is to migrate the Ambari user and its associated credentials from the PostgreSQL 9.2 database to the new external PostgreSQL database.

```

Hostname (localhost):
Port (5432):
Database name (ambari):
Postgres schema (ambari):
Username (ambari):
Enter Database Password (ambari):

```

Enter the fully-qualified domain **Hostname**, **Port**, **Database name**, **Postgres schema**, **Username**, and **Database Password**.



Important: When prompted to run a script to create your Ambari database schema, **do not run the script at this time**.

This is because you will want to re-use your newly-configured PostgreSQL database, now migrated from the embedded 9.2 database in Ambari.

15. Open a new console for the Ambari node in your cluster and, at the command prompt, run the following script to ensure that the Ambari server can connect to the new PostgreSQL 10 external database, and to copy the database driver .jar file.

```
sudo ambari-server setup --jdbc-db=postgres --jdbc-driver=/usr/share/java/postgresql-42.2.2.jar
```

16. When the Ambari Server setup is complete, close this second Ambari node console.
17. In your original Ambari node console, type **y** to proceed with configuring the remote database connection properties.

The Ambari services will now be restarted. To monitor the progress, log in to Ambari at **https://<ambari_node_fqdn>:8080** and then click the **Current Operations** button. All the service restarts are listed, along with their status.

18. In the Ambari node console, when prompted, enter the password for the new PostgreSQL 10.4 database user that you created in Step 6 above.
19. Enter, and then confirm, a new password for the new PostgreSQL registry database that will now be created.



Notes:

- If you plan to use the new schema registry introduced with Interset5.9.2, you will install and configure the schema registry and then update the dependant components only after you complete the upgrade.
- The post-upgrade instructions to implement the schema registry feature are detailed in "[Install and Configure the Schema Registry](#)" on page 32.

Upgrade Elasticsearch to Version 6.8

Upgrading Elasticsearch to version 6.8.1 includes the following tasks:

- [removing the Elasticsearch plug-ins](#)
- [upgrading Elasticsearch to version 6.8.1](#)
- [upgrading Kibana to version 6.8.1](#)

Remove the Elasticsearch Plug-ins

Because the current Elasticsearch plug-ins will not be compatible with the new Elasticsearch version 6.8.1, you must remove them. Following the upgrade, you can install their correct versions.

1. On each Search node, and on the Reporting node in your Interset cluster, run the following command to identify all the installed Elasticsearch plug-ins:

```
sudo /usr/share/elasticsearch/bin/elasticsearch-plugin list
```

If a list of plug-ins is returned, copy this list to a text file. You will re-install the plug-ins following the Elasticsearch upgrade.



Tip: It is possible that no plug-ins will be returned. In this case, skip Step 2 and proceed directly to the next section.

2. For each plug-in returned in Step 1, run the following command to remove it:

```
sudo /usr/share/elasticsearch/bin/elasticsearch-plugin remove <plug-in_Name>
```

Upgrade Elasticsearch to Version 6.8.1

Execute the following steps on each Search node, and on the Reporting node in your Interset cluster.

1. As the **interset** user, SSH to a Search node in your Interset cluster.

```
ssh interset@<search_node_fqdn>
```

2. Run the following command to disable shard allocation.

```
curl -X PUT "<reporting_node_fqdn>:9200/_cluster/settings" -H 'Content-Type: application/json' -d '{
  "persistent": {
    "cluster.routing.allocation.enable": "primaries"
  }
}'
```

3. Run the following command to stop all non-essential Elasticsearch indexing and perform a synchronized flush.

```
curl -X POST "<search_node_fqdn>:9200/_flush/synced"
```

4. On the Reporting node in your Interset cluster, run the following commands to stop Reporting and Kibana.

```
sudo systemctl stop reporting
```

```
sudo systemctl stop kibana
```

5. On the Reporting node in your Intersect cluster, run the following command to stop Elasticsearch.

```
sudo systemctl stop elasticsearch
```

6. Still on the Reporting node, run the following yum command to upgrade Elasticsearch.

```
sudo yum upgrade -y elasticsearch-6.8.1
```

7. Repeat Steps 5 and 6 above on each node in your Intersect cluster where Elasticsearch is installed to stop and upgrade Elasticsearch.



Tip: This normally includes all Search and Reporting nodes.

8. On each node where Elasticsearch is installed, do the following:

- Open the **elasticsearch.yml** file, and then add the parameters and values listed below.

```
xpack.monitoring.enabled: false
```

```
xpack.security.enabled: false
```

```
xpack.watcher.enabled: false
```

```
xpack.ml.enabled: false
```

If these fields are already present, edit their values to match what is listed above unless your particular setup requires you to use these features.

- For each plug-in that you removed earlier, run the following command to re-install the plug-in:

```
sudo /usr/share/elasticsearch/bin/elasticsearch-plugin install file:///path/to/plugin.zip
```

- Start Elasticsearch using the following command:

```
sudo systemctl start elasticsearch.service
```

9. On any node where Elasticsearch is installed, run the following command to re-enable the shard allocation.

```
curl -X PUT "<reporting_node_fqdn>:9200/_cluster/settings" -H 'Content-Type: application/json' -d'
{
  "persistent": {
    "cluster.routing.allocation.enable": null
  }
}'
```

10. Wait until the cluster is fully recovered and displaying a green state.



Tip: If your Intersect cluster has only one Search node, the **'status'** will always be **"yellow"**. In this situation, you can proceed directly to the next section.

You can monitor the cluster recovery using the following cURL requests:

- To retrieve the cluster state (color), run the following command:

```
curl -X GET "<reporting_node_fqdn>:9200/_cluster/health?pretty"
```

The response will be similar to the following:

```
{
  "cluster_name": "Interset",
  "status": "green",
  "number_of_nodes": 3,
  "active_primary_shards": 3,
  "active_shards": 2843,
  "relocating_shards": 0,
  "initializing_shards": 2,
  "unassigned_shards": 1220,
  "active_shards_percent_as_number": 69.93849938499385
}
```

- For a more detailed view of the cluster recovery, run the following command:

```
curl -X GET "<reporting_node_fqdn>:9200/_cat/recovery"
```

Upgrade Kibana to Version 6.8.1

1. As the **interset** user, ssh to the Reporting node in your Interset cluster.

```
ssh interset@<reporting_node_fqdn>
```

2. Upgrade Kibana to version 6.8.1 using the following command.

```
sudo yum upgrade -y kibana-6.8.1
```

3. Run the following command to reload the Kibana daemon.

```
sudo systemctl daemon-reload
```

4. Start Kibana, using the following command:

```
sudo systemctl start kibana
```

Upgrade Interset Reporting to Version 5.9.2

To upgrade the Interset Reporting component to version 5.9.2, do the following:

1. As the **interset** user, ssh to the Reporting node in your Interset cluster.

```
ssh interset@<reporting_node_fqdn>
```

2. If you did not previously stop the Reporting service, run the following commands to stop Reporting.

```
sudo systemctl stop reporting
```

3. Run the following command to upgrade Interset Reporting.

```
sudo yum upgrade -y interset-reporting
```

4. Update the Reporting configuration by running the following commands to update the **investigator.yml** file.

```
sudo sed -i 's| domain:| #domain:|' /opt/interset/etc/investigator.yml
```

```
sudo sed -i 's|url: jdbc:phoenix.*|url: jdbc:phoenix:thin:url=http://<master_node_fqdn>:8765;serialization=PROTOBUF|' /opt/interaset/etc/investigator.yml
```

```
sudo sed -i 's|esIPAddr:.*|esIPAddr: <reporting_node_fqdn>|' /opt/interaset/etc/investigator.yml
```

```
sudo sed -i 's|esClusterName:.*|esClusterName: "interaset"|' /opt/interaset/etc/investigator.yml
```

5. On the Reporting node, run the following commands to upgrade the Nginx configuration:

```
sudo cp -vp /opt/interaset/etc/nginx.conf /opt/interaset/etc/nginx.conf.backup.`date +%s`
```

```
sudo mv /opt/interaset/etc/nginx.conf.rpmnew /opt/interaset/etc/nginx.conf
```

```
sudo sed -i 's|server_name server.company.com;|server_name <reporting_node_fqdn>;|' /opt/interaset/etc/nginx.conf
```

```
sudo /opt/interaset/bin/sysprep/scripts/configure_nginx_reporting
```

6. If you customized nginx to use a local TLS certificate, do the following:

- On the Reporting node, navigate to **/opt/interaset/etc**, and then open the **nginx.conf** configuration file.
- Locate the **ssl_certificate** and **ssl_certificate_key** parameters.
- Enter the path for each of these parameters.

For example,

```
ssl_certificate /etc/nginx/ssl/interaset.crt;
ssl_certificate_key /etc/nginx/ssl/interaset.key;
```



Tip: You can obtain the paths for your local TLS certificate and certificate key from the **nginx.conf.backup** file created in the previous step.

- Save and close the **nginx.conf** file.

7. Run the following commands to restart all Reporting services:

```
sudo systemctl restart nginx
```

```
sudo systemctl restart reporting
```

8. Wait for Reporting to start.



Tip: You can run the following command to monitor the status of the Reporting component.

```
(sudo tail -f -n0 '/opt/interaset/log/reporting.log' & ) | grep -q
'org.eclipse.jetty.server.Server: Started' && echo | sudo tee -a
/opt/interaset/log/reporting.log
```

If no response is returned, it is because the Reporting component has already started.

Upgrade the InterSet Exports Service

Complete the following steps to upgrade the Exports service, which is the PDF report and data extraction service:

1. As the **interSet** user, ssh to the Reporting node in your InterSet cluster.

```
ssh interSet@<reporting_node_fqdn>
```

2. Run the following commands to upgrade the InterSet Exports service.

```
sudo yum install -y nodejs-8.11.3 --disablerepo="*" --enablerepo="interSetepel"
```

```
sudo yum install -y interSet-exports
```

Upgrade the InterSet Analytics Component to Version 5.9.2

Upgrading Analytics brings the InterSet Analytics component to version 5.9.2 and updates the Analytics (**interSet.conf**) configuration file. Execute the following steps on each Master node in your InterSet Analytics cluster.

