



Interaset

© 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	8
Supported Environments	8
Supported Data Sources	9
Intended Audience	10
Installation	10
Additional Support	11
Administer the Interaset System	12
Manage Tenants and Users	12
Create a New Tenant	13
Remove a Tenant	14
Create a New User	14
Modify a User's Role	15
Configure Multi-tenant Authentication	15
Configure a New Data Source	16
Configure an SSL Context Service	17
Create a Process Group	19
Configure the Schema Registry	19
Enable the Controller Services	20
Configure the Data Flow	20
Start the Data Flow	25
Configure Analytics	27
Enable Windowed Analytics	28
Configure the 'Peek-Back' Window for Windowed Analytics	29
Configure Search	29
Configure New Action Mappings for Existing Data Sources	34
Resolve Multiple Entities in Source Data	35
Configure Interaset to Ingest Universal Third-party Alerts	35
Use the API to Define and Map the Universal Alert Columns	36

Use Interset Workflow to Create Violations with Risk Weights	38
Update the Universal Alert Data Ingest Settings	39
Administer Interset for End Users	40
Create Workflows	40
Verify the Workflow Engine is Running	40
Configure Workflow for Different Data Sources	41
Configure Workflow for a Custom Schema	45
Workflow Scenarios	46
Workflow Notifications and Outcomes	46
Scale Your Workflows	48
Manage Workflow Lists	49
Workflow Example	50
Manage Bots and Bot-like Users	52
Tune the Analytics	55
Use Interset	58
Users and Other Entities	58
Behaviors	58
Accumulating Risk	58
Overall Risk Page	59
Entities Page	60
User-Defined Tags	61
Explore Page	62
Matrix of Anomalies & Violations	63
Apply Filters to Entity and Anomaly Data	64
Anomalies & Violations Panel	65
Anomaly and Violation Flags	66
Add Comments	68
Entity Details Panel	68
Authentications Panel	69
Top Risky Panel	70

Most Exits By User Panel	71
Top Screen Capture Produces Panel	71
Top Users To Trigger Violations Panel	71
CSV Reports	71
PDF Reports	72
Explore Raw Events	72
Advanced Features	73
Workflows	73
Bots and Bot-like Users	74
Advanced Configuration Options	75
Custom Theme	75
Reset Theme to Default	77
Add Custom Text to the Interface Banner and Footer	77
Configuring Delimiters	80
User Accounts and Permissions	80
Interset Accounts	80
Ambari and HDP Accounts	80
Other Third-party Component Accounts	81
Appendix A: Interset Cluster Components	82
Interset Components	82
Third-party Components	82
Component Distribution	85
Interset Configuration	85
Appendix B: Explore the Analytics Using Sample Data	90
Create the Samples Tenant	90
Create an Administrator User for the Samples Tenant	90
Copy the Interset Sample Datasets	91
Configure the Sample Authentication Data Source	91
Run Analytics	91
Appendix C: Security Best Practices	92

Changing Default Account Passwords	92
Ambari	92
Investigator/Workflow	94
Interset User Interface X-Frame-Options	98
Firewall Configuration for Interset Servers	98
Network Topology Recommendations	99
Enabling SSL for Interset Reporting	100
Operating System User Permissions and Patch Levels	100
MD5 / SHA1 / SHA256 Checksums	100
Appendix D: Modify Kerberos Ticket	101
Modifying Kerberos Ticket Lifetime in Ambari	101
Modifying Kerberos Ticket Lifetime in Cloudera Manager	101
Appendix E: Change Passwords for a Secure Environment	102
Changing TLS Certificates Password	102
Re-Configure CDH/TLS by Changing Passwords	102
Create Key Tables	103
Changing Schema Registry Passwords	105
Changing Analytics Password	107
Reinstalling Kafka Topic	108
Changing Elasticsearch Passwords	109
Reconfiguring the Reporting Node	112
Changing ZooKeeper Configurations and Stopping Storm	112
Changing Workflow Passwords	113
Changing Nifi Passwords	116
System Configuration Files where Passwords Change	119
Updating NiFi Controller Services and Processors Configuration After System Password Rotation	121
Updating Interset Analytics Configuration After System Password Rotation	122
Appendix F: Migrate an Existing External PostgreSQL Server Database to Another Server	124
Backing Up and Restoring PostgreSQL Databases	124
Transferring the Database to Another Server	125

Index130

Introduction

This guide describes how to administer and configure the Intersect 5.9.2 system and its users. It also provides an overview of the Intersect cluster requirements, dependencies, components, and best practices information.

Intersect uses data science and advanced analytics to identify the top risky entities and behaviors occurring in your organization. Using your organization's data, Intersect establishes the *normal* behavior for your organizational entities and then, using advanced analytics, identifies the *anomalous* behaviors that constitute potential risks such as compromised accounts, insider threats, or other cyber threats.

Intersect's innovative user experience, true machine learning, and big data platform easily identify and prioritize high risk anomalies, allowing your security practitioners to instantly explore the underlying raw event data. The Intersect analytical models apply risk scores to individual users to provide security teams with relevant, prioritized information quickly enough to stop the activity before data loss occurs.

Intersect is a server-based product that is deployed in a clustered configuration. This means that the software is distributed across multiple machines, where each machine (which can be a physical machine or a virtual machine running on a VM server such as VMware ESX) is called a node. The distribution of load and responsibilities across multiple nodes is what makes the Intersect solution a scalable system that can handle large amounts of data: the more nodes in your deployment, the more data Intersect can handle.

Supported Environments

Intersect Analytics 5.9.2 is supported in the following x86_64 environments:

- CentOS 7.6
- Red Hat Enterprise 7.6

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

- Oracle OpenJDK 8u201/211
- Elasticsearch 6.8.1

Intersect 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 Customer Support at <https://softwaresupport.softwaregrp.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.

Installation

The Interset installation has several parts:

- installing and configuring components common to all machines in the Interaset configuration
- installing and configuring the software on the
 - Monitoring node
Where we refer to the Monitoring node in this document, we are referring to the node in your Interaset cluster where Ambari or Cloudera Manager is installed.
 - Master node(s)
 - Stream node(s)
 - Compute node(s)
 - Search node(s)
 - Reporting node
- tenant configuration

Additional Support

Should you experience unexpected results or identify issues that are not addressed in this document, please contact Micro Focus Customer Support at <https://softwaresupport.softwaregrp.com/>.

Administer the Interset System

This section describes how to administer Interset to ensure optimal system performance for end users. Topics include:

- [managing tenants and users](#)
- [configuring local multi-tenant authentication](#)
- [configuring new data sources](#)
- [configuring Analytics](#)
- [configuring the Search functionality](#)
- [configuring new data source action mappings](#)
- [resolving multiple entities in source data](#)
- [configuring the cluster to ingest universal third-party alerts](#)

Manage Tenants and Users

When you install Interset, the first tenant, tenant **0** (ID zero) is created by default. The users **admin** and **user** are created in tenant **0**, and configured with the login credentials **admin/password** and **user/password**.

You can create new users in this default Tenant **0**, or you can create a new tenant and add new users to that tenant. In either case, you will use the appropriate tenant ID in the sections below.

For users to log into the Interset user interfaces, they must have a user ID and password that can authenticate to the Interset system.



Tip: If you have established tenants from an earlier version of Interset, you must maintain the character case (upper case, lower case) of those tenant IDs (TIDs). You do not need to do this for new tenants.

You create both tenant records and user records on the Reporting node, as the Interset **root** user.

You can manage tenants and users on the **Tenants** page.



Tip: If you have LDAP or SAML configured, you must log in to Interset as the LDAP or SAML **"rootUser"** to administer tenants or users, instead of the Interset **root** user, in the steps that follow. For more information, please see "Configure LDAP Authentication" in the Interset 5.9.2 *Installation and Configuration Guide*.

Create a New Tenant

You can create a single tenant or import a CSV file to create multiple tenants.

To create a single new tenant:

1. Log in to Interset as **root** (the default password is **root**).
2. Click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, click **New**.
4. In the **Create a new Tenant** dialog, enter a new **Tenant ID** and **Tenant Name**.
5. Click **Create Tenant**. The new tenant appears in the tenant list.

To create multiple tenants:

1. Create and save a CSV file containing the list of tenants you want to create. The file must have the following format:

```
tenantId,name  
555,TenantName1  
556,Another New Tenant  
557,tenant3
```

2. Log in to Interaset as **root** (the default password is **root**).
3. On the **Tenants** page, click **CSV Import**.
4. In the **Upload a list of tenants** dialog, either drag-and-drop or browse to select the file you want to import.
5. Click **Upload**. The new tenants appear in the tenant list.



Note: Any time you configure a tenant to include a new data source, you must also configure Search. For more information, please see ["Configure a New Data Source" on page 16](#) and ["Configure Search" on page 29](#).

Remove a Tenant

1. Log in to Interaset as **root** (the default password is **root**).
2. Click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, select a tenant from the list on the left.
4. Click **Delete Tenant**.



Note: If you delete a tenant, you will also delete all related users.

5. Click **OK** to delete the tenant.

Create a New User

To create a new user:

1. Log in to Interaset as **root** (the default password is **root**).
2. If you are not presented with the **Tenants** page, click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, select a tenant from the list on the left, and then click **New User**.

4. In the **Create a User for <tenant name>** dialog, do one of the following:
 - For a local authentication environment, enter the **Name**, **Username**, **Role**, and **Password** for the new user.
 - For a SAML environment:
 - To create a local user, select the **Create a local user** checkbox. Enter the **Name**, **Username**, **Role**, and **Password** for the new user.
 - To create a SAML user, ensure the **Create a local user** checkbox is cleared. Enter the **Name**, **Username**, and **Role** for the new user. The **Username** must exactly match an existing SAML user.
5. Click **Create User**. The new user appears in the user list.



Ensure that the passwords you select for local users are very secure. These passwords are used as API keys, and there is no second factor authentication, so these passwords must be strong.

Modify a User's Role

Any new user who logs in to Intersect is created automatically in the Investigator server database. By default, any new user is assigned the **user** role, with the fewest privileges. You can modify a user's role on the **Tenants** page.

Steps

1. Log in to Intersect as **root** (the default password is **root**).
2. Click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, select a tenant from the list on the left. The users for the selected tenant are displayed in the main panel.
4. Double-click the **Role** for the desired user, and then click the down arrow to see the list of roles.
5. Click the required role to select it for the user.
6. Press **Enter** or click somewhere else on the page to save the selection.

Configure Multi-tenant Authentication

To enable user access to more than one tenant in the Intersect cluster, you can attach an existing user to another tenant.

To attach a user:

1. Log in to Interset as **root** (the default password is **root**).
2. If you are not presented with the **Tenants** page, click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, from the list on the left, select the tenant to which you want to attach the user, and then click **Attach User**.
4. In the **Attaching User for <tenant name>** dialog, enter the **Username** for the existing user from the previous tenant, and select the **Role** for this new tenant.
If the **Username** you enter does not exist in any tenant, an error message will appear.
5. Click **Attach User**. The new user appears in the user list, and now has access to both tenants.



Tips:


- If a user has access to the data for only one tenant, when they log in to Inter-set, they will see the data for that tenant.
- If a user has access to the data for more than one tenant, when they log in to Inter-set, they will see the data for the first tenant configured for their user profile. The tenants to which the user has access will appear in a dropdown menu to the right of the **Apps** field. To switch between tenants, the user simply selects the desired tenant from the dropdown menu.



- If a user with an **administrator** role switches to a tenant for which no data is available, the user is redirected to the **Settings** page.
- If a user with a **user** role switches to a tenant for which no data is available, the user is logged out.

Configure a New Data Source

With Interset 5.9.2, the data ingest process uses Apache NiFi to extract, transform, and load data. The following diagram illustrates the data flow in Interset 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 Customer Support at <https://softwaresupport.softwaregrp.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/`.

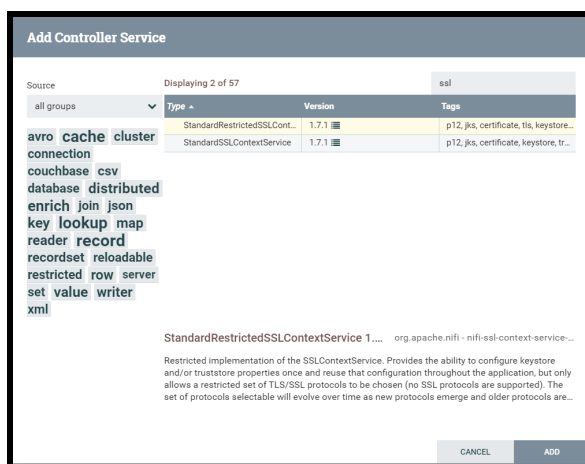


In a secure cluster, the NiFi URL is `https://<nifi_node_fqdn>:8085/nifi/`

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/interset/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/interset/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 **Interset 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 **Interset 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 **Interset 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 **Interset 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 - Interset 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/sslcert/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/sslcert/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 `interset_ad_events_0_0`
5. If you are running Interset on a secure cluster, set the properties required for Kerberos as follows:
 - **Security Protocol:** Select SSL protocol to communicate with brokers. Corresponds to Kafka's `security.protocol` property.
 - **Kerberos Credentials Service:** No value required.
 - **Kerberos Service Name:** No value required.
 - **Kerberos Principal:** No value required.
 - **Kerberos Keytab:** No value required.
 - **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
9. In a Web browser, navigate to `http://<ambari_fqdn>:8080`, and log in to Ambari using your login credentials.
10. Select the **Flume** service, and then click the **Configs** tab.

11. Click the **Group** dropdown arrow, and then select your Flume Configuration Group.
12. In the Flume agent **Configs** text box, paste the Data Ingest configuration information you created previously, and then click **Save**.



You may have to click **Override** to the right of the **Configs** text box to gain write access to the text box.

The new Data Ingest configuration contains all of the configuration information required to load your data from Kafka to Elasticsearch and to HBase.

13. Scroll to the top of the Flume page and then, in the upper right, click **Restart**, and then click **Restart All Affected**.
14. When prompted, click **Confirm Restart All**.

Configure Analytics

An **interaset.conf** file must be created for each tenant in your configuration.

Questions about the setting of these values and the impacts that are not covered in the content of this document can be directed to Micro Focus Customer Support at <https://softwaresupport.softwaregrp.com/>.

- **Analytics**

In the Analytics section of the **interaset.conf** file, you can configure each tenantID to reference a specific ZooKeeper instance.

```
tenantID = 0
# Update 'localhost' below to correct ZooKeeper URL
zkPhoenix = localhost:2181:/hbase-unsecure
```

- **Elasticsearch**

In the Elasticsearch section of the **interaset.conf** file, you can configure each Elasticsearch cluster to reference a specific Elasticsearch host.

```
# Elasticsearch cluster name
esClusterName=interaset
# Elasticsearch host
esHost=localhost
```

- If your Interset installation includes the 5.9.2 Endpoint Server, update the tenantID in the Elasticsearch index as follows:.

```
# If using Endpoint, update tenantID in index pattern below. Otherwise unused.
esIndexName=interaset_wdc_rawdata_0-*
```

- **Spark Runtime**

The Spark Runtime controls the tuning settings that are applied to Spark for the Analytics jobs. These settings should not be changed unless you are advised to do so by Micro Focus Interset Support.

```
# Number of spark tasks to split tasks by. Should be 2x - 3x cluster cores
parallelism = 32
# How many executors to request
numExecutors = 8
# How much RAM to request
executorMem = 4g
# How many cores per executor
executorCores = 2
```

Enable Windowed Analytics

After your InterSet cluster is installed and configured, your historical data has been ingested and analyzed, and data is appearing in the InterSet user interface, we recommend that you optimize the ongoing processing of data by enabling **Windowed Analytics**.

By default, InterSet is configured in batch mode, which means that when new data is ingested, Analytics is run on both the new and the existing data. Although this process is beneficial when you first install and configure your InterSet cluster (for testing and validation purposes), running Analytics on the entirety of your data on an ongoing basis unnecessarily uses system resources.

When you enable Windowed Analytics, you configure InterSet to run Analytics only on newly ingested data, as determined by the date of the last Analytics run and the timestamp of the data. InterSet identifies the data it has already analyzed, and then runs Analytics only on the new data. These results are then aggregated with the existing results to produce updated, current Analytics results for the entire data set.

Windowed Analytics has a positive impact on performance and stability because it allows the system to analyze and aggregate smaller, more consistently sized quantities of data than batch mode, particularly as the total amount of data in your system continues to grow.



Note: After you have validated the initial data ingest and Analytics run for your InterSet cluster, you might need to ingest and analyze historical data. In this scenario, you must continue to run Analytics in batch mode to ensure all the data is included.

Steps

1. Navigate to `/opt/interSet/analytics/conf/` and open the **interSet.conf** file.
2. In the **Analytics configuration** section, edit the **batchProcessing** parameter as follows:

Change

```
batchProcessing = true
```

to

```
batchProcessing = false
```

Windowed Analytics is now enabled, and will begin with the next Analytics run.



The first Windowed Analytics run for a tenant performs a full batch run to establish the baseline for the system going forward. The second and subsequent runs occur as Windowed Analytics.

Configure the 'Peek-Back' Window for Windowed Analytics

The 'peek-back' window is a best-effort buffer that ensures that delayed or out-of-order data is not missed between Windowed Analytics runs.

The value for the peek-back window is specified in milliseconds. The default value for the peek-back window is 24 hours (86400000.0 milliseconds).

This means that when a Windowed Analytics job runs, it loads data starting from the end of the last run, minus 24 hours. For example, if the previous windowed run started at midnight and completed successfully at noon today, the next run loads data beginning from noon yesterday. The intent of this overlap is to include data that was ingested after midnight, but that has event timestamps between noon and midnight of the previous day.

The peek-back window can be adjusted as follows:

1. As the **interaset** user, log in to the **Master** node where Interset is installed.
2. At the command console, run the following command:

```
/opt/interaset/analytics/bin/sql.sh --action console
```



Note: If you are using Kerberos, you must run `kinit` and specify a keytab/principal with appropriate permissions in HBase before you can successfully run this command.

3. In the Phoenix console, enter the following commands, replacing `VALUE_IN_MS_HERE` with the desired value of the peek-back window, specified in milliseconds (for example, 86400000.0):

```
UPSERT INTO PARAMETERS(TID, NAME, VAL) VALUES('0', 'SCORE WINDOW BUFFER', VALUE_IN_MS_HERE);
UPSERT INTO PARAMETERS(TID, NAME, VAL) VALUES('0', 'AGGREGATE WINDOW BUFFER', VALUE_IN_MS_HERE);
```

4. Type `!quit` to exit the console.

Configure Search

Any time you configure a tenant to include a new data source, you must also configure Search.



Important: You cannot configure Search until data has been ingested and Analytics has been run.

After running Analytics for a tenant for the first time, the search indexes must be set up. This step needs to be done only once per tenant; however, it must be done after the first Analytics run.

Setting up the Kibana indexes will enable the Search feature for the different data types.


For this section, you need to know:

- the tenant ID (TID, the default is 0)
- the data sources being used

The data sources will always include violations (from Workflow), and one or more of Active Directory event logs, repository data, printer logs, and so on.

Steps

- Open a Web browser, and go to **`https://<reporting_fqdn>/search/<tid>`**
- In the Kibana home page, click **Connect to your Elasticsearch index**.



- On the **Create index pattern** page, under **Step 1 of 2: Define index pattern** in the **Index pattern** box, enter the index name from the table below, replacing **<tid>** with your tenant ID.

Although earlier releases allowed uppercase characters in the **<tid>**, the **<tid>** must be lower-case characters.

Data Type	Index Name	Timestamp Field
All	intersect_*_rawdata_<tid>	timestamp
Access	intersect_access_rawdata_<tid>	timestamp
Access, Active Directory, Authentication auditD, VPN	intersect_auth_rawdata_<tid>	timestamp
auditD	intersect_auditd_rawdata_<tid>	timestamp
Authentication	Intersect_authentication_rawdata_<tid>	timestamp
Email	intersect_email_rawdata_<tid>	timestamp
Finance Expense	intersect_expense_rawdata_<tid>	timestamp
Intersect EDR Server	intersect_sensor_rawdata_<tid>	timestamp
NetFlow	intersect_netflow_rawdata_<tid>	timestamp
Printer	intersect_printer_rawdata_<tid>	timestamp

Data Type	Index Name	Timestamp Field
Repository	interiset_repo_rawdata_<tid>	timestamp
Universal Alert	interiset_violations_<tid>	timestamp
VPN	Interiset_vpn_rawdata_<tid>	timestamp
Web Proxy	interiset_webproxy_rawdata_<tid>	timestamp
Workflow Violations	interiset_violations_<tid>	timestamp

**Notes:**

- The Workflow Violations **interiset_violations_<tid>** index will not exist unless a violation was triggered by a Workflow.
- You will always need to add the Workflow index, and one or more of the other data types.
- If Kibana does not recognize the index pattern, do not create it at this time.
- When you configure the **All** data type **interiset_*_rawdata_<tid>** index, you are creating a search index for queries that run against all ingested data types. This type of search index is useful when, for example, you want to see all the Analytics data for one specific entity (user, machine, controller, and so on) across all data types.



Tip: When you enter your index pattern name, Kibana provides an indication whether your index name matches existing indexes.

4. Click **Next step**.

Create index pattern

Kibana uses index patterns to retrieve data from Elasticsearch indices for things like visualizations.

☐ Include system indices

Step 1 of 2: Define index pattern

Index pattern

You can use a * as a wildcard in your index pattern.
You can't use spaces or the characters \, /, ?, ", <, >, |.

> Next step

✓ Success! Your index pattern matches 4 indices.

interaset_ad_rawdata_tst-2017-03-w13
interaset_ad_rawdata_tst-2017-04-w14
interaset_ad_rawdata_tst-2017-04-w15
interaset_ad_rawdata_tst-2017-04-w16

Rows per page: 10 ▾

5. Under **Step 2 of 2: Configure settings**, in the **Time Filter field name** box, click the dropdown arrow, and then select the appropriate **Timestamp Field** value from the table above.

6. Click **Create index pattern**.

Create index pattern

Kibana uses index patterns to retrieve data from Elasticsearch indices for things like visualizations.

☐ Include system indices

Step 2 of 2: Configure settings

You've defined `interiset_ad_rawdata_tst` as your index pattern. Now you can specify some settings before we create it.

Time Filter field name

Refresh

timestamp


The Time Filter will use this field to filter your data by time. You can choose not to have a time field, but you will not be able to narrow down your data by a time range.

[Show advanced options](#)

[< Back](#)[Create index pattern](#)

Kibana displays a summary of the new index pattern, including a list of all of the data fields, each field's core type, and whether the field is searchable and aggregatable. To the left of the summary table, Kibana provides a list of all configured index patterns.

7. For each subsequent data type, click **+Create Index Pattern** to the left of the index pattern summary table, and then repeat Steps 3 through 7.
8. After creating index patterns and thus enabling Search for one or more data types, click **Discover** in the Kibana sidebar menu to begin searching.



