Micro Focus Security ArcSight Intelligence

Software Version: 6.3.0

Release Notes

Document Release Date: May 2021
Software Release Date: May 2021

Legal Notices

Micro Focus
The Lawn
22-30 Old Bath Road
Newbury, Berkshire RG14 1QN
UK

https://www.microfocus.com

Copyright Notice

Trademark Notices

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number
- Document Release Date, which changes each time the document is updated
- Software Release Date, which indicates the release date of this version of the software

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

https://www.microfocus.com/support-and-services/documentation

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://softwaresupport.softwaregrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwaregrp.com/
ArcSight Product Documentation	https://community.softwaregrp.com/t5/ArcSight-Product-Documentation/ct-p/productdocs

Contents

Introduction

This release notes document for Intelligence 6.3 includes the following:

- What's New?
- Known Issues
- Technical Requirements
- Downloading Intelligence
- Installing Intelligence
- Upgrading Intelligence
- Licensing Information
- Contacting Micro Focus

What's New?

The following sections outline the key features and the issues resolved in this release:

- FIPS 140-2 Compliance
- Custom Model Support
- Fusion User Management
- Cloud-native Deployment
- Policy for Retaining Events
- Integration with Fusion
- Resolved Issues

FIPS 140-2 Compliance

ArcSight Intelligence is now FIPS 140-2 compliant. Currently, all sub-components in the Intelligence architecture operate in the FIPS 140-2 mode. For all sub-components, the FIPS 140-2 mode is always enabled and you do not have the option to disable it.

For more information, see Using ArcSight Platform and Products in FIPS Mode in the Administrator's Guide for ArcSight Platform.

Custom Model Support

Intelligence provides support for custom machine learning (ML) models. This feature enables you to enhance Intelligence with models that provide analytics tailored to your unique environments. It also provides a method for extending Intelligence analytics to address new use cases such as the detection of new patterns of unusual behavior. You can also customize the alert template associated with a custom model. An alert template provides a way to describe an anomaly in the Intelligence UI by using the textual information provided as part of the alert template's metadata.

For more information, see Enabling Custom Model Support in the Administrator's Guide for ArcSight Platform.

Fusion User Management

You now use Fusion to create users and assign relevant roles and permissions for the users to access Intelligence.

Cloud-native Deployment

You can now deploy and configure Intelligence in the following cloud environments:

- **Azure** - Leverages the capabilities of the Microsoft Azure cloud platform.
- **Amazon Web Services (AWS)** - Leverages its cloud-native services and capabilities.

For more information, see Setting Up Your Azure Deployment Architecture and Setting Up Your Deployment Architecture (Amazon Web Services) in the Administrator's Guide for ArcSight Platform.

Policy for Retaining Events

In Intelligence, you can now configure a policy to retain raw events in the database for a specified period of time, after which the raw events are removed automatically.

For more information, see the Configuring the Policy for Retaining Events section in the Administrator's Guide for ArcSight Platform.

Integration with Fusion

The Fusion UI now includes a link called **INSIGHTS** which allows direct navigation to the Intelligence UI. However, if you have deployed Intelligence with Recon, then click **INSIGHTS** > **ENTITIES AT RISK** in the Fusion UI to navigate to the Intelligence UI. The Intelligence UI also includes the **Dashboard** link that takes the users back to the Fusion UI.

Resolved Issues

- Search Query Is Not Populated in Recon for the Anomaly Types 605 and 606 of the Web Proxy Data
- Intelligence Installation Using the ArcSight Installer Results in Intelligence Analytics Failure when the Analytics Pods Start for the First Time
- Multi-node Installation of Intelligence Using the ArcSight Installer Results in Analytics Failure Because of HDFS NameNode Mismatch
- Analytics Might Fail or ENTITYID Might Truncate if the Ingested Data has Entities More Than 255 Characters
- Links in the PDF Reports Do Not Work
- Cannot Log in to the Intelligence UI if the URL Encoding Option for a Multi-node Setup is Set as Hash
- Recon Search Fails to Work for Fileshare and Resource Data if the Values of the File Name and File Path Fields Contain "\"
- When Custom SSL Chain Certificates are Used, Logstash Pods Fail to Communicate with Transformation Hub
- Filtering Does Not Work for the '-' Character in Intelligence UI
- Bad Message 413 reason: Request Entity Too Large

Search Query Is Not Populated in Recon for the Anomaly Types 605 and 606 of the Web Proxy Data

Issue: In the Intelligence UI **>Explore** page, when you filter anomalies with 605 or 606, anomalies are displayed in the **Anomalies & Violation** panel based on the filter you provided. When you click an anomaly, a dialog box provides context about the anomaly or violation. When you click **View Events>Search in Recon**, the search query is not populated in Recon. [FT-21921]

Fix: The search query is now populated in Recon.

Intelligence Installation Using the ArcSight Installer Results in Intelligence Analytics Failure when the Analytics Pods Start for the First Time

Issue: If you install Intelligence by using the ArcSight Platform Installer and execute the install script, the analytics pods move to the Running state even before the postinstall is complete. Intelligence analytics runs for the first time but it fails and you get an error that "Investigation Schema does not exist". This happens because when the analytics pods come up after the install command is complete, they look for the analytics schema. However, the analytics schema is created in the database only after you execute the postinstall script and the postinstall is complete. [FT-21922]

Fix: Now, when you use the ArcSight Platform Installer, Intelligence analytics runs fine when the analytics pods start for the first time.