1. As the **interSet** user, ssh to the Master node where Analytics is installed in your cluster.

```
ssh interSet@<master_node_fqdn>
```

2. Run the following commands to upgrade Analytics:

```
sudo yum upgrade -y interSet-analytics
```

```
cd /opt/interSet/analytics/bin; ./sql.sh --dbServer <zookeeper_fqdn>:2181:/hbase-unsecure --action migrate
```

3. Navigate to the **/opt/interSet/analytics/conf/** directory, and locate the **interSet.conf.rpmnew** file.



Notes:

- The **interSet.conf.rpmnew** file will be present in the **/opt/interSet/analytics/conf/** directory only in situations where the installed 5.7.x **interSet.conf** file was edited following the initial InterSet 5.7.x installation.
- If your **/opt/interSet/analytics/conf/** directory does not contain an **interSet.conf.rpmnew** file, proceed directly to Step 6 below.
- If your InterSet cluster has multiple tenants configured and, therefore, more than one **interSet.conf** file (eg. **interSet.conf.1**, **interSet.conf.2**, and so on), create as many copies of the **interSet.conf.rpmnew** file as required so that you have one for each configured tenant in your upgraded 5.9.2 cluster. Having sufficient copies of the **interSet.conf.rpmnew** file will allow you to create 5.9.2 versions of each **interSet.conf** file in your cluster.

4. If the **interaset.conf.rpmnew** file is present, do the following to update your Analytics configuration (**interaset.conf**) file(s):

- Open the **interaset.conf** and **interaset.conf.rpmnew** file(s).
- Copy the modified values from your 5.7.x **interaset.conf** file(s) to the **interaset.conf.rpmnew** file(s), and then save the modified **interaset.conf.rpmnew** file(s).



Tip: Two settings in the 5.7.x **interaset.conf** file are no longer used and do not appear in the 5.9.2 **interaset.conf.rpmnew** file. These settings are **storiesFolder** and **entityStoriesFolder**, located in the **EntityScores output** section.

- Run the following commands, using your Analytics 5.7.x configuration (**interaset.conf**) and Analytics 5.7.x (**interaset.conf.rpmnew**) file names, to back up your 5.7.x **interaset.conf** file and rename the updated 5.9.2 **interaset.conf.rpmnew** file to **interaset.conf**:

```
cp -vp /opt/interaset/analytics/conf/<interaset.conf>
/opt/interaset/analytics/conf/<interaset.conf>.backup.`date +%s`
```

```
mv -v /opt/interaset/analytics/conf/<interaset.conf>.rpmnew /opt/interaset/analytics/conf/<interaset.conf>
```

- When prompted to confirm that you want to overwrite the existing **interaset.conf** file with the modified **interaset.conf.rpmnew** file, type **y** and press **Enter**.



Tips:

- If you are running InterSet Analytics against WDC data, the **esHost** variable must be populated and must match the value in the previous version of the **interaset.conf** file.
- The **zkPhoenix** value must be populated with the same value used in the previous version of the **interaset.conf** file. In addition, the **parallelism**, **numExecutors**, **executorMem**, **executorCores**, and **driverMem** values should match the values in the previous version of the **interaset.conf** file.

5. Repeat Steps 4 (b), (c), and (d) as many times as required, substituting your Analytics 5.7.x configuration (**interaset.conf**) and Analytics 5.9.2 (**interaset.conf.rpmnew**) file names, to create a new Analytics 5.9.2 **interaset.conf** file for each configured tenant in your InterSet cluster.
6. If the **interaset.conf.rpmnew** file is not present, do the following to update your Analytics configuration (**interaset.conf**) file(s):
- Create as many copies of the new 5.9.2 **interaset.conf** file as required so that you have one for each configured tenant in your upgraded 5.9.2 cluster.
 - Back up each of your remaining 5.7.x **interaset.conf** configuration files.
 - Copy the modified values from each of your backed-up 5.7.x configuration (**interaset.conf**) files to a new 5.9.2 **interaset.conf** file.
 - Save the new 5.9.2 **interaset.conf** configuration files, using your 5.7.x file names.

Upgrade the Intersect EDR Server to Version 5.9.2

1. As the **intersect** user, ssh to an EDR server node in your Intersect cluster.

```
ssh intersect@<edr_server_node_fqdn>
```

2. Run the following command to stop the EDR server.

```
sudo systemctl stop edr-server
```

3. Repeat the first two steps on each EDR server in your Intersect cluster.
4. As the **intersect** user, ssh to the primary EDR server in your cluster.



Tips:

- The primary EDR server is the EDR server node within a tenant where the **nodeID=0**.
- You will upgrade the primary EDR server node first.
- To identify which of the EDR server nodes is the primary EDR server in the cluster, run the following command.

```
$ grep com.intersect.endpoint.flux.nodeId /opt/intersect/flux/conf/system.conf  
com.intersect.endpoint.flux.nodeId=0
```

- If your cluster includes multiple tenants, you will have one primary EDR server node per tenant.

5. On each primary EDR server node(s) in the cluster, run the following command to upgrade the EDR server.

```
sudo yum upgrade intersect-edr-server -y
```



Tip: If you have multiple tenants in your Intersect cluster and therefore multiple primary EDR servers, upgrade each primary EDR server in the cluster before moving on to the others.

6. On the Reporting node, navigate to the **/opt/intersect/etc** directory and open the **investigator.yml** file.
7. In the **investigator.yml** file, locate the **flux-proxy** section, and then do the following:
 - Uncomment the lines as indicated in the code snippet below.
 - Enter the **<TID>** for the tenant for which the primary EDR server is configured.
The default tenant is tenant **0**.
 - Update the **hostname** parameter with the fully-qualified domain name of the primary EDR server node.
 - Leave the **enableSSL** parameter set to **true**.

- For the **username** parameter, enter the default username: **admin**
- For the **password** parameter, enter the default password: **4serv!**

```
flux-proxy:
0:
  hostname: <EDR-server_node_fqdn>
  enableSSL: true
  username: admin
  password: 4serv!
```

**Tips:**

- You need only complete the configuration information in the **investigator.yml** file for the primary EDR server node(s).
- If you have multiple tenants in your cluster, you will have a primary EDR server node for each tenant.

8. Save and close the **investigator.yml** file.

9. Start the EDR service on the primary EDR server node(s) using the following command:

```
sudo systemctl start edr-server
```

**Tips:**

- You can verify the EDR Server configuration(s) in the **flux.log** file(s). The **flux.log** file(s), are located in the **/var/log/intersect** directory on each EDR Server node.
- For information about managing the EDR Sensors in your environment, please see **Manage Your EDR Sensors** in the Intersect 5.9.2 *Administrator and User Guide*.

10. On the Reporting node in your Intersect cluster, run the following command to restart the Reporting service.

```
sudo systemctl restart reporting
```

To view the **Intersect Sensors** page in the Intersect user interface, log in to Intersect using your administrator credentials, and from the **Settings** dropdown menu, select **Intersect Sensors**.

11. After you are satisfied that the upgrade completed successfully on each primary EDR server node, on each additional EDR server node in the cluster, run the following commands to upgrade and then start the remaining EDR servers.

```
sudo yum upgrade intersect-edr-server -y
```

```
sudo systemctl start edr-server
```


Upgrade Interset Data Ingest to Version 5.9.2

After the Interset Analytics component is successfully upgraded, you can upgrade the Data Ingest components.

Upgrade Stream Nodes



Tips:

- The Analytics component must be upgraded before the Data Ingest components, as new ingest .jar files are upgraded by the Analytics rpm package.
- If prompted for a password, enter **interset**.

1. As the **interset** user, SSH to the Ambari node in your Interset cluster.

```
ssh interset@<ambari_node_fqdn>
```

2. Run the following command to create a directory named **interset_upgrade** under the **/tmp** directory:

```
mkdir -vp /tmp/interset_upgrade/
```

3. Copy the ingest .jar files from the Master node where Analytics is installed, using the following command:

```
scp interset@<master_node_fqdn>:/opt/interset/ingest/*.jar /tmp/interset_upgrade/
```

4. Run the following command for each Stream node in your Interset cluster:

```
scp -r /tmp/interset_upgrade/ interset@<stream_node_fqdn>:/tmp 2>&1
```

5. Run the following command:

```
sudo rm -rfv /tmp/interset_upgrade
```

6. As the **interset** user, SSH to a Stream node in your Interset cluster.

```
ssh interset@<stream_node_fqdn>
```

7. Run the following command to remove obsolete Flume .jar files:

```
sudo rm -rf /usr/hdp/current/flume-server/plugins.d/interset/lib/*
```

8. Run the following commands to copy new 5.9.2 Flume .jar files to the Interset Flume **lib** directory:

```
sudo chown -R interset:interset /usr/hdp/current/flume-server/plugins.d/interset/lib
```

```
mv /tmp/interset_upgrade/*.jar /usr/hdp/current/flume-server/plugins.d/interset/lib
```

9. Run the following commands to link the Phoenix .jar files:

```
sudo ln -vs /usr/hdp/current/phoenix-client/lib/phoenix-flume-*.jar /usr/hdp/current/flume-server/plugins.d/interset/lib/
```

```
sudo ln -vs /usr/hdp/current/phoenix-client/lib/phoenix-core-*.jar /usr/hdp/current/flume-server/plugins.d/interset/lib/
```

10. Repeat Steps 5 through 8 on each Stream node in your Interset cluster.

Upgrade NiFi to Version 1.7.1

Where the use of NiFi for data ingest was optional in previous releases, InterSet 5.9.2 now uses NiFi 1.7.1 for the data extraction and transformation portion of the data ingest process. As a result, you must upgrade NiFi to version 1.7.1.



Important: You must stop all existing NiFi services, and then upgrade each NiFi service to version 1.7.1 before you restart any of the NiFi services. You cannot upgrade the NiFi services one at a time; doing so will result in incompatible versions running concurrently in the cluster.

1. On all Stream nodes in the InterSet cluster, run the following command to stop the existing NiFi service.

```
sudo service nifi stop
```

2. Navigate to **`http://<mirror_fqdn>/5.9.2/nifi`**, and then run the following command to copy the **nifi-1.7.1-bin.tar.gz** archive to the **/opt/nifi** directory on each Stream node in your InterSet cluster.

```
sudo scp 5.7.0/nifi//nifi-1.7.1-bin.tar.gz interset@<stream_node_fqdn>:/opt/nifi
```

3. On each Stream node in the cluster, do the following:

- Navigate to **/opt/nifi**, and then run the following command to extract the **nifi-1.7.1-bin.tar.gz** archive.

```
sudo tar -xvzf nifi-1.7.1-bin.tar.gz
```

- Run the following command to change the ownership of the **/opt/nifi** directory and all sub-directories to **nifi:hadoop**.

```
sudo chown -R nifi:hadoop /opt/nifi/
```

- Run the following commands to unlink the current softlink and create a new softlink pointing to the new NiFi installation directory.

```
sudo unlink current
```

```
sudo ln -s nifi-1.7.1 current
```

- Navigate to **<previous_nifi_version_directory>/conf**, and for each of the following files, copy any customized settings to the same files in the new version 1.7.1 **/opt/nifi/conf** directory:
 - **nifi.properties**
 - **bootstrap.conf**
 - **logback.xml**

**Tips:**

- If no new properties were added in these files, you can copy them directly from the previous version's directory to the new version 1.7.1 directory. If you copy the **nifi.properties** file, ensure that you update the version number in the file.
- When you are configuring the paths to your existing repos, and the path to the **flow.xml.gz** file in the NiFi 1.7.1 **nifi.properties** file, ensure that all information is accurate and contains no typos.
- Ensure that, when you configure the **bootstrap.conf** file, you configure the same **Run As** user that was used in the previous version.