Configure New Action Mappings for Existing Data Sources

1. Open a Web browser, and sign in to the Interset server as an Administrator.
2. On the **Overall Risk** page, click the **Settings** icon  in the upper-right corner to open the **Settings** page.
3. On the **Settings** page, click **Map Actions** to launch the **Action Mapping** wizard.
4. On the **Action Mapping** page, under **Add New Actions**, click the **Data Type** dropdown arrow to select the data source type.
5. Under **Action**, enter the new data source action value you want to map.
6. Under **Mapping**, click the dropdown arrow to select the schema action value to which the new data source action will be mapped.
7. Click **Add**  to continue adding and mapping new data source action values until your list is complete.



Tip: To remove an existing mapped action, click **Remove** .

8. Click **Add**  to save the new action mappings.

For all future ingests of the selected data type, the new action mappings will be applied.



Tip: You can update the data source action mappings by following these same steps.

Resolve Multiple Entities in Source Data

This section describes how to resolve entities when one user entity appears with more than one `<username>` value in your source data. For example, if John Smith registers in your source data as both `jsmith` and `jsmith@yourcompany.com`, you will want to have the analytics for John Smith consolidated as one complete story, rather than as two separate ones.

The resolution of these multiple entities representing the same user must occur prior to data ingest. Therefore, redefining these entities and mapping one to the other is carried out using an expression in the extract processors in NiFi to match the entities as required.

For example to change the `user` field to remove the email domain, do the following:

1. Identify the NiFi processor you wish to modify (you might need to enter a process group to find the required processor).
2. Right-click the processor, and select **Configure**.
3. Select the **Properties** tab, and click the plus sign to add a new property. The **Add Property** dialog appears.
4. In the **Property Name** text box, enter a forward slash followed by the name of the field you want to update. In this example, we enter `/user`.
5. Enter the expression as required to modify the field, and then click **OK**. In this case, we enter `${field.value:substringBefore('@')}`. This expression removes all the text in the `user` field after (and including) the '@'.
6. In the **Configure Processor** dialog, click **Apply** to save the new property.

The next time your source data is ingested, each user with an email domain will be converted to just the user name portion. In our example `jsmith@yourcompany.com` becomes `jsmith`.

For more information, see the Intersect 5.9.2 *Data Ingest Guide*.

Configure Intersect to Ingest Universal Third-party Alerts

When you ingest third-party Universal Alert data as violations, these alerts can then be analyzed by Intersect in the context of the entirety of your organizational data. These Universal Alert violations are similar to other Intersect violations, forming part of the hourly risk information and appearing as anomalies in the Intersect user interface.

Examples of supported third-party alert sources include DLP products, malware alerts, networking alerts, and so on. The third-party Universal Alert data must be in .csv format.

Ingesting third-party Universal Alert data as violations involves:

- using the API to first define the columns in your third-party Universal Alert data, and then to map those columns to the schema;
- using Workflow to create violations based on conditions you define; and
- updating the Universal Alert Data Ingest Configuration information and ingest the third-party alerts.



Tip: If you ingest third-party Universal Alert data using NiFi, you must ensure that the field names of the schema specified in the NiFi CSV reader match the column headings of the source data. The field names in the NiFi schema must also match the Workflow column definitions, either set by default or modified using the `/rules/{tid}/config` endpoint in the API.



Important: When Universal Alert data is initially ingested, these alerts have a risk score of zero (0). To see these new Universal Alert violations and for them to reflect a risk score other than zero, you must create one or more Workflows to define new violations with an assigned severity.

Use the API to Define and Map the Universal Alert Columns

After your Universal Alert data has been exported to .csv format, you can use the API to identify the columns in your data; and then to map those columns to the schema upon ingest.

Interset provides a number of predefined column definitions in the API, as identified below. If you require column definitions beyond those provided by Interset, you have the flexibility to add those definitions directly in the API.

Column Heading	Value Type	Display	Description
user	string	user	The user involved in the event.
source	string	source	The hostname, machine, and/or system from which the event was generated.
eventUUID	string	event UUID	Unique identifier of the alert event.
destination	string	destination	The resource the client is trying to access (e.g., /assets/cat.gif).
eventType	string	event type	The type of event.
application	string	application	The application associated with the event.
action	string	action	The outcome of the event (e.g., success, failure).
message	string	message	The text of the message in the event.
duration	string	duration	The time spent servicing the event, in milliseconds
filename	string	filename	The name of the file associated with the event.
projectname	string	project name	The name of the project associated with the generated alert event.
severity	string	severity	The severity of the associated event.
signature	string	signature	The Windows event code associated with the event (e.g., 4624).
size	number	size	Total size of the file associated with the event, in bytes.
tags	string	tags	Metadata associated with events or alerts.
vendorProduct	string	vendor product	The reporting software. For example, McAfee DLP.

Column Heading	Value Type	Display	Description
status	string	status	The resolution status of the alert (for example, resolved, under investigation, false positive, etc.).
sourceIp	string	source IP	The client IP associated with the event.
destinationIp	string	destination IP	The destination IP the client is trying to access.

Steps

1. In a Web browser, log in to the Reporting server as an Administrator.
2. On the **Overall Risk** page, click **Settings**  and then, in the dropdown list, select **API**.
Swagger opens in a new browser window.
3. In Swagger, expand **rules{tenantId}**, and then expand **GET /rules/{tid}/config**.
4. In the **Parameters** section, enter a tenant ID, and then click **Try it out!**.
When the response appears, you can explore the predefined Universal Alert data column definitions in the **Response Body** section.
5. To edit or extend the predefined column heading parameters to align them with the column headings of your organization's Universal Alerts, do the following:
 - copy all of the text in the **Response Body** section to a text editor;
 - in the text editor, update the parameter information as required;
For example, if your Universal Alert data column heading is "username" instead of "user" as currently defined in the API, replace "user" with "username".
 - If you have Universal Alert data column headings that do not appear in the predefined parameters, you can create a new mapping by creating a new JSON block. For example:

```
{
  "name": "bus_unit",
  "type": "string",
  "display": "Business Unit"
}
```



Tips:

- The **"name"** represents the value of your Universal Alert data column heading as it is defined in your ingest configuration file.
- The **"type"** is either **'string'** or **'number'**.
- The **"display"** value is the name of your of your Universal Alert data column heading in the Workflow user interface.

6. When you have finished updating the column heading parameters, while still in the text editor, scroll down to **"ViolationsMapping"** section, and update the schema mapping for each parameter that you modified.

For example, if you changed "user" to "username", update the **"userId"** schema parameter from "user" to "username".

- copy all of the text from the text editor;
 - in Swagger, expand the **PUT /rules/{tid}/config** section, and paste the copied text in the **body** text box in the bottom half of the page;
 - enter the tenant ID in the **tid** box; and
 - Click **Try it out!** to save your changes.
7. To verify your changes, expand the **GET /rules/{tid}/config** section, enter the tenant ID in the **tid** box, and then click **Try it out!** to view the updated Universal Alert column headings.

Use Interset Workflow to Create Violations with Risk Weights

1. In the **Overall Risk** page, click **Settings**  and then, in the dropdown list, select **Workflows**.
2. At the top right of the **Current Workflows** page, click **Create Workflow** to open the **Create a Workflow** page.
3. On the **Create a Workflow** page, in the **Data Source** section, select **Universal Alert**.



Note: The Universal Alert has only one **Detect** option and one **Trigger** option available to choose. Select these options to proceed to the **Conditions** section.

4. In the **Conditions** section, select your conditions and then click **Save**.



Tip: The selection list contains the column headings you defined in the **PUT /rules/{tid}/config** section of the API in Swagger.

5. Confirm your selected conditions. To add more conditions, click **add a new condition....**
When you are satisfied with your conditions selection, choose **and** or **or** next to each condition to combine them appropriately.
6. In the **Outcomes** section, select **flag as a violation**.
7. In the **Set severity to** field, select the severity level you want applied to the violation from the drop down menu (**none**, **low**, **medium**, **high**, or **extreme**) then, click **Save**.
This severity level will be attached to the risk score of the new violations.
8. To add more outcomes, click **add a new outcome....**
9. Review your Workflow and make any changes necessary.
10. When you are satisfied with your Workflow, enter the name of your Workflow in the **Workflow Title** field at the top of the page, then click **Save Changes**.
The Workflow appears in the **Current Workflows** list on the **Workflows** page.
11. To activate the Workflow, on the **Workflows** page, in the **Activate** column, click the slider that corresponds to your Workflow.



Tip: You can sort your Workflows using the arrows at the top of the **Name**, **Creation Date**, **Last Modified**, or **Activate** columns.

Update the Universal Alert Data Ingest Settings

Intersect uses Apache NiFi for data extraction, transformation, and loading.

If you need to change any settings related to ingesting Universal Alert data, follow the instructions in the "Ingesting Data with Apache NiFi" section of the *Intersect 5.9.2 Data Ingest Guide*, or contact Intersect Customer Support for assistance.

Administer Interset for End Users

There are many tasks that you can perform as the Interset Administrator to ensure that the Analytics end users have access to the information they need, when they need it.

These tasks include:

- [creating Workflows](#) for notifications, violations, and other outcomes
- [managing bots and bot-like users](#)
- [tuning analytics](#)

Create Workflows

Workflows are an extremely valuable tool for ensuring that the Interset Analytics consumers, who rely on the data to identify the top risky entities and behaviors occurring in the organization, can be immediately aware of specific areas of potential concern and that appropriate, proactive measures may be undertaken in a timely manner. With Workflows, Interset Administrators can build custom use cases to highlight specific information for their end users, based on a wide range of criteria.

When Workflows are configured and include email or SMS notifications, and REST API calls, these notifications and corrective actions occur instantly, in real time. When Workflows include the creation of violations or changes in severity level, these outcomes are available the next time Analytics runs on the data.

For example, a Security team may wish to be immediately notified, and have initial corrective measures put in place any time an individual inappropriately attempts to access a restricted project repository. By creating Workflows based on specific conditions, the Security team can easily configure real time email notifications, and even control access privileges using REST APIs when those conditions are met.

In the Workflow page, dropdown menus and dynamic options facilitate building the Workflow definitions.

Only users with Administrator privileges can create workflows.

When creating Workflows, there are a number of things you will want to learn:

- that the [Workflow Engine is running](#)
- how to [configure Workflow for different data sources](#)
- the [criteria and/or conditions](#) that warrant Workflows
- [how to alert your end users](#) to risky entities and behaviors
- how to [scale your Workflows](#)

At the end of this section is a [Workflow example](#) that will guide you as you create your own Workflows.

Verify the Workflow Engine is Running

Before Administrators begin working with Workflows, you should confirm that the Workflow created for that tenant is running.

1. On the Master node, run the following status command as the **interaset** user:

```
/opt/interaset/rules/bin/workflow.sh --status /opt/interaset/rules/conf/rules.conf
```

The expected response includes information on Topology_name, STATUS, Num tasks, Num_workers, and Uptime-secs.

2. On the Master node, run the following validation command as the **interaset** user:

```
/opt/interaset/rules/bin/workflow.sh --validate /opt/interaset/rules/conf/rules.conf  
/opt/interaset/rules/conf/kafka-config.properties
```

Additional ways to verify that Workflows are processing data include:

- Creating a Workflow that includes setting a violation that is sure to be triggered, ingesting the data, and then running the following command on the Data node to see if messages are showing up in the Kafka topics:

```
/usr/hdp/current/kafka-broker/bin/kafka-console-consumer.sh --from-beginning --max-messages 10 --zookeeper  
localhost:2181 --topic interaset_violations_<TID>
```

- Using the Phoenix sqlline utility to confirm that data is being written to HBase:

```
phoenix-sqlline
```

and then running the following command:

```
SELECT * FROM VIOLATIONS WHERE TID='<TID>'
```



Tip: Workflows are configured on a per-tenant basis. When your configuration includes multiple tenants, each tenant must have its own Workflow engine and associated Workflow configuration file (for example, **rules-tid-0.conf**).

Configure Workflow for Different Data Sources

To create Workflows for data sources other than Active Directory authentication, you will need to configure the data source in the Workflow **rules.conf** configuration file.



Important: Do not delete any of the lines in the **rules.conf** files, even if they are not relevant to your tenant or Workflow.

To make a new data source available in Kafka (for use in Workflow), you must configure a NiFi processor to ingest the data and create the Kafka topic for that data source on the Stream node. The process groups in the Interaset Marketplace template each include a processor that writes data to Kafka for use in Workflow. For more information, see the *Data Ingest Guide*.

In addition, because Workflow runs against data as it is ingested, you cannot create Workflows on historical data. To ensure the Workflow results you want, ensure that the Kafka topics exist, and the Kafka Spout topics are enabled, before ingesting your data. The Kafka and Kafka Spout topics required by Interaset, along with the default tenant and data source IDs (tid, did), are listed in the table below.

Data Source Type	Kafka Topic	Kafka Spout Topic	Default did	Default tid
Authentication events	interiset_auth_events_<did>_<tid>	#KafkaSpoutAdTopics = <kafka_topic>	0	0
Repository events	interiset_repo_events_<did>_<tid>	#KafkaSpoutRepoTopics = <kafka_topic>	0	0
AuditD events	interiset_auditd_events_<did>_<tid>	#KafkaSpouAuditdTopics = <kafka_topic>	0	0
Web Proxy events	interiset_webproxy_events_<did>_<tid>	#KafkaSpoutProxyTopics = <kafka_topic>	0	0
Windows Printer events	interiset_printer_events_<did>_<tid>	#KafkaSpoutWindowsPrinterTopics = <kafka_topic>	0	0
InterSet Violations	interiset_violations_<tid>	n/a	n/a	0
NetFlow	interiset_netflow_events_<did>_<tid>	#KafkaSpoutNetFlowTopics = <kafka_topic>	0	0
Finance Expense	interiset_expense_events_<did>_<tid>	#KafkaSpoutExpenseTopics = <kafka_topic>	0	0
Email	interiset_email_events_<did>_<tid>	#KafkaSpoutEmailTopics = <kafka_topic>	0	0
Universal Alert	interiset_alert_events_<did>_<tid>	#KafkaSpoutGenericAlertsTopics = <kafka_topic>	0	0

Steps

1. Navigate to the **Stream** node, and run the following command to list the Kafka topics created for your InterSet cluster:

```
/usr/hdp/current/kafka-broker/bin/kafka-topics.sh --zookeeper <master_node_fqdn>:2181 --list
```

The response returned will include all the existing Kafka topics. If the Kafka topic for your data source is returned, proceed directly to Step 3.

2. If the required Kafka topic is not returned, run the following command to create the Kafka topic:

```
/usr/hdp/current/kafka-broker/bin/kafka-topics.sh --zookeeper <master_node_fqdn>:2181 --create --topic \"<kafka_topic>\" --partitions 8 --replication-factor 1
```

- Run the following command, this time to ensure the topic was created:

```
/usr/hdp/current/kafka-broker/bin/kafka-topics.sh --zookeeper <master_node_fqdn>:2181 --list
```

The new Kafka topic will appear in the list.

3. On the Master node, navigate to the **/opt/interiset/rules/conf** directory, and open the **rules.conf** configuration file.
4. If you plan to use the schema registry, do the following:

- In the **Service Connections** section of the **rules.conf** file, under **Schema Registry**, verify that the **ConfluentRegistry** parameter is not commented out.

The **ConfluentRegistry** parameter must be enabled 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, 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. If you do not plan to use the schema registry, do the following:

- In the the **Service Connections** section of the **rules.conf** file, under **Schema Registry**, comment out the **ConfluentRegistry** parameter.
- In the **Event Processing** section, update the **non-Schema Registry** path parameters by:

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

- removing the hash (#) symbol from the **TopicEvents** parameter 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

Data Source Type	<archType>
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

- To configure Workflow for an authentication data source other than Active Directory, do the following:

- Locate the KafkaSpoutTopics parameter.
- Edit the **KafkaSpoutTopics** parameter to include the spout topics for the other authentication data source as a comma-separated list.

For example, to enable auditD and VPN, add the following spout topics:

```
KafkaSpoutTopics = interaset_auditd_events_0_0, interaset_vpn_events_0_0
```

and then save the **rules.conf** file.

- Still on the Master node, run the following commands to kill the currently-running topology, and then redeploy it:

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

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

Use these same steps to enable other datasources such as repository, WebProxy, and so on.

- To configure Workflow email notifications, do the following:

- In the **#Notification Settings** section, change the value of the **EmailOutputEnabled** parameter to **true**.

```
EmailOutputEnabled = true
```

- Enter the appropriate email configuration information:

```
EmailServer = <email_server_name>
EmailServerPort = <email_server_SMTP_port>
EmailSetSSLOnConnect = true
EmailUser = <email_login>
EmailPassword = <email_password>
EmailFromAddress = <your_FROM_address>
```

- To receive text message notifications via your smart phone, enter your credentials as follows:



Tip: You must have a Twilio account (www.twilio.com).

```
SmsSid = <your_twilio_SID>
SmsToken = <your_twilio_API_token>
SmsFromNumber = <your_twilio_FROM_number>
```

8. Save the updated **rules.conf** file.
9. To deploy the updated Workflow configuration, run the following command:

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

10. To ensure that the current topology is running, run the status command:

```
/opt/interaset/rules/bin/workflow.sh --status /opt/interaset/rules/conf/rules.conf
```

The expected response includes information on **Topology_name**, **STATUS**, **Num_tasks**, **Num_workers**, and **Uptime-secs**.

11. To validate the currently running topology, run the following command:

```
/opt/interaset/rules/bin/workflow.sh --validate /opt/interaset/rules/conf/kafka-config.properties
/opt/interaset/rules/conf/rules.conf
```

Use these same steps to enable other data sources for Workflow, such as repository, WebProxy, and so on.

Configure Workflow for a Custom Schema

If you have defined a custom schema in the Confluent Schema Registry, you can reference that schema in Workflow by updating the **rules.conf** file.

To refer to a custom schema in **rules.conf**:

1. Determine the name of the custom schema, as defined in the **name** field of the Avro schema. For example in the email schema, the schema name is **EmailRecord**, as shown below:

```
{
  "namespace": "com.interaset.schema",
  "name": "EmailRecord",
  "type": "record",
  "fields": [
    {
      "name": "timestamp",
      "type": "string",
      "doc": "ISO 8601 Timestamp, with millisecond resolution. Eg '2016-01-01T13:00:00.000-05:00'."
    },
    {
      "name": "senderEmailID",
      "type": "string",
      "doc": "An identifier for the user who sent an email."
    },
    . . .
  ]
}
```

2. Determine the appropriate archetype for the data you plan to ingest. The archTypes are listed in the table in ["Configure Workflow for Different Data Sources" on page 41](#). For example, for the email schema, the archType is **email**.
3. Open the **rules.conf** file for editing, and add a new line with the following format:

```
SchemaSubjectEvents.<schema_name>=<archType>
```

4. For example, if you created a custom repository schema with the name **MyRepoRecord**, using the

repo archType, you would add the following line:

```
SchemaSubjectEvents.MyRepoRecord=repo
```

Workflow Scenarios

Workflow offers a full range of criteria that you can apply to define the conditions to trigger a notification, a violation, or a REST API call to implement a corrective action.

Working together with your organization's Security specialists, you define the circumstances that require Workflows:

- Workflows can be created and applied to the following types of data: Active Directory log, Linux Audit (auditd), Web Proxy, Intersect Analytics, Printer, Repository log, NetFlow, Finance Expense, and Universal Alert.
For example, specific project repositories may be highly sensitive and require that the Security team be alerted to activity by users not belonging to a particular work group. Or, Security team members may wish to be made aware of any user risk score above an identified threshold. These are ideal circumstances to implement a Workflow.
- Depending on the data source, a Workflow can be applied to the following entities: a user, a user's login, a machine, or a project.
Isolating activity to a specific entity in a Workflow virtually allows your Security team to be alerted to, and corrective actions implemented when the defined criteria relative to that entity are met.
- Workflow action Triggers are specific to each data source.
For example, for repository log data, you can select the Trigger **retrieving data.**; for Sensor data, you can select **attaching a file to an Outlook email** as a Trigger.

For more information about defining the data source, entity, and triggers in Workflows, see ["Workflow Example" on page 50](#).

Workflow Notifications and Outcomes

Working together with your organization's Security specialists, you can configure a number of different notifications and outcomes to occur when the criteria defined in the Workflow are met:

- **Flag as a violation**
Violations are special alerts designed to identify specific risky behaviors. When you include **Flag as a violation** as a Workflow outcome, you also define the associated severity of that violation. The violation and associated severity are then factored into the Analytics and appear in the **Matrix of Anomalies & Violations** the next time Analytics is run.
Flagging as a violation is not a supported outcome when the Analytics is the data source. This is because it could inappropriately create a circular dependency.
To flag a Workflow outcome as a violation, on the **Actions** page, select **flag as a violation** and then set the violation severity level: **low**, **medium**, **high**, or **extreme**.

- **Change the importance**

Changing the importance of an entity is designed to make that entity more sensitive to anomalies. When two entities trigger the exact same anomaly, the entity with higher importance will end up with a higher risk score and higher risky hours scores than the entity with low importance.

Changing the importance is not a supported outcome when Interset Analytics is the data source. This is because it could inappropriately create a circular dependency.

To create a Workflow outcome that changes the importance of the entity, on the **Actions** page, select **change the importance** radio button. Then, specify the entity and the new level of importance.

- **Send an email**

Configure an email to be sent to alert one or more recipients when the Workflow criteria are met.

To create a Workflow outcome that sends an email to one or more recipients, on the **Actions** page, select **send an email** radio button. Then, complete the required fields to provide the email information.

- **Send an SMS**

Send a text message to alert one or more recipients when the Workflow criteria are met.

To create a Workflow outcome that sends a text message, on the **Actions** page, select **send an SMS**. Then, complete the required fields to provide the SMS information.

Note: When you complete the fields for SMS or email notification, any backslash gets converted to a front slash when the fields get converted to outputs. This is to ensure that the formatting is maintained in the output.

- **Call a REST API**

Call a REST API to implement an action when the Workflow criteria are met.

For example, configure a Workflow to restrict access permissions when the risk score for any user with Administrator privileges is higher than 50.

To create a Workflow outcome that calls a REST API, on the **Actions** page, select **call a REST API**. Then, complete the required fields to call the REST API.

- **Send to McAfee DXL**

Send a notification to the McAfee Data Exchange Layer (DXL) to alert systems and/or applications when the Workflow criteria are met.

To create a Workflow outcome that sends a notification to the McAfee DXL, on the **Actions** page, select **send to DXL**. Then, complete the required fields to send the notification to DXL.

- **Export as CEF messages to an ArcSight syslog server**

Send a notification as a CEF export to ArcSight syslog servers to inform users when the Workflow criteria are met.

To create a Workflow outcome that exports a CEF message to an ArcSight syslog server, on the **Actions** page, select **send to CEF syslog**. Then, complete the required fields to export the CEF message.

- **Send to Splunk**

Send a notification to Splunk to alert systems and/or applications when the Workflow criteria are met.