Multi-node Installation of Intelligence Using the ArcSight Installer Results in Analytics Failure Because of HDFS NameNode Mismatch

Issue: When you perform a multi-node installation of Intelligence by using the ArcSight Platform Installer, the postinstall script creates the HDFS configuration files on all the database nodes. By default, it adds the corresponding database hostname or IP address instead of the HDFS NameNode hostname or IP address in the **core-site.xml** HDFS configuration file. This leads to Intelligence analytics failure when it is run for the first time and you get an error that could not connect to the HDFS URL: hdfs://<HDFS NameNode hostname or IP address>:30820. [FT-21924]

Fix: Intelligence analytics now runs fine as the ArcSight Platform Installer creates the **core-site.xml** file with the correct hostname of the node.

Analytics Might Fail or ENTITYID Might Truncate if the Ingested Data has Entities More Than 255 Characters

Issue: If the ingested data has entities more than 255 characters, either of the following may occur:

- The value of ENTITYID in the analytics schema might truncate after running analytics.
- Analytics might fail with the following error messages:
  - Unable to create/insert into target table: UEBA.ENTITIES_RISK with SaveMode: Append. ERROR MESSAGE: ERROR: java.sql.SQLIntegrityConstraintViolationException: [Vertica][VJDBC](6745) ERROR: Duplicate key values: 'TID=0,ENTITYTYPE=usr,ENTITYID=<entity_id_ displayed>,SCORETYPE=vul,"TIMESTAMP"=2020-09-21 22:00:00,TIME_ BUCKET=hourly' – violates constraint 'UEBA.ENTITIES_RISK.C_PRIMARY'
  - Exception in thread "main" java.sql.SQLIntegrityConstraintViolationException: [Vertica]VJDBC ERROR: Duplicate MERGE key detected in join [(UEBA.PREVIOUS_ ENTITIES_RISK x UEBA.PREVIOUS_ENTITIES_RISK_STAGING) using PREVIOUS_ ENTITIES_RISK_super and subquery (PATH ID: 1)]; value

[FT-21970]

Fix: Now, if the ingested data has entities more than 255 characters, analytics runs fine and ENTITYID in the analytics schema does not truncate.

Links in the PDF Reports Do Not Work

Issue: When you click a link in a PDF report downloaded from the Intelligence UI, you must be directed to the Intelligence UI. Instead, you are directed to the **https://interset-api-svc:9090/interset/0/** URL and you get a site cannot be reached error.

For example, for a PDF Report which gives the top risky users' information, when you click a top risky user listed in the PDF, instead of being directed to the Intelligence UI, you are directed to **https://interset-api-svc:9090/interset/0/**. [FT-21873]

Fix: Now, when you click a link in a PDF report downloaded from the Intelligence UI, you are directed to the Intelligence UI.

Cannot Log in to the Intelligence UI if the URL Encoding Option for a Multi-node Setup is Set as Hash

Issue: In the Kubernetes Dashboard of CDF Management Portal, for the **investigator-default-yaml** file in the arcsight-installer namespace, when you specify the URL encoding option as Hash, you cannot log in to the Intelligence UI when you try to log in again. [FT-21920]

Fix: Now, if you specify the URL Encoding Option as Hash, you can log in to the Intelligence UI when you try to log in again.

Recon Search Fails to Work for Fileshare and Resource Data if the Values of the File Name and File Path Fields Contain “\”

Issue: In the Intelligence UI **> Explore** page **> Anomalies & Violation** panel based on the filter you provided. When you click an anomaly related to FileShare or Resource, a dialog box provides context about the anomaly or violation. When you click **View Events > Search in Recon**, the Recon search fails to work for Fileshare and Resource data if the values of the **File Name** and **File Path** fields contain “\”. [FT-21997]

Fix: Now, the Recon search works fine for Fileshare and Resource data if the values of the **File Name** and **File Path** fields contain “\”.

When Custom SSL Chain Certificates are Used, Logstash Pods Fail to Communicate with Transformation Hub

Issue: When configured with custom SSL chain certificates, the Logstash pods fail to communicate with Transformation Hub and enter into the CrashLoopBackOff stage. This issue occurs when you use the intermediate CA certificates. The issue does not occur when you use the default CA certificates. [FT-21957]

Fix: Now, the Logstash pods do not fail to communicate with Transformation Hub when configured with custom SSL chain certificates.

Filtering Does Not Work for the '-' Character in Intelligence UI

Issue: When filtering using the search filter in the Intelligence UI, if the search includes the '-' character, no results are displayed. [FT-20825]

Fix: Now, filtering works fine if the search includes the '-' character.

Bad Message 413 reason: Request Entity Too Large

Issue: While logging to the Intelligence UI, a bad message **413** is encountered. [FT-20164]

Fix: Clear the cookies for the site and log in again.