- In each Stream node **nifi.properties** file, do the following.

- Set the following property values:

```
nifi.web.http.port=<node_port>
nifi.cluster.is.node=true
nifi.cluster.node.address=<node_fqdn>
nifi.cluster.node.protocol.port=<node_protocol_port>
nifi.state.management.embedded.zookeeper.start=true
nifi.state.management.provider.cluster=<zk_provider>
nifi.state.management.embedded.zookeeper.properties=./conf/zookeeper.properties
nifi.zookeeper.connect.string=<comma-separated_list_of_host:port_ZooKeeper_connections>
```



Tip: The **nifi.zookeeper.connect.string** value would appear as, for example,

```
nifi.zookeeper.connect.string=zk_server1:2181,zk_server2:2181,zk_
server3:2183
```

- Add the following properties to ensure that the cluster coordinator and other nodes have the time required to select the correct data flow:

```
nifi.cluster.flow.election.max.wait.time=5 mins
nifi.cluster.flow.election.max.candidates=<number_of_nodes_with_NiFi_in_the_cluster>
```

4. Copy the custom NiFi archive (.nar) files to each Stream node in your cluster.

- Navigate to http://<mirror_fqdn>/5.9.2/interaset, and then run the following command to copy the **nifi-interaset-5.9.2.tar.gz** archive to a temporary directory on each Stream node in your InterSet cluster (typically **/tmp**).

```
sudo scp 5.9.2/interaset/nifi-interaset-5.9.2.tar.gz interaset@<stream_node_fqdn>:/tmp
```

- Extract the **nifi-interaset-5.9.2.tar.gz** archive.
- Copy the **nifi-interaset-nar-5.9.2-SNAPSHOT.nar** file to the **/opt/nifi/meta-data/customNars** directory on each Stream node(s).

5. On each Stream node in the Interset cluster, run the following command to start the new NiFi 1.7.1 service.

```
sudo service nifi start
```

**Tips:**

- Following the upgrade, you can review the **nifi-app.log** file for errors, as well as for any nodes that fail to join the upgraded cluster.
- If you previously configured NiFi to run as a service, navigate to the **/etc/init.d** directory, open the **nifi** file, and then edit the **NIFI_HOME** property as follows to ensure that any path or links for that service are updated to point to the newly-installed version executables:

```
NIFI_HOME=/opt/nifi/nifi-1.7.1
```

- You may need to set path to Java_Home in the **nifi-env.sh** file on each Stream node. To do this, navigate to the **/opt/nifi/nifi-1.7.1/bin** directory, and then open the **nifi-env.sh** file to edit the Java_Home property.

To obtain the path to Java_Home, type the following:

```
echo $JAVA_HOME
```

Update Flume Data Ingest Configuration Information

After the upgrade is complete, Reporting is running, and you can log in to Interset, you will need to update any Flume Data Ingest configuration information.

To do this, Micro Focus Interset Support will assist you to create new Interset 5.9.2 Data Ingest configurations for your Flume Configuration Group in Ambari. Assistance from Micro Focus Interset Support is necessary in this release due to significant updates and improvements in the Flume Data Ingest configuration files.

For the detailed changes implemented in the 5.9.2 Flume Data Ingest configurations for all data source types, please see ["Flume Configuration Changes" on page 48](#) in the *Flume Configuration Guide*.

Upgrade Interset Elasticsearch Index Data, Mappings, and Templates

After you upgrade Analytics, you must execute the following commands to upgrade your Elasticsearch index data, mappings, and templates.



Tip: You must have completed the Analytics upgrade before you can update these indexes.

1. As the **interaset** user, ssh into the Master node where Analytics is installed.

```
ssh interaset@<master_node_fqdn>
```

2. Run the following command to update the templates.

```
/opt/interaset/bin/elasticsearch/update_templates.sh --esHost <search_node_fqdn>
```

3. Run the following command to update the all indices to use the latest mappings.

```
/opt/interaset/bin/elasticsearch/reindex.sh -esHost <search_node_fqdn> --belowVersion 6.8.1 --ignorePipeline
```

4. As the **spark** user, run the following command to upgrade your data.

```
sudo -u spark bash -c "cd /opt/interaset/analytics/bin; ./upgrade_5.7_to_5.9.2.sh
```

5. As the **interaset** user, ssh into the Reporting node.

```
ssh interaset@<reporting_node_fqdn>
```

6. Run the following command to restart reporting and verify that the latest mappings were correctly applied.

```
sudo systemctl restart reporting
```

Upgrade Workflow to Version 5.9.2

Before you upgrade Workflow, you must suspend all active Workflows by killing the active Storm topologies. Now that the Analytics and Reporting components have been upgraded, you can proceed to upgrade and redeploy the Workflow component.



Tip: Analytics must be upgraded before Workflow, as new Workflow .jar files are upgraded by the Analytics rpm.

Suspend the Active Workflows

1. On the Master node where Analytics is installed, run the following command to stop the Workflow component:

```
/opt/interaset/rules/bin/workflow.sh --kill /opt/interaset/rules/conf/<rules.conf>
```



Tip: If your Workflow **rules.conf** configuration file is not named rules.conf, substitute your Workflow configuration file name where **rules.conf** appears in the above command. If you have multiple Workflow configuration files, run this command for each of them.

Upgrade Workflow

1. As the **interaset** user, ssh to the Master node where Analytics is installed in your Interset cluster.

```
ssh interaset@<master_node_fqdn>
```

2. Navigate to the `/opt/interset/rules/conf/` directory, and locate the existing **rules.conf** files for your cluster.



Tip: If you configured Workflow for each tenant in your cluster, you will have multiple **rules.conf** configuration files.

3. Run the following command to create a backup of each of your existing 5.7.x **rules.conf** files.

```
sudo cp -vp /opt/interset/rules/conf/<rules.conf> /opt/interset/rules/conf/<rules.conf>.backup.`date +%s`
```

4. Run the following command, using the 5.9.2 **rules.conf.rpmnew** file to create new **rules.conf** configuration files.

```
sudo cp -vp /opt/interset/rules/conf/rules.conf.rpmnew /opt/interset/rules/conf/<rules.conf>
```



Tips

- If your Interset cluster has Workflow configured for multiple tenants, create as many copies of the **rules.conf.rpmnew** file as required so that you have one for each configured tenant in your upgraded 5.9.2 cluster. Having sufficient copies of the **rules.conf.rpmnew** file will allow you to create 5.9.2 versions of each **rules.conf** file in your cluster.
- Save the copies of the **rules.conf.rpmnew** file with the previous 5.7.x **rules.conf** filenames.

5. After you have created the new 5.9.2 **rules.conf** files for your cluster, do the following:
 - Open a 5.7.x **rules.conf.backup** file, and the corresponding 5.9.2 **rules.conf** file.
 - Copy the configured values from your configured **rules.conf.backup** file to the **rules.conf** file.
 - In the **Service Connections** section of the new 5.9.2 **rules.conf** file, comment out the **ConfluentRegistry** parameter.
This parameter is required only when you configure Interset to use the schema registry. For more information about the schema registry, please contact Micro Focus Interset Support.
 - In the **Event Processing** section of the new 5.9.2 **rules.conf** file, update the **non-Schema Registry path** parameters by:

```
#TopicEvents.<kafkaTopicName1> = <archType1>
```

- removing the hash (#) symbol to enable this setting,
- creating as many **#TopicEvents.<kafkaTopicName1>** entries as necessary for your tenant and Workflow data source types,

- entering the appropriate **<kafkaTopicName>** for the data source type, and
- entering the relevant **<archType>** value from the table below.

Data Source Type	<archType>
Access	access
Active Directory	ad, active_directory, activedirectory
Email	email
Endpoint	endpoint, wdc
Expense	expense
Alert	interaset-standard, interasetstandard, alert, uaf
Linux AuditD	auditd, linuxauditd
NetFlow	netflow
Printer	printer
Repository	repo, repository
Sensor	sensor
VPN	vpn
Web Proxy	proxy, webproxy, webproxy
Windows Printer	winprinter, windows_printer, windowsprinter

- Save the modified **rules.conf** file.
 - Repeat these steps for each configured **rules.conf** file in your cluster.
6. When all required Workflow **rules.conf** configuration files are updated and saved, run the following command to deploy each of your updated 5.9.2 Workflow configurations.

```
/opt/interaset/rules/bin/workflow.sh --deploy /opt/interaset/rules/conf/<rules.conf>
```

Restart Ambari Services

When the InterSet upgrade is complete, some Ambari services may need to be restarted. To restart these services:

1. In a Web browser, navigate to **http://<ambari_fqdn>:8080**, and log in to Ambari using administrator credentials.
2. In the left pane, click the **Actions** dropdown list, and then from the dropdown menu select **Restart All Required**.
3. When prompted, click **Confirm Restart All**.

Your InterSet 5.9.2 upgrade is now complete.

Appendix A: Install and Configure the Schema Registry

This Appendix provides instructions for installing and configuring the schema registry in your upgraded Interset 5.9.2 cluster.