The Workflow engine must have been configured to send notifications to Splunk. For more information, in the "Configure Workflow" chapter of the Interset 5.9.2 *Installation and Configuration Guide*, see "Configure Splunk".

To create a Workflow outcome that sends a notification to Splunk, in the **Outcomes** section, select **send to Splunk**. Then, complete the required fields to send the notification to Splunk.

- **Send to Phantom**

Send a notification to Phantom to alert systems and/or applications when the Workflow criteria are met.

The Workflow engine must have been configured to send notifications to Phantom. For more information, in the "Configure Workflow" chapter of the Interset 5.9.2 *Installation and Configuration Guide*, see "Configure Phantom".

To create a Workflow outcome that sends a notification to Phantom, on the **Actions** page, select **send to Phantom**. Then, complete the required fields to send the notification to Phantom.

For more information about defining outcomes in Workflows, see ["Workflow Example" on page 50](#).

Scale Your Workflows

To facilitate the implementation of Workflows that apply to multiple entities, Interset allows you to define multiple entities of the same type using a separate file. For example, rather than having 20 individual Workflows to highlight each time one of 20 administrator users fails to authenticate, you can now have one Workflow that uses the data uploaded from a text file that, in turn, lists the 20 administrator users.

Similarly, you may want to create a Workflow in which a condition may have multiple values. For example, you want to be alerted whenever a specific user account accesses one of several destination IP addresses via Web proxy. Rather than having separate Workflows to identify each individual destination IP address, you can now have one Workflow that uses the destination IP addresses uploaded from one text file.

Each of these Workflow list files must:

- be tenant-specific;
- reside in a location accessible to the Interset cluster;
- have no more than one entry per line;
- use file names that do not contain spaces or any of the following characters: `/`, `:`, `.`, or `\\`;
- be in .txt. format; and
- contain data values that respect the format of the ingested data.

In addition, the Workflow list file can:

- contain an unlimited number of entries (for example, users); and
- be updated with new information as required.

You must give your Workflow list a name. Interset uses this name internally to store the list. The Workflow list name:

- must be tenant-specific;
- must not contain spaces or any of the following characters: `/`, `:`, `.`, or `\\`; and
- does not have to match the list file name.



Important: When you configure a Workflow to use a list file, you must type the list name into the **Enter a list name** field. As a result, it is important that you maintain a record of your Workflow lists.

Manage Workflow Lists

To manage your Workflow lists, you can run the following commands from the `/opt/interiset/rules/bin` directory on the **Master node** where Analytics is installed:

- To upload a Workflow list file for use in your Workflow:

```
./workflow.sh --createList /opt/interiset/rules/conf/rules.conf <ListName> <path/to/List/List_file_name.txt>
```

- To update an existing Workflow from a Workflow list file:

```
./workflow.sh --updateList /opt/interiset/rules/conf/rules.conf <ListName> <path/to/List/List_file_name.txt>
```

- To show all the Workflow lists for the tenant:

```
./workflow.sh --showLists /opt/interiset/rules/conf/rules.conf
```

- To show the contents of a Workflow list:

```
./workflow.sh --catList /opt/interiset/rules/conf/rules.conf <ListName>
```

- To delete a Workflow list:

```
./workflow.sh --deletelist /opt/interiset/rules/conf/rules.conf <ListName>
```

- To upload all the Workflow list files in the specified directory (defaults to `/opt/interiset/rules/lists/`):

```
./workflow.sh --createAllLists /opt/interiset/rules/conf/rules.conf <path/to/List_directory>
```



Workflow uses the filenames (up to the `'.'` character) in the specified directory as the names of the Workflow lists it creates. Therefore, the file names must comply to the naming requirements for Workflow lists.

- To update all the Workflow lists corresponding to the files in the specified directory (defaults to `/opt/interiset/rules/lists/`):

```
./workflow.sh --updateAllLists /opt/interaset/rules/conf/rules.conf <path/to/List_directory>
```



Workflow uses the filenames (up to the '.' character) in the specified directory to identify the Workflow lists to update.

- To delete the (previously uploaded) Workflow lists corresponding to the files in the specified directory (defaults to **/opt/interaset/rules/lists/**):

```
./workflow.sh --deleteAllLists /opt/interaset/rules/conf/rules.conf <path/to/List_directory>
```



Workflow uses the filenames (up to the '.' character) in the specified directory to identify the Workflow lists to delete.

Workflow Example

The following example is provided to guide you as you create your own Workflows.

Workflow Example: Highlight Risk on a High Priority Project

ABC Company has a highly sensitive, high priority project entitled Pilgrim. The Pilgrim project files are stored in a Perforce repository with access restricted to four (4) ABC Company employees. Because of the extreme sensitivity of the Pilgrim project, the ABC Company Security team has decided to create a Workflow to ensure that, on an ongoing basis, they are advised if any entity other than the four authorized users accesses the Pilgrim project files. In addition, they decide to include Workflow outcomes that result in the creation of a violation, with a severity level of Extreme when an unauthorized entity attempts to access Pilgrim files.

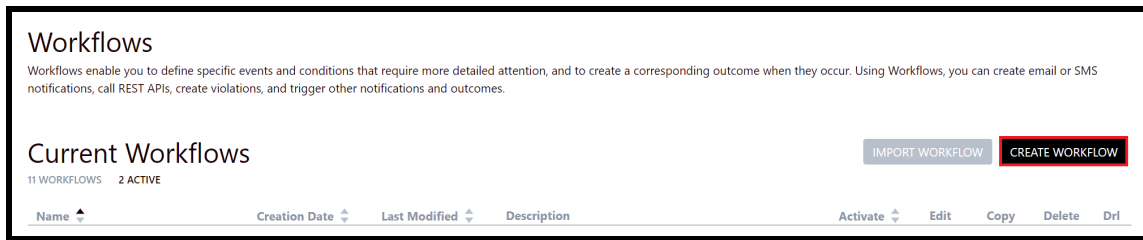


Tips:

- From the **Current Workflows** page, on the right side, you can click the **DRL** button to reveal the Workflow code of a saved Workflow.
- All string comparison operators except **contains** are case sensitive.
- When you create a Workflow containing a backslash ('\'), you must escape this character for the condition to execute properly. To escape the backslash character, enter two backslashes instead of one.

Steps

1. On the right of the **Workflows** page, click **Create Workflow** to open the page to create a Workflow.



2. At the top of the **Workflow** page, in the **Workflow Title** area, type a name for the new workflow.
3. In the **Workflow Description** area, type a description of the new workflow.
4. In the **Data Source** section, from the **Events** drop-down menu, select **Repository**.
5. In the **Detect** section, from the **Detect** drop-down menu, select **someone**.
6. In the **Trigger** section, from the **Trigger** drop-down menu, select **done anything**.
7. In the **Conditions** section, click **add a new condition**, and then from two subsequent **Condition** drop-down menus, select **project** and then select **contains**.
8. From the **List** drop-down menu, select **string**.
9. In the **Enter a substring...** field, type **Pilgrim**, then click **Save**.
The condition **project contains "Pilgrim"** appears in the **Conditions** section.
10. In the **Conditions** section, click **add a new condition**, and then from two subsequent **Condition** drop-down menus, select **user** and **does not contain**.
11. From the **List** drop-down menu, select **string**.
12. In the **Enter a substring...** field, type **Chelsea**, then click **Save**.
The condition **user does not contain "Chelsea"** now also appears in the **Conditions** section, with a new drop-down menu to **and** or **or** the multiple conditions.



Tip: Chelsea is one of the four ABC Company employees authorized to access the Pilgrim project files.

13. From the **or** drop-down menu, select **or**.
14. Repeat Steps 8 though 10 for the remaining three (3) ABC Company employees authorized to access the Pilgrim project files: **Andrew**, **Satoru**, and **Priya**.
The Workflow is now configured to identify any user that accesses the Pilgrim files who is not authorized to do so.
15. In the **Outcomes** section, click **add a new outcome...**, and then from the **Outcome** drop-down menu, select **send an email**.
 - In the **To** field, type **security@abccompany.com**
 - In the **Subject** field, type **Pilgrim**
 - In the **Email Body** field,

- type **Unauthorized activity in Pilgrim project**:
 - type a { character to insert a property and then, from the drop-down list, select **user**
 - In the **Choose the frequency** drop-down list, select **daily**.
16. In the **Outcomes** section, click **add a new outcome....**
 17. From the drop-down list, select **flag as a violation**, and then set the severity to **extreme**.



Important: Violations that have a severity level of extreme will, on their own, create high risk hours in the Analytics. Setting too many violations at this severity level may create more high risk hours than can be consumed.

18. Click **Save** to save the new outcome.
19. (Optional) Click **Test Run** to validate the workflow. Test Run helps you ensure that violations and other outcomes trigger as expected. It indicates an **X** if the event does not trigger the workflow. In such cases, review the event and conditions to fix any errors.
20. The Workflow appears in the **Current Workflows** list on the **Workflows** page.
21. To activate the Workflow, on the **Workflows** page, in the **Activate** column, drag the slider that corresponds to your Workflow to the right.



Tip: You can sort your Workflows using the arrows at the top of the **Name**, **Creation Date**, **Last Modified**, or **Activate** columns.

Result

One month following the implementation of the **Pilgrim** Workflow, the following occurred:

- **security@abccompany.com** received an email with the Subject **Unauthorized activity in Pilgrim project**. In the body of the email, the username **Brittany** appeared.
- The Security team immediately explored the risky hours and associated log files in Intersect, and identified Brittany Smith as the user who had triggered the email notification. Further investigation revealed that Brittany Smith was the user who had accessed the Perforce Pilgrim repository.
- The Security team placed immediate restrictions on Brittany Smith's user access, and began a further review of her activity:
 - In the **Explore** page, the Security team was able to review all of Brittany's past activity and identify some low to medium risky behaviors, including some unusual work hours.
 - The Security team was able to establish, in consultation with Human Resources, that Brittany Smith was a disgruntled employee. Shortly following this event, Brittany Smith's employment at ABC Company was terminated.

Manage Bots and Bot-like Users

Internet bots, or Web robots, are software applications that run automated tasks. If your organization has system bot activity, this activity – because of the exceptional speed with which the activity occurs – will

likely generate Interset Risky Hours in your Analytics. Interset Analytics identifies those system users it deems to be bots, and strips them from the **Matrix of Anomalies & Violations**.

There is often very real difficulty identifying those system users that are bots and those that are live humans, based on the user activity alone. Your Security team should work with you to identify those system users that are truly bots, and those that are not.

After the true bots are identified, you can configure Interset to remove these bots from the Analytics. Similarly, if bot-like users have been stripped from the Analytics but are not bots, you can configure Interset to ensure that these users remain in the Analytics.

Steps


1. Log in to Interset as an Administrator.
2. In the **Overall Risk** page, click **Settings**  and then, in the dropdown list, select **API** to open the API in Swagger.
3. Ensure that you are authenticated as the **root** user. The default password is **root**.

Note: If you configured LDAP or SAML authentication for Interset , you must log in to Swagger as the user you configured as the LDAP or SAML "**rootUser**". For more information, please see "Configure Authentication" in the Interset 5.9.2 *Installation and Configuration Guide*.

- In the Swagger header, click **Authorize**.



- In the **Available authorizations** dialog box, under **Select OAuth2.0 Scopes**, click **Authorize**.



If you are already authenticated as the **root**, this button will display **Logout**. If the **Authorize** button displays **Logout**, you can cancel out of this step.

- In the log in dialog box, enter your user credentials and then click **Sign in**.
5. Expand the **tuning{tenantId}** section.
 6. Scroll down, and then expand the row where **PUT /tuning/{tid}/tags/{tag}** appears.
 7. Under **Parameters**, click the **entityType** dropdown box and select **users** from the list of available parameters.
 8. In the **entityId** box, type the username, or ID, for the user.
To ensure that the user **entityId** is accurate, get this ID from the **Explore** page (if currently a user shown in the Analytics) or from the PDF Report (if currently reported as a bot).
 9. In the **tag** box, do one of the following:
 - If the entity is currently a user shown in the Analytics and you want to tag this entity as a bot, type **FORCEBOT**.

- If the entity is currently shown as a bot in the PDF Report and you want to tag this entity as a user, type **NOTBOT**.

8. Click **Try it out!** to enable the new entity tag.

Tune the Analytics

After you have had the opportunity to explore the Interset Analytics and investigate the leads identified in the Interset Dashboard, you may want to fine-tune the importance applied by the Analytics to the events in your source data.

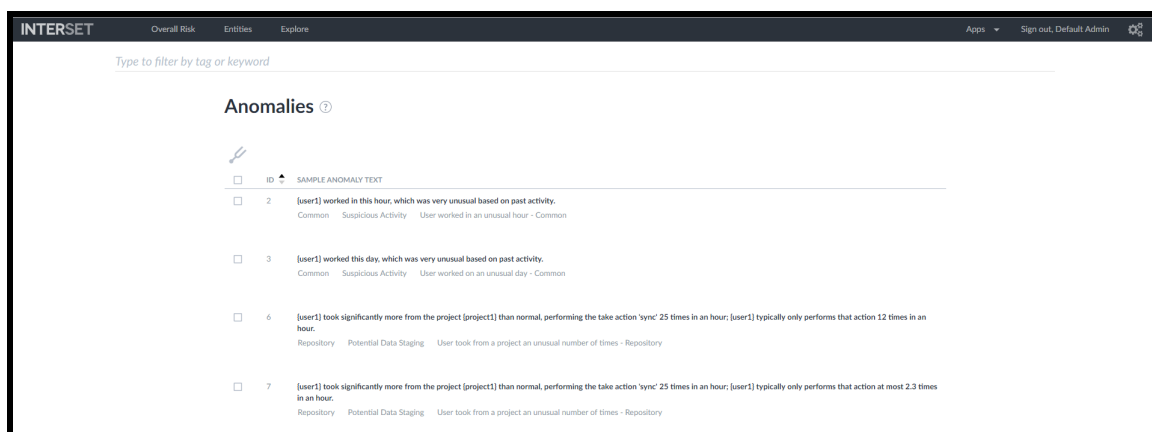
For example, perhaps due to the nature of your business, your employees have never – not once – accessed the corporate information systems outside of the standard 9:00 am to 5:00 pm work hours. In this scenario, should an employee one day access your corporate information system outside of the standard work hours, the potential for that access to be a risk to your organization could be much more significant than it would be in an organization in which employees routinely access the corporate systems at any hour. As a result, you might want to increase the importance of the group of anomalies in the anomaly family, **User worked in an unusual hour**. When you increase the importance of this anomaly family, anomalies of this type that are identified in the analytics will have a higher risk score than they would have using the default importance level.

You fine-tune the importance of individual anomalies, or grouped anomaly families, on the **Anomalies** page of the Interset user interface.

Steps

1. In a Web browser, log in to Interset as an Administrator.
2. In the Interset **Overall Risk** page, click **Settings** and then, in the dropdown menu, select **Anomalies**.


The **Anomalies** page opens, listing all of the anomalies triggered on your source data by the Inter-set models. Each anomaly appears with the following information: the Interset model **ID**, as well as the model **Data Type**, **Threat Type**, and **Family Type**.



3. To change the importance of an anomaly **Family Type**, do the following:

- At the top of the **Anomalies** page, click in the **Type to filter by tag or keyword** field.

A new dialog box opens, displaying the **Data Type**, **Threat Type**, and **Family Type** for the anomalies identified.

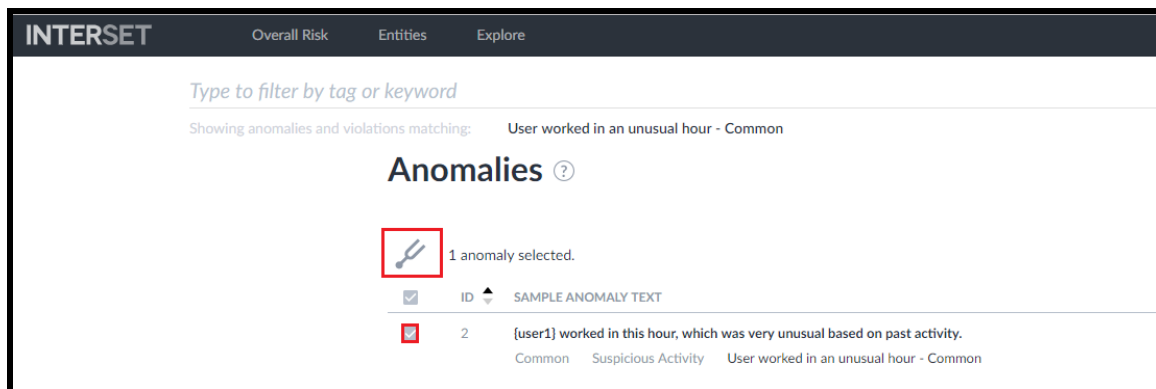


3. Select a **Data Type**, **Threat Type**, or **Family Type**.

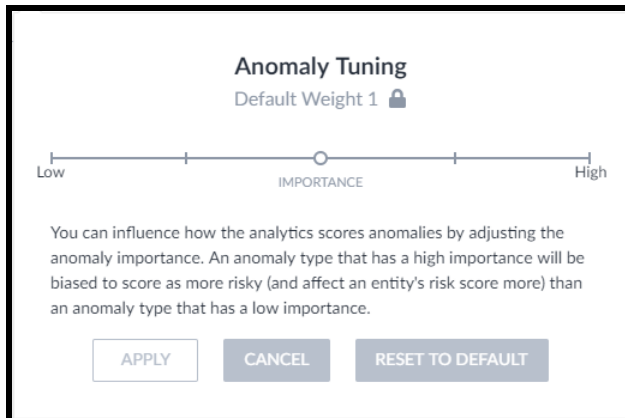
Following the example discussed above, under **Family Type**, you would select the **User worked in an unusual hour** anomaly family type.

The anomalies of that type triggered by the Analytics are now isolated in the **Anomalies** page.

4. Select an anomaly, and then click **Tuning**.



5. In the **Anomaly Tuning** dialog box, click one of the available values on the horizontal rule.

The image shows a dialog box titled "Anomaly Tuning". Below the title, it says "Default Weight 1" followed by a lock icon. There is a horizontal slider with tick marks at "Low", "IMPORTANCE", and "High". The slider handle is positioned at the "IMPORTANCE" mark. Below the slider, there is a paragraph of text: "You can influence how the analytics scores anomalies by adjusting the anomaly importance. An anomaly type that has a high importance will be biased to score as more risky (and affect an entity's risk score more) than an anomaly type that has a low importance." At the bottom of the dialog box, there are three buttons: "APPLY", "CANCEL", and "RESET TO DEFAULT".

Continuing with the example above, you would click the **High** value on the horizontal rule to increase the importance of these anomalies to the highest amount available.



Note: The **Default Weight** of the anomaly is indicated at the top of the **Anomaly Tuning** dialog box, with a lock symbol. InterSet strongly recommends that you avoid changing any anomaly weight unless instructed to do so by Micro Focus InterSet Support.

4. Click **Apply**.
5. Repeat Steps 5 through 7 for the remaining anomalies.

The next time Analytics is run, the new **Importance** value will be applied to the anomalies.



Tip: To return the anomaly importance to the default setting at any time, select the anomaly, click **Tuning**, and then in the **Anomaly Tuning** dialog box, click **Reset to Default**.

Use Intersect

Intersect uses advanced analytical models to measure behavior and to quantify risks. These models range from cluster models, which group together users and assets based on specific behavioral vectors, to volumetric anomaly models, rare activity models, and other higher-order models. Many different behavioral vectors are tracked and measured, which reduces the ability for malicious users or compromised accounts to "fake" normal behavior.

The Intersect models are true advanced behavioral models: they don't rely on binary rules or arbitrary thresholds. Rather, these models measure the probability that an observed action is truly anomalous and represents a true potential risk. Using this type of approach leads to a continuous, prioritized list of risks, and helps improve the efficiency of IT security teams and their tools.

The use of Intersect machine learning models means that you are not required to perform any additional configuration for the analytical models to execute. Through observation, Intersect learns what constitutes normal behavior for the entities within your organization, and immediately begins to quantify abnormal behavior. There are no thresholds to set, no rules to author, and no configurations to undertake.

Intersect displays the results of the Analytics in an interface that provides at-a-glance actionable information on current risk, and flexible multi-entity historical data exploration.

Users and Other Entities

Entities are the foundation of Intersect Analytics. Entities are the objects involved in behaviors. For example, if a user **Philip** accesses **Fileshare A**, then the event contains, at minimum, one behavior, and two entities. **Philip's account** and **Fileshare A** are the two entities, and the access is the behavior.

Behaviors

Behaviors are often thought of as single events. In the previous example, the access can be captured in one single event. If that event happens to be a malicious action, finding that one malicious event is virtually impossible. This is because there can be billions of these events, and the overwhelming majority of events are perfectly legitimate and normal behaviors.

Accumulating Risk

As behaviors occur, Intersect processes these events and calculates that which is normal from dozens of behavioral perspectives. For example, Intersect will count how many times Philip accesses **Fileshare A** each hour, how often his authentication attempts fail on **Fileshare A**, at what time of day, or which day of week he is normally active, etc.

These metrics are all calculated using unsupervised machine learning. This means that the system identifies what is normal, rather than organizational security practitioners setting thresholds which may be reasonable for some, but completely inappropriate for others.

As new observed behaviors occur, Intersect determines whether the behaviors are normal, or unusual. When unusual, Intersect calculates how unusual the behavior is. The more unusual the behavior, the higher the significance of the anomaly. When anomalies are identified, these anomalies influence the risk score

of the entities that are involved in the behavior. The more an entity is involved in significant anomalies, the higher that entity's risk score. For example, if Philip accesses **Fileshare A** 100 times in an hour, and accesses 100 other fileshares that he's never accessed before, his risk score will spike, because the behavior simulates internal recon or lateral movement. In addition, because **Fileshare A** was involved in a significant set of anomalies, its risk score will also spike.

This comprehensive reporting allows practitioners to explore the anomalies from different perspectives. In cases where multiple user accounts are accessing **Fileshare A** in an abnormal manner, the user behavior may not appear abnormal and therefore the risk scores may not spike significantly, however, **Fileshare A** would have a significant spike in its risk score, providing a signal to security practitioners that **Fileshare A** requires attention.

As entities are involved in risky behaviors, their risk scores increase. The riskier the entity's behavior the more the risk increases. When the entity is not engaging in any activity, the risk score decays downward towards zero; as a result, when the entity goes a long time without registering any suspicious activities, its risk score will trend toward zero.


Overall Risk Page

When you first log in to Interset, you are taken to the **Overall Risk** page. This page allows you to see, at a glance, the overall risk status of your organization. For example, in the screenshot below, you immediately see:

- just over 4 million events were analyzed
- about 836 thousand anomalies and violations were found
- 22 active risky entities were identified
- the overall risk is medium and increasing rapidly
- the threat of Potential Lateral Movement is contributing 22% to the overall risk
- the various streams of the graph indicate the potential threat types involved
- the types of entities involved and their risk counts
- the top five risky users

When you click an entity type, the **Entities** page opens, where additional information for the selected entity type is displayed.

