



# **ArcSight User Behavior Analytics**

Software Version: 6.10

## Installation Guide

4/13/2018

Powered by  **SECURONIX**

# Legal Notices

## Warranty

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.

The information contained herein is subject to change without notice.

The network information used in the examples in this document (including IP addresses and hostnames) is for illustration purposes only.

Micro Focus ArcSight products are highly flexible and function as you configure them. The accessibility, integrity, and confidentiality of your data is your responsibility. Implement a comprehensive security strategy and follow good security practices.

This document is confidential.

## Restricted Rights Legend

Confidential computer software. Valid license from Micro Focus required for possession, use or copying. 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.

## Copyright Notice

© Copyright 2018 Micro Focus or one of its affiliates.

Follow this link to see a complete statement of copyrights and acknowledgments: <https://community.softwaregrp.com/t5/ArcSight-Discussions/Third-Party-Copyright-Notices-and-License-Terms/td-p/1589228>

## Legal Notice for Open Source Code

vLGPLv3, LGPLv2, EPL 1.0, CCDL

This product includes code licensed under the LGPLv3 licensed-code, LGPLv2 licensed-code, Eclipse Public License 1.0, CCDL-licensed code, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Micro Focus.

To obtain such source code on CD, send a check or money order in the amount of US \$10.00 (for shipping and handling) to:

Micro Focus

Attn: Gordon Lee

1140 Enterprise Way

Sunnyvale, CA 94089

USA

Please specify the product and version for which you are requesting the source code.

## Support

### Contact Information

Phone	A list of phone numbers is available on the Micro Focus ArcSight Technical Support Page: <a href="https://softwaresupport.softwaregrp.com/documents/10180/14684/esp-support-">https://softwaresupport.softwaregrp.com/documents/10180/14684/esp-support-</a>
-------	--

**Contact Information, continued**

	<a href="#">contact-list</a>
Support Web Site	<a href="https://softwaresupport.softwaregrp.com/">https://softwaresupport.softwaregrp.com/</a>
Protect 724 Community	<a href="https://community.softwaregrp.com/t5/Protect724/ct-p/Protect724">https://community.softwaregrp.com/t5/Protect724/ct-p/Protect724</a>

# Contents

---

<b>Introduction</b>	<b>6</b>
About ArcSight User Behavior Analytics	6
Who Should Read this Document	6
ArcSight UBA Architecture	7
ArcSight UBA Reference Architecture	8
Installation Components	9
Installation Flow	10
Installation Checklist	10
Prerequisites	11
Supported Operating Systems	11
Supported Hadoop Deployment	11
Certification Matrix	12
Required Communication Ports	13
Supported Browsers	13
Java	14
Tomcat	14
<b>Installation Prerequisites</b>	<b>15</b>
MySQL Installation	15
Recommended Best Practice	18
Increase OS User Process Limit	18
<b>Installation Options</b>	<b>20</b>
Linux GUI Installer	20
Running the Linux GUI Installer	20
Command Line Installer	30
Running the Command Line Installer	31
<b>Post-Installation Procedures</b>	<b>45</b>
Running the Spark Application Utility Script	45
Configuring Redis	52
<b>Hadoop Integration</b>	<b>54</b>
Configure Hadoop Initial Settings	54
Configure Hadoop Services for ArcSight UBA	55



---

<b>Uninstall ArcSight UBA .....</b>	<b>60</b>
<b>Appendix A: Spark Job Properties .....</b>	<b>61</b>

# Introduction

Use the Installation Guide to learn how to install ArcSight User Behavior Analytics. In this guide, you can find the following information:

1. About ArcSight User Behavior Analytics
2. Reference Architecture
3. Installation components
4. Installation options
5. Post-installation procedures

## About ArcSight User Behavior Analytics

ArcSight UBA is a big data security analytics platform built on Hadoop that utilizes ArcSight UBA machine learning-based anomaly detection techniques and threat models to detect sophisticated cyber and insider attacks. ArcSight UBA uses Hadoop both as its distributed security analytics engine and long-term data retention engine. Hadoop nodes can be added as needed, allowing the solution to scale horizontally to support hundreds of thousands of events per second (EPS).

Features:

- Supports a rich variety of security data including security event logs, user identity data, access privileges, threat intelligence, asset metadata, and netflow data.
- Normalizes, indexes, and correlates security event logs, network flows, and application transactions.
- Utilizes machine learning-based anomaly detection techniques, including behavior profiling, peer group analytics, pattern analysis, and event rarity to detect advanced threats.
- Provides out-of-the-box threat and risk models for detection and prioritization of insider threat, cyber threat, and fraud.
- Risk-ranks entities involved in threats to enable an entity-centric (user or devices) approach to mitigating threats.
- Provides Spotter, a blazing-fast search feature with normalized search syntax that enables investigators to investigate today's threats and track advanced persistent threats over long periods of time, with all data available at all times.
- Provides the Investigation Workbench to detect links across disparate datasets to enable quick investigations and hunting for cyber threats.

## Who Should Read this Document

This document is intended for system administrators, system integrators, and deployment teams who need to install the application.

## ArcSight UBA Architecture

The ArcSight UBA solution utilizes services in a Hadoop cluster. ArcSight UBA provides deployment on an existing Hadoop cluster (see the Installation Guide for supported Hadoop distributions and versions).

ArcSight UBA provides integration with many different systems. The solution includes the following nodes that integrate with the Hadoop services:

- ArcSight UBA ConsoleNode: These nodes are Edge nodes in a Hadoop cluster that are used for the ArcSight UBA user interface and the configuration repository for all components used by the solution. Each Console Node performs the following tasks:
  - Provide visualizations for monitoring events, threat management dashboards, investigations and incident response
  - Build custom dashboards with visualizations for viewing violation and event data
  - Configure all ingestion jobs - user identities, access privileges, threat intelligence, security events and others
  - Administration interface for application support, personnel and administrators
  - Configure all policies and analytics, including behavior-based anomaly detection, peer-based analytics, threat modeling and risk analytics