Install the Schema Registry

1. On the monitoring node, navigate to `/tmp/interset_installer` and then run the following command to launch the menu to install the schema registry.

```
./upgrade.sh
```

2. In the Interset 5.7.0 menu, enter **2** to run the schema registry installation.

```

#####      888      888      8888888888      8888888888      .d8888b.
888      888      888      888      888      888      d88P  d88P  Y88b
888      888      888      888      888      888      d88P  888  888
888 88888b. 888888 .d88b. 8888888 .d8888b. d88b. 888888 8888888b. d88P  888  888
888 888 "88b 888 d8P Y8b 888P" 88K d8P Y8b 888 "Y88b 88888888 888 888
888 888 888 888 888888888 888 "Y8888b. 88888888 888      888 d88P  888 888
888 888 888 Y88b. Y8b. 888      X88 Y8b. Y88b. Y88b d88P d8b d88P d8b Y88b d88P
8888888 888 888 "Y888 "Y888 888 88888P' "Y8888 "Y888 "Y8888P" Y8P d88P Y8P "Y8888P"

[ 1] Update Configs, Upgrade Database
[ 4] Schema Registry Setup
[ 6] Stream node(s) Installation
[ c] Re-transfer config file to all servers. Use this if server names have been updated in etc/config!
[ l] Copy node log files to current host
[ q] Exit the installer
  
```

3. When prompted to upload Interset schemas to the schema registry, type **y** and then press Enter. After the schema upload to the Master node in the Interset cluster is complete, the expected result is:


```
Registry Setup Is Complete!
```


Configure the ConfluentSchemaRegistry Controller Service

1. If you haven't already, open a browser and navigate to `<nifi_node_fqdn>:8085/nifi/`.
2. At the top-most level of the NiFi user interface, right-click on the canvas and select **Configure**.
3. In the **NiFi Flow Configuration** dialog, select the **Controller Services** tab.
4. Scroll through the list to find an item with **Type** field value **ConfluentSchemaRegistry <version>**.

For example, in the **Interset Marketplace** template, locate the **ConfluentSchemaRegistry - Interset Avro Schemas** controller service.

If you don't find a **ConfluentSchemaRegistry** entry:

- a. Click **Create a new controller service** .
- b. Start typing "confluent" into the **Filter** text box, until the **ConfluentSchemaRegistry** controller service appears in the list.

- c. Select the **ConfluentSchemaRegistry** controller service and click **Add**.
- d. If required, rename the controller service you just added by clicking **Configure**, selecting the **Settings** tab, entering a new name, and clicking **Apply**.
5. In the right-most column of the **ConfluentSchemaRegistry** controller service, click **Configure** .
6. Select the **Properties** tab, and in the **Value** text box next to the **Schema Registry URLs** property, type the fully-qualified domain name (FQDN) and port for each of your schema registry data-bases (separated by commas).



The default location of the schema registry is **https://<master_node_fqdn>:9190**

7. Click **Apply**.

For more information, see the NiFi documentation for the [ConfluentSchemaRegistry](#) controller service.



You can access the schema registry's user interface (for example, to make modifications to a schema) at the following URL:

- **http://<master_node_fqdn>:9190/ui/**

Configure Readers to Use the ConfluentSchemaRegistry Controller Service

Any processors in your NiFi flow that read or write record-based information must be set up to retrieve the schema from the schema registry. There are many built-in controller services in NiFi that perform record reading or writing. Some examples are the following (where *<version>* is the NiFi version and build, such as 1.7.1.0.1.0.0-123):



- AvroReader *<version>*
- CSVReader *<version>*
- AvroRecordSetWriter *<version>*
- GrokReader *<version>*
- JsonRecordSetWriter *<version>*

To configure these readers to use the schema registry, do the following:

1. From either the NiFi canvas or from the appropriate processor properties, click **Configure**.
2. Select the **Controller Services** tab.
3. Locate the required reader (for example, in the **Interaset Marketplace** template, locate the **CSVReader - By Schema Name** controller service).

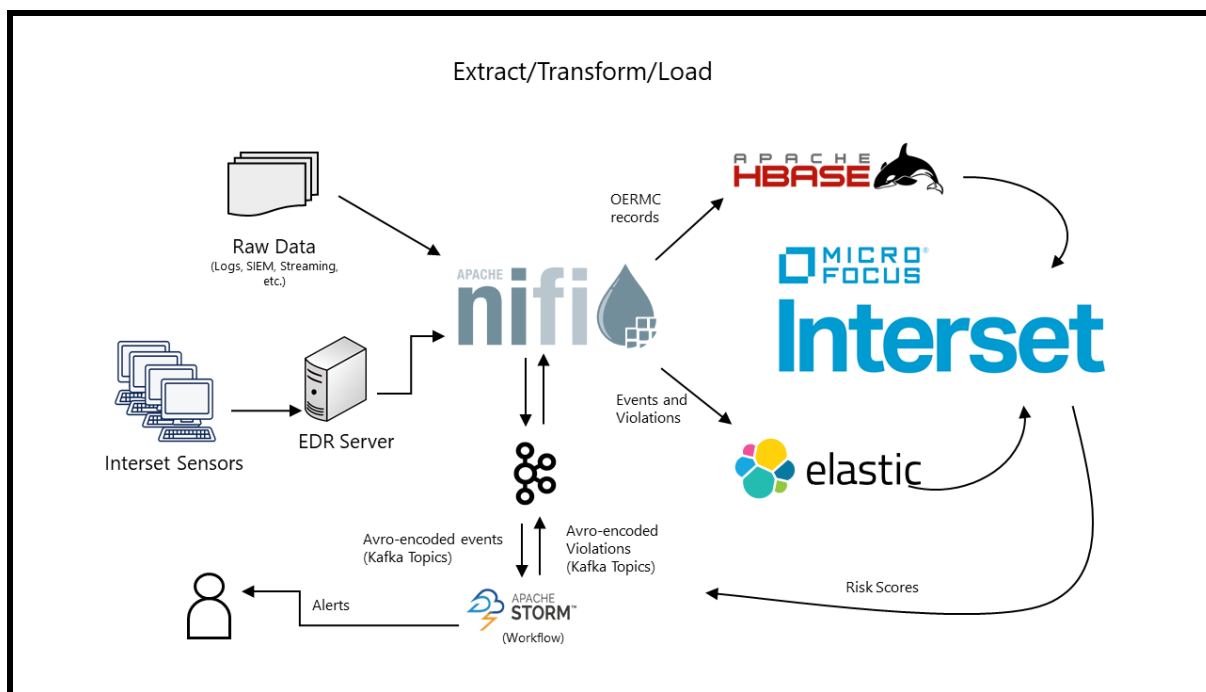


In the **Interaset Marketplace** template, you can find all the relevant processors by looking at the **Settings** tab of the **Configure Controller Service** dialog for the **ConfluentSchemaRegistry - Interaset Avro Schemas** controller service.

4. Click **Configure** .
5. Set the **Schema Registry** field to the name of the **ConfluentSchemaRegistry** controller service you created above.
If you are using the **Interaset Marketplace** template, enter **ConfluentSchemaRegistry - Interaset Avro Schemas**.
6. Set the **Schema Access Strategy** property to **Use Schema Name Property**.
7. Set the **Schema Name** field to the name of the required schema, or to a FlowFile attribute containing the name.
For example, the **Interaset Marketplace** template contains an **UpdateAttribute** processor for each data type that sets the correct Avro schema name for that data type.
8. Click **Apply**.
9. In the right-most column of the **ConfluentSchemaRegistry** controller service, click **Enable**  to activate the controller service.
10. In the **Enable Controller Service** dialog, click **Enable**, and then click **Close**.

Configure a New Data Source

With Interaset 5.9.2, the data ingest process uses Apache NiFi to extract, transform, and load data. The following diagram illustrates the data flow in Interaset 5.9.2.



When you configure data ingest in Interset 5.9.2, you:

- configure an SSL controller service, which is required by some of the processors in the the **Interset Marketplace** template
- create a process group to separate the new data flow from the **Interset Marketplace** template
- specify the location of the Interset schema registry
- enable the controller services
- configure the processors in the NiFi flow that are responsible for extracting, transforming, and loading the data
- start the data flow

For this data ingest, we use the sample AD dataset (**ad_sample.csv.gz**) provided by Interset.

If you don't already have the sample authentication dataset, you can obtain it from the Interset online customer repository. If you require access credentials, please contact Micro Focus Interset Support at interset.support@microfocus.com.



Ensure that you copy the sample data to a directory accessible to NiFi. Make a note of this location; you will need it later.


Configure an SSL Context Service

Some of the processors in the **Interset Marketplace** template require an SSL context service to be specified. To configure an SSL context service, do the following:

1. Open a browser and navigate to `<nifi_node_fqdn>:8085/nifi/`.

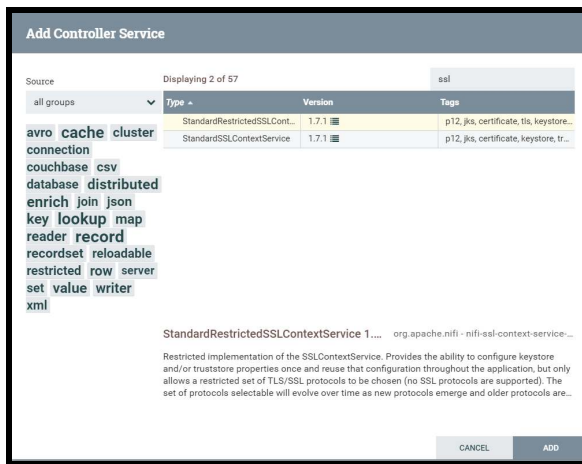



In a secure cluster, the NiFi URL is `https://<nifi_node_fqdn>:8085/nifi/`

2. Right-click the canvas, then click **Configure**.
3. In the **NiFi Flow Configuration** dialog, select the **Controller Services** tab.
4. Click **Create a New Controller Service** . The **Add Controller Service** dialog opens.
5. In the **Filter** text box (above the **Tags** column), type `ssl` to filter the list of services.

You should see the following services:

- `StandardRestrictedSSLContextService`
- `SSLContextService`



6. Select `StandardRestrictedSSLContextService`, and then click **Add**.
The service is added to the list of Controller Services in the **NiFi Flow Configuration** dialog.
7. In the right-most column of the Controller Services list, click **Configure** .
8. In the **Configure Controller Services** dialog, set the properties as follows:
 - **Keystore Filename:** The path to the Keystore file. The default location is `/etc/security/interaset/sslcert/localhost-keystore.jks`.
 - **Keystore Password:** The password for the Keystore.
 - **Key Password:** The password for the key. If this is not specified, but the Keystore Filename, Password, and Type are specified, then the Keystore Password will be assumed to be the same as the Key Password.
 - **Keystore Type:** JKS
 - **Truststore Filename:** The path to the trust store file. The default location is `/etc/security/interaset/sslcert/all-truststore.jks`.
 - **Truststore Password:** The password for the Truststore.

- **Truststore Type:** JKS
- **TLS Protocol:** TLS

9. Click **Apply**.

Create a Process Group



We strongly recommend that you create a new process group to build your data flow. This ensures that the processors in the **Interaset Marketplace** template will not be unintentionally started.