When you click one of the **Top 5 Riskiest Users**, the **Explore** page opens, with the selected user's name applied to the **anomalies and violations** filter.



Entities Page

The **Entities** page provides the entity risk scores, sorts the entities and their risk scores in descending order, and then also provides the trending information, the entity name, the potential threat type, and the most relevant anomaly identified by InterSet. Potential threat types are determined by the most relevant risky activity in the system.

At the top of the **Entities** page, you can use the different entity tabs to explore the riskiest entities grouped by their type, such as **Users**, **Projects** or **Controllers**, for example. Tabs in bold text represent entities that are present in the data. Typically, you will explore your list of users first.



Note: You may see a difference between the number of entities that exist in your data, and the number of entities that appear in the InterSet Analytics user interface. This is because:

- The entities that appear in the **Entities** page include only
 - those entities with a current risk score greater than zero (0); and
 - those entities that have a current risk score of zero (0), but for which anomalies were identified during the selected time period.
- Entities identified as BOTs do not appear in the user interface.

The screenshot shows the INTERSET 'Entities' page. At the top, there's a search bar and a navigation bar with 'Overall Risk', 'Entities', and 'Explore'. Below the navigation bar, there's a row of icons representing different entity types: Users (2,278), Machines (63), Controllers (53), IP Addresses (104), Printers (10), Projects (43), Resources (20), Shares (136), Websites (8), Files (0), and Servers (0). The main section is titled 'Top Risky Entities' and displays a table of entities with their risk scores and potential threats.

#	Current Risk	Entity	Potential Threat	Riskiest Anomaly/Violation
1	98	joshua.newman	Suspicious Activity	NOVEMBER 2ND 2016, 3:29 AM Active 3-4 AM Nov 2, 2016 (5% anomalous)
2	98	vap3lad3.ljlt.com	Data Theft	NOVEMBER 1ST 2016, 1:00 AM Sent 767MB of data to vap3lad3.ljlt.com on first access.
3	98	QA-WIN-7-CONN.local	Data Theft	NOVEMBER 2ND 2016, 1:00 AM Sent 1.16TB of data to 204.212.170.143 in an hour.
4	97	shiela.mathis	Internal Recon	NOVEMBER 2ND 2016, 10:00 AM Attempted to access inactive shared drives /archives/Designs, /Network Shares/ITFM.
5	95	benjamin.mitchell	Data Staging	NOVEMBER 2ND 2016, 1:00 AM Accessed files totalling 1.75TB.

Potential threat types are determined by the riskiest activity identified by Intersect for that entity. For example, if the riskiest alert results from behaviors in which a user account is accessing unusual locations or assets, the potential threat type will appear as **Potential Lateral Movement**, and a summarized description of the anomaly will be shown on the right of the page. This provides immediate context for security practitioners, and enables them to more quickly determine whether further investigation is required.

User-Defined Tags

On the **Entities** page, you can associate **User-defined tags** with individual entities or many entities at the same time. You can also create new tags and delete tags you don't need any more.



Important: Do not use **bot**, **forcebot**, or **notbot** as names for a **User-defined tag**.

To manage tags:

1. Choose one or more entities from the **Top Risky Entities** list by selecting the checkbox(es) in the left-most column.

2. Click the **Tag Management** icon ().

The **Tag Management** dialog shows the tags associated with the selected entities. Tags that have checkmarks are associated with all the selected entities. Tags with a stroke through the middle of the checkbox are associated with one or more (but not all) of the selected entities. Tags with no checkmarks are not associated with any of the selected entities.


- Do one of the following:
 - To associate tags with the selected entities, select the checkbox next to one or more tags and click **Apply**.
 - To remove the association between a tag and the selected entities, clear the checkbox next to the tag and click **Apply**. The **Total entities tagged** field is updated to show the number of entities associated with the tag.
 - To create a new tag, click **Create a new tag?** Enter the new name in the **tag-name** field, and click **Create**. The new tag is created and associated with the selected entities.
 - To rename a tag, select the tag by clicking its name, then click the **tag-name** field and type the new name. Click **Save Changes** to save the new name.
 - To delete a tag, select the tag by clicking its name, then click **Delete**. Click **Yes** to confirm the deletion.

You can also create **User-defined tags** on the **Explore** page. See ["Use Intersect" on page 58](#).

Explore Page


When you select an entity, the **Explore** page opens, where the entity's name is filtered. Here, all **Anomalies & Violations** associated with that entity are shown within the established time range. To find or filter another entity, use the search filter at the top of the **Explore** page.

The **Explore** page information allows you to use to determine the types of risky activities that are occurring within your organization. The **Explore** page features the **Matrix of Anomalies & Violations**, the **Contribution to Risky by Threat** graph, and the **Top Risky Users** and **Anomalies & Violations** panels, which are displayed by default.



Matrix of Anomalies & Violations

The **Matrix of Anomalies & Violations** is a visual representation of the **Anomalies & Violations** in your data set, displayed as squares, color-coded to reflect their severity.



You can change the time window for the **Matrix of Anomalies & Violations** to reflect a time period of specific interest. You can choose the following time periods: **24 Hours**, **7 Days**, **30 Days**, **Year**, or you can set the time period to include **All Data**. To zoom in on a specific area of the matrix, click the **+** icon and then click and drag your cursor across the area of the matrix where you want to zoom in. To zoom out, click the **-** icon, or select one of the predefined time windows. To pan across the time window, click and drag your cursor across the matrix (zoom must not be enabled). As you zoom or pan, all aspects of the user interface update dynamically and accordingly.


You can use the slider to the left of the matrix to filter alerts based on their risk level. This enables you to reduce the number of alerts displayed in a gradual manner, and as appropriate. You can also click one of the **Risk** squares below the graph to set the slider filter to that risk level. For example, if you wanted to view **Medium Risk** and above, you would click the yellow **Medium Risk** square. This would filter out all low risk alerts, as shown in the example below.

In the **Matrix of Anomalies & Violations** timeline, you can filter the analytics on the associated entities displayed in **Anomalies & Violations**. For example, setting the project filter to: **dev/rel3/kanga** and setting the user filter to **nell.bernstein**, displays only **Anomalies & Violations** involving both the **dev/rel3/kanga** project and **nell.bernstein**.

Periods of Risky Activity features an **Overall Risk Trend** which displays a baseline within the graph. When you add an entity filter to the **Periods of Risky Activity** graph, a new **Risk Trend** line based on that entity is created. This custom **Risk Trend** displays a baseline based on that entity's activity. You can have multiple **Risk Trends** displayed at once. You can also hide and show the **Risk Trends** by selecting the name of the **Risk Trend**.

Contribution to Risk by Threat

Below the **Matrix of Anomalies & Violations** is the **Contribution to Risk by Threat** graph. This graph organizes and displays potential threat types by their percentage of the overall risks. You can filter the graph by threat type by selecting the threat type name or square in the graph. For example, if you wanted to highlight the percentage that **Potential Internal Recon** represents in the graph, you would select the **Potential Internal Recon** name or square underneath the graph. To reveal/hide the **Contribution to Risk by Threat** graph, click the **Contribution to Risk by Threat** heading.



Apply Filters to Entity and Anomaly Data

At the top of the **Entities** and **Explore** pages is a field where you can select filters to apply to the data that is displayed.

- On the **Entities** page, the filter field is labeled **Type to filter entities by name, tag, or type...**
- On the **Explore** page, the filter field is labeled **Type to filter anomalies and violations...**

From the filter field you can choose a filter from the dropdown menu, or you can search for filters by typing the filter name. Depending on your data set, you can apply filters on many aspects of your data, including **Users**, **Entity Types**, **Projects**, **Controllers**, and **User-defined tags**.

When you select filters, the filters appear in a list under the filter field:

- On the **Entities** page, the filter list is labeled **Showing entities matching:**
- On the **Explore** page, the filter list is labeled **Showing anomalies and violations matching:**

To disable a filter:


- In the filter list, hover your cursor over the filter name and then click the checkbox on the left. The filter is still displayed but is no longer applied to the data. Repeat this process to enable a disabled filter.

To remove a filter from the filter list:

- In the filter list, hover your cursor over the filter name and click the **X** on the right. The filter is removed from the list, but is still available to use if you select it again.

Anomalies & Violations Panel

The **Anomalies & Violations** panel displays triggered activities in the form of a list. Each **Anomaly** or **Violation** has a time stamp, risk color, description, potential threat type, and associated entities attached to it. The **Anomalies & Violations** list can be sorted by **Time** (default) or by **Risk**.



To sort the list:

- At the top left of the **Anomalies & Violations** panel, click the dropdown menu and then select **Sort by time** or **Sort by risk**.

To apply filters based on an **Anomaly** or **Violation**:

- Below the description of the **Anomalies** or **Violation**, click the tags you wish to apply to the filter.

To disable a filter:

- At the top left of the page, under **Type to filter anomalies and violations...**, hover your cursor over the filter name and then click the checkbox on the left. Repeat this process to enable a disabled filter.

To delete a filter:

- At the top left of the page, under **Type to filter anomalies and violations...**, hover your cursor over the filter name and click the **X** on the right.

When you click in an **Anomaly** or **Violation** box, a visualization is provided to enhance context and includes a description of the activity. From here you can choose to explore the raw events that triggered the **Anomaly** or **Violation**, For more information on exploring raw events, please see the **Explore Raw Events** section.



Note: When you configure a Workflow to send a resulting email, and you added new lines of text in the email body when creating the Workflow, the email notification received by the recipient will include a semicolon at the end of each text line. These semicolons do not impact Workflow functionality in any way, and can be ignored.

Anomaly and Violation Flags

Interset provides five (5) possible flags that you can use to characterize, or mark individual anomalies and violations within the Analytics. These five flags are represented by the following symbols:



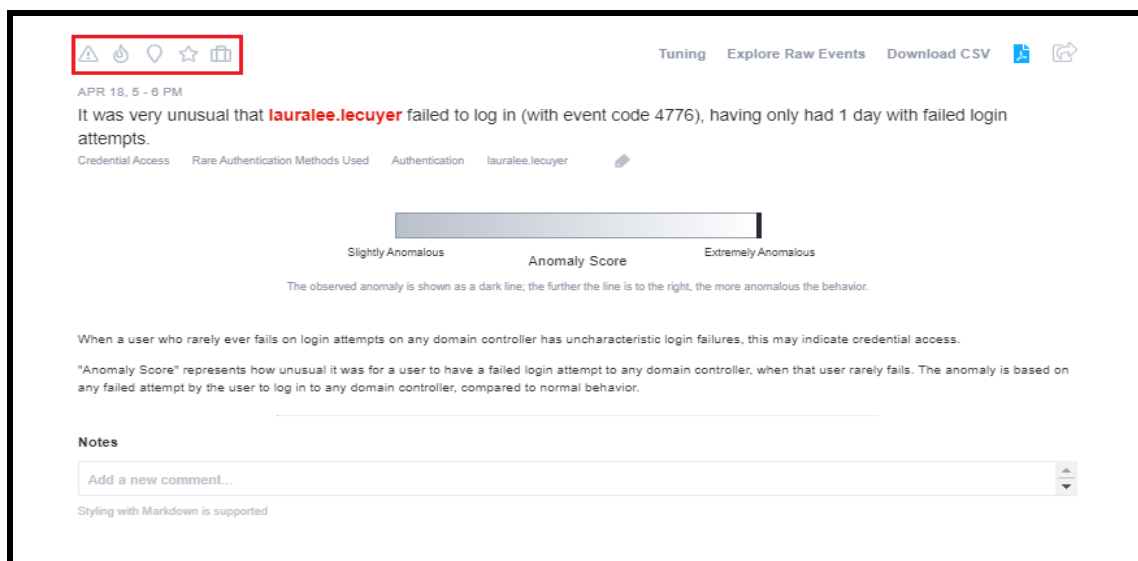
These five flags, or symbols, have no established definitions; as a result, your organization can determine the appropriate meaning for each symbol within the context of the anomalies and violations that you want to highlight in your data.

For example, you may decide to use one of these symbols to identify anomalies and violations resulting from failed access or log in attempts. You choose which of the flags you will use for this purpose and then, in the Anomalies & Violations panel, you mark the individual anomalies accordingly. When you have finished marking the anomalies and violations, you have only to select the flag as a filter to produce a list of all failed access and log in attempts.

To create flags:

1. In the **Anomalies & Violations** panel, click in an anomaly for which you want to set a flag.

The anomaly opens, displaying the flags in the upper left corner.

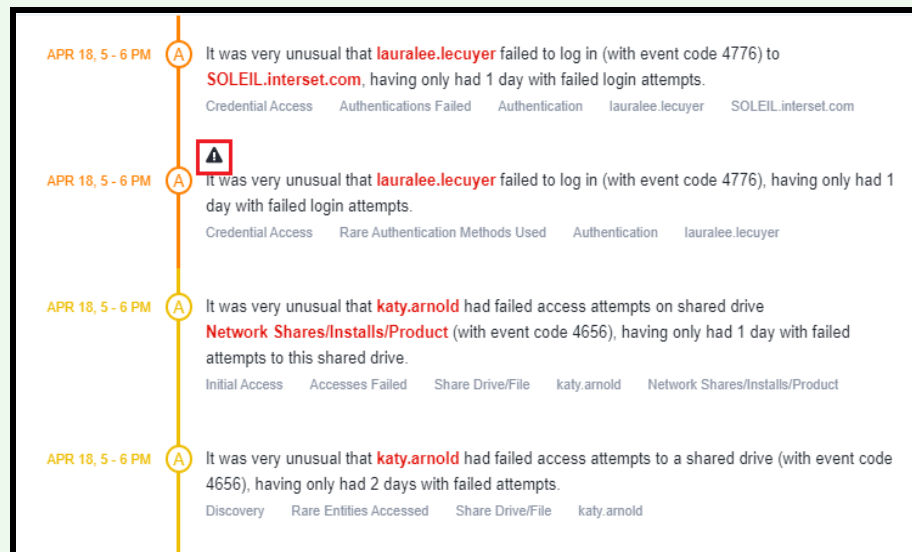


2. Click on a flag symbol to enable it for the anomaly.




**Tips:**

- When a flag is enabled for an anomaly or violation, the symbol changes from light to dark.
- When you close the anomaly and return to the list of anomalies and violations, the enabled flag appears within the anomaly and provides a visual cue.




- When you hover over the flag, the date and time it was created, and the account that created it, are displayed.

3. To view all the anomalies and violations flagged with a specific symbol, click in the **Type to filter anomalies and violations** field.
4. In the filter drop-down menu, under **Flags**, select the flag on which you want to filter the anomalies and violations.



The Explore page now displays only the anomalies and violations that match the defined filter, which is displayed directly below the **Type to filter anomalies and violations** field.



Tips:

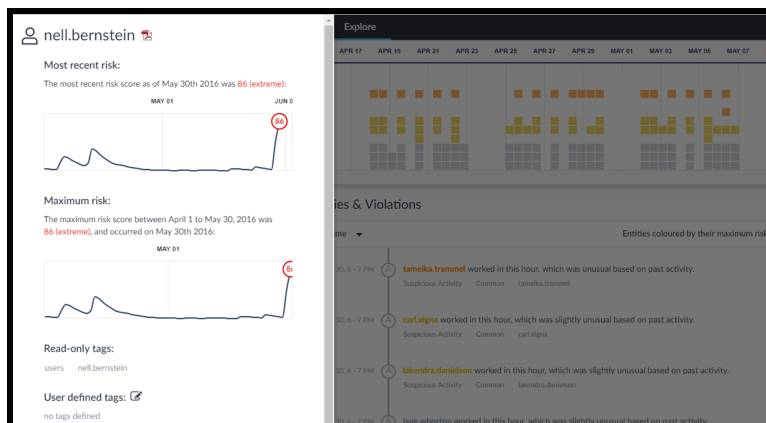
- To return to the unfiltered view of anomalies and violations, click the **X** beside the filter.
- To remove a flag from an anomaly or violation, open the individual anomaly or violation, and then click the flag to disable it.

Add Comments

Based on your investigation of an anomaly or a violation, you may want to add your observations so that other members of your team can leverage the information you have gathered so far. You can add a comment by clicking on an item in the Anomalies & Violations panel, and typing in the **Notes** field.

Entity Details Panel

When you select an entity, an **Entity Details** panel containing additional information on the entity you clicked opens. If you selected a **User** entity, for example, the **Entity Details** panel might display information regarding the **Most Recent Risk Score**, **Maximum Risk** score within the time frame, Read-only tags, User-defined tags, Typical working hours, and Typical weekly activity. To download a PDF report on the entity, click the PDF icon beside the entity name.



From the **Entity Details** you can create and apply **User-defined tags**.



Important: Do not use **bot**, **forcebot**, or **notbot** as names for a **User-defined tag**.

To create a tag:

1. On the right side of **User-defined tags**, click the  icon. A + appears below the **User-defined tags** section.
2. Click the +.
3. In the dialog box, enter the name of the tag you want to create.
4. On the right side of **User-defined tags**, click **done** to save the tag.

To delete a tag:

1. On the right side of **User-defined tags**, click the  icon.
2. Click a tag to highlight it.
3. Press your **Delete** or **Backspace** key to delete the tag.
4. On the right side of **User-defined tags**, click **done** to save your changes.



Note: If you use the same tag for multiple entity types, the results of filtering may also return entities that are associated with entities of that tag. For example, filtering on a tag of "Boston" which has been applied to users and controllers located in Boston may return users outside of Boston that have interacted with the controllers with that tag.

Authentications Panel

The **Authentications** panel displays the total number of successful and failed authentication attempts, sorted by entities with the most failed attempts in descending order.

To add the **Authentications** panel:

- Click the **+** symbol and then select **Authentications**.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

Most Accessed Panel

On the **Explore** page, you can view the **Most Accessed** entities of your whole dataset, or of specific entities.

To add a **Most Accessed** panel:

- At the bottom of the page beside the leftmost tab, click the **+** symbol, select **Most Accessed** and then select a filter.

Each **Most Accessed** filter displays a list of entities that have been interacted with, sorted in descending order. You can further explore the **Most Accessed** entities by selecting an entity to open the **Entity Details** panel.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

Top Risky Panel

On the **Explore** page, the **Top Risky** panel provides a list of the top risky entities by type, displaying the **Top Risky Users** by default. You can change the filter to display a different entity type by clicking **Top Risky Users**, selecting **Top Risky**, and then selecting an entity type. The **Top Risky** list is sorted by maximum risk.

Note: You may see a difference between the number of entities that exist in your data, and the number of entities that appear in the Intersect Analytics user interface. This is because:

- Only entities with anomalies appear in the Top Risky list; entities without identified anomalies within the selected time range are filtered out. In addition, entities identified as BOTs do not appear in the user interface.
- When you select the all data timeframe, all entities that have ever had at least one identified anomaly will be shown.

To add a new **Top Risky** panel:

- At the bottom of the page beside the leftmost panel, click the **+** symbol, select **Top Risky** and then select an entity type.

You can further explore the **Top Risky** entities by click in an entity box to open the **Entity Details** panel.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

Most Exits By User Panel

On the **Explore Page**, you can view the **Most Exits By User**. This panel provides a table list of users with the most 'Exit' activity, sorted in descending order. 'Exit' activity includes saves to USB, file uploads, print actions, and more.

To add the **Most Exits By User** panel:

- At the bottom of the page beside the leftmost panel, click the **+** symbol and then select **Most Exits By User**.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

Top Screen Capture Produces Panel

On the **Explore Page**, you can view the **Top Screen Capture Produces**. This panel provides a table list of the users who produce the highest number of screen captures, sorted in descending order. Screen captures includes the standard methods of copying the active image, such as **PrtScn**, **Alt+PrtScn**, and the Windows Snipping Tool (PC); **command+shift+3**, **command+shift+4**, **control+command+shift+3**, and **control+command+shift+4** (Mac); as well as Snagit, a third-party screen capture application.

To add the **Top Screen Capture Produces** panel:

- At the bottom of the page beside the leftmost panel, click the **+** symbol and then select **Top Screen Capture Produces**.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

Top Users To Trigger Violations Panel

On the **Explore Page**, you can view the **Top Users To Trigger Violations**. This panel provides a table list of the top users who have triggered Workflow violations, sorted in descending order.

To add the **Top Users To Trigger Violations** panel:

- At the bottom of the page beside the leftmost panel, click the **+** symbol and then select **Top Users To Trigger Violations**.

To remove the panel:

- Click the **X** symbol at the top right of the panel.

CSV Reports

CSV reports provide you with the raw data of the **Anomalies & Violations**. A CSV Report can provide you with further insight on how an entity is behaving. For example, a user entity CSV Report may contain

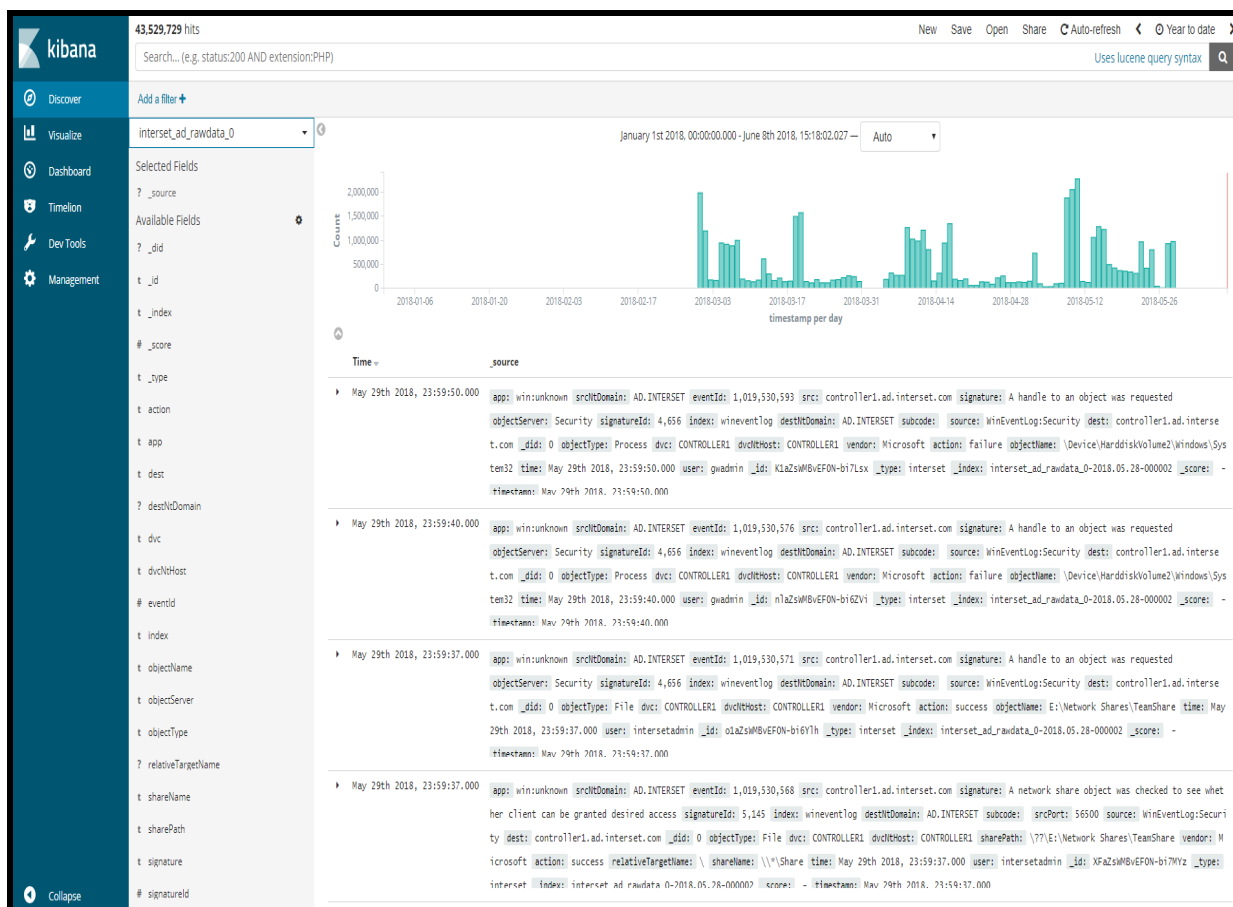
information regarding country of origin, actions taken, username and object type. To download a CSV report, click in the box of a **Anomaly** or **Violations** and then, on the right side of the visualization, click **Download CSV**.