Known Issues

Micro Focus strives to ensure that our products provide quality solutions for your enterprise software needs. The following issues are currently being researched. If you need further assistance with any issue, contact Micro Focus Customer Support.

- Elasticsearch and Logstash Pods Fail in an AWS Deployment Because of Permission Issues
- Elasticsearch and Logstash Pods Fail in an Azure Deployment Because of Permission Issues
- Logstash Pod Fails on Data Ingestion in Azure and AWS Deployments When Using Self-Signed Certificates
- Either Intelligence Search API or Log In to Intelligence UI Fails with a Timeout Error for Large Data Sets in the Database
- Intelligence Search API Fails with a Timeout Error for Large Data Sets in the Database
- Init Container for HDFS NameNode and HDFS DataNodes Fails in Multiple NIC Environments
- Intelligence Analytics Fails After Upgrading to Intelligence 6.3.0 Because of Database Schema Migration Failure
- Special Characters for Database Credentials
- Filtering Using the Hand Icon in the Matrix of Anomalies & Violations Does Not Update the Top Risky Users Accordingly
- CSV Reports Do Not Have Timestamps in the Date and Time Format
- Changing a BOT User to a NOTBOT User Has No Effect on Inactive Projects
- Daylight Savings Time
- Repartition Percentage Threshold
- Changing the HDFS NameNode Does Not Terminate the Previous Instance of the HDFS NameNode Container
- Certificate Warnings in Logstash Logs

Elasticsearch and Logstash Pods Fail in an AWS Deployment Because of Permission Issues

Issue: When configuring the EFS for deploying Intelligence in AWS, even after setting the permissions in the arcsight-volume folder to 1999:1999, the Elasticsearch and Logstash pods enter into a CrashLoopBackOff state from a Running state.

Workaround: If the pods enter into the CrashLoopBackOff state, perform the following steps:

1. Log in to the bastion host.
2. Navigate to the following directory and set the permissions to 1999:1999 again:

   ```
   cd /mnt/efs/<parent_folder_name>
   chown -R 1999:1999 arcsight-volume
   ```

3. Wait for the Elasticsearch and Logstash pods to come up.
4. If the pods enter into a Running state and then into a CrashLoopBackOff state, keep repeating steps 2 and 3 till the pods are stable, that is, they do not move from the Running state to the CrashLoopBackOff state.

Elasticsearch and Logstash Pods Fail in an Azure Deployment Because of Permission Issues

Issue: When preparing the NFS server for deploying Intelligence in Azure, even after setting the permissions in the arcsight-volume folder to 1999:1999, the Elasticsearch and Logstash pods enter into a CrashLoopBackOff state from a Running state.

Workaround: If the pods enter into the CrashLoopBackOff state, perform the following steps:

1. (Conditional) If the NFS server is not the Azure NetApp Files server, do the following:
   a. From your jump host, SSH to the NFS VM using its private IP address.
   b. Log in to the NFS VM.
   c. Become root.
   d. Navigate to the following directory and set the permissions to 1999:1999 again:

   ```
   cd /nfs
   chown -R 1999:1999 arcsight-volume
   ```

2. (Conditional) If the NFS server is the Azure NetApp Files server, do the following:
   a. From your jump host, become root.
   b. Execute the following command to retrieve the directory on which the Azure NetApp Files server is mounted:

   ```
   df -h
   ```

   The directory corresponding to <IP address of the NetApp Files server>/volume is the directory on which the Azure NetApp Files server is mounted.
   c. Navigate to the directory retrieved in the previous step and set the permissions to 1999:1999 again:

   ```
   cd /<Azure NetApp Files server directory>
   chown -R 1999:1999 arcsight-volume
   ```

3. Wait for the Elasticsearch and Logstash pods to come up.
4. If the pods enter into a Running state and then into a CrashLoopBackOff state, keep repeating steps 4 and 5 till the pods are stable, that is, they do not move from the Running state to the CrashLoopBackOff state.

Logstash Pod Fails on Data Ingestion in Azure and AWS Deployments When Using Self-Signed Certificates

Issue: In an Azure or AWS deployment of Intelligence, when data is ingested, the Logstash pod enters into a CrashLoopBackOff state from a Running state. This issue occurs if you have configured CDF in the cloud (Azure or AWS) environments with self-signed certificates.

Workaround: Perform the following steps:

1. Do one of the following to connect to the cluster:
   - For Azure, connect to the jump host.
   - For AWS, connect to the bastion.
2. Execute the following command to scale down the Logstash nodes:

   ```
   kubectl -n $(kubectl get namespaces | grep arcsight | cut -d ' ' -f1)
   scale statefulset interset-logstash --replicas=0
   ```

3. Execute the following command to modify the logstash-config-pipeline configmap:

   ```
   kubectl -n $(kubectl get namespaces | grep arcsight | cut -d ' ' -f1)
   edit configmaps logstash-config-pipeline
   ```

4. Update the value of the **verify_mode** field from "verify_peer" to "verify_none".
5. Save the configmap.
6. Execute the following command to scale up the Logstash nodes:

   ```
   kubectl -n $(kubectl get namespaces | grep arcsight | cut -d ' ' -f1)
   scale statefulset interset-logstash --replicas=<number_of_replicas>
   ```

Init Container for HDFS NameNode and HDFS DataNodes Fails in Multiple NIC Environments

Issue: In an environment having multiple network interface controllers (NICs), the init container of the HDFS NameNode and HDFS DataNodes fails. Though this issue will be fixed in the upcoming 6.3.0.1 patch, follow the workaround for now.