To create a new process group:

1. If you haven't already, open a browser and navigate to `<nifi_node_fqdn>:8085/nifi/`.





In a secure cluster, the NiFi URL is `https://<nifi_node_fqdn>:8085/nifi/`

2. In NiFi, double-click the **Interaset Marketplace** template to open it.
3. In the **Additional Deliverables - For ingesting data into HBase and Elasticsearch** (yellow) canvas group, right-click the **AD Loading (Generic CSV) Configuration** process group and click **Copy**.
4. Right-click the **Interaset Marketplace** canvas and select **Leave Group**.
5. Right-click the **NiFi Flow** canvas and select **Paste**. A copy of the **AD Loading (Generic CSV) Configuration** process group appears on your canvas next to **Interaset Marketplace** template.

You can optionally rename the new process group:

1. Right-click the **AD Loading (Generic CSV) Configuration** process group you just created, and then click **Configure**.
2. On the **General** tab, enter a new name in the **Process Group Name** field.
3. Click **Apply**, and click **OK** when the change is confirmed.
4. Close the configuration dialog to return to the main canvas.

Configure the Schema Registry

1. Right-click the process group you just created, and select **Configure**.
2. On the **Controller Services** tab, scroll down to the **ConfluentSchemaRegistry - Interaset Avro Schemas** controller service, and click anywhere in the row to select it.
3. In the right-most column of the table, in the selected row, click the right-arrow , then click **View Configuration** .
4. The **Configure Controller Service** dialog opens.

5. On the **Properties** tab, do the following:

- In the **Schema Registry URLs** field, enter the URL and port of the Interaset schema registry.



The default location is `https://<master_node_fqdn>:9190`

- In the **SSL Context Service** field, select the SSL context service that you set up in the [Configure an SSL Context Service](#) section.



6. Click **Apply**.

7. Close the **NiFi Flow Configuration** dialog to return to the main canvas.

Enable the Controller Services

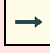
Before you can run any processors in NiFi, you must enable the controller services for the process group you just created.

To enable controller services:

1. Right-click the process group, and then click **Configure**.
2. Select the **Controller Services** tab.
3. For each controller service displayed, ensure that **Enabled** is displayed in the **State** column.
4. For any disabled controller services, click the right-arrow , and then click **Enable** . The **Enable Controller Service** dialog is displayed.
5. Click **Enable**, and then click **Close**.
6. Repeat for all disabled controller services, and then close the configuration dialog.



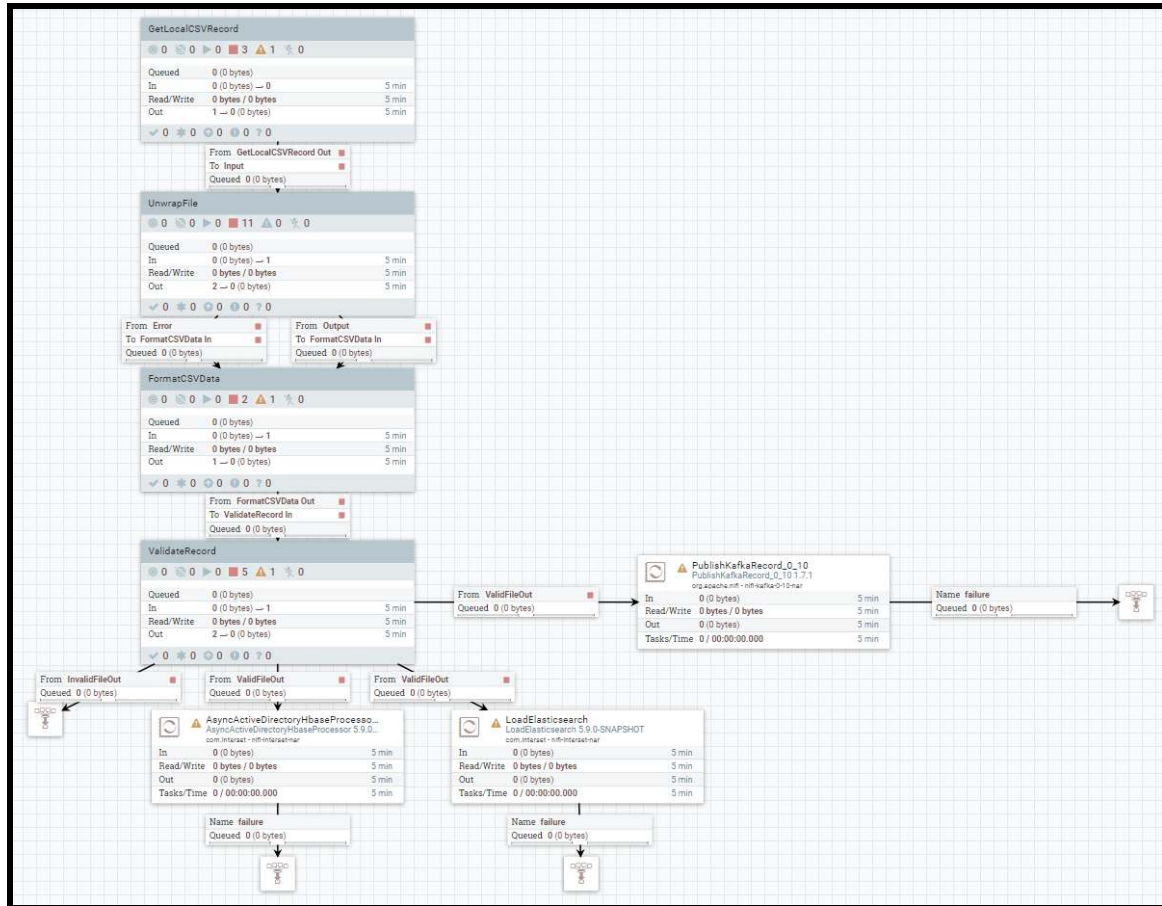
Important: It may happen that some services won't start because of dependencies.

If that happens, following the right-arrow  to dependent service to **Start** or **Configure** as required.

Configure the Data Flow

In this section, you configure various processors in the NiFi flow that are responsible for extracting, transforming, and loading the data into Interaset.

Enter the process group you just created by double-clicking it. You should see the following:



To configure the data flow, you edit properties in the following process groups:

- GetLocalCSVRecord
- FormatCSVData
- ValidateRecord

You must also set properties in the following processors:

- AsyncActiveDirectoryHbaseProcessor
- LoadElasticsearch
- PublishKafkaRecord_0_10

Configure the GetLocalCSVRecord Process Group

1. Double-click the **GetLocalCSVRecord** process group to enter it.
2. Right-click the **ListFile** processor and select **Configure**.
3. On the **Properties** tab, next to **Input Directory**, enter the path to the directory where your data is stored (for example, `/opt/interaset/data/`), and then click **Apply**.
4. Right-click the canvas and select **Leave Group** to return to the top level of the flow.

Configure the FormatCSVData Process Group

1. Double-click the **FormatCSVData** process group to enter it.
2. Right-click the **UpdateRecord - Format CSV Fields** processor and select **Configure**.
3. On the **Properties** tab, verify that the expression in the **/time** field contains a date format for the incoming records that matches the date format in the incoming file. For example, the format `"yyyy/MM/dd HH:mm:ssXXX"` matches dates like `2017/04/01 08:01:15-05:00`, whereas the format `"yyyy-MM-ddTHH:mm:ssXXX"` matches date like `2017-04-01T08:01:15-05:00`.
4. If you made changes to the date format, click **Apply**; otherwise, click **Cancel**.
5. Right-click the canvas and click **Leave Group** to return to the top level of the flow.

Configure the AsyncActiveDirectoryHbaseProcessor

To set up this Hbase processor, you must:

- configure the NiFi node environment
- configure the NiFi processor properties

Configure the NiFi Node Environment

1. Run the following command on the **Nifi** node (replace `<KERBEROS_DOMAIN.COM>` with the realm you set when you configured Kerberos) to ensure that the NiFi user has an active authentication ticket through Kerberos:

```
sudo -u nifi kinit -kt /etc/security/interset/keytab/i_nifi.service.keytab i_nifi/<nifi_node_fqdn>@<REALM>
```

Verify the Validity of the ticket by running the following command. This will provide you with the ticket details about the default principal, start and Expiry date along with the due date for next renewal.

```
sudo -u nifi klist
```

2. For CDH, run the following command on the **Compute** node to determine the path to the Region-server Key tab and initialize the key :

```
sudo -u hbase kinit -kt /var/run/cloudera-scm-agent/process/`sudo ls -lrt /var/run/cloudera-scm-agent/process/ | awk '{print $9}' | grep 'hbase-MASTER$|hbase-REGIONSERVER$' | tail -1`/hbase.keytab hbase/<compute_node_fqdn>@<REALM>
```

Verify the Validity of the ticket by running the following command. This will provide you with the ticket details about the default principal, start and Expiry date along with the due date for next renewal.

```
sudo -u hbase klist
```

3. For HDP, run the following command to initialize the key:

```
sudo -u hbase kinit hbase/<compute_node_fqdn>@<REALM> -kt /etc/security/keytabs/hbase.service.keytab
```


4. Run the following command on the **NiFi** node to restart NiFi:

```
service nifi restart
```



Important: For ticket renewal use the following commands within the crontab file. These commands will ensure logging for any errors during the renewal of the ticket.

Run on the **Compute node(s)**

```
sudo -u hbase sh -c "kinit -R > /tmp/hbase_kinit.log 2>&1"
```

Run on the **Nifi node(s)**

```
sudo -u nifi sh -c "kinit -R > /var/log/nifi/kinit_cron.log 2>&1"
```

Configure the AsyncActiveDirectoryHbaseProcessor Properties

1. Right-click the **AsyncActiveDirectoryHbaseProcessor** processor and select **Configure**.
2. Select the **Properties** tab.
3. In text box next to the **Zookeeper Quorum** property, enter the name of your Zookeeper hosts. The default location is **<master_node_fqdn>**, for example, **master.interset.com**.
4. Ensure the **Zookeeper Node Parent** property is correctly set:
 - for HDP in an unsecured environment: **/hbase-unsecure**
 - for HDP in a Kerberos environment: **/hbase-secure**
 - for CDH: **/hbase**
5. Set value of the **Tenant ID (tid)** property to the tenant for which you are ingesting data.
6. Set the value of the **Data Instance ID (did)** to the appropriate Data Instance ID. The default is 0.
7. If you are running InterSet on a secure cluster, set the properties required for Kerberos on the **Compute** node:
 - **Security Auth Enabled:** **true**
 - **Security Authentication:** **kerberos** (to enable regular Kerberos authentication) or **protected** (to enable encryption on RPC payloads).
 - **Security Kerberos Principal:** The name of the Hbase principal to use for Kerberos authentication.

```
hbase/compute_node_fqdn@<REALM>
```
 - **Security RPC Protection:** This property indicates whether to encrypt RPCs while they are traversing the network. Possible values are **authentication** (no encryption), **integrity** (no encryption) or **privacy** (encrypted). Requires authentication to be enabled. The default is **privacy**.