PDF Reports

After an investigation has sufficient evidence to warrant an escalation, information can be exported to a PDF format so that incident response can begin immediately. To generate a PDF report for your organizational risk, on the **Overall Risk** page, next to the date at the top of the page, click the PDF icon. To generate a report for a user entity, from the **Explore** page, click a user entity name to open the **Entity Details** panel, and then click the PDF icon beside the entity name to download a PDF report. With this report, you can quickly share the findings of the investigation without having to manually create any additional documents. The report helps provide an understanding of what constitutes a risky or normal behavior for any entity.

Explore Raw Events

When you click an item in the **Anomalies & Violations** panel, a dialog box appears that provides additional context about the anomaly or violation. To see the events that triggered the risky activity, in the top right of the dialog box, click **Explore Raw Events**. This launches a pre-populated query in Kibana, where you can further explore the events.



Kibana provides security practitioners with a quick way to explore the context around the raw events that triggered the anomaly. This can include expanding the time range, changing filter options, or any other faceted search.

Advanced Features

In the Intersec user interface, you can take advantage of a number of advanced features that allow you to manage the security of your organization more effectively. For example, you can:

- work with your Intersec Administrator to create [Workflows](#) (real-time notifications and violations)
- work with your Intersec Administrator to [manage bots and bot-like users](#)

Workflows

Workflows are an extremely valuable tool for ensuring that those who rely on Intersec Analytics to identify the top risky entities and behaviors occurring in the organization, can be immediately aware of specific areas of potential concern and that appropriate, proactive measures may be undertaken in a timely manner. With Workflows, Administrators can build custom use cases to highlight specific information for their end users, based on a wide range of criteria.

When Workflows are configured and include email or SMS notifications, and REST API calls, these notifications and corrective actions occur instantaneously, in real time. When Workflows include the creation of violations or changes in severity level, these outcomes are available the next time Intersect Analytics runs on the data.

For example, a Security team may wish to be immediately notified, and have initial corrective measures put in place any time an individual inappropriately attempts to access a restricted project repository. By creating Workflows based on specific conditions, the Security team can easily configure real time email notifications, and even control access privileges using REST APIs when those conditions are met.

Only users with Administrator privileges can create Workflows. To learn more about creating Workflows, see ["Create Workflows" on page 40](#).

Bots and Bot-like Users

Internet bots, or Web robots, are software applications that run automated tasks. If your organization has system bot activity, this activity – because of the exceptional speed with which the activity occurs – will likely generate anomalies in your Analytics. Intersect Analytics identifies those system users it deems to be bots, and strips them from the Investigator matrix view.

There is often very real difficulty identifying those system users that are bots and those that are live humans, based on the user activity alone. Your Intersect Administrator will work with you to identify those system users that are truly bots, and those that are not.

After the true bots are identified, your Administrator can configure Intersect to remove these bots from the Analytics. Similarly, if bot-like users have been incorrectly tagged and stripped from the analytics, your administrator can configure Intersect to ensure that these users remain in the Analytics.

For more information about bots and bot-like users, see ["Manage Bots and Bot-like Users" on page 52](#).

Advanced Configuration Options

After installing and configuring Interset in your environment, the following advanced configuration information may be useful.

Custom Theme

The custom branding feature allows you to change the logos and the navigation bar colors of the Interset interface to reflect that of your business and organizational needs. Text references to Interset can also be changed to those of your company. You can reset the interface to the Interset default at any time.

These changes are implemented on the Reporting node. After changes are saved, they will be applied to all tenants on the node.


Note: Changes made to the **NavBar Logo**, **Company Name**, **NavBar Color**, **Accent Color**, **Log In Page Gradients**, and **Font Color** during this process are shown in the **Sample Preview** window. Changes to the **Log In Logo** and **Enable Powered by Interset message** cannot be seen until saved. Changes will not be implemented until you save them.

To edit the theme of the Interset interface:

1. Open a Web browser, go to **https://<reporting_node_fqdn>/dashboard** and log in as an administrator.
2. On the **Overall Risk** page, at the top right of the page, click the **Settings**  to open the **Settings** page.




3. On the **Settings** page, click **CHANGE THEME**.



The **Theme Panel Page** opens.

4. To change the **Log In Logo**:

- a. Next to the **Log In Logo** input field, click **Browse**.


The screenshot shows the 'Theme Panel' interface. At the top, there's a dark header with 'INTERSET' and navigation links 'Overall Risk', 'Current Risk', and 'Explore'. Below the header, the 'Theme Panel' title is followed by three input fields. The first field is 'Log In Logo', which has a red circle around its 'BROWSE' button. Below it is a hint: 'Log in logo must be at least 560px x 100px'. The second field is 'NavBar Logo', with its 'BROWSE' button also visible. Below it is a hint: 'NavBar logo must be at least 300px x 80px'. The third field is 'Company Name', which currently displays 'InterSet'.

- b. Navigate to the directory where the image is located, select the graphic file to upload as your **NavBar Logo**, and then click **Open**.

The **Log In Logo** must be in the .png format, have a maximum dimension of 560px by 100px, and be no larger than 2MB in size.

5. To change the **NavBar Logo**:


- a. Next to the **NavBar Logo** input field, click **Browse**.

The screenshot shows the 'Theme Panel' interface. The 'Log In Logo' field now displays 'YOURLOGO1.png' and its 'BROWSE' button. The 'NavBar Logo' field has a red circle around its 'BROWSE' button. The 'Company Name' field still displays 'InterSet'.

- b. Navigate to the directory where the image is located, select the graphic file to upload as your **NavBar Logo**, and then click **Open**.

The **NavBar Logo** must be in the .png format, have a maximum dimension of 300px by 80px, and be no larger than 2MB in size.

6. To change the **Company Name**, in the **Company Name** text box, enter the company name to be displayed.



This changes/removes any text instances of Interset and changes them to the company name you entered.

7. To change the **Log In Page** gradients and **NavBar** colors, click the color box next to the field you want to change.

You can change the colors of the following:

- **Log In Page** Gradient 1
- **Log In Page** Gradient 2
- **NavBar** Color
- Accent Color

Font colors will change automatically based on your color choices.

8. To save your modified theme, at the bottom right of the page, click **Save Theme**.

Reset Theme to Default

You can reset your theme to the default at any time by completing the following steps:

1. Open a Web browser and go to **https://<reporting_node_fqdn>/dashboard** and log in as an administrator.
2. On the **Overall Risk** page, at the top right of the page, click **Settings**.
3. On the **Settings** page, click **CHANGE THEME**.
4. On the **Theme Panel Page**, at the bottom right of the page click **Reset Theme**.

Add Custom Text to the Interface Banner and Footer


To further customize the appearance of the Interset interface for your organization, you can add your own custom banner text to appear over the **NavBar Logo** on the login page, as well as custom footer text to appear on the bottom right of the login page.

To add custom text, do the following:

1. Open a Web browser, go to **https://<reporting_node_fqdn>**, and then log in as an administrator.
2. On the **Overall Risk** page, at the top right of the page, click **Settings**  to open the **Settings** page.



3. On the **Settings** page, click **CHANGE THEME**.




The **Theme Panel Page** opens.

4. To change the banner label:

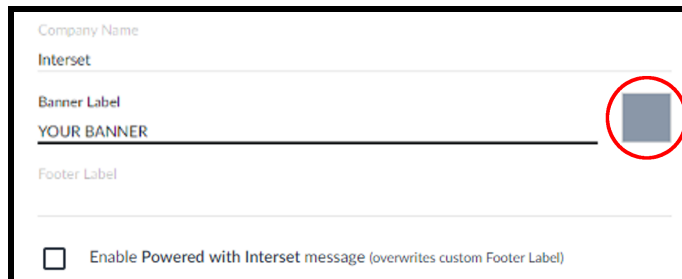
- a. In the **Banner Label** text box, type the banner text to be displayed.

The banner label appears at the top left of the page, above the NavBar logo.



The screenshot shows the 'Theme Panel' interface. It includes fields for 'Log In Logo', 'NavBar Logo', 'Company Name', 'Banner Label', and 'Footer Label'. The 'Banner Label' field is highlighted with a red circle and contains the text 'YOUR BANNER'. To the right of the 'Banner Label' field is a small gray square color box. Below the 'Banner Label' field is the 'Footer Label' field, which contains the text 'YOUR FOOTER'. At the bottom of the panel is a checkbox labeled 'Enable Powered with Interiset message (overwrites custom Footer Label)'.

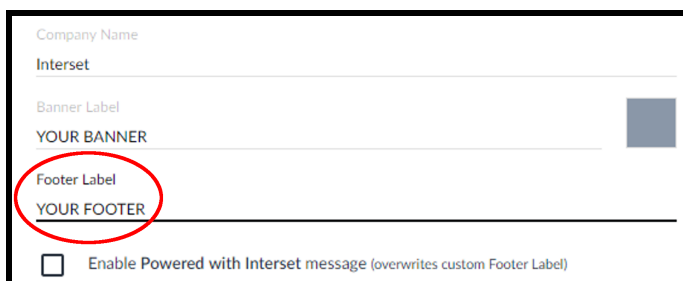
- b. To change the color of the banner label text, to the right of the **Banner Label** text box, click the color box and then select a new color.



This screenshot is a closer view of the 'Banner Label' section of the Theme Panel. The 'Banner Label' field contains 'YOUR BANNER'. To its right is a small gray square color box, which is highlighted with a red circle. Below this is the 'Footer Label' field with 'YOUR FOOTER' and the same 'Enable Powered with Interiset message' checkbox at the bottom.

5. To change the footer label, in the **Footer Label** text box, type the footer text to be displayed.

The footer label appears at the bottom right of the login page.



This screenshot shows the 'Footer Label' field in the Theme Panel, which is highlighted with a red circle and contains the text 'YOUR FOOTER'. The 'Banner Label' field above it contains 'YOUR BANNER'. The 'Enable Powered with Interiset message' checkbox is visible at the bottom.

6. To display **Powered by Interiset** at the bottom of the login page, select the **Enable Powered by Interiset message** checkbox.

When enabled, this replaces any custom text in the **Footer Label** textbox to **Powered by Interset**.

7. To save your changes, at the bottom right of the page, click **Save Theme**.

Configuring Delimiters

For data sources identified as .csv, the default delimiter is the comma. You can specify a variety of other delimiters, including the semicolon, pipe, tab, and space.

If your data source uses a different delimiter, you can manually configure that delimiter by changing either:

- the `Value Separator` field in the **CSVReader - <schema>** Flume controller services, OR
- the `agent.sources.dirSource.deserializer.csvSeparatorCharacter` parameter in the Flume configuration file.

User Accounts and Permissions

The user accounts that appear in the tables below are created during the normal installation of Interset, Ambari, and HDP. None of these user accounts are required to run as root.

Interset Accounts

Service	Component(s)	Default User Account
Interset	Interset	interset
Admin	Interset	admin

Ambari and HDP Accounts

Service	Component(s)	Default User Account
Accumulo	Accumulo Tracer, Accumulo Monitor, Accumulo GC, Accumulo Master	accumulo (HDP 2.2 or later)
Ambari	Ambari Server, Ambari Agent	interset
Ambari Metrics	Metrics Collector, Metrics Monitor	ams
Ambari Infra	Infra Solr instance	infra-solr
Atlas	Atlas Metadata Server	atlas (HDP 2.3 or later)
Falcon	Falcon Server	falcon
HBase	MasterServer, RegionServer	hbase
HDFS	NameNode, SecondaryNameNode, DataNode	hdfs
Hive	Hive Metastore, HiveServer2	hive
HUE	HUE	hue
Kafka	Kafka Broker	kafka
Knox	Knox Gateway	knox
Mahout	Mahout clients	mahout (HDP 2.2 or later)
MapReduce2	HistoryServer	mapred
Oozie	Oozie Server	oozie
PostgreSQL	PostgreSQL (with Ambari Server)	postgres (Created as part of installing the default

Service	Component(s)	Default User Account
		PostgreSQL database with Ambari Server. If you are not using the Ambari PostgreSQL database, this user is not needed.)
Ranger	Ranger Admin, Ranger Usersync	ranger (HDP 2.2 or later)
Ranger KMS	Ranger KMS Server	kms (HDP 2.3 or later)
Slider	Slider clients	slider
SmartSense	HST Server, HST Agent, Activity Analyzer, Activity Explorer	<same as ambari agent>
Spark2	Livy Servers	livy
Spark2	Spark History Server	spark (HDP 2.2 or later)
Sqoop	Sqoop	sqoop
Storm	Master (Nimbus, DRPC Server, Storm REST API, Server, Storm UI Server) Slave (Supervisors, Logviewers); Slave (Supervisors, Logviewers)	storm
Tez	Tez clients	tez
WebHCat	WebHCat Server	hcat
YARN	NodeManager, ResourceManager	yarn
Zeppelin Notebook	Zeppelin Notebook	zeppelin
ZooKeeper	ZooKeeper	zookeeper

For more information about the Ambari and HDP service accounts, please go to <https://docs.hortonworks.com/> and then use the Search box to locate the relevant topic.

Other Third-party Component Accounts

Service	Component(s)	Default User Account
Cassandra	Cassandra	cassandra
Elasticsearch	Elasticsearch	elasticsearch

Appendix A: Interset Cluster Components

This section provides information about:

- the Interset components
- the third-party components
- the recommended distribution for the Single-instance Node and Multi-instance Node production installations described in this Guide

Interset Components

The Interset components, which will be installed on different nodes in the recommended configurations, include:

- **Interset Analytics**

This component performs the vital task of determining individual behavioral baselines, and then discovering and ranking deviations from those baselines.

Interset Analytics is installed on the Master node(s).

- **Interset Reporting**

This component provides the REST API, as well as the rich user interface that allows the analytics results and raw data to be explored visually.

Interset Reporting is installed on the Reporting node.

- **Interset Workflow**

This component applies user-defined rules to highlight specific events and trigger follow-up actions. These user-defined events contribute to the analytics.

Workflow is installed on the Master node(s).

Third-party Components

The Interset cluster third-party components, also distributed among multiple nodes in the recommended configurations, include:

- **Apache Ambari Server**

The Apache Ambari project simplifies Apache Hadoop management with the development of software for provisioning, managing, and monitoring Hadoop clusters. Ambari provides an intuitive, central Hadoop management user interface backed by its REST APIs.

Apache Ambari server is installed on the Monitoring node.

- **Apache Ambari Metrics**

The Ambari Metrics System (AMS) collects, aggregates, and serves up Hadoop and cluster metrics in Ambari-managed clusters.

Apache Ambari Metrics is installed on the Ambari, Compute, Master, and Stream nodes.

- **Apache Ambari Client**

An Ambari client is the node in the cluster that provides the client libraries for any services managed by Ambari, and supports installed client applications (such as Interaset Analytics).

Apache Ambari Client is installed on the Stream node(s).

- **Cloudera Manager Server**

Cloudera Manager is an end-to-end application for managing CDH clusters. Cloudera Manager provides granular visibility into and control over every part of the CDH cluster—empowering operators to improve performance, enhance quality of service, increase compliance, and reduce administrative costs. With Cloudera Manager, you can easily deploy and centrally operate the complete CDH stack and other managed services.

Cloudera Manager Server is installed on the Monitoring node.

- **Cloudera Manager Agents**

The Cloudera Manager Agent is a Cloudera Manager component that works with the Cloudera Manager Server to manage the processes that map to role instances.

Cloudera Manager Agents are installed on the Monitoring, Compute, Master, and Stream nodes.

- **Apache HDFS™**

The Hadoop Distributed File System (HDFS) is a distributed file system that provides high-throughput access to application data. All Interaset Analytics data, residing in the HBase database, is stored in HDFS.

Apache HDFS is installed on the Data node.

- **Apache HBase™**

HBase is a scalable, distributed database that supports structured data storage for large tables. HBase stores the Analytics data for the Interaset cluster.

Apache HBase is installed across the Compute and Master node(s).

- **Apache Storm™**

Apache Storm reliably processes unbounded streams of data, doing for real-time data processing what Hadoop does for batch processing.

In a Single Instance Node configuration, Apache Storm is installed on the Master and Stream node.

In a Multi-instance Node configuration, Apache Storm is installed on multiple Master nodes and the Stream node.

- **Apache Spark2™**

Spark2 is a fast, general computing engine for Hadoop data. Spark2 executes the Analytics, providing a simple and expressive programming model to support a wide range of applications, including ETL, machine learning, stream processing, and graph computation.

Apache Spark2 is installed on the Master node.

- **Apache NiFi**

Apache NiFi is an easy to use, powerful, and reliable system to process and distribute data. It supports powerful and scalable directed graphs of data routing, transformation, and system mediation logic. NiFi is used for the Interaset Analytics data ingest.

Apache NiFi is installed on the NiFi node.

- **Apache Kafka**

Apache Kafka is a distributed publish-subscribe messaging system that is designed to be fast, scalable, and durable. In the InterSet cluster, Kafka is used for data transport to Storm.

Apache Kafka server is installed on the Stream node.

- **Apache ZooKeeper™**

ZooKeeper is a high-performance coordination service for distributed applications. In the InterSet cluster, ZooKeeper manages the coordination of the various component configurations.

ZooKeeper is installed on the Master node(s).

- **Elasticsearch**

Elasticsearch is an open source, broadly-distributable and easily-scalable enterprise-grade search engine. Elasticsearch houses all of the InterSet Analytics raw events, and provides all of the data that drives the user interface.

Elasticsearch is installed on the Search node.

- **Kibana**

Kibana is an open source data visualization plug-in for Elasticsearch. Kibana serves as the user interface and data exploration mechanism for Elasticsearch.

Kibana is installed on the Reporting node.

- **Nginx**

Nginx is a free, open-source, high-performance HTTP and reverse proxy server, as well as an IMAP/POP3 proxy server. Nginx is recognized for its high performance, stability, rich feature set, simple configuration, and low resource consumption.

Nginx is installed on the Reporting nodes.

- **FreeType**

FreeType is a public software library for rendering fonts. InterSet uses FreeType when rendering PDF Reports.

FreeType is installed on the Reporting node.

- **Chromium**

Chromium is an open-source Web browser with PDF generation capability in headless mode.

Chromium is installed on the Reporting node.

Component Distribution

The Interset production installation distributes the components across separate machines, or nodes, identified as follows:

- **Monitoring node**

The Monitoring node is where the cluster management server (Ambari or Cloudera Manager) is installed. Ambari and Cloudera Manager provide convenient Web user interfaces that simplify the deployment and management of the different components that make up the Interset solution.

- **Master node(s)**

The Master node(s) is used for various infrastructure components, and for starting the Interset Analytics process.

- **Search node(s)**

The Search node(s) is used for the Elasticsearch cluster and, in turn, are used by the Interset Reporting components.

- **Stream node(s)**

The Stream node is used for moving data to Storm for determining violations.

- **NiFi node(s)**

The NiFi node is used for ingesting data, moving data to the compute nodes for the Analytics, and later to the Search node(s) for Interset Reporting.

- **Compute node(s)**

The compute node(s) is used for running the Interset machine learning algorithms on the ingested data to detect anomalous behaviour and score entities. This process uses big-data components for analysis and storage.

- **Reporting node(s)**

The Reporting node provides a Web interface for Interset Reporting, and for further exploring and investigating anomalies identified by Interset Analytics.

Interset Configuration

Your Interset configuration will depend primarily on the amount of data to be analyzed. Interset recommends two basic configurations:

- Single-instance Node Configuration

In this configuration, there is only one instance of each node type.

- Multi-instance Node Configuration

In this configuration, there are multiple instances of various node types, depending on your data volumes.

Single-Instance Node Configuration

In a single-instance node configuration, the Interset and third-party components are distributed as follows:

Monitoring node	Master Node	Stream Node	Compute Node	Search Node	NiFi Node	Reporting Node			
Monitoring Server	ZooKeeper	Metrics Monitor	Metrics Monitor	Elasticsearch	Apache NiFi	Nginx			
	MetricsMonitor	Kafka Broker	Yarn Node Manager			Interset Reporting (Kibana, Nginx, Reporting Config, Chromium)			
	Metrics Collector	Hadoop Client Components	HBase RegionServer						
	Yarn App Timeline Server	Ingest .jar files for Streaming	HDFS DataNode						
	Yarn Resource Manager		Storm Supervisor						
	Yarn History Server								
	HBase Master								
	HDFS (S)NameNode								
	Spark History Server								
	Hadoop Client Components								
	Storm DRPC Server								
	Storm Nimbus								
	Storm UI Server								
	Interset Analytics (including Workflow)								

Multi-Instance Node Configuration

To maximize redundancy in the infrastructure and performance of the overall cluster, InterSet recommends the Multi-instance Node configuration illustrated below. This configuration can also be set up as a high availability (HA) system.

In this multi-instance node configuration, the InterSet and third-party components are distributed as follows:

Monitoring node	Master Node 1	Master Node 2	Master Node 3	NiFi Node (s)	Stream Node	Compute Node	Search Node	Reporting Node						
Monitoring Server	ZooKeeper	ZooKeeper	ZooKeeper	Apache NiFi	Metrics Monitor	Metrics Monitor	Elasticsearch	Nginx						
	Metrics Monitor	Metrics Monitor	Metrics Monitor		Kafka Broker	Yarn Node Manager		Interaset Reporting (Kibana, Nginx, Reporting Config, Chromium)						
	Metrics Collector	Hadoop Client Components	Hadoop Client Components		Hadoop Client Components	HBase RegionServer								
	Hadoop Client Components	Yarn App Timeline Server	Yarn Resource Manager		Interaset Ingest .jar files for Streaming	HDFS DataNode								
	HBase Master	Yarn Resource Manager	HBase Master			Storm Supervisor								
	HDFS NameNode	Yarn MapReduce2 History Server	HDFS JournalNode											
	HDFS JournalNode	HDFS JournalNode	Storm Passive Nimbus											
	ZooKeeperFC	HDFS NameNode												
	Storm DRPC Server	ZooKeeperFC												
	Storm Active Nimbus	Spark History Server												
	Storm UI Server													
	Interaset Analytics (including Workflow)													

For information and assistance calculating the optimal InterSet topology for your organization, please contact Micro Focus Customer Support at <https://softwaresupport.softwaregrp.com/>.