Workaround: Perform the following steps:

1. Launch a terminal session and log in to any of the Kubernetes nodes as a root user.
2. Create the **/tmp/ft23676.json** directory with the following content:
   {"spec":{"template": {"spec":{"initContainers": [{"name": "install","env": [{"name": "POD_NAME","valueFrom":{ "fieldRef":{ "apiVersion": "v1", "fieldPath": "metadata.name" }}}]}}}}}

   Important: Ensure that you use the exact closing characters in the content.
3. Execute the following commands to retrieve the namespace and apply patch to the HDFS NameNode and HDFS DataNode pods:

   ```
   NAMESPACE=$(kubectl get namespaces | grep arcsight-installer | awk '{ print $1}')
   ```

   ```
   kubectl patch sts hdfs-namenode -n $NAMESPACE --patch "$(cat /tmp/ft23676.json)"
   ```

   ```
   kubectl patch ds hdfs-datanode -n $NAMESPACE --patch "$(cat /tmp/ft23676.json)"
   ```
4. Execute the following commands to scale down and scale up the HDFS NameNode pods:

   ```
   kubectl -n $NAMESPACE scale statefulset hdfs-namenode-0 --replicas=0
   kubectl -n $NAMESPACE scale statefulset hdfs-namenode-0 --replicas=1
   ```
   The pods then start.
5. Execute the following command to verify the status of the pods:

   ```
   kubectl get pods -n $NAMESPACE
   ```

[FT-23676]

Intelligence Analytics Fails After Upgrading to Intelligence 6.3.0 Because of Database Schema Migration Failure

Issue: In a high availability cluster, after you upgrade from a previous version (6.1.0 or 6.2.0) of Intelligence to the 6.3.0 version, Intelligence analytics might fail. On verifying the analytics pod status, you might notice two analytics pods in a CrashLoopBackOff state instead of just one analytics pod running. This happens because while upgrading the database, the migration script - **V10__fix_entity_id_column_length.sql** fails as a result of

the primary key constraint missing in the **ENTITIES_RISK** table. This in turn leads to the database schema migration failure, and hence Intelligence analytics fails.

Workaround: Perform the following steps after you upgrade the database:

1. Launch a terminal session and log in to the master node as a root user.
2. Execute the following commands to scale down the analytics pod:

```
export NS=$(kubectl get namespaces | grep arcsight | cut -d ' ' -f1)
kubectl scale deployment interset-analytics -n $NS --replicas=0
```

3. Execute the following command to verify that the analytics pod is not running:

```
kubectl get pods -n $NS | grep analytics
```

4. Log in to a database node as a root user.
5. Log in to fusiondb as the appadmin user:

```
/opt/vertica/bin/vsql -h `hostname -f` fusiondb appadmin
[password prompt]
```

6. Execute the following Data Definition Language (DDL) command to create the **default_secops_intelligence.ENTITIES_RISK_SUPPORT_TABLE** table if it does not already exist:

Important: If you are doing a rolling upgrade from 6.1.0 to 6.2.0 to 6.3.0, then look up the **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID** column length, and in the following command, replace 1024 in ENTITYID VARCHAR(1024) with the column length of **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID**.

```
CREATE TABLE IF NOT EXISTS default_secops_intelligence.ENTITIES_RISK_
SUPPORT_TABLE
(
  TID           VARCHAR(3)   NOT NULL,
  ENTITYTYPE    VARCHAR(3)   NOT NULL,
  ENTITYID      VARCHAR(1024) NOT NULL,
  SCORETYPE     VARCHAR(3)   NOT NULL,
  TIMESTAMP     TIMESTAMP    NOT NULL,
  TIME_BUCKET   VARCHAR(100),
  LONG_SUM      DOUBLE PRECISION,
  SHORT_SUM     DOUBLE PRECISION,
  PRIMARY KEY (TID, ENTITYTYPE, ENTITYID, SCORETYPE, TIMESTAMP, TIME_
BUCKET) ENABLED
) ORDER BY TID, ENTITYTYPE, ENTITYID, SCORETYPE, TIMESTAMP, TIME_
BUCKET
  PARTITION BY DATE_TRUNC('DAY',TO_TIMESTAMP_TZ(EXTRACT(EPOCH FROM
```

```
TIMESTAMP)) AT TIME ZONE 'UTC') ACTIVEPARTITIONCOUNT
2;
```

7. Execute the following Data Manipulation Language (DML) command to insert data from the **ENTITIES_RISK** table to the **ENTITIES_RISK_SUPPORT_TABLE** table:

```
INSERT INTO default_secops_intelligence.ENTITIES_RISK_SUPPORT_TABLE
select * from default_secops_intelligence.ENTITIES_RISK;
COMMIT;
```