- **Security SASL Client Config:** Set this property to the name of the section (for example, Client) in the JAAS configuration file you created in the previous section. This name is used to look up the configuration when NiFi is authenticating a user against a region server.
- **Append Classpath:** This property extends the classpath that the processor will run with, allowing the processor to access the **hbase-site.xml** client configuration file to establish the client connection with HBase. You can set this by using the following command:

```
CDH: /etc/hbase/conf:/etc/hadoop/conf:/opt/cloudera/parcels/CDH/lib/hbase/lib/*:/opt/interaset/phoenix_lib_links/*
```

```
HDP: /etc/hbase/conf:/etc/hadoop/conf:/usr/hdp/current/hbase-client/lib/*:/opt/interaset/phoenix_lib_links/*
```

8. Click **Apply** to close the processor properties and return to the main flow.

Configure the LoadElasticsearch Processor

1. Right-click the **LoadElasticsearch** processor and select **Configure**.
2. Select the **Properties** tab.
3. In text box next to the **Elasticsearch Hostnames** property, enter the name(s) and port(s) of your Elasticsearch hosts.



The default location is `<search_node_fqdn>:9200`.

4. Edit the **Index Name** property to include the correct tenant ID. For example, if your tenant is 234, the value of the **Index Name** property must be set to `interaset_ad_rawdata_234`.
5. Set value of the **Tenant ID (tid)** property to the tenant for which you are ingesting data.
6. Set the value of the **Data Instance ID (did)** to the appropriate Data Instance ID. The default is 0.
7. If you are running Interaset on a secure cluster, set the properties required for Kerberos as follows:
 - **Security User:** The Elasticsearch user and password. Must be in the format `<user>:<password>`. The default user is `elastic`.
 - **SSL Enabled:** `true`
 - **Truststore Path:** The path to the trust store file. The default location is `/etc/security/interaset/sslcrt/all-truststore.jks`.
 - **Truststore Password:** The password to the trust store.
 - **Keystore Path:** The path to the keystore file. The default location is `/etc/security/interaset/sslcrt/localhost-keystore.jks`.
 - **Keystore Password:** The password to the keystore.
8. Click **Apply** to close the processor properties and return to the main flow.

Configure the PublishKafkaRecord_0_10 Processor



Important: If you are not using Workflow in your Interaset installation, delete the connection leading to the **PublishKafkaRecord_0_10** processor. Otherwise, as the data ingest progresses, the processor queue will reach its maximum capacity and then apply back-pressure to the upstream processors, which will cause the entire ingest to stall.

1. Right-click the **PublishKafkaRecord_0_10** processor and select **Configure**.
2. Select the **Properties** tab.
3. In text box next to the **Kafka Brokers** property, enter the name(s) and port(s) of your Kafka Broker(s).



- The default location in Ambari is `<stream_node_fqdn>:6667`
- In a secured cluster in Cloudera, the default location is `<stream_node_fqdn>:9093`
- In an unsecured cluster in Cloudera, the default location is `<stream_node_fqdn>:9092`
- In case of multiple stream nodes, each stream and port needs to be specified and separated with a coma. `< stream_ node_fqdn>:port,<stream_node_fqdn>:port...etc`

4. Edit the Topic Name property to include the correct tenant ID and data source ID. For example, for the default tenant ID (0) and data source ID (0), **Topic Name** would be set to `interaset_ad_events_0_0`
5. If you are running Interaset on a secure cluster, set the properties required for Kerberos as follows:
 - **Security Protocol:** Select the protocol used to communicate with brokers. Corresponds to Kafka's `security.protocol` property.
 - **Kerberos Credentials Service:** Specify the Kerberos Credentials Controller Service that should be used for authenticating with Kerberos.
 - **Kerberos Service Name:** Enter the Kerberos principal name that Kafka runs as. This can be defined either in Kafka's JAAS config or in Kafka's config. Corresponds to Kafka's `security.protocol` property. It is ignored unless one of the SASL options of the **Security Protocol** property are selected.
 - **Kerberos Principal:** Enter the Kerberos principal that will be used to connect to brokers. If not set, there must be a JAAS configuration file set in the JVM properties JVM properties defined in NiFi's `bootstrap.conf` file. This principal will be set into Kafka's `sasl.jaas.config` property.

- **Kerberos Keytab:** Enter the Kerberos keytab that will be used to connect to brokers. If not set, there must be a JAAS configuration file set in the JVM properties defined in NiFi's **bootstrap.conf** file.



Tip: These values can be copied from previously run commands:

```
sudo -u nifi kinit -kt /etc/security/interaset/keytab/i_nifi.service.keytab i_nifi/<stream_node_fqdn>@<REALM>
```

For Kerberos Keytab: /etc/security/interaset/keytab/i_nifi.service.keytab

For Kerberos Principal: i_nifi/<stream_node_fqdn>@<REALM>

- **SSL Context Service:** Select the SSL context service to use for communicating with Kafka (the service that you configured in the [Configure an SSL Context Service](#) section).
6. Click **Apply** to close the processor properties and return to the main flow.

Start the Data Flow



To start the data flow:

1. In a web browser, return to the NiFi user interface.
2. If you haven't already done so, enable the controller services in your top-level process group.



Tip:

To enable controller services:

1. Right-click the process group, and then click **Configure**.
2. Select the **Controller Services** tab.
3. For each controller service displayed, ensure that **Enabled** is displayed in the **State** column.
4. For any disabled controller services, click the right-arrow , and then click Enable . The **Enable Controller Service** dialog is displayed.
5. Click **Enable**, and then click **Close**.
6. Repeat for all disabled controller services, and then close the configuration dialog.

3. Right-click the **AD Loading (Generic CSV) Configuration** group and select **Start**.

Your process flow is now running, and you should see the byte counts incrementing at each stage of the flow. You might need to refresh your view periodically by right-clicking the canvas and selecting **Refresh**.

Configure Violations Loading

1. If you haven't already, open a browser and navigate to `<nifi_node_fqdn>:8085/nifi/`.



In a secure cluster, the NiFi URL is `https://<nifi_node_fqdn>:8085/nifi/`

2. In NiFi, double-click the **InterSet Marketplace** template to open it.
3. In the **Additional Deliverables - For ingesting data into HBase and Elasticsearch** (yellow) canvas group, right-click the **Violations Loading Configuration** process group and click **Copy**.
4. Right-click the **InterSet Marketplace** canvas and select **Leave Group**.
5. Right-click the **NiFi Flow** canvas and select **Paste**. A copy of the **Violations Loading Configuration** process group appears on your canvas next to **InterSet Marketplace** template.
6. Right-click the **Violation Loading** processor and select **Enter group**.
7. To configure **Violation to Hbase**, right-click on the process and click on configure.
8. Select the properties tab

Update Workflow for Schema Registry

1. On the Master node in your InterSet cluster, navigate to `/opt/interSet/rules/conf`, and open the `rules.conf` file.
2. In the **Service Connections** section, in the **Non Schema Registry** section, add a hash (#) symbol to all **TopicEvents** parameters to comment them out.



Tip: The **TopicEvents** parameters are required only when not using the schema registry.

3. In the **Service Connections** section, in the **Schema Registry** section, remove the hash (#) symbol from the **ConfluentRegistry** parameter, and then enter the following value, where `<schema_registry_fqdn>` is the fully-qualified domain name of the machine where the schema registry is installed.

```
ConfluentRegistry = <schema_registry_fqdn>:9190/api/v1/confluent/
```

4. In the **Event Processing** section, do the following:
 - Uncomment and then enter the required values for the **SchemaSubjectEvents** parameter (s), using the values defined in the table below.

```
#SchemaSubjectEvents.<schemaName1> = <archType1>
```

For example, for repository data, the **SchemaSubjectEvents** parameter is entered as:

```
SchemaSubjectEvents.RepositoryRecord = repo
```


Tips:

- The **<schemaName>** is the value of the name field from within the Avro schema itself, for example, **AccessRecord**.
- Enter as many **SchemaSubjectEvents** parameters as required for your data source types.

Data Source Type	<archType>
Access	access
Active Directory	ad, active_directory, activedirectory
Email	email
Endpoint	endpoint, wdc
Expense	expense
Alert	interaset-standard, interasetstandard, alert, uaf
Linux AuditD	auditd, linuxauditd
NetFlow	netflow
Printer	printer
Repository	repo, repository
Sensor	sensor
VPN	vpn
Web Proxy	proxy, webproxy, webproxy
Windows Printer	winprinter, windows_printer, windowsprinter

5. Save and close the **rules.conf** file.
6. Repeat these steps for each configured Workflow **rules.conf** file in your cluster.
7. When all required Workflow **rules.conf** configuration files are updated and saved, run the following command to deploy each of your updated 5.9.2 Workflow configurations.

```
/opt/interaset/rules/bin/workflow.sh --deploy /opt/interaset/rules/conf/<rules.conf>
```

Update the EDR Server for Schema Registry

To enable the schema registry for the EDR server, you must make the following changes to the EDR server system configuration file.

1. On each EDR server node in your Interset cluster, do the following:
 - As the **interaset** user, ssh to an EDR server node in your Interset cluster.

```
ssh interaset@<edr_server_node_fqdn>
```

- Run the following command to stop the EDR server.

```
sudo systemctl stop edr-server
```

- Navigate to the **/opt/interaset/flux/conf** directory, and then open the **system.conf** file.
- In the **Endpoint Server (Flux) Properties** section, add the following line to the **com.interaset.endpoint.flux.output** grouping:

```
com.interset.endpoint.flux.output.<schema_registry_fqdn>:9190/api/v1/confluent/
```

- Save and close the **system.conf** file.
- Run the following command to start the EDR server.

```
sudo systemctl start edr-server
```

Appendix B: Flume Configuration Changes

The following tables list the changes in the Flume configuration files between Interaset 5.7.x and Interaset 5.9.2.

You must make the changes listed to allow your Interaset 5.7.x Flume configurations to work with Interaset 5.9.2.

If you need to upgrade your existing configurations, please contact Micro Focus Interaset Support at interaset.support@microfocus.com for assistance.

Extract Configurations

The extract configurations are located in the `/conf-templates-2.0/extract` directory. The table below lists the subdirectory for each configuration.