Appendix B: Explore the Analytics Using Sample Data

This appendix provides an end-to-end introduction to Intersect, using sample datasets. These tasks should only be performed after Intersect has been installed and configured in your environment.

For information about installing and configuring Intersect, please see the Intersect 5.9.2 *Installation and Configuration Guide*.

When you configure the sample data, you will:

- create a dedicated tenant for the sample data;
- configure NiFi to ingest sample data from universal Active Directory logs in .csv format;
- run Intersect Analytics; and
- review the top risky users and their associated most significant anomalies and underlying events.

The sample dataset includes four different data types (authentication, repository, Web proxy, and NetFlow) to demonstrate that Intersect provides an integrated view of risk across multiple datasets. As Intersect Analytics is designed to highlight the top risky entities, IT security team members can therefore prioritize their efforts and investigations.

The sample datasets are small, to reduce any impact on the performance or storage requirements of your Intersect production installation.

For the purposes of this example, we will only be ingesting sample authentication data.

Create the Samples Tenant

By creating a separate Samples tenant, you reduce any impact to your deployment. This Guide uses **int** as the tenant ID (TID) but you may choose an alternate tenant ID.

1. Log in to Intersect as **root** (the default password is **root**).
2. Click **Settings**  and then, in the dropdown list, select **Tenants**.
3. On the **Tenants** page, click **New**.
4. In the **Create a new Tenant** dialog, enter a new **Tenant ID** (in this case, **int**) and **Tenant Name**.
5. Click **Create Tenant**. The new tenant appears in the tenant list.

Create an Administrator User for the Samples Tenant

For users to log in to Intersect, they must have a username and password.

1. Log in to Intersect as **root** (the default password is **root**).
2. Click **Settings**  and then, in the dropdown list, select **Tenants**.

3. On the **Tenants** page, select the tenant you just created from the list on the left, and then click **New User**.
4. In the **Create a User for <tenant name>** dialog, enter the **Name**, **Username**, **Role**, and **Password** for the new user.

For the purposes of this exercise, select the **admin** role. The **admin** role can perform tasks such as configuring data sources, creating Workflows, and accessing the REST API, while the **user** role cannot.

5. Click **Create User**. The new user appears in the user list.

Copy the Interset Sample Datasets

The sample datasets are in the **sampledata** directory of the Interset repository.

The sample dataset includes three data types, each of which requires a separate data source configuration. In the Interset repository, locate the three sample datasets and copy them to a location that is accessible to the stream node(s). This location should have three directories, one for each dataset:

```
/opt/interset/sampledata/authentication
```

```
/opt/interset/sampledata/repository
```

```
/opt/interset/sampledata/webproxy
```

If you cannot access the repository, please contact Micro Focus Customer Support at <https://software.support.softwaregrp.com/>.

Configure the Sample Authentication Data Source

Follow the instructions in "[Configure a New Data Source](#)" on page 16 to configure NiFi and ingest the sample data.

Run Analytics

To run Analytics, run the following command on the **Master node** as the **Spark** user:

```
/opt/interset/analytics/bin/analytics.sh /opt/interset/analytics/conf/interset.conf
```

When Analytics has completed its run, log into the Interset user interface to view the data, using the new tenant and user credentials.

Appendix C: Security Best Practices

This appendix will provide you with security information regarding:

- Changing Default Account Passwords
- Firewall Configuration for Interset Servers
- Network Topology Recommendations
- Enabling SSL for Interset Reporting
- OS User Permissions & Patch Levels
- MD5 / SHA1 / SHA256 Checksums

Changing Default Account Passwords

The Interset system contains default accounts on Ambari, Investigator/Workflow. The following section will outline the accounts, location they are configured, use, and the impact of changing their credentials.


Ambari

By default, the Ambari management **admin** account username and password is set to **admin**.


You can change the password for the **admin** account using Ambari, or using a shell script. Changing the Ambari password will have no impact on other parts of the product.

To change the password using Ambari:


1. Log in to Ambari as **Admin** (or another administrator user if one is configured for the cluster).
2. Click **admin** to reveal the dropdown menu, and then select **Manage Ambari**.




- On the Welcome Page, click **Users**.



- On the **Users** page, click the **admin** user.



- On the **Users/admin** page, click **Change Password**.



- In the **Your Password** box enter the current password.
The default password is **admin**.
- In the **New User Password** box, enter your new password, and then confirm the password.
- Click **OK** to save the new password.

To change the password using a shell script, on the **Ambari node**:

- Run the following command:

```
/opt/interaset/bin/sysprep/scripts/update_passwords.sh
```

- Choose **option 1**.

3. Enter the Ambari Server URI, including the port (for example, `https://ambarinode.company.com:8443`).
4. Enter the existing admin password.
5. Enter the new password for the admin user.

When the password change is successful the following response will appear:

```
Password successfully changed!
```



Note: Ambari can also integrate with your existing LDAP. For more information about integrating Ambari with LDAP, please go to <https://docs.hortonworks.com/> and then use the Search box to locate the relevant topic.

Investigator/Workflow

Multiple Investigator/Workflow accounts are configured by default throughout the course of the Interset installation. These accounts are:

- root
- admin
- user
- Workflow_0

The Workflow_0 tenant is created to maintain a non-expiring session token, which allows the Workflow engine to communicate with the Interset Reporting server without having to log in.

Any of these accounts can be managed through the **update_passwords.sh** tool, or through **Swagger** as outlined in this guide.

To update these accounts via **update_passwords.sh**, do the following:

1. On the **Ambari node**, run the following command:

```
/opt/interset/bin/sysprep/scripts/update_passwords.sh
```

2. Choose **option 2**.
3. Enter the Reporting Server URI, including protocol (for example, `https://reporting.company.com`).
4. Enter the existing root password.
5. Enter the Tenant ID of the user you wish to update.



Note: root exists in the **adm** tenant.

6. Enter the name of the user for which you wish the change the password.

7. Enter the new password for the user noted in Step 6.

If the password change is successful the following response will appear:

```
Password successfully changed!
```




Note: If you have configured a multi-tenant environment, ensure that the user IDs and updated passwords are the same for each tenant. Please refer to the ["Configure Multi-tenant Authentication" on page 15](#) section for the instructions on multi-tenant environment configuration.


By modifying the Workflow_0 account, a new authorization token will need to be generated and used in the `/opt/interaset/rules/conf/rules.conf` file on the **Master node** where Analytics is installed.

To get a new authorization token:

1. In an Incognito window, log in to Interaset as an Administrator.
2. Click **Settings**  and then, in the dropdown list, select **API** to open the API in Swagger.
3. In the Swagger header, click **Authorize**.



4. In the **Available authorizations** dialog box, under **Select OAuth2.0 Scopes**, click **Authorize**.



The 'Available authorizations' dialog box has a title bar. Below the title, it says 'Select OAuth2.0 Scopes'. A paragraph explains that scopes grant different levels of access and provides a link 'Learn how to use'. Another paragraph states that the API requires specific scopes and asks the user to select which ones to grant to Swagger UI. It shows the 'Authorization URL' as '/login' and the 'flow' as 'implicit'. There are two 'Authorize' buttons and one 'Cancel' button at the bottom.

Available authorizations

Select OAuth2.0 Scopes

Scopes are used to grant an application different levels of access to data on behalf of the end user. Each API may declare one or more scopes. [Learn how to use](#)

API requires the following scopes. Select which ones you want to grant to Swagger UI.

Authorization URL: /login

flow: implicit

Authorize

Api key authorization

name: access_token


in: header

value:

Authorize

Cancel


5. In the login dialog box, enter the user credentials with the updated Workflow user password from the previous set of steps.
6. Click **info** to expand the section.



A table titled 'Interset Reporting API' with columns for actions and their corresponding operations. The 'info' row is highlighted with a red circle.

actions	Show/Hide	List Operations	Expand Operations
info	Show/Hide	List Operations	Expand Operations
rawEvents(tenantId)	Show/Hide	List Operations	Expand Operations
reports(tenantId)	Show/Hide	List Operations	Expand Operations
rules(tenantId)	Show/Hide	List Operations	Expand Operations
search(tenantId)	Show/Hide	List Operations	Expand Operations
search(tenantId)tags	Show/Hide	List Operations	Expand Operations
tenants	Show/Hide	List Operations	Expand Operations
theme	Show/Hide	List Operations	Expand Operations
tuning(tenantId)	Show/Hide	List Operations	Expand Operations

7. Click **GET /info/session** to expand the section.



The 'info' section is expanded, showing a list of endpoints. The 'GET /info/session' endpoint is highlighted with a red circle.

info	Show/Hide	List Operations	Expand Operations
GET /info/auth			Get authentication provider information
GET /info/build			Get version information
GET /info/session			Get session information

8. Click **Try it out!**

If the action is successful, **200** will appear in the **Response Code** box.

9. In the **Response Body**, copy the value associated with **"access_token"**, and make a note of listed tenants that need a token to be updated in their associated **rule.conf** files.

The **Response Body** should look similar to the following:

```
{
  "userId": "Workflow_0",
  "userDisplayName": "Workflows Admin",
  "persistentSessions": true,
  "accessToken": "4rpB5hG7KXBq41C8CvQqZgWUz0tq",
  "roles": [
    {
      "userId": "Workflow_0",
      "tenantId": "0",
      "role": "admin",
      "tenantName": "Default On-Premise Tenant",
      "hasSensorProxy": false,
      "features": [
        "showThemePanel"
      ]
    },
    {
      "userId": "Workflow_0",
      "tenantId": "ps1",
      "role": "admin",
      "tenantName": "tenant2",
      "hasSensorProxy": false,
      "features": [
        "showThemePanel"
      ]
    }
  ],
}
```



Important: Close the Incognito Window without logging out of Swagger. This will ensure that **"access_token"** remains valid for the Workflow user.

10. On the **Master node** where Analytics is installed, navigate to **/opt/interaset/rules/conf/** and edit the **rules.conf** file.
11. Copy the **"accessToken"** value from the response body generated in Step 8, and then paste that value into **"RulesApiServerAuthorizationToken"** section for each (required) tenant listed in the response body from Step 8, to update the value.



Note: The following command examples use the default tenant "0". The **rules.conf** file defaults to tenant **0** on an initial installation.

12. Validate that the change is valid by executing the following command:

```
/opt/interaset/rules/bin/workflow.sh --validate /opt/interaset/rules/conf/rules.conf
```

The expected output for token validation example is:

```
Testing config: RulesApiServer AuthorizationToken.....
```

```
SUCCESS: Testing RulesApiServer AuthorizationToken
```

13. Run the following command to kill an existing Workflow deployment:

```
/opt/interset/rules/bin/workflow.sh --kill /opt/interset/rules/conf/rules.conf
```

14. Redeploy Workflow using the following command:

```
/opt/interset/rules/bin/workflow.sh --deploy /opt/interset/rules/conf/rules.conf
```



Important: Depending on your configuration, the **Workflow** user may have a different account for every tenant in your environment; alternatively, multiple tenants may be consolidated under one **Workflow** user. For a single tenant configuration, these steps should be repeated for each tenant, if not originally configured with their own password(s).

Interset User Interface X-Frame-Options

To prevent Web browser framing attacks, Interset has configured the user interface Web page X-Frame-Options to prevent browsers from displaying any content in frames not originating with Interset.

Should your organization require the ability to embed the Interset Dashboard or Kibana in an inline frame element (iframe), you can disable the default X-Frame-Options.

1. On the Reporting node in your Interset cluster, navigate to the **/opt/interset/etc** directory, and open the **nginx.conf** file.
2. Locate, and then uncomment, the following line:

```
#add_header X-Frame-Options "SAMEORIGIN";
```

3. Save and close the file.

Firewall Configuration for Interset Servers

By default, the Interset installation requires iptables/firewalld to be disabled to prevent blockers in the installation. If these services (iptables/firewalld) are required, we recommend that you configure and start them after the installation is complete.

Sample configurations for the firewall rules, listed on a node-role basis, are available at <https://interset.zendesk.com/hc/en-us/articles/115001728148-How-To-Configure-iptables-firewalld-for-Interset-Servers> for both iptables and firewalld.

**Important:**

- These documents are provided as-is, and are for reference only. Depending on your environment, your Interset cluster may require configuration beyond the information provided in these documents. Interset is not responsible for any issues that arise from using this information.
- These configurations assume that all outgoing communication is permitted from each server.
- Spark executors use randomized port numbers; as a result, when your firewall is configured in this fashion, you will not be able to view stdout/stderr for executors (and possibly drivers) unless you whitelist specific hosts to have access to a massive port range. We suggest that you use **yarn logs -applicationId <application ID>** to retrieve logs wherever this firewall configuration is used.

Network Topology Recommendations

Interset recommends the use of a dedicated VLAN or network segment specifically for the Interset environment. This will allow for more-easily configurable intra-cluster communication, while explicitly restricting external access to the cluster, with the exception of those explicit allowances required for your business needs.

On Amazon Web Services (AWS), we suggest that this be achieved through the use of Security Groups that have the following configurations:

- All Interset servers are members of the same Security Group;
- The Security Group contains an exception for itself;
(For reference, please see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/security-group-rules-reference.html>)
- External access is restricted to port 22 as a management network (i.e., a corporate subnet), and to port 443 to a broader audience for access to the Reporting Server.

**Notes:**

- Through the use of SSH tunneling, port 22 can provide access to all other ports required for management purposes.
- If a more open approach is preferred, **All TCP** can be enabled for the management subnet.

These concepts can be applied to an on-premise installation; however, you will need to modify the concepts above to best fit with your network security requirements.

Enabling SSL for Interaset Reporting

Interaset supports the use of TLS encryption, which can be implemented with a FIPS 140-2 validated (NIST Certified with CMVP certificate) encryption system using the AES 256 algorithm.

By default, Interaset Reporting is configured to use 2048-bit TLS encryption with a self-signed certificate. We strongly suggest the use of a properly signed certificate in a production scenario, which would be provided by your company.

To update the configuration to use your specific certificate, do the following:

1. On the **Reporting node**, modify `/etc/nginx/conf.d/interaset.conf` to change the `ssl_certificate` and `ssl_certificate_key` values to point to your certificate and key.
2. Restart Nginx using the following command:

```
sudo systemctl restart nginx
```

Operating System User Permissions and Patch Levels

Interaset uses multiple local accounts on the machines in the cluster as a result of its own requirements and those of the supporting software (for example, HDP, Elasticsearch, Cassandra, and so on).

The Interaset 5.9.2 **interaset** user has a default password assigned during installation. (For information about the Interaset default user passwords, please see "[User Accounts and Permissions](#)" on page 80.) All other service accounts do not have set passwords, and must be accessed using **su** rather than by direct login.

By default, the **interaset** user is given unlimited sudo access, whereas the other accounts are restricted to managing only their own processes and directories. For more information about HDP services, please go to <https://docs.hortonworks.com/> and then use the **Search** feature to locate the relevant topic.

If direct root access via SSH was used during the Interaset installation, Interaset recommends that this direct root access via SSH be disabled when the installation is complete.

For Operating System processes, Interaset recommends that all servers be kept up to date with the latest supported security patches. Generally, you can achieve this by running **yum update**; however, your organization may have different processes in place.

Information regarding Interaset supported environments can be found in the Interaset 5.9.2 *Installation and Configuration Guide*.

MD5 / SHA1 / SHA256 Checksums


Interaset provides MD5, SHA1, and SHA 256 checksums of every file available on our repository. We strongly encourage that the individual performing the installations and upgrades of Interaset software validate the download archives or executables against the provided checksums (for example, for mirroring the repository) to ensure that the download was received intact and unchanged.

Appendix D: Modify Kerberos Ticket

Modifying Kerberos Ticket Lifetime in Ambari

The default Kerberos ticket lifetime is 24 hours (1 day). The following section details about how to modify Kerberos configurations in Ambari and increase the lifetime for the ticket.

1. Log in to Ambari as **Admin** (or another administrator user if one is configured for the cluster).
2. In the left panel, click **Kerberos** and then **CONFIGS**.
3. Scroll down on the config options and select **Advanced krb5-conf**. In the properties, modify the `ticket_lifetime`, which is set to 24h by default.



4. Click **Save**.

Modifying Kerberos Ticket Lifetime in Cloudera Manager

The default Kerberos ticket lifetime is 24 hours (1 day). The following section details about how to modify Kerberos configurations in Cloudera Manager and increase the lifetime for the ticket.

1. Log in to the Cloudera Manager as **Admin** (or another administrator user if one is configured for the cluster).
2. Click the **Administration** tab on the menu panel and select the **security** option.
3. Select **Kerberos Credentials** and click the **Configuration** tab. This will list all the current kerberos settings.
4. Scroll down to **Kerberos Ticket Lifetime** and modify the age of the ticket as per your requirement.

Appendix E: Change Passwords for a Secure Environment

This section primarily outlines the ways to change passwords for various components within a secure InterSet environment.

Changing TLS Certificates Password

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the InterSet installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **4** to select **4] Generate TLS Certificates**, and then press **Enter**.

```
Please select an installation option from the list above: 4
```

3. In the next steps, the installer will prompt you to enter new passwords for `truststore_password` and `keystore_password`. You will see a similar output once you run the option from the installer.

```
===== In function main_create_tls_cert: FUNCNAME=main_create_tls_cert main =====
===== In function init_tls_cert_gen: FUNCNAME=init_tls_cert_gen main_create_tls_cert main =====

Please enter new password for tls:keystore_password
new password: newpassword
Please re-enter password for tls:keystore_password
re-enter password: newpassword

Please enter new password for tls:truststore_password
new password: newpassword
Please re-enter password for tls:truststore_password
re-enter password: newpassword
...
...
Certificate is to be certified until Apr 12 19:17:42 2022 GMT (824 days)

Write out database with 1 new entries
Data Base Updated
Importing keystore /opt/interaset/installer/tmp/sslcrt/certout/out/ <reporting_node_fqn>-keystore.p12 to
/opt/interaset/installer/tmp/sslcrt/certout/out/ <reporting_node_fqn>-keystore.jks...
[Storing /opt/interaset/installer/tmp/sslcrt/certout/out/ <reporting_node_fqn>-keystore.jks]
...
...
Certificate was added to keystore
===== function main_create_tls_cert: FUNCNAME=main_create_tls_cert main Complete! =====
```

Re-Configure CDH/TLS by Changing Passwords

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the InterSet installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

- From the available installation options, type **6** to select **[6] Configure CDH/TLS**, and then press **Enter**.

```
Please select an installation option from the list above: 6
```

- In the next steps, the installer will prompt you to enter new passwords for cloudera_manager and kerberos. Following is an example of the output which you should see once you run the option:

```
Please select an installation option from the list above: 6

Please enter your existing password for cloudera_manager:admin
admin password: cloudera

Please enter new password for cloudera_manager:admin: newpassword
Please re-enter password for cloudera_manager:admin: newpassword
...
...
Please enter new password for kerberos:root
new password: newpassword
Please re-enter password for kerberos:root
re-enter password: newpassword
...
...
===== In function set_cm_tls_status: FUNCNAME=set_cm_tls_status main_cdh_tls_setup main_configure_secure_
cdh main =====
ADDING USE_TLS=true to /opt/interaset/installer/etc/config
Config file .interaset_profile transferred to <monitoring_node_fqn>
Config file .interaset_profile transferred to <computing_node_fqn>
Config file .interaset_profile transferred to <master_node_fqn>
Config file .interaset_profile transferred to <reporting_node_fqn>
Config file .interaset_profile transferred to <stream_node_fqn>
...
...
cm_restart complete!
===== In function main_configure_secure_cdh: FUNCNAME=main_configure_secure_cdh main Complete! =====
```

- Open a new browser using the **incognito window** and verify the password change by logging into the CDH manager `http://<monitoring_node_fqn>:7180` using your **new password**.