8. Execute the following queries to verify that the row count of both the tables match:

```
Select count(*) from default_secops_intelligence.ENTITIES_RISK;
Select count(*) from default_secops_intelligence.ENTITIES_RISK_
SUPPORT_TABLE;
```

9. Execute the following DDL command to alter the **ENTITIES_RISK** table to rename it to **ENTITIES_RISK_ORIGINAL**:

```
ALTER TABLE default_secops_intelligence.ENTITIES_RISK RENAME TO
ENTITIES_RISK_ORIGINAL;
```

10. Execute the following DDL command to alter the **ENTITIES_RISK_SUPPORT_TABLE** table to rename it to **ENTITIES_RISK**:

```
ALTER TABLE default_secops_intelligence.ENTITIES_RISK_SUPPORT_TABLE
RENAME TO ENTITIES_RISK;
```

11. Execute the following DDL command to create the **default_secops_intelligence.PREVIOUS_ENTITIES_RISK_SUPPORT_TABLE** table:

> **Important**: If you are doing a rolling upgrade from 6.1.0 to 6.2.0 to 6.3.0, then look up the **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID** column length, and in the following command, replace 1024 in ENTITYID VARCHAR(1024) with the column length of **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID**.

```
CREATE TABLE IF NOT EXISTS default_secops_intelligence.PREVIOUS_
ENTITIES_RISK_SUPPORT_TABLE
(
  TID           VARCHAR(3)   NOT NULL,
  ENTITYTYPE    VARCHAR(3)   NOT NULL,
  ENTITYID      VARCHAR(1024) NOT NULL,
  SCORETYPE     VARCHAR(3)   NOT NULL,
  TIMESTAMP     TIMESTAMP    NOT NULL,
  TIME_BUCKET   VARCHAR(100),
  LONG_SUM      DOUBLE PRECISION,
  SHORT_SUM     DOUBLE PRECISION,
```

```
  PRIMARY KEY (TID, ENTITYTYPE, ENTITYID, SCORETYPE, TIMESTAMP, TIME_
BUCKET) ENABLED
) ORDER BY TID, ENTITYTYPE, ENTITYID, SCORETYPE, TIMESTAMP, TIME_
BUCKET
PARTITION BY DATE_TRUNC('DAY',TO_TIMESTAMP_TZ(EXTRACT(EPOCH FROM
TIMESTAMP)) AT TIME ZONE 'UTC') ACTIVEPARTITIONCOUNT
2;
```

12. Execute the following DML command to insert data from the **PREVIOUS_ENTITIES_RISK** table to the **PREVIOUS_ENTITIES_RISK_SUPPORT_TABLE** table:

```
INSERT INTO default_secops_intelligence.PREVIOUS_ENTITIES_RISK_
SUPPORT_TABLE select * from default_secops_intelligence.PREVIOUS_
ENTITIES_RISK;
COMMIT;
```

13. Execute the following queries to verify that the row count of both the tables match:

```
Select count(*) from default_secops_intelligence.PREVIOUS_ENTITIES_
RISK;
Select count(*) from default_secops_intelligence.PREVIOUS_ENTITIES_
RISK_SUPPORT_TABLE;
```

14. Execute the following DDL command to alter the **PREVIOUS_ENTITIES_RISK** table to rename it to **PREVIOUS_ENTITIES_RISK_ORIGINAL**:

```
ALTER TABLE default_secops_intelligence.PREVIOUS_ENTITIES_RISK RENAME
TO PREVIOUS_ENTITIES_RISK_ORIGINAL;
```

15. Execute the following DDL command to alter the **PREVIOUS_ENTITIES_RISK_SUPPORT_TABLE** table to rename it to **PREVIOUS_ENTITIES_RISK**:

```
ALTER TABLE default_secops_intelligence.PREVIOUS_ENTITIES_RISK_
SUPPORT_TABLE RENAME TO PREVIOUS_ENTITIES_RISK;
```

16. Execute the following DDL command to alter the length of the ENTITYID column of the **PREVIOUS_ENTITIES_RISK_STAGING** table:

Important: If you are doing a rolling upgrade from 6.1.0 to 6.2.0 to 6.3.0, then look up the **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID** column length, and in the following command, replace 1024 in the command with the column length of **OBSERVED_ENTITY_RELATION_MINUTELY_COUNTS.ENTITYID**.

```
ALTER TABLE default_secops_intelligence.PREVIOUS_ENTITIES_RISK_STAGING
alter column ENTITYID SET DATA TYPE VARCHAR(1024);
```

17. Execute the following DML command to update the success column of the **SCHEMA_VERSION** table:

```
update default_secops_intelligence.schema_version set success =true
where version='10' and script = '6.3.0/V10__fix_entity_id_column_
length.sql' and success = false;
COMMIT;
```

18. Execute the following query to verify that the success column in the **SCHEMA_ VERSION** table has been set to true:

```
select * from default_secops_intelligence.schema_version where
version='10' and script = '6.3.0/V10__fix_entity_id_column_
length.sql';
```

19. Log in to the master node as a root user and execute the following commands to scale up the analytics pod:

```
export NS=$(kubectl get namespaces | grep arcsight | cut -d ' ' -f1)
kubectl scale deployment interset-analytics -n $NS --replicas=1
```

20. Execute the following command to verify the analytics pod is running:

```
kubectl get pods -n $NS | grep analytics
```

21. Repeat steps 4 and 5, then execute the following DDL commands to drop the backed up tables:

```
DROP TABLE default_secops_intelligence.ENTITIES_RISK_ORIGINAL CASCADE;
DROP TABLE default_secops_intelligence.PREVIOUS_ENTITIES_RISK_ORIGINAL
CASCADE;
```

22. Proceed to run analytics to sync Elasticsearch with the previously ingested and analyzed data. For more information on running analytics, see the the *Running Analytics on Demand* section in the Administrator's Guide for ArcSight Platform.

Special Characters for the Database Credentials

Issue: Intelligence Analytics does not support the following characters when you specify the database credentials:

- Whitespace
- Single quotes

Workaround: There is no workaround at this time.

[FT-21600]

Filtering Using the Hand Icon in the Matrix of Anomalies & Violations Does Not Update the Top Risky Users Accordingly

Issue: In the Intelligence UI **> Explore** page **> Matrix of Anomalies & Violations**, when you click the hand icon, then click and drag the cursor across the matrix to see the risky users and anomalies or violations for a specific time interval, the users in the **Top Risky Users** panel are not updated accordingly.

For example, if there are no risky users and the associated anomalies or violations in the selected time interval (for example, between 4:30 pm and 4:35 pm on a day), the **Top Risky Users** panel still displays users.

Workaround: There is no workaround at this time.

[FT-20451]

CSV Reports Do Not Have Timestamps in the Date and Time Format

Issue: The CSV Reports have timestamps in the Unix Epoch format instead of the date and time format.

Workaround: There is no workaround at this time.

[FT-20978]

Changing a BOT User to a NOTBOT User Has No Effect on Inactive Projects

Issue: When anomalies are identified because few users access a specific project, and one or more of the users are flagged as bots, changing the BOT users to NOTBOT users — and therefore increasing the number of non-bot users accessing the project — will not impact the project's identification as 'inactive'. Anomalies will therefore continue to be identified when the project is accessed, even though more non-bot users are now regularly accessing the project.

Workaround: There is no workaround at this time.

[FT-8934]

Daylight Savings Time

Issue: During the weeks immediately following Daylight Savings Time (DST) clock changes, you may observe an increase in reported Normal Working Hours anomalies. These anomalies, which are due to automatic software clock changes, will usually have risk scores of zero (0), and are reflective of the perceived Normal Working Hours pattern shift.

Workaround: There is no workaround needed.

[FT-8601]

Repartition Percentage Threshold

Issue: In the **CDF Management Portal > Configure/Deploy** page **> Intelligence**, when you specify a value for the **Repartition Percentage Threshold** field, the installer does not validate the value. However, Intelligence Analytics fails if the value is not set between 0.7 and 1.0 as stated in the tooltip.

Workaround: Ensure that you set a value between 0.7 and 1.0.

[FT-20011]

Changing the HDFS NameNode Does Not Terminate the Previous Instance of the HDFS NameNode Container

Issue: In the **CDF Management Portal > Configure/Deploy** page **> Intelligence**, when you change the value of the **HDFS NameNode** field to deploy the HDFS NameNode container on another worker node, the older instance of the HDFS NameNode container goes into a pending state instead of being terminated.

Workaround: Perform the following steps after changing the value in the field:

1. In the CDF Management Portal, click **Cluster > Nodes**.
2. Click the [-] icon for the **intelligence-namenode:yes** label present on the worker node.

3. From **Predefined Labels**, drag and drop the **intelligence-namenode:yes** label to the worker node to which you want to add it. Ensure the worker node matches the new value you specified in the **HDFS NameNode** field.
4. Configure the database with HDFS. For more information, see the "Configuring the Database with HDFS for Intelligence" section in Administrator's Guide for ArcSight Platform.
5. Restart the HDFS DataNodes. Do the following:
   a. Launch a terminal session and log in to a worker node where an HDFS DataNode is deployed.
   b. Execute the following commands:

```
NAMESPACE=$(kubectl get namespaces | grep arcsight-installer | awk '{ print $1}')
```

```
kubectl get pods -n $NAMESPACE | grep -e 'hdfs\|interset-analytics' | awk '{print $1}' | xargs kubectl delete pod -n $NAMESPACE --force --grace-period=0
```

[FT-20019]

Certificate Warnings in Logstash Logs

Issue: When you view the Logstash logs, you might come across the following warnings:

- ** WARNING ** Detected UNSAFE options in elasticsearch output configuration!
- ** WARNING ** You have enabled encryption but disabled certificate verification.
- ** WARNING **To make sure your data is secure change :ssl_certificate_vertification to true

Workaround: There is no workaround needed. You can ignore these warnings as there is no impact in the functionality.

[FT-20038]

Error Message Related to Keystore Format in Intelligence Analytics Logs

Issue: When you view the Intelligence Analytics logs, you might come across multiple instances of the following error message:

```
Invalid keystore format
```

Workaround: There is no workaround needed. You can ignore such instances of the error message because it does not affect the functioning of Intelligence Analytics.

[FT-22885]

Installer Does Not Validate the Value You Specify for Elasticsearch Data Retention Period

Issue: In the **CDF Management Portal** > **Configure/Deploy** page > **Intelligence** > **Elasticsearch Configuration** section, the installer does not validate the value you specify for the **Elasticsearch Data Retention Period** field. The tool-tip for the **Elasticsearch Data Retention Period** field suggests that you should specify a value greater than 30 for indices retention. However, there is no validation preventing you from entering a value that is less than 30. If you specify a value that is less than 30, the value for **Elasticsearch Data Retention Period** will be set to the minimum default value of 30 days.

Workaround: There is no workaround at this time.

[FT-23129]

Intelligence Analytics Fails After Upgrading to Intelligence 6.3

Issue: Intelligence Analytics fails after you upgrade to Intelligence 6.3.0.

Workaround: After upgrading to Intelligence 6.3.0, irrespective of whether you have enabled Kerberos Authentication (to secure HDFS) or not, you must update the `hdfs-site.xml` and `core-site.xml` files on the database nodes.

[FT-23151]

Uninstalling Intelligence Does Not Delete All Files

Issue: When you uninstall Intelligence, some files are not deleted from the `/opt/arcsight/k8s-hostpath-volume/interset` directory of all the worker nodes. Therefore, when you install Intelligence again, the intelligence pods stay in Init state.

Workaround: Before installing Intelligence again, manually delete the remaining files from the `/opt/arcsight/k8s-hostpath-volume/interset` directory of all the worker nodes. If you have modified the value of the **Elasticsearch Node Data Path** field in the **Intelligence** tab of the CDF Management Portal, check and manually delete the remaining

files from the directory you have specified for the **Elasticsearch Node Data Path** field for all the worker nodes.

[FT-23241]

Unable to Retrieve Elasticsearch Indices

Issue: When your Elasticsearch Cluster is not stable and you run the reindex jobs, the jobs run successfully but display the following error message in the job details:

```
Error occurred while getting all ES indices: Request cannot be executed;
I/O reactor status: STOPPED
```

Workaround: You must restart the Elasticsearch cluster to refresh the Elasticsearch environment.

[FT-23359]

Anomaly Tuning Does Not Update the Value of the IMPORTANCE field When You Modify Both Weight and IMPORTANCE Fields

Issue: In the Intelligence UI > **Explore** page > **Anomalies & Violation** panel > click an anomaly or violation > **Tuning** > **ADVANCED TUNING** > select an anomaly and click the  icon, when you modify the values of both the **Weight** and **IMPORTANCE** fields, Anomaly Tuning updates only the value of the **Weight** field and does not update the value of the **IMPORTANCE** field. However, when you modify the value of only the **IMPORTANCE** field, Anomaly Tuning updates the value of the **IMPORTANCE** field.

Workaround: Do not modify the value of the **IMPORTANCE** field together with the value of the **Weight** field. You must modify the value of the **IMPORTANCE** field separately.

[FT-23438]

The Entities at Risk Option Does Not Appear After Upgrading Intelligence From 6.1 to 6.2 to 6.3

Issue: After upgrading Intelligence from 6.1.0 to 6.2.0 to 6.3.0, on the Fusion Dashboard page > **INSIGHTS**, the **Entities at Risk** option does not appear.

Workaround: You should scale down and then scale up the `intelligence-arcsightconnector-api` pod.

[FT-23378]

Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Special Character

Issue: In the **CDF Management Portal** > **Configure/Deploy** page > **Intelligence** > **KeyStores** section > **KeyStore Password** field, if you specify a password that starts with a special character, most pods enter into the CrashLoopBackOff state.

Workaround: For the **KeyStore Password** field, do not specify a password that starts with a special character.

[FT-23474]

HTTP Status 400 - Bad Request

Issue: If the cookie request size exceeds the cookie size limit, your screen displays a **HTTP Status 400 - Bad Request** message when you try to open the CDF Management Portal.

Workaround: Perform the following steps:

1. Open a certified web browser.
2. Login to the Management portal as the administrator.
   `https://<virtual_FQDN>:5443`
3. Click **CLUSTER** > **Dashboard**. You will be redirected to the **Kubernetes Dashboard**.
4. Under **Namespace**, search and select the `arcsight-installer-xxxx` namespace.
5. Under **Config and Storage**, click **Config Maps**.
6. Click the filter icon, and search for `investigator-default-yaml`.
7. Click  and select **Edit**.
8. In the **YAML** tab, under the `interset-cookie` section, add the following:

   ```
   path: /interset;SameSite=Lax
   ```

9. Click **Update**.
10. To apply the changes, restart the `interset-api` pods by either deleting the `interset-api` pods or scaling down the `interset-api` deployments using the

following commands:

```
kubectl delete pods -n <arcsight-installer-namespace> <interset-api-pod-1> <interset-api-pod-2>
```

OR

```
kubectl scale deployment -n <arcsight-installer-namespace> interset-api --replicas=0
kubectl scale deployment -n <arcsight-installer-namespace> interset-api --replicas=2
```

11. Log in to Intelligence or other application user interfaces available for this domain such as the CDF Management Portal or the Fusion dashboard.
12. Using the **Developer tools** option in your browser, ensure that the **INTERSET_SESSION** cookie is only available to request with **/interset** in the path.

    

    To verify the information about cookies passed in each request, in the **Developer tools** option of your browser, click **Network** > **Cookies**.

[FT-23189]

Intelligence Search API Fails with a Timeout Error for Large Data Sets in the Database

Issue: Intelligence Search API fails with the esSocketTimeout exception while querying a large data set (approximately 4 billion records) in the database, along with ingestion and analytics running simultaneously.

Workaround: Perform the following steps:

1. Open a certified web browser.
2. Log in to the CDF Management portal as the administrator.

   `https://<virtual_FQDN>:5443`
3. Click **CLUSTER** > **Dashboard**. You are redirected to the **Kubernetes Dashboard**.
4. In **Namespace**, search and select the `arcsight-installer-xxxx` namespace.
5. In **Config and Storage**, click **Config Maps**.
6. Click the filter icon, then search for `investigator-default-yaml`.
7. In the **db-elasticsearch** section of the YAML tab, modify the **esSocketTimeout** value based on the data size.

For example, if the Intelligence search API takes 150 seconds to retrieve data from the database, then ensure that you set the **esSocketTimeout** value to more than 150 seconds to avoid the exception.

Ensure that you set the **esSocketTimeout** value in milliseconds.

8. Click **Update**.
9. Restart the interset-api pods:
   a. Launch a terminal session and log in to the master or worker node.
   b. Execute the following command to retrieve the namespace:

```
export NS=$(kubectl get namespaces | grep arcsight|cut -d ' ' -f1)
```

   c. Execute the following commands to restart the interset-api pods:

```
kubectl -n $NS scale deployment interset-api --replicas=0
```

```
kubectl -n $NS scale deployment interset-api --replicas=2
```

[342464]

Either Intelligence Search API or Log In to Intelligence UI Fails with a Timeout Error for Large Data Sets in the Database

Issue: Either Intelligence Search API or Log in to Intelligence UI fails with the IOException: Listener Timeout after waiting for 30 seconds while querying a large data set (approximately 2 billion records) in the database.

Workaround: Perform the following steps:

1. Open a certified web browser.
2. Log in to the CDF Management portal as the administrator.
   `https://<virtual_FQDN>:5443`
3. Click **CLUSTER** > **Dashboard**. You are redirected to the **Kubernetes Dashboard**.
4. In **Namespace**, search and select the `arcsight-installer-xxxx` namespace.
5. In **Config and Storage**, click **Config Maps**.
6. Click the filter icon, then search for `investigator-default-yaml`.

7. In the **db-elasticsearch** section of the YAML tab, modify the **esListenerTimeout** value based on the data size.

   For example, if the Intelligence search API takes 150 seconds to retrieve data from the database, then ensure that you set the **esListenerTimeout** value to more than 150 seconds to avoid the exception.

   

   Ensure that you set the **esListenerTimeout** value in milliseconds.

8. Click **Update**.
9. Restart the interset-api pods:
   a. Launch a terminal session and log in to the master or worker node.
   b. Execute the following command to retrieve the namespace:

   ```
   export NS=$(kubectl get namespaces | grep arcsight|cut -d ' ' -f1)
   ```

   c. Execute the following commands to restart the interset-api pods:

   ```
   kubectl -n $NS scale deployment interset-api --replicas=0
   ```

   ```
   kubectl -n $NS scale deployment interset-api --replicas=2
   ```

[314858]

Technical Requirements

For information about the software and hardware requirements for your deployment and performance tuning guidelines, see the Technical Requirements for ArcSight Platform.

Downloading Intelligence

Before you begin installing Intelligence, you must download necessary product installation packages. The installation package also includes the respective signature file for validating that the downloaded software is authentic and not tampered by a third party.

To review the list of the files and versions to download for this release, see the Release Notes for ArcSight Platform.

Installing Intelligence

Micro Focus provides several options for deploying your Intelligence environment. For more information, see the Administrator's Guide for ArcSight Platform.

Upgrading Intelligence

You can upgrade only from 6.2.0 to 6.3.0 version of Intelligence. For more information about upgrading Intelligence, see the Administrator's Guide for ArcSight Platform.

The **Investigation** schema is now **default_secops_adm** and the **UEBA** schema is now **default_secops_intelligence**.

Licensing Information

For information about activating a new license, see the Administrator’s Guide for ArcSight Platform.

Contacting Micro Focus

For specific product issues, contact Micro Focus Support at https://softwaresupport.softwaregrp.com/.

Additional technical information or advice is available from several sources:

- Product documentation, Knowledge Base articles, and videos: https://softwaresupport.softwaregrp.com/
- The Micro Focus Community pages: https://community.microfocus.com/

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this computer, click the link above and an email window opens with the following information in the subject line:

Feedback on Release Notes (Intelligence 6.3.0)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to arcsight_doc@microfocus.com.

We appreciate your feedback!