Subdirectory	File	Changes (from 5.6.0 to 5.7.0)
/arcsight-logger-ad-authentication	arcsight-logger-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interaset_auth_events_{{ did }}_{{ tid }}_arcSightLogger to: interaset_ad_events_{{ did }}_{{ tid }}_arcSightLogger Channel Kafka topic name changed from: interaset_auth_raw_arcsight_logger_{{ did }}_{{ tid }} to: interaset_ad_raw_arcsight_logger_{{ did }}_{{ tid }}
	arcsight-logger-historical-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interaset_auth_events_{{ did }}_{{ tid }}_arcSightLoggerh to: interaset_ad_events_{{ did }}_{{ tid }}_arcSightLoggerh Channel Kafka topic name changed from: interaset_auth_raw_arcsight_logger_{{ did }}_{{ tid }} to: interaset_ad_raw_arcsight_logger_{{ did }}_{{ tid }}

Subdirectory	File	Changes (from 5.6.0 to 5.7.0)
/generic-files	file-line-by-line-to-kafka.conf	None.
	file-multiline-csv-to-kafka.conf	None.
/mcafee-esm-ad-authentication	mcafee-esm-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_mcafeeEsm to: interset_ad_events_{{ did }}_{{ tid }}_mcafeeEsm Channel Kafka topic name changed from: interset_auth_raw_mcafee_esm_{{ did }}_{{ tid }} to: interset_ad_raw_mcafee_esm_{{ did }}_{{ tid }}
	mcafee-esm-historical-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_mcafeeEsmh to: interset_ad_events_{{ did }}_{{ tid }}_mcafeeEsmh Channel Kafka topic name changed from: interset_auth_raw_mcafee_esm_{{ did }}_{{ tid }} to: interset_ad_raw_mcafee_esm_{{ did }}_{{ tid }}
/qradar	qradar-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_qRadar to: interset_ad_events_{{ did }}_{{ tid }}_qRadar Channel Kafka topic name changed from: interset_auth_raw_qradar_{{ did }}_{{ tid }} to: interset_ad_raw_qradar_{{ did }}_{{ tid }}
	qradar-historical-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_qRadarh to: interset_ar_events_{{ did }}_{{ tid }}_qRadarh

Subdirectory	File	Changes (from 5.6.0 to 5.7.0)
		<ul style="list-style-type: none"> Channel Kafka topic name changed from: <code>interaset_auth_raw_qradar_{{ did }}_{{ tid }}</code> to: <code>interaset_ar_raw_qradar_{{ did }}_{{ tid }}</code>
/splunk-ad	splunk-ad-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>interaset_auth_events_{{ did }}_{{ tid }}_splunk</code> to: <code>interaset_ad_events_{{ did }}_{{ tid }}_splunk</code> Channel Kafka topic name changed from: <code>interaset_auth_raw_splunk_{{ did }}_{{ tid }}</code> to: <code>interaset_ad_raw_splunk_{{ did }}_{{ tid }}</code>
	splunk-historical-ad-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>interaset_auth_events_{{ did }}_{{ tid }}_splunkh</code> to: <code>interaset_ad_events_{{ did }}_{{ tid }}_splunkh</code> Channel Kafka topic name changed from: <code>interaset_auth_raw_splunk_{{ did }}_{{ tid }}</code> to: <code>interaset_ad_raw_splunk_{{ did }}_{{ tid }}</code>
	splunk-historical-windowsprinter-to-kafka.conf	Channel Kafka topic name changed from: <code>interaset_printer_raw_splunk_{{ did }}_{{ tid }}</code> to: <code>interaset_winprinter_raw_splunk_{{ did }}_{{ tid }}</code>
	splunk-windowsprinter-to-kafka.conf	Channel Kafka topic name changed from: <code>interaset_printer_raw_splunk_{{ did }}_{{ tid }}</code> to: <code>interaset_winprinter_raw_splunk_{{ did }}_{{ tid }}</code>
/syslog	syslog-to-kafka.conf	None.

Transform Configurations

The extract configurations are located in the `/conf-templates-2.0/transform` directory. The table below lists the subdirectory for each configuration.

Directory	File	Changes (from 5.6.0 to 5.7.0)
/arcsight-logger-ad-authentication	kafka-arcsight-logger-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>interaset_auth_events_{{ did }}_{{ tid }}_arcsightLogger_transform</code> to: <code>interaset_ad_events_{{ did }}_{{ tid }}_arcsightLogger_transform</code> Source Kafka topic name changed from: <code>interaset_auth_raw_arcsight_logger_{{ did }}_{{ tid }}</code> to: <code>interaset_ad_raw_arcsight_logger_{{ did }}_{{ tid }}</code> Source Kafka consumer group ID changed from: <code>interaset_auth_raw_arcsight_logger_{{ did }}_{{ tid }}_transform_group</code> to: <code>interaset_ad_raw_arcsight_logger_{{ did }}_{{ tid }}_transform_group</code>
/auditd	kafka-auditd-to-kafka.conf	None.
/generic-alerts	kafka-csv-alert-to-kafka.conf	None.
/generic-authentication	kafka-csv-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>interaset_auth_events_{{ did }}_{{ tid }}_csv_transform</code> to: <code>interaset_ad_events_{{ did }}_{{ tid }}_csv_transform</code> Source Kafka topic name changed from: <code>interaset_auth_raw_csv_{{ did }}_{{ tid }}</code> to: <code>interaset_ad_raw_csv_{{ did }}_{{ tid }}</code> Source Kafka consumer group ID changed from: <code>interaset_auth_raw_csv_{{ did }}_{{ tid }}_transform_group</code> to:

Directory	File	Changes (from 5.6.0 to 5.7.0)
		<p>interiset_ad_raw_csv_{{ did }}_{{ tid }}_transform_group</p> <ul style="list-style-type: none"> Channel Kafka topic name changed from: interiset_auth_events_{{ did }}_{{ tid }} to: interiset_ad_events_{{ did }}_{{ tid }}
/generic-printer	kafka-csv-printer-to-kafka.conf	<ul style="list-style-type: none"> Example CSV format (in comments) changed from: timestamp,eventCode,user,_,_,action,printer,fileName,size,pages,_ to: timestamp,signatureId,user,_,_,action,printer,fileName,size,pages,_ Record type for dePrToAvro interceptor (interceptors.dePrToAvro.recordType) changed from: com.interset.schema.WindowsPrinterRecord to: com.interset.schema.PrinterRecord
/generic-repository	kafka-csv-repo-to-kafka.conf	None.
/generic-webproxy	kafka-csv-webproxy-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interiset_webproxy_events_{{ did }}_{{ tid }}_csv_transform to: interiset_proxy_events_{{ did }}_{{ tid }}_csv_transform Source Kafka topic name changed from: interiset_webproxy_raw_csv_{{ did }}_{{ tid }} to: interiset_proxy_raw_csv_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interiset_webproxy_raw_csv_{{ did }}_{{ tid }}_transform_group to: interiset_proxy_raw_csv_{{ did }}_{{ tid }}_transform_group Channel Kafka topic name changed from: interiset_webproxy_events_{{ did }}_{{ tid }} to: nterset_proxy_events_{{ did }}_{{ tid }}

Directory	File	Changes (from 5.6.0 to 5.7.0)
/github	kafka-github-repo-to-kafka.conf	None
/mcafee-esm-ad-authentication	kafka-mcafee-esm-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>intersect_auth_events_{{ did }}_{{ tid }}_mcafeeEsm_transform</code> to: <code>intersect_ad_events_{{ did }}_{{ tid }}_mcafeeEsm_transform</code> Source Kafka topic name changed from: <code>intersect_auth_raw_mcafee_esm_{{ did }}_{{ tid }}</code> to: <code>intersect_ad_raw_mcafee_esm_{{ did }}_{{ tid }}</code> Source Kafka consumer group ID changed from: <code>intersect_auth_raw_mcafee_esm_{{ did }}_{{ tid }}_transform_group</code> to: <code>intersect_ad_raw_mcafee_esm_{{ did }}_{{ tid }}_transform_group</code> Channel Kafka topic name changed from: <code>intersect_auth_events_{{ did }}_{{ tid }}</code> to: <code>intersect_ad_events_{{ did }}_{{ tid }}</code>
/netflow	kafka-csv-netflowV10-to-kafka.conf	New optional interceptors added to specify custom internal IPv4 address ranges for source and destination.
	kafka-csv-netflowV5-to-kafka.conf	
	kafka-csv-netflowV9-to-kafka.conf	
/perforce	kafka-perforce-repo-to-kafka.conf	None.
	kafka-perforce-structured-repo-to-kafka.conf	None.
/qradar	kafka-qradar-auth-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: <code>intersect_auth_events_{{ did }}_{{ tid }}_qRadar_transform</code> to: <code>intersect_ad_events_{{ did }}_{{ tid }}_qRadar_transform</code> Source Kafka topic name changed from: <code>intersect_auth_raw_qradar_{{ did }}_{{ tid }}</code> to: <code>intersect_ad_raw_qradar_{{ did }}_{{ tid }}</code>

Directory	File	Changes (from 5.6.0 to 5.7.0)
		<ul style="list-style-type: none"> Source Kafka consumer group ID changed from: interaset_auth_raw_qradar_{{ did }}_{{ tid }}_transform_group to: interaset_ad_raw_qradar_{{ did }}_{{ tid }}_transform_group Channel Kafka topic name changed from: interaset_auth_events_{{ did }}_{{ tid }} to: interaset_ad_events_{{ did }}_{{ tid }}
/splunk-ad	kafka-splunk-ad-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interaset_auth_events_{{ did }}_{{ tid }}_splunk_transform to: interaset_ad_events_{{ did }}_{{ tid }}_splunk_transform Source Kafka topic name changed from: interaset_auth_raw_splunk_{{ did }}_{{ tid }} to: interaset_ad_raw_splunk_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interaset_auth_raw_splunk_{{ did }}_{{ tid }}_transform_group to: interaset_ad_raw_splunk_{{ did }}_{{ tid }}_transform_group Channel Kafka topic name changed from: interaset_auth_events_{{ did }}_{{ tid }} to: interaset_ad_events_{{ did }}_{{ tid }}
	kafka-splunk-windowsprinter-to-kafka.conf	<ul style="list-style-type: none"> Source Kafka topic name changed from: interaset_printer_raw_splunk_{{ did }}_{{ tid }} to: interaset_winprinter_raw_splunk_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interaset_printer_raw_splunk_{{ did }}_{{ tid }}_transform_group

Directory	File	Changes (from 5.6.0 to 5.7.0)
		to: interset_winprinter_raw_splunk_{{ did }}_{{ tid }}_transform_group • Channel Kafka topic name changed from: interset_printer_events_{{ did }}_{{ tid }} to: interset_winprinter_events_{{ did }}_{{ tid }}
/windows-event-log-csv-export	kafka-interaset-wel-csv-ad-to-kafka.conf	• Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_interaset_wel_csv_transform to: interset_ad_events_{{ did }}_{{ tid }}_interaset_wel_csv_transform • Source Kafka topic name changed from: interset_auth_raw_interaset_wel_csv_{{ did }}_{{ tid }} to: interset_ad_raw_interaset_wel_csv_{{ did }}_{{ tid }} • Source Kafka consumer group ID changed from: interset_auth_raw_interaset_wel_csv_{{ did }}_{{ tid }}_transform_group to: interset_ad_raw_interaset_wel_csv_{{ did }}_{{ tid }}_transform_group • Channel Kafka topic name changed from: interset_auth_events_{{ did }}_{{ tid }} to: interset_ad_events_{{ did }}_{{ tid }}
	kafka-wel-csv-ad-to-kafka.conf	• Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_wel_csv_transform. to: interset_ad_events_{{ did }}_{{ tid }}_wel_csv_transform • Source Kafka topic name changed from: interset_auth_raw_wel_csv_{{ did }}_{{ tid }} to: interset_ad_raw_wel_csv_{{ did }}_{{ tid }}

Directory	File	Changes (from 5.6.0 to 5.7.0)
		<ul style="list-style-type: none"> Source Kafka consumer group ID changed from: interaset_auth_raw_wel_csv_{{ did }}_{{ tid }}_transform_group to: interaset_ad_raw_wel_csv_{{ did }}_{{ tid }}_transform_group Channel Kafka topic name changed from: interaset_auth_events_{{ did }}_{{ tid }} to: interaset_ad_events_{{ did }}_{{ tid }}
	kafka-wel-csv-printer-to-kafka.conf	<ul style="list-style-type: none"> Source Kafka topic name changed from: interaset_printer_raw_wel_csv_{{ did }}_{{ tid }} to: interaset_winprinter_raw_wel_csv_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interaset_printer_raw_wel_csv_{{ did }}_{{ tid }}_transform_group to: interaset_winprinter_raw_wel_csv_{{ did }}_{{ tid }}_transform_group Channel Kafka topic name changed from: interaset_printer_events_{{ did }}_{{ tid }} to: interaset_winprinter_events_{{ did }}_{{ tid }}
	kafka-wel-viewer-csv-ad-to-kafka.conf	<ul style="list-style-type: none"> Agent name changed from: interaset_auth_events_{{ did }}_{{ tid }}_wel_viewer_csv_transform to: interaset_ad_events_{{ did }}_{{ tid }}_wel_viewer_csv_transform Source Kafka topic name changed from: interaset_auth_raw_wel_viewer_csv_{{ did }}_{{ tid }} to: interaset_ad_raw_wel_viewer_csv_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interaset_auth_raw_wel_viewer_csv_{{ did }}_{{ tid }}_transform_group

Directory	File	Changes (from 5.6.0 to 5.7.0)
		to: interset_ad_raw_wel_viewer_csv_{{ did }}_{{ tid }}_transform_group

Loading Configurations

The loading configurations are located in the **/conf-templates-server/kafka** directory.

File	Changes (from 5.6.0 to 5.7.0)
kafka-ad-to-es.conf	<ul style="list-style-type: none"> Renamed from kafka-auth-to-es.conf. Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_es to: interset_ad_events_{{ did }}_{{ tid }}_es Source Kafka topic name changed from: interset_auth_events_{{ did }}_{{ tid }} to: interset_ad_events_{{ did }}_{{ tid }} Source Kafka consumer group ID changed from: interset_auth_events_{{ did }}_{{ tid }}_es_group to: interset_ad_events_{{ did }}_{{ tid }}_es_group New interceptor to handle encode/decode with schema registry.
kafka-ad-to-hbase.conf	<ul style="list-style-type: none"> Renamed from kafka-auth-to-hbase.conf. Agent name changed from: interset_auth_events_{{ did }}_{{ tid }}_hbase to: interset_ad_events_{{ did }}_{{ tid }}_hbase Source Kafka consumer group ID changed from: interset_auth_events_{{ did }}_{{ tid }}_hbase_group to:

File	Changes (from 5.6.0 to 5.7.0)
	<p>interaset_ad_events_{{ did }}_{{ tid }}_hbase_group</p> <ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New settings to control <code>objectName</code> and <code>sharePath</code> directory depth.
kafka-auditd-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-auditd-to-hbase.conf	New interceptors to handle encode/decode with schema registry.
kafka-email-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-email-to-hbase.conf	New interceptors to handle encode/decode with schema registry.
kafka-expense-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-expense-to-hbase.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for HBase sink serializer to specify schema registry URL
kafka-netflow-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-netflow-to-hbase.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for HBase sink serializer to specify schema registry URL New setting to turn on the generation of NetFlow packet relations.
kafka-repo-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-repo-to-hbase.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for HBase sink serializer to specify schema registry URL
kafka-sensor-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-sensor-to-hbase.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for HBase sink serializer to specify schema registry URL
kafka-violations-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-violations-to-phoenix.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for Phoenix sink serializer to specify schema registry URL
kafka-wdc-to-es.conf	New interceptor to handle encode/decode with schema registry.
kafka-wdc-to-hbase.conf	<ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for Hbase sink serializer to specify schema registry URL
kafka-webproxy-to-es.conf	<ul style="list-style-type: none"> Agent name changed from: interaset_webproxy_events_{{ did }}_{{ tid }}_es to: interaset_proxy_events_{{ did }}_{{ tid }}_es Source Kafka topic name changed from:

File	Changes (from 5.6.0 to 5.7.0)
	<p>interiset_webproxy_events_{{ did }}_{{ tid }}</p> <p>to:</p> <p>interiset_proxy_events_{{ did }}_{{ tid }}</p> <ul style="list-style-type: none"> Source Kafka consumer group ID changed from: interiset_webproxy_events_{{ did }}_{{ tid }}_es_group <p>to:</p> <p>interiset_proxy_events_{{ did }}_{{ tid }}_hbase_group</p> <ul style="list-style-type: none"> New interceptor to handle encode/decode with schema registry.
kafka-webproxy-to-hbase.conf	<ul style="list-style-type: none"> Agent name changed from: interiset_webproxy_events_{{ did }}_{{ tid }}_hbase <p>to:</p> <p>interiset_proxy_events_{{ did }}_{{ tid }}_hbase</p> <ul style="list-style-type: none"> Source Kafka topic name changed from: interiset_webproxy_events_{{ did }}_{{ tid }} <p>to:</p> <p>interiset_proxy_events_{{ did }}_{{ tid }}</p> <ul style="list-style-type: none"> Source Kafka consumer group ID changed from: interiset_webproxy_events_{{ did }}_{{ tid }}_es_group <p>to:</p> <p>interiset_proxy_events_{{ did }}_{{ tid }}_hbase_group</p> <ul style="list-style-type: none"> New interceptors to handle encode/decode with schema registry. New setting for Hbase sink serializer to specify schema registry URL. New setting to specify tracking relations by <code>clientId</code> field instead of <code>clientName</code> field.
kafka-windowsprinter-to-es.conf	<ul style="list-style-type: none"> Agent name changed from: interiset_printer_events_{{ did }}_{{ tid }}_es <p>to:</p> <p>interiset_printer_events_{{ did }}_{{ tid }}_es</p> <ul style="list-style-type: none"> Source Kafka topic name changed from: interiset_printer_events_{{ did }}_{{ tid }} <p>to:</p>

File	Changes (from 5.6.0 to 5.7.0)
	<p>interiset_winprinter_events_{{ did }}_{{ tid }}</p> <ul style="list-style-type: none">Source Kafka consumer group ID changed from: interiset_printer_events_{{ did }}_{{ tid }}_es_group to: interiset_winprinter_events_{{ did }}_{{ tid }}_es_groupNew interceptor to handle encode/decode with schema registry.
kafka-windowsprinter-to-hbase.conf	<ul style="list-style-type: none">Agent name changed from: interiset_printer_events_{{ did }}_{{ tid }}_hbase to: interiset_winprinter_events_{{ did }}_{{ tid }}_hbaseSource Kafka topic name changed from: interiset_printer_events_{{ did }}_{{ tid }} to: interiset_winprinter_events_{{ did }}_{{ tid }}Source Kafka consumer group ID changed from: interiset_printer_events_{{ did }}_{{ tid }}_hbase_group to: interiset_winprinter_events_{{ did }}_{{ tid }}_hbase_groupNew setting for Hbase sink serializer to specify schema registry URL.

Index

Analytics

- upgrading 21

audience

- intended 7

configurations

- optional 32

- updating Data Ingest 28

configuring

- a new data source 34

- postgreSQL 10 external database 14

Data Ingest

- updating configurations 28

- upgrading 25

data sources

- configuring 34

downloading

- Interset installer 9

EDR Server

- update for schema registry 46

- upgrading 23

Elasticsearch

- removing plug-ins 17

- upgrading 17

- upgrading index data 28

- upgrading mappings 28

- upgrading templates 28

extracting

- Interset installer 9

Flume

- stopping 13

installations

- optional 32

installer

- downloading 9

- extracting 9

installing

- postgresql 10 external database 14

intended

- audience 7

Interaset 5.5.x

- upgrading 13

Interaset 5.6

- preparing to upgrade 9

Interaset installer

- downloading 9

- extracting 9

interaset user

- upgrading 14

Introduction 5

Kerberos

- configuring in NiFi 41-43

Kibana

- upgrading 19

mirror server 9

NiFi

- configure SSL controller service 35

- configuring Kerberos 41-43

- configuring the example flow 38

- copying a process group 37
- enabling controller services 38
- specifying schema registry location 37
- starting the data flow 44
- upgrading 26
- offline
 - upgrade 9
- optional
 - configurations 32
 - installations 32
- postgresql
 - configuring 14
 - installing external database 14
- preparing
 - upgrade Interset 5.6 9
- removing
 - Elasticsearch plug-ins 17
- Reporting
 - upgrading 19
- schema registry
 - specify location in NiFi 35, 37
 - update EDR server 46
 - update Workflow 45
- server
 - mirror 9
- stopping
 - Flume service 13
- updating
 - Data Ingest configurations 28

upgrading 13

Analytics 21

Data Ingest 25

EDR Server 23

Elasticsearch 17

Elasticsearch index data 28

Elasticsearch mappings 28

Elasticsearch templates 28

intersect user 14

Kibana 19

NiFi 26

offline 9

preparation 9

Reporting 19

to Intersect 5.6 13

Workflow 29

Workflow

update for schema registry 45

upgrading 29