Create Key Tables

- In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

- From the available installation options, type **7** to select **[7] Create the Key tables**, and then press **Enter**.

```
Please select an installation option from the list above: 7
```

You should see the following output:

```
WARNING: no policy specified for i_spark@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_spark@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
Authenticating as principal root/admin@INTERSET.COM with password.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type aes256-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type aes128-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type des3-cbc-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type arcfour-hmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type camellia256-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type camellia128-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type des-hmac-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Entry for principal i_spark@INTERSET.COM with kvno 3, encryption type des-cbc-md5 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_spark.headless.keytab.
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_spark@INTERSET.COM" modified.
WARNING: no policy specified for i_nifi@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_nifi@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
Authenticating as principal root/admin@INTERSET.COM with password.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type aes256-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type aes128-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type des3-cbc-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type arcfour-hmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type camellia256-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type camellia128-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type des-hmac-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Entry for principal i_nifi@INTERSET.COM with kvno 3, encryption type des-cbc-md5 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.headless.keytab.
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_nifi@INTERSET.COM" modified.
WARNING: no policy specified for i_flume@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_flume@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
Authenticating as principal root/admin@INTERSET.COM with password.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type aes256-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type aes128-cts-hmac-sha1-96 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type des3-cbc-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type arcfour-hmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type camellia256-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type camellia128-cts-cmac added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type des-hmac-sha1 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Entry for principal i_flume@INTERSET.COM with kvno 3, encryption type des-cbc-md5 added to keytab
WRFFILE:/opt/interaset/installer/tmp/keytabs/i_flume.headless.keytab.
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_flume@INTERSET.COM" modified.
```

```

WARNING: no policy specified for i_reporting@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_reporting@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
Authenticating as principal root/admin@INTERSET.COM with password.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type aes256-cts-hmac-sha1-96 added to
keytab WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type aes128-cts-hmac-sha1-96 added to
keytab WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type des3-cbc-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type arcfour-hmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type camellia256-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type camellia128-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type des-hmac-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Entry for principal i_reporting@INTERSET.COM with kvno 3, encryption type des-cbc-md5 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_reporting.headless.keytab.
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_reporting@INTERSET.COM" modified.
WARNING: no policy specified for i_storm@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_storm@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
Authenticating as principal root/admin@INTERSET.COM with password.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type aes256-cts-hmac-sha1-96 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type aes128-cts-hmac-sha1-96 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type des3-cbc-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type arcfour-hmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type camellia256-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type camellia128-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type des-hmac-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Entry for principal i_storm@INTERSET.COM with kvno 3, encryption type des-cbc-md5 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_storm.headless.keytab.
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_storm@INTERSET.COM" modified.
===== main_create_keytabs Complete! =====

```

Changing Schema Registry Passwords

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **8** to select **[8] Schema Registry Setup**, and then press **Enter**.

```
Please select an installation option from the list above: 8
```

You should see the following output:.

```
Please select an installation option from the list above: 8
Registry DB setup starting.
Registry database installation was previously successful. Skipping. Complete!
If rerunning install on <monitoring_node_fqn>, delete /opt/interaset/log/dbs_registry_success.log first and try again.
```

3. At this point, quit the installer and on the **Monitoring node**, delete the `/opt/interaset/log/dbs_registry_success.log` file mentioned in the code above and start the installer again from Step 1. Once you select option 8 again, you will be asked to re-enter your existing password for **DB user**.

```
Please enter your existing password for DB user - Postgres: current_password
```

4. To further verify your current Postgres password, open a second terminal where Postgres has been setup for your environment and run the following command:

```
psql -U postgres scm
```

5. Enter your current password to go to the Postgres shell. Once logged in successfully, exit out of Postgres shell.
6. Return to the installer terminal from **step 4** and enter the new passwords as prompted.

```
must be alphanumeric upper/lowercase
Please enter new password for TLS Keystore password (must be alphanumeric upper/lowercase): newpassword
Please re-enter password for TLS Keystore password: newpassword
must be alphanumeric upper/lowercase
Please enter new password for TLS Truststore password (must be alphanumeric upper/lowercase): newpassword
Please re-enter password for TLS Truststore password: newpassword
...
...
Install of registry on <master_node_fqn> Complete!
You can optionally upload the Interaset schemas to the schema registry. If this is a fresh installation, this is
strongly recommended, however it will overwrite any local changes to existing schema versions.
Do you want to upload Interaset schemas to the schema registry? [Y]:
...
...
Performing Schema upload on <master_node_fqdn>
...
...
Registry Setup Is Complete!
```

7. Perform the following steps **manually** on a new **Monitoring node** terminal.

8. Change the database user registry password:

```

su - postgres
Last login: Thu Nov 28 12:06:35 EST 2019

-bash-4.2$ psql schema_registry registry
Password for user registry: old_password
psql (10.4)
Type "help" for help.

schema_registry=> ALTER USER registry WITH PASSWORD 'newpassword';
ALTER ROLE
schema_registry=> \q

-bash-4.2$ psql schema_registry registry
Password for user registry: new_password
psql (10.4)
Type "help" for help.

schema_registry=> \q

-bash-4.2$ exit
logout

```

9. On the MASTER node edit the `/opt/interaset/registry/conf/registry.yaml` file to modify the following configurations:

```

# BEGIN DB CONFIG
# jdbc provider configuration is:
storageProviderConfiguration:
  providerClass: "com.hortonworks.registries.storage.impl.jdbc.JdbcStorageManager"
  properties:
    db.type: "postgresql"
    queryTimeoutInSecs: 30
    db.properties:
      dataSourceClassName: "org.postgresql.ds.PGSimpleDataSource"
      dataSource.url: "jdbc:postgresql://<monitoring_node_fqn>/schema_registry"
      dataSource.user: "registry"
      dataSource.password: "new_password"

```

10. Restart schema registry on the MASTER node:

```

systemctl restart hwx-registry

systemctl status hwx-registry

• hwx-registry.service - Hortonworks Schema Registry
  Loaded: loaded (/usr/lib/systemd/system/hwx-registry.service; disabled; vendor preset: disabled)
  Active: active (running) since Thu 2020-01-09 13:14:58 EST; 8s ago
  Main PID: 21447 (java)
  CGroup: /system.slice/hwx-registry.service
          └─21447 java -Xmx1G -Xms1G -server -XX:+UseParNewGC -XX:+UseConcMarkSweepGC -XX:+CMSClassUnloadingEnabled
          -XX:+CMSScavengeBeforeRemark -XX:+DisableExplicitGC -Djava.awt.headless

```

11. Verify whether the changes took effect by accessing the registry server URL in incognito window mode: `https://<master_node_fqdn>:9190/ui/#/`

Changing Analytics Password

1. In a command console on the Monitoring node, run the following commands to navigate to the `/opt/interaset/installer` directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **9** to select **[9] Analytics installation**, and then press **Enter**.

```
Please select an installation option from the list above: 9
```

The following steps will prompt you to change the Analytics password.

```
Please select an installation option from the list above: 9
```

```
...
...
Configuring of Phoenix Jars Complete!
...
...
Starting analytics installation
...
...
must be alphanumeric upper/lowercase
Please enter new password for Elastic password (must be alphanumeric upper/lowercase): newpassword
Please re-enter password for Elastic password: newpassword
...
...
Install of analytics on master_node_fqdn Complete!

Analytics Node Configuration Complete!
```

Reinstalling Kafka Topic

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **10** to select **[10] Stream node(s) installation**, and then press **Enter**.

```
Please select an installation option from the list above: 10
```

The following steps will confirm Stream node configurations.

```
Please select an installation option from the list above: 10
```

```
Starting Stream Node installation
...
...
Install of kafka topics on stream_node_fqdn Complete!

Stream Node Configuration Complete!
```

Changing Elasticsearch Passwords

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interSet/installer** directory and launch the InterSet installation menu.

```
cd /opt/interSet/installer
```

```
./install.sh
```

2. From the available installation options, type **11** to select **[11] Search node(s) installation**, and then press **Enter**.

```
Please select an installation option from the list above: 11
```

The following steps will prompt you to change the Elasticsearch passwords:

```
Please select an installation option from the list above: 11
```

```
Starting Search Node installation
```

```
must be alphanumeric upper/lowercase
```

```
Please enter new password for Elastic Super password (must be alphanumeric upper/lowercase): newpassword
```

```
Please re-enter password for Elastic Super password: newpassword
```

```
Please enter new password for Elastic InterSet password (must be alphanumeric upper/lowercase): newpassword
```

```
Please re-enter password for Elastic InterSet password: newpassword
```

```
Please enter new password for Elastic Kibana password (must be alphanumeric upper/lowercase): newpassword
```

```
Please re-enter password for Elastic Kibana password: newpassword
```

```
Please enter new password for Logstash password (must be alphanumeric upper/lowercase): newpassword
```

```
Please re-enter password for Logstash password: newpassword
```

```
Installing and Configuring Elasticsearch on reporting_node_fqdn
```

```
...
```

```
...
```

```
An elasticsearch keystore already exists. Overwrite? [y/N]Created elasticsearch keystore in /etc/elasticsearch
```

```
...
```

```
...
```

```
Install of search on search_node_fqdn Complete!
```

```
Search Node Configuration Complete!
```

3. Run the following command to reset the Elasticsearch super user "elastic" password **manually**

```
curl -k -u elastic 'https://search_node_fqdn:9200/_xpack/security/_authenticate?pretty'
```

```

Enter host password for user 'elastic': new_password
{
  "error" : {
    "root_cause" : [
      {
        "type" : "security_exception",
        "reason" : "failed to authenticate user [elastic]",
        "header" : {
          "WWW-Authenticate" : [
            "ApiKey",
            "Basic realm=\"security\" charset=\"UTF-8\""
          ]
        }
      }
    ],
    "type" : "security_exception",
    "reason" : "failed to authenticate user [elastic]",
    "header" : {
      "WWW-Authenticate" : [
        "ApiKey",
        "Basic realm=\"security\" charset=\"UTF-8\""
      ]
    }
  },
  "status" : 401
}

```

The output shows that the password for user "elastic" is incorrect.

4. Now run the following command using the **old** Elasticsearch password:

```

curl -k -u elastic 'https://search_node_fqdn:9200/_xpack/security/_authenticate?pretty'
Enter host password for user 'elastic': old_password
{
  "username" : "elastic",
  "roles" : [
    "superuser"
  ],
  "full_name" : null,
  "email" : null,
  "metadata" : {
    "_reserved" : true
  },
  "enabled" : true,
  "authentication_realm" : {
    "name" : "reserved",
    "type" : "reserved"
  },
  "lookup_realm" : {
    "name" : "reserved",
    "type" : "reserved"
  }
}

```

The output shows that user "elastic" still use an old password!

5. Reset the password for the **elastic user**. To do this, create an alternate superuser and then authenticate as that user in order to change the password.

Follow the steps below to perform the change.

1. Stop your Elasticsearch node.

```
systemctl stop elasticsearch
systemctl status elasticsearch
• elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor preset: disabled)
   Active: inactive (dead) since Thu 2020-01-09 15:02:38 EST; 7s ago
```

2. Ensure that the file realm is available on your Elasticsearch node. If you are using a default X-Pack configuration for authentication, then the file realm is available and you don't need to modify anything. However, if you have explicitly configured the authentication realms in your `elasticsearch.yml` file, then you need to add a file realm to the list of realms.

3. Use the `bin/x-pack/users` command to create a new file-based superuser:

```
cd /usr/share/elasticsearch/bin
./elasticsearch-users useradd tmp_admin -p tmp_password -r superuser
```

4. Start Elasticsearch node

```
systemctl start elasticsearch
systemctl status elasticsearch
• elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor preset: disabled)
   Active: active (running) since Thu 2020-01-09 15:04:24 EST; 7s ago
```

5. Using `curl` command, reset the password for the elastic user:

```
curl -k -u tmp_admin -XPUT 'https://search_node_fqdn:9200/_xpack/security/user/elastic/_password?pretty'
-H 'Content-Type:
application/json' -d'
{
    "password" : "newpassword"
}
'
```

Enter host password for user 'tmp_admin': tmp_password

6. Verify the new password

```
curl -k -u elastic 'https://search_node_fqdn:9200/_xpack/security/_authenticate?pretty'
Enter host password for user 'elastic': newpassword
{
  "username" : "elastic",
  "roles" : [
    "superuser"
  ],
  "full_name" : null,
  "email" : null,
  "metadata" : {
    "_reserved" : true
  },
  "enabled" : true,
  "authentication_realm" : {
    "name" : "reserved",
    "type" : "reserved"
  },
  "lookup_realm" : {
    "name" : "reserved",
    "type" : "reserved"
  }
}
```

The password has changed now.

7. At this step if required , stop Elasticsearch and then remove the file realm from your elasticsearch.yml and remove the temp_admin user from the file realm.
6. You can now verify the changed password by accessing Elasticsearch indices. To do so open an incognito window and login to: **https://<search_node_fqn>:9200/_cat/indices/** as user "elastic" using the new password. The output should look as the following:

```
yellow open anomalies_t01_gen2_2017.19 SMLy2FegTNaSS4AkB9rUTw 5 1 50706 0 14.6mb 14.6mb
yellow open anomalies_0_gen3_2017.15 CA1zet8-TIeULbVS-DJ61A 5 1 61327 0 17.5mb 17.5mb
yellow open interset_ad_rawdata_0-2017-04-w14 eBORTtUdTJy3gfQto3tmzQ 5 1 41991 0 3.5mb 3.5mb
yellow open working_hours_0_2019-11-28_20:32:42 UYjvawbURLSsPIxscTws9w 5 1 4848 0 906.9kb 906.9kb
yellow open entity_relation_counts_t01_gen2_2017.13 fVL1aJU1QyCR7jJJeETHww 5 1 1223 0 512.2kb 512.2kb
yellow open risk_scores_0_gen1_2017.16 2c_9k8MuTdePUTjULHmRrg 5 1 20865 0 3.9mb 3.9mb
yellow open entity_relation_counts_0_gen1_2017.20 syCnReXiSfqcXsy62u1WJw 5 1 5898 0 1.3mb 1.3mb
yellow open entity_relation_counts_0_gen1_2017.19 t8BYUYVRNTMmoDIzAxwuiPQ 5 1 5973 0 1.4mb 1.4mb
yellow open anomalies_0_gen3_2017.19 iZjqbYiWRI0eB5gLD0ioVA 5 1 61908 0 17.9mb 17.9mb
yellow open interset_ad_rawdata_0-2017-05-w21 1gW8y9NmSISumNGRXQF95w 5 1 42130 0 3.9mb 3.9mb
green open .security-6 CorhltSQTDm20p85dXAR0g 1 0 5 0 22.2kb 22.2kb
...
...
```

Reconfiguring the Reporting Node

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interset/installer** directory and launch the InterSet installation menu.

```
cd /opt/interset/installer
```

```
./install.sh
```

2. From the available installation options, type **12** to select **[12] Reporting node(s) installation**, and then press **Enter**.

```
Please select an installation option from the list above: 12
```

You should see the following output:

```
Starting reporting installation
...
...
Install of reporting on reporting_node_fqdn Complete!
...
...
Reporting Node Configuration Complete!
```

3. To verify the changes, open a browser in incognito mode and log in to **http://<reporting_node_fqn>**.

Changing ZooKeeper Configurations and Stopping Storm

Before running Step 13 for Workflow configurations, you must run the following steps **manually**:

1. Modify ZooKeeper configuration:
 1. Log in to CDH Manager.
 2. Select **Zookeeper** and click **Configurations**.
 3. Select **Scope Server**.

4. Search for the **Java Configuration Options for Zookeeper Server** property.
 5. Add value **-Dzookeeper.skipACL=yes** into an empty field.
 6. Save the changes and restart ZooKeeper services.
2. Stop all Storm services:

```
a) On the MASTER node run commands:

sudo systemctl stop storm-nimbus
sudo systemctl stop storm-drpc
sudo systemctl stop storm-ui

sudo systemctl status storm-nimbus
sudo systemctl status storm-drpc
sudo systemctl status storm-ui

b) On the COMPUTE node run command:

sudo systemctl stop storm-supervisor
sudo systemctl status storm-supervisor

c) On the MASTER node run command:

/opt/cloudera/parcels/CDH/bin/zookeeper-client -server `hostname -f`:2181

[zk: <master_node_fqn>:2181(CONNECTED) 0] ls /
[cluster, controller, brokers, zookeeper, admin, isr_change_notification, log_dir_event_notification,
controller_epoch, consumers, latest_producer_id_block, config, storm-1, hbase]
[zk: <master_node_fqn>:2181(CONNECTED) 1] rmr /storm-1
[zk: <master_node_fqn>:2181(CONNECTED) 2] ls /
[cluster, controller, brokers, zookeeper, admin, isr_change_notification, log_dir_event_notification,
controller_epoch, consumers, latest_producer_id_block, config, hbase]
[zk: <master_node_fqn>:2181(CONNECTED) 3] quit
```

3. Start Storm services:

```
a) On the COMPUTE node run commands:

sudo systemctl start storm-supervisor
sudo systemctl status storm-supervisor

b) On the MASTER node run commands:

sudo systemctl start storm-nimbus
sudo systemctl start storm-drpc
sudo systemctl start storm-ui

sudo systemctl status storm-nimbus
sudo systemctl status storm-drpc
sudo systemctl status storm-ui
```

Changing Workflow Passwords

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **13** to select **[13] Workflow configuration**, and then press **Enter**.

Please select an installation option from the list above: **13**

3. After selecting Workflow Configuration, you will see the following output. Enter passwords as prompted.

```

Please select an installation option from the list above: 13

Starting Storm installation

-----
Please enter new password for TLS Truststore password:
Please re-enter password for TLS Truststore password:
Please enter new password for TLS Keystore password:
Please re-enter password for TLS Keystore password:
Installing and Starting Storm on <compute_node_fqn>
Authenticating as principal root/admin@INTERSET.COM with password.
WARNING: no policy specified for i_nimbus/<compute_node_fqn> defaulting to no policy
add_principal: Principal or policy already exists while creating "i_nimbus/<compute_node_fqn>".
Authenticating as principal root/admin@INTERSET.COM with password.
...
...
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_nimbus/<compute_node_fqn>" modified.
2020-01-09 12:32:52: createServiceKeyTab ... SUCCESS
...
...
Authenticating as principal root/admin@INTERSET.COM with password.
Principal "HTTP/<compute_node_fqn>" modified.
2020-01-09 12:32:53: createServiceKeyTab ... SUCCESS
...
...
Install of Storm on <compute_node_fqn> Complete!
Install of Storm on <master_node_fqn> Complete!

Storm Installation Complete!

Starting Workflow installation
...
Configuring and Starting Workflow on <compute_node_fqn>
configure_workflow.sh          100% 13KB 9.5MB/s 00:00
registry_config                100% 301 658.1KB/s 00:00
common_functions.sh            100% 25KB 7.6MB/s 00:00
dependencies.version           100% 179 133.3KB/s 00:00
secure.properties              100% 1205 282.5KB/s 00:00
i_storm.headless.keytab        100% 522 1.0MB/s 00:00
Configuring and Starting Workflow on <master_node_fqn>
configure_workflow.sh          100% 13KB 6.1MB/s 00:00
registry_config                100% 301 297.2KB/s 00:00
common_functions.sh            100% 25KB 24.0MB/s 00:00
dependencies.version           100% 179 453.3KB/s 00:00
secure.properties              100% 1205 2.6MB/s 00:00
i_storm.headless.keytab        100% 522 444.9KB/s 00:00
Install of workflow on <compute_node_fqn> Complete!

Waiting for installation to complete on all hosts...(test 27 of 50)
If necessary, check progress in /opt/interSet/log/interSetlog_configWorkflowLog_install.txt on <master_node_fqn>
For convenience, here are the last 5 lines from that log file:

=====
10295 [main] WARN o.a.s.u.Utils - STORM-VERSION new 1.2.1 old 1.2.1
11082 [main] INFO o.a.s.StormSubmitter - Finished submitting topology: Workflow_0

2020-01-09T12:35:31-0500 INFO - ==== VERIFY WORKFLOW ENGINE IS RUNNING ====
Retrieving status of Workflow: Workflow_0 ...

Waiting for installation to complete on all hosts...(test 28 of 50)
If necessary, check progress in /opt/interSet/log/interSetlog_configWorkflowLog_install.txt on <master_node_fqn>
For convenience, here are the last 5 lines from that log file:

```

```

=====
| TESTING: RulesApiServer AuthorizationToken...
|
*-----

SUCCESS: RulesApiServer AuthorizationToken is correctly configured.
...
...
Configuration for Email output is disabled, skipping test.

Configuration for SMS output is disabled, skipping test.

Configuration for Splunk output is disabled, skipping test.

Install of workflow on master_node_fqdn Complete!

Workflow Configuration Complete!

```

4. To verify the changes, you can log in to the reporting server UI and create a new Workflow as a test.

Changing Nifi Passwords

Before running installer step [14] for NiFi installation, perform the following steps **manually**.

Use the **kadmin** tool to change principal password for **i_nifi user** and run the following command on the **Monitoring Node**.

```

kadmin.local
Authenticating as principal root/admin@INTERSET.COM with password.
kadmin.local: list_principals
...
i_nifi/<stream_node_fqn>
i_nifi@INTERSET.COM
i_nifi_user@INTERSET.COM
...

kadmin.local: change_password i_nifi_user@INTERSET.COM
Enter password for principal "i_nifi_user@INTERSET.COM": newpassword
Re-enter password for principal "i_nifi_user@INTERSET.COM": newpassword
Password for "i_nifi_user@INTERSET.COM" changed.
kadmin.local: quit

```



Notes: You must specify a password that is not in the principal's password history.

Run the installer steps:

1. In a command console on the Monitoring node, run the following commands to navigate to the **/opt/interaset/installer** directory and launch the Interaset installation menu.

```
cd /opt/interaset/installer
```

```
./install.sh
```

2. From the available installation options, type **14** to select **[14] Nifi installation on Nifi node(s)**, and then press **Enter**.

Please select an installation option from the list above: **14**

3. Once the option is selected, change the passwords as prompted.

```

Please select an installation option from the list above: 14

Starting Nifi installation

-----

Please enter new password for TLS Keystore password: newpassword
Please re-enter password for TLS Keystore password: newpassword
Please enter new password for TLS Truststore password: newpassword
Please re-enter password for TLS Truststore password: newpassword
Please enter new password for Nifi Admin User (i_nifi_user): newpassword
Please re-enter password for Nifi Admin User (i_nifi_user): newpassword

Authenticating as principal root/admin@INTERSET.COM with password.
WARNING: no policy specified for i_nifi_user@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_nifi_user@INTERSET.COM".
Authenticating as principal root/admin@INTERSET.COM with password.
WARNING: no policy specified for i_nifi_user@INTERSET.COM; defaulting to no policy
add_principal: Principal or policy already exists while creating "i_nifi_user@INTERSET.COM".
Installing and Configuring NiFi on <stream_node_fqn>
Authenticating as principal root/admin@INTERSET.COM with password.
WARNING: no policy specified for i_nifi/<stream_node_fqn> defaulting to no policy
add_principal: Principal or policy already exists while creating "i_nifi/<stream_node_fqn>".
Authenticating as principal root/admin@INTERSET.COM with password.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type aes256-cts-hmac-sha1-96 added to
keytab WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type aes128-cts-hmac-sha1-96 added to
keytab WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type des3-cbc-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type arcfour-hmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type camellia256-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type camellia128-cts-cmac added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type des-hmac-sha1 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Entry for principal i_nifi/<stream_node_fqn> with kvno 7, encryption type des-cbc-md5 added to keytab
WRFILE:/opt/interaset/installer/tmp/keytabs/i_nifi.<stream_node_fqn>.service.keytab.

Authenticating as principal root/admin@INTERSET.COM with password.
Principal "i_nifi/<stream_node_fqn>" modified.
2020-01-08 14:02:46: createServiceKeyTab ... SUCCESS
nifi_config                                100% 1442    484.5KB/s    00:00
registry_config                           100% 301    464.5KB/s    00:00
common_functions.sh                       100% 25KB    22.6MB/s    00:00
install_nifi.sh                           100% 30KB    12.3MB/s    00:00
secure.properties                         100% 1205     3.1MB/s    00:00
jq-linux64                               100% 3861KB   78.5MB/s    00:00
i_nifi.<stream_node_fqn>.service.keytab

Install of nifi on <stream_node_fqn> Complete!

NiFi Node(s\ ) Configuration Complete!
PLEASE OPEN WEB BROWSER AND NAVIGATE TO: https://<stream_node_fqn>:8085/nifi
AND LOGIN AS THE USER: i_nifi_user
You may need to wait up to a minute for the cluster to vote on master election and fully start.

```

```
Verify that NiFi admin web UI is available open up browser incognito window and login to https://nifi_node_
fqdn:8085/nifi/ as i_nifi_user using new password.
```

System Configuration Files where Passwords Change

The following section lists all the system files where password changes are reflected:

1. On the NiFi node: `/opt/nifi/current/conf/nifi.properties`

```
# security properties #
nifi.sensitive.props.key=
nifi.sensitive.props.key.protected=
nifi.sensitive.props.algorithm=PBEWITHMD5AND256BITAES-CBC-OPENSSL
nifi.sensitive.props.provider=BC
nifi.sensitive.props.additional.keys=

nifi.security.keystore=/etc/security/interset/sslcert/localhost-keystore.jks
nifi.security.keystoreType=JKS
nifi.security.keystorePasswd=newpassword
nifi.security.keyPasswd=
nifi.security.truststore=/etc/security/interset/sslcert/all-truststore.jks
nifi.security.truststoreType=JKS
nifi.security.truststorePasswd=newpassword
nifi.security.needClientAuth=false
nifi.security.user.authorizer=managed-authorizer
nifi.security.user.login.identity.provider=kerberos-provider
nifi.security.ocsp.responder.url=
nifi.security.ocsp.responder.certificate=
```

2. On the Master node: `/opt/interset/registry/conf/registry.yaml`

```
# BEGIN DB CONFIG
# jdbc provider configuration is:
storageProviderConfiguration:
providerClass: "com.hortonworks.registries.storage.impl.jdbc.JdbcStorageManager"
properties:
db.type: "postgresql"
queryTimeoutInSecs: 30
db.properties:
  dataSourceClassName: "org.postgresql.ds.PGSimpleDataSource"
  dataSource.url: "jdbc:postgresql://<monitoring_node_fqn>/schema_registry"
  dataSource.user: "registry"
  dataSource.password: "newpassword"
# END DB CONFIG

#swagger configuration
swagger:
resourcePackage: com.hortonworks.registries.schemaregistry.webservice

#enable CORS, may want to disable in production
enableCors: true
#Set below 3 properties if server needs a proxy to connect. Useful to download mysql jar
#httpProxyUrl: "http://proxyHost:port"
#httpProxyUsername: "username"
#httpProxyPassword: "password"

# BEGIN CONNECTOR
server:
  applicationConnectors:
  - type: https
    port: 9190
    keyStorePath: /etc/security/interaset/sslcert/localhost-keystore.jks
    keyStorePassword: newpassword
    trustStorePath: /etc/security/interaset/sslcert/all-truststore.jks
    trustStorePassword: newpassword
    needClientAuth: false
    validateCerts: false
    validatePeers: false
  adminConnectors:
  - type: https
    port: 9191
    keyStorePath: /etc/security/interaset/sslcert/localhost-keystore.jks
    keyStorePassword: newpassword
    trustStorePath: /etc/security/interaset/sslcert/all-truststore.jks
    trustStorePassword: newpassword
    needClientAuth: false
    validateCerts: false
    validatePeers: false
# END CONNECTOR
```

3. On the Master node: /opt/interaset/analytics/conf/interaset.conf

```
# Elasticsearch SSL configuration parameters
esXPackUser=elastic:new_password
keystorePath=/etc/security/interaset/sslcert/localhost-keystore.jks
keystorePassword=new_password
sslEnabled=true
selfSignedAllowed=true
truststorePath="/etc/security/interaset/sslcert/all-truststore.jks"
truststorePassword = new_password
```

4. On the Search (and Reporting) nodes: /etc/kibana/kibana.yml

```
elasticsearch.username: "kibana"
elasticsearch.password: "new_password"
```

5. On the Search (and Reporting) nodes: /etc/elasticsearch/elasticsearch.yml

```
xpack.ssl.keystore.password: interset
xpack.security.http.ssl.keystore.password: new_password
xpack.security.http.ssl.truststore.password: new_password
xpack.security.transport.ssl.keystore.password: new_password
xpack.security.transport.ssl.truststore.password: new_password
```

6. On the Search (and Reporting) nodes: /opt/interaset/etc/investigator.yml

```
# KB User Begin
kibana:
  username: kibana
  password: new_password
# KB User End
#####
#
# Elasticsearch
#
#####
#search-db: elasticsearch
db-elasticsearch:
  esHosts: <reporting_node_fqn>

# ES Secure configuration Begin
esUser: elastic
esPassword: new_password
esKeystorePassword: new_password
esKeystorePath: /etc/security/interaset/sslcert/localhost-keystore.jks
esEnableSSL: true
# ES Secure configuration End
```

Updating NiFi Controller Services and Processors Configuration After System Password Rotation

Before running ingest and after completing the password rotation steps, you have to perform the following steps **manually**

Open a new browser and navigate to the NiFi UI **<nifi_node_fqdn>:8085/nifi/**. Log in as the **i_nifi_user** user with the new password.

1. Right-click on the main canvas and select **Configure**.

On the NiFi Flow Configuration page, select the **CONTROLLER SERVICES** tab. NiFi global Controller Services will open up with a scope **NiFi Flow**. Fix the **invalid Controller Services** by updating their properties and enable all disabled Controller Services.

2. Open the processor group where your ingest to Kafka, HBase, and Elasticsearch processors reside. Right-click on the canvas and select **Configure**.

On the **Processor Group Name Configuration** page, select the **CONTROLLER SERVICES** tab.

NiFi Controller Services will open up with a scope of your **Processor Group Name**, along with the **global NiFi Flow**.

Fix **invalid** Controller Services by updating their properties and enable all disabled Controller Services.

3. If you are ingesting into a new tenant, open the **PublishKafkaRecord** processor and update the **Topic Name** field value to point to a new tenant Kafka topic.
4. Open the **AsyncActiveDirectoryHbaseProcessor** processor and update the **Tenant ID (tid)** field value if a new tenant is used.
5. Open the **LoadElasticsearch** processor and update the **Security User**, **Truststore Password**, and the **Keystore Password** fields with new passwords. Also, update the **Index Name**, **Tenant ID (tid)**, and **Data Instance ID (did)** values if a new tenant is being used.
6. If you are ingesting violations, open the relevant processor group.
7. Open the **Violations to ES** processor and update the **Security User**, **Truststore Password** and **Keystore Password** values with new passwords. Also, update **Index Name**, **Tenant ID (tid)**, and **Data Instance ID (did)** values for the new tenant.
8. Open the **Violations to HBase** processor and update the **Tenant ID (tid)** field value if a new tenant is being used.

At this point, the NiFi configuration will be up-to-date and ready to run data ingest.

Updating Interaset Analytics Configuration After System Password Rotation

Before running analytics and after ingest, perform the following steps manually:

1. 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 klist
klist: No credentials cache found (filename: /tmp/krb5cc_982)

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@INTERSET.COM
```

Verify the result:

```
sudo -u hbase klist
Ticket cache: FILE:/tmp/krb5cc_982
Default principal: hbase/<compute_node_fqn>

Valid starting    Expires          Service principal
01/13/2020 10:44:42 01/14/2020 10:44:42  krbtgt/INTERSET.COM@INTERSET.COM
renew until 01/20/2020 10:44:42
```

2. On the **Compute node**, verify that the user hbase can access **Phoenix database server** by running the following command:

```
sudo -u hbase /opt/cloudera/parcels/APACHE_PHOENIX/bin/phoenix-sqlline.py
```

Use Ctrl+D to quit from the command prompt.

3. Log in to the **Master node** and verify that ticket is valid for **i_spark** user by running the following command:

```
sudo -u i_spark klist
```

```
Ticket cache: FILE:/tmp/krb5cc_1001
Default principal: i_spark@INTERSET.COM

Valid starting      Expires              Service principal
01/13/2020 14:15:33 01/14/2020 14:15:33  krbtgt/INTERSET.COM@INTERSET.COM
renew until 01/20/2020 14:15:33
```

If the ticket has expired, re-generate it by using the following command:

```
sudo -u i_spark kinit i_spark@INTERSET.COM -kt /etc/security/interaset/keytab/i_spark.headless.keytab
```

4. On the **Master node**, verify that the **i_spark** user can access Phoenix database server by running the following command:

```
sudo -u i_spark /opt/cloudera/parcels/APACHE_PHOENIX/bin/phoenix-sqlline.py
```

5. Re-running step [9] **Analytics installation** of the installer will update passwords only in the default **/opt/interaset/analytics/conf/interaset.conf** file. Therefore, if applicable you may need to manually update tenant-specific configuration file(s).

Edit tenant-specific **/opt/interaset/analytics/conf/tenant_specific_name.conf** file(s) and update **Elasticsearch SSL configuration parameters** section with new passwords.

```
# Elasticsearch SSL configuration parameters

esXPackUser=elastic:new password
keystorePath=/etc/security/interaset/sslcert/localhost-keystore.jks
keystorePassword=new password
sslEnabled=true
selfSignedAllowed=true
truststorePath="/etc/security/interaset/sslcert/all-truststore.jks"
truststorePassword = new password
```

6. After completing all the steps for analytics verification, start analytics by running the following command:

```
sudo -u i_spark /opt/interaset/analytics/bin/analytics.sh /opt/interaset/analytics/conf/interaset.conf
```

Appendix F: Migrate an Existing External PostgreSQL Server Database to Another Server

This section primarily outlines the ways to migrate an existing PostgreSQL server database from one host another host .

Backing Up and Restoring PostgreSQL Databases

Database backup and restore can be done for each database. In this case, you have to create backups for **metastore**, **schema_registry**, and **scm** databases.

- **For a Single Database:**

Use the following commands to back up and restore:

For Example:

For Backup

```
pg_dump -U postgres -W -C -d metastore > /var/tmp/metastore_db.sql
```

Password:

For Restore:

```
psql -U postgres -W -d metastore -f /var/tmp/metastore_db_ddl.sql
```

- **For All databases:**

You can back up and restore all databases hosted on the PostgreSQL server at once. To do so, follow the steps below:

1. Create a full backup using the following command:

Example:

```
pg_dumpall -U postgres -W -f /var/tmp/all_databases_backup.sql
```

Password:

Password:

Password:

Password:

Password:

Password:



Notes: You will be asked to enter the password several times.

2. Verify the backed up files by using the following commands.

```
ls -al /var/tmp/ | grep .sql
```

```
-rw-rw-r--. 1 interset interset 38031439 Jan 21 19:08 all_databases_backup.sql
```

3. To restore a full backup use the following command: Example:

```
psql -U postgres -W -f /var/tmp/all_databases_backup.sql
```

Transferring the Database to Another Server

1. Based on your environment, use one of the backup methods described above to back up databases.
2. **Stop all cluster services** on the Cloudera Manager.
3. Navigate to the `/opt/interaset/log` directory on the **Monitoring node** and **delete (or rename)** `interasetlog_postgres_serverLog.txt` and `postgres_server_success.log` files.



Notes: Do not delete `dbs_CDH_success.log` and `interasetlog_install_dbs_CDHLog.txt`. Otherwise, the installer will attempt to create scm and metastore databases. However, if the logs are deleted accidentally in the process and the database was created, you will need to drop them before running the restore command.



Notes: The **database roles**, should not be dropped at this point.

4. Open the `/opt/interaset/installer/etc/config` file and edit the Postgres host by providing a new host-name.
5. Run the `install.sh` script:

```
cd /opt/interaset/installer
```

```
./install.sh
```

- Select option c to transfer the updated configuration to all hosts.
- Run installer step: [1] Initial server config (set up SSH, `offline_setup` script, databases...)

After performing host InterSet baseline installation, the installer will come to a point of installing the database server on a new host.



Notes: Type in the **same password** used previously for the PostgreSQL database server installation.

```

...
...
Install Postgres Server? [y/n]: y

Postgres server setup starting.

must be alphanumeric upper/lowercase
Please enter new password for DB user - Postgres (must be alphanumeric upper/lowercase):
Please re-enter password for DB user - Postgres:

SSH Successful and install postgres server script started on rvoyer-cdh-endpoint-1.ad.interset.com!
Postgres server install on rvoyer-cdh-endpoint-1.ad.interset.com successful!
Postgres server install Complete!

Postgres client setup starting.

SSH Successful and install postgres client script started on rvoyer-cdh-ambari-1.ad.interset.com!
Postgres client install on rvoyer-cdh-ambari-1.ad.interset.com successful!
Postgres client install was previously successful on rvoyer-cdh-master-1.ad.interset.com. Skipping.
If rerunning install on rvoyer-cdh-master-1.ad.interset.com, delete /opt/interset/log/postgres_client_
success.log first and try again.
Postgres client install was previously successful on rvoyer-cdh-master-2.ad.interset.com. Skipping.
If rerunning install on rvoyer-cdh-master-2.ad.interset.com, delete /opt/interset/log/postgres_client_
success.log first and try again.
Postgres client install Complete!

CDH Monitoring DB setup starting.
CDH database install was previously successful. Skipping. Complete!
If rerunning install on rvoyer-cdh-ambari-1.ad.interset.com, delete /opt/interset/log/dbs_CDH_success.log first
and try again.

```

- Verify the status of the server to ensure that it is now running on the new PostgreSQL database server host:

```

systemctl status postgresql-10
• postgresql-10.service - PostgreSQL 10 database server
Loaded: loaded (/usr/lib/systemd/system/postgresql-10.service; enabled; vendor preset: disabled)
Active: active (running) since Thu 2020-01-16 20:00:21 UTC; 4min 21s ago

```

- Log in to the Postgres server installed on the new host. Now drop the **metastore** and **scm data-bases**:

```

sudo su postgres
bash-4.2$ psql -l
could not change directory to "/home/interset": Permission denied
Password:
List of databases
Name      | Owner   | Encoding | Collate  | Ctype    | Access privileges
-----+-----+-----+-----+-----+-----
postgres  | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | 
template0 | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/postgres      +
template1 | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | postgres=CTc/postgres
(3 rows)

```

- Copy your **database backup(s)** from the original Postgres server host to a new Postgres server host:

```

scp all_databases.sql interset@rvoyer-cdh-1xd-ambari-1.ad.interset.com:/tmp

```

9. Restore all database(s) by running the following command:

```
psql -U postgres -W -f /tmp/all_databases_backup.sql
```

Note: Ignore "role < role_name> already exists" error(s). This is expected.

10. Verify all the databases are restored back:

```
psql -l -U postgres
List of databases
Name          | Owner   | Encoding | Collate  | Ctype    | Access privileges
-----+-----+-----+-----+-----+-----
metastore     | metastore | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
postgres     | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
schema_registry | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =Tc/postgres      +
               |          |          |          |          | postgres=CTc/postgres+
               |          |          |          |          | registry=CTc/postgres
scm           | scm      | UTF8     | en_US.UTF-8 | en_US.UTF-8 |
template0     | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | =c/postgres       +
               |          |          |          |          | postgres=CTc/postgres
template1     | postgres | UTF8     | en_US.UTF-8 | en_US.UTF-8 | postgres=CTc/postgres+
               |          |          |          |          | =c/postgres
(6 rows)
```

11. Verify the list of database users:

```
sudo su postgres

bash-4.2$ psql postgres postgres
Password for user postgres:
psql (10.4)
Type "help" for help.

postgres=# \du
List of roles
Role name | Attributes | Member of
-----+-----+-----
metastore |            | {}
postgres | Superuser, Create role, Create DB, Replication, Bypass RLS | {}
registry  |            | {}
scm       |            | {}
```

12. Verify the database users passwords:

```
bash-4.2$ psql metastore metastore
Password for user metastore:
psql (10.4)
Type "help" for help.

metastore=> \q

bash-4.2$ psql scm scm
Password for user scm:
psql (10.4)
Type "help" for help.

scm=> \q

bash-4.2$ psql schema_registry registry
Password for user registry:
psql (10.4)
Type "help" for help.

schema_registry=> \q
```

13. Update the schema registry database host.

- Open a new browser and navigate to the **Schema Registry UI**: https://<master_fqdn>:9190/ui/#/ and verify if its running.
- On the **Master node**, **stop the registry service**:

```
sudo systemctl stop hwx-registry
```

- Refresh the browser window to ensure it is not connected.
- Edit **/opt/interaset/registry/conf/registry.yaml** and change the database server in **DB CONFIG** section:

```
# BEGIN DB CONFIG
# jdbc provider configuration is:
storageProviderConfiguration:
providerClass: "com.hortonworks.registries.storage.impl.jdbc.JdbcStorageManager"
properties:
db.type: "postgresql"
queryTimeoutInSecs: 30
db.properties:
dataSourceClassName: "org.postgresql.ds.PGSimpleDataSource"
dataSource.url: "jdbc:postgresql://rvoyer-cdh-lxd-endpoint-1.ad.interaset.com/schema_registry"
dataSource.user: "registry"
dataSource.password: "interaset"
# END DB CONFIG
```

- Stop the PostgreSQL server on the original node and verify the service status by running the following commands:

```
sudo systemctl stop postgresql-10
sudo systemctl status postgresql-10
```

- Start the schema registry server on the **Master node** by using the following command:

```
sudo systemctl start hwx-registry
```

- Refresh your browser window to verify that the server is up and running.

14. On the **Monitoring node**, stop the **Cloudera server**:

```
sudo systemctl stop cloudera-scm-server
```

15. Edit **com.cloudera.cmf.db.host** section of the **/etc/cloudera-scm-server/db.properties** configuration file. Example:

```
# Auto-generated by scm_prepare_database.sh on Wed Jan 15 19:00:11 UTC 2020
#
# For information describing how to configure the Cloudera Manager Server
# to connect to databases, see the "Cloudera Manager Installation Guide."
#
com.cloudera.cmf.db.type=postgresql
com.cloudera.cmf.db.host=rvoyer-cdh-endpoint-1.ad.interaset.com
com.cloudera.cmf.db.name=scm
com.cloudera.cmf.db.user=scm
com.cloudera.cmf.db.setupType=EXTERNAL
com.cloudera.cmf.db.password=interaset
```

16. Start the Cloudera server and verify its status by running the following commands:

```
sudo systemctl start cloudera-scm-server

sudo systemctl status cloudera-scm-server
● cloudera-scm-server.service - Cloudera CM Server Service
Loaded: loaded (/usr/lib/systemd/system/cloudera-scm-server.service; enabled; vendor preset: disabled)
Active: active (running) since Thu 2020-01-16 21:28:18 UTC; 8s ago
```

17. Log in to the CDH Manager UI and update the Hive configuration:
 - Select **Hive** and click **Configuration**.
 - Under Category, select **Hive Metastore Database**.
 - Verify **Hive Metastore Database Type** is set to **PostgreSQL**.
 - Verify **Hive Metastore Database Name** is set to **metastore**.
 - Verify **Hive Metastore Database Port** is set to **5432**.
 - Verify **Hive Metastore Database User** is set to **metastore**.
 - Verify **Hive Metastore Database Password** is set to the one configured.
18. On the CDH manager, stop all cluster services.
19. Open a new terminal for each registered host on the cluster and restart the **Cloudera Agent**, using the following command:

```
sudo systemctl restart cloudera-scm-agent
```
20. Once all agents have restarted, open a terminal for the **Monitoring node** and restart the **Cloudera Server**:

```
sudo systemctl restart cloudera-scm-server
```
21. Log in to the CDH Manager Web UI and restart the **Cloudera Management Service**.
22. Once the Cloudera Management Service has restarted, restart all the cluster services.

Index

accounts

- user 80

action mappings

- configuring new 34

- updating 34

administering

- Interset 12

- Interset for users 40

advanced

- configuration options 75

- features 73

alerts

- configuring malware 35

- configuring networking 35

- configuring third-party 35

- configuring universal 35

analytics

- enabling windowed 28

- tuning 55

Apache Ambari Client 83

Apache Ambari Metrics 82

Apache Ambari Server 82

Apache HBase 83

Apache HDFS 83

Apache Kafka 84

Apache NiFi 83

Apache Spark 83

Apache Storm 83

Apache ZooKeeper 84

attaching

- users from other tenants 16

audience

- intended 10

authentication

- configuring multi-tenant 15

banners

- custom text 77

bots

- managing 52

branding

- custom 75

Change Password 102

changing

- user role 15

Chromium 84

Cloudera Manager Agents 83

Cloudera Manager Server 83

cluster

- components 82

components

- distribution 85

- Interset 82

- third-party 82

Compute

- node(s) 85

configuration

- advanced 75

- configuration file
 - editing Workflow 41
- configuring
 - "peek-back window" 29
 - CSV delimiters 80
 - DLP alerts 35
 - interset.conf 27
 - malware alerts 35
 - multi-tenant authentication 15
 - networking alerts 35
 - new action mappings 34
 - new data sources 16
 - rules.conf 41
 - sample data 90
 - third-party alerts 35
 - universal alerts 35
 - value separators 80
- creating
 - new tenants 13
 - new users 14
 - Workflows 40
- CSV files
 - value separator 80
- custom
 - banner text 77
 - themes 75
- custom schema
 - configuring with Workflow 45
- data
 - resolving multiple entities 35
- data sources
 - configuring new 16
- deleting
 - tenants 14
- delimiters
 - configuring 80
- disabling
 - X-Frame-Options 98
- distribution
 - components 85
- DLP alerts
 - configuring 35
- editing
 - Workflow configuration file 41
- Elasticsearch 84
- enabling
 - windowed analytics 28
- entities
 - resolving multiple in source data 35
- examples
 - Workflow 50
- features
 - advanced 73
- files
 - configuring interset.conf 27
- flags 66
- FreeType 84
- installing
 - sample data 90

- intended
 - audience 10
- Interset
 - administering 12
 - administering for users 40
 - components 82
 - using 58
- interset.conf
 - configuring 27
- Introduction 8
- Kerberos
 - configuring in NiFi 23-25
- Kerberos Ticket Lifetime 101
- Kibana 84
- malware
 - configuring alerts 35
- managing
 - BOT-like users 52
 - bots 52
 - tenants 12
 - users 12
- mappings
 - configuring new action 34
- Master
 - node(s) 85
- Migrating External PostgreSQL 124
- modifying
 - user role 15
- Monitoring
 - node 85
- networking
 - configuring alerts 35
- new
 - creating tenants 13
- new data sources
 - configuring 16
- Nginx 84
- NiFi
 - configure SSL controller service 17
 - configuring Kerberos 23-25
 - configuring the example flow 20
 - copying a process group 19
 - enabling controller services 20
 - node(s) 85
 - specifying schema registry location 19
 - starting the data flow 25
- nodes
 - Compute 85
 - Master 85
 - Monitoring 85
 - NiFi 85
 - Reporting 85
 - Search 85
 - Stream 85
- notifications
 - Workflow 46
- options
 - advanced configuration 75
- outcomes
 - Workflow 46

permissions	tenants
user 80	creating new 13
removing	deleting 14
tenants 14	managing 12
Reporting	removing 14
node 85	text
resolving	custom banners 77
multiple entities in source data 35	themes
roles	custom 75
modifying user 15	third-party
rules.conf	components 82
configuring 41	configuring alerts 35
sample data	third-party alerts
configuring 90	configuring 35
installing 90	tuning
scaling	analytics 55
Workflows 48	universal alerts
scenarios	configuring 35
Workflow 46	updating
schema registry	action mappings 34
specify location in NiFi 17, 19	user
Search	accounts 80
node 85	permissions 80
source data	user-defined tags 61
resolving multiple entities 35	users
Stream	creating 14, 16
node(s) 85	managing bot-like 52
tags	modifying role 15
user-defined 61	using
	Analytics 58

- Intersect 58
- Workflows 73
- verifying
 - Workflows 40
- windowed analytics
 - configuring "peek-back window" 29
 - enabling 28
- Workflow
 - configuring custom schema 45
 - creating 40, 48
- Workflow configuration
 - editing 41
- Workflows
 - examples 50
 - notifications and outcomes 46
 - scaling 48
 - scenarios 46
 - using 73