
Micro Focus Security ArcSight Intelligence

Software Version: 6.4.0

Release Notes

Document Release Date: Feb 2022

Software Release Date: Feb 2022



Legal Notices

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

<https://www.microfocus.com>

Copyright Notice

© Copyright 2022 Micro Focus or one of its affiliates

Confidential computer software. Valid license from Micro Focus required for possession, use or copying. The information contained herein is subject to change without notice.

The only warranties for Micro Focus products and services are 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.

No portion of this product's documentation may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser's internal use, without the express written permission of Micro Focus.

Notwithstanding anything to the contrary in your license agreement for Micro Focus ArcSight software, you may reverse engineer and modify certain open source components of the software in accordance with the license terms for those particular components. See below for the applicable terms.

U.S. Governmental Rights. For purposes of your license to Micro Focus ArcSight software, "commercial computer software" is defined at FAR 2.101. If acquired by or on behalf of a civilian agency, the U.S. Government acquires this commercial computer software and/or commercial computer software documentation and other technical data subject to the terms of the Agreement as specified in 48 C.F.R. 12.212 (Computer Software) and 12.211 (Technical Data) of the Federal Acquisition Regulation ("FAR") and its successors. If acquired by or on behalf of any agency within the Department of Defense ("DOD"), the U.S. Government acquires this commercial computer software and/or commercial computer software documentation subject to the terms of the Agreement as specified in 48 C.F.R. 227.7202-3 of the DOD FAR Supplement ("DFARS") and its successors. This U.S. Government Rights Section 18.11 is in lieu of, and supersedes, any other FAR, DFARS, or other clause or provision that addresses government rights in computer software or technical data.

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.microfocus.com/about/legal/>.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

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://www.microfocus.com/documentation/argsight/

Contents

- Introduction 4
- What's New? 5
 - MSSP Licensing 5
 - Non-MSSP Licensing Enhancements 5
 - Overall Risk Page Updates 5
 - Common ArcSight Masthead 6
 - Resolved Issues 6
 - Intelligence Analytics Fails After Upgrading to Intelligence 6.3 6
 - Error Message Related to Keystore Format in Intelligence Analytics Logs 6
 - Anomaly Tuning Does Not Update the Value of the IMPORTANCE field When You Modify Both Weight and IMPORTANCE Fields 6
- Known Issues 7
 - Either the Intelligence Search API or Login to the Intelligence UI or both Fail with a Timeout Error (IOException: Listener Timeout) for Large Data Sets in the Database 8
 - Intelligence Search API Fails with a Timeout Error (esSocketTimeout exception) for Large Data Sets in the Database 9
 - Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Space Character 10
 - Elasticsearch and Logstash Pods Fail in an AWS Deployment Because of Permission Issues 10
 - Elasticsearch and Logstash Pods Fail in an Azure Deployment Because of Permission Issues 11
 - Logstash Pod Fails on Data Ingestion in Azure and AWS Deployments When Using Self-Signed Certificates 12
 - Erroneous Warning about Recon License 12

Intelligence Analytics Fails After Upgrading to Intelligence 6.4.0 Because of Database Schema Migration Failure	13
Special Characters for the Database Credentials	17
Filtering Using the Hand Icon in the Matrix of Anomalies & Violations Does Not Update the Top Risky Users Accordingly	17
CSV Reports Do Not Have Timestamps in the Date and Time Format	17
Changing a BOT User to a NOTBOT User Has No Effect on Inactive Projects	18
Daylight Savings Time	18
Repertition Percentage Threshold	18
Changing the HDFS NameNode Does Not Terminate the Previous Instance of the HDFS NameNode Container	18
Certificate Warnings in Logstash Logs	19
Installer Does Not Validate the Value You Specify for Elasticsearch Data Retention Period	20
Uninstalling Intelligence Does Not Delete All Files	20
Unable to Retrieve Elasticsearch Indices	20
The Entities at Risk Option Does Not Appear After Upgrading Intelligence From 6.1 to 6.2 to 6.3	21
Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Special Character	21
HTTP Status 400 - Bad Request	21
Technical Requirements	22
Downloading Intelligence	22
Installing Intelligence	23
Upgrading Intelligence	23
Licensing Information	23
Contacting Micro Focus	23
Send Documentation Feedback	24

Introduction

This release notes document for Intelligence 6.4.0 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:

- [MSSP Licensing](#)
- [Non-MSSP Licensing Enhancements](#)
- [Overall Risk Page Updates](#)
- [Common ArcSight Masthead](#)
- [Resolved Issues](#)

MSSP Licensing

In addition to the existing licensing, Intelligence also supports licensing to accommodate Managed Security Service Provider (MSSP) deployments. MSSP licenses take precedence over existing Non-MSSP licenses when licenses of both types are in use. When the validity of the MSSP license expires, the Non-MSSP license comes into effect.

Non-MSSP Licensing Enhancements

Enhancements have been made to differentiate between human users and machine users. The existing Non-MSSP license is now measured against the number of human users that Intelligence runs analytics on. The license information in the Intelligence UI now includes the number of analyzed human users and the total number of analyzed users (human and machine), in addition to other license details.

Overall Risk Page Updates

The **Overall Risk** page no longer exists in the Intelligence UI. Instead, the functionality is part of the **Entity Risk** dashboard, which is an out-of-the-box dashboard in the Fusion UI.

The Fusion UI includes the new widgets that make up the **Entity Risk** dashboard.

Common ArcSight Masthead

The masthead of the Intelligence UI has been updated to align with the ArcSight masthead.

Resolved Issues

- [Intelligence Analytics Fails After Upgrading to Intelligence 6.3](#)
- [Error Message Related to Keystore Format in Intelligence Analytics Logs](#)
- [Anomaly Tuning Does Not Update the Value of the IMPORTANCE field When You Modify Both Weight and IMPORTANCE Fields](#)

Intelligence Analytics Fails After Upgrading to Intelligence 6.3

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

Fix: The issue has been fixed now.

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

Fix: This error message is not being displayed now.

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.

Fix: The issue has been fixed now.

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](#).

- [Either the Intelligence Search API or Login to the Intelligence UI or both Fail with a Timeout Error \(IOException: Listener Timeout\) for Large Data Sets in the Database](#)
- [Intelligence Search API Fails with a Timeout Error \(esSocketTimeout exception\) for Large Data Sets in the Database](#)
- [Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Space Character](#)
- [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](#)
- [Erroneous Warning about Recon License](#)
- [Intelligence Analytics Fails After Upgrading to Intelligence 6.4.0 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](#)

- [Installer Does Not Validate the Value You Specify for Elasticsearch Data Retention Period](#)
- [Uninstalling Intelligence Does Not Delete All Files](#)
- [Unable to Retrieve Elasticsearch Indices](#)
- [The Entities At Risk Option Does Not Appear After Upgrading Intelligence From 6.1 to 6.2 to 6.3](#)
- [Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Special Character](#)
- [HTTP Status 400 - Bad Request](#)

Either the Intelligence Search API or Login to the Intelligence UI or both Fail with a Timeout Error (**IOException: Listener Timeout**) for Large Data Sets in the Database

Issue: Either the Intelligence Search API or login to the Intelligence UI or both fail 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.



Note: 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
```

Intelligence Search API Fails with a Timeout Error (**esSocketTimeout exception**) 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.



Note: 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
```

Most Pods Enter into the CrashLoopBackOff State if the KeyStore Password Starts with a Space 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 space character, most pods enter into the CrashLoopBackOff state.

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

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>
```

Erroneous Warning about Recon License

Issue: In an ArcSight Platform deployment that has Intelligence with an MSSP license, you will receive the usual notifications that the licenses are about to expire. However, if the MSSP license expires, the Platform erroneously displays a warning that the Recon license has expired even though Recon is not deployed. This issue does not occur when Recon is deployed, with or without the MSSP license.

Workaround: There is no workaround for this issue.

Intelligence Analytics Fails After Upgrading to Intelligence 6.4.0 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 6.4.0 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 6.3.0 to 6.4.0, 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;
```

- 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;
```

- 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;
```

- 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;
```

- 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;
```

- 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 6.3.0 to 6.4.0, 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;
```

- 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;
```

- 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;
```

- 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;
```

- 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;
```

- 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 6.3.0 to 6.4.0, 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);
```

- 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;
```

- 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';
```

- 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
```

- Execute the following command to verify the analytics pod is running:

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

- 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 *Running Analytics on Demand* section in the [Administrator's Guide for ArcSight Platform](#).

Special Characters for the Database Credentials

Issue: The following characters are not supported for the database credentials:

- Whitespace
- Single quotes

Workaround: There is no workaround at this time.

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.

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.

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.

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.

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.

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 the [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
```

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_verification to true

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

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.

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.

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.

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.

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.

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**.

Technical Requirements

For information about the software and hardware requirements for your deployment and performance tuning guidelines, see the [Technical Requirements Guide](#).

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 Intelligence only in succession. For example, if you are using the 6.3.0 version of Intelligence and you need to upgrade to 6.4.0, you must upgrade in the following order: 6.3.0 > 6.3.1 > 6.4.0.

This release does not support upgrade of the ArcSight Database from the previous version. Instead, you must remove the previous version of the ArcSight Database and perform a fresh installation of the 22.1 (current) version.



Micro Focus does not recommend that ArcSight Intelligence customers attempt to upgrade without consulting with Micro Focus as upgrading at this time will result in loss of baseline data. We are aware that this is inconvenient for customers, and we are currently working on addressing this limitation in an future release.

For more information about upgrading Intelligence, see the [Administrator's Guide for ArcSight Platform](#).

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 6.4.0)

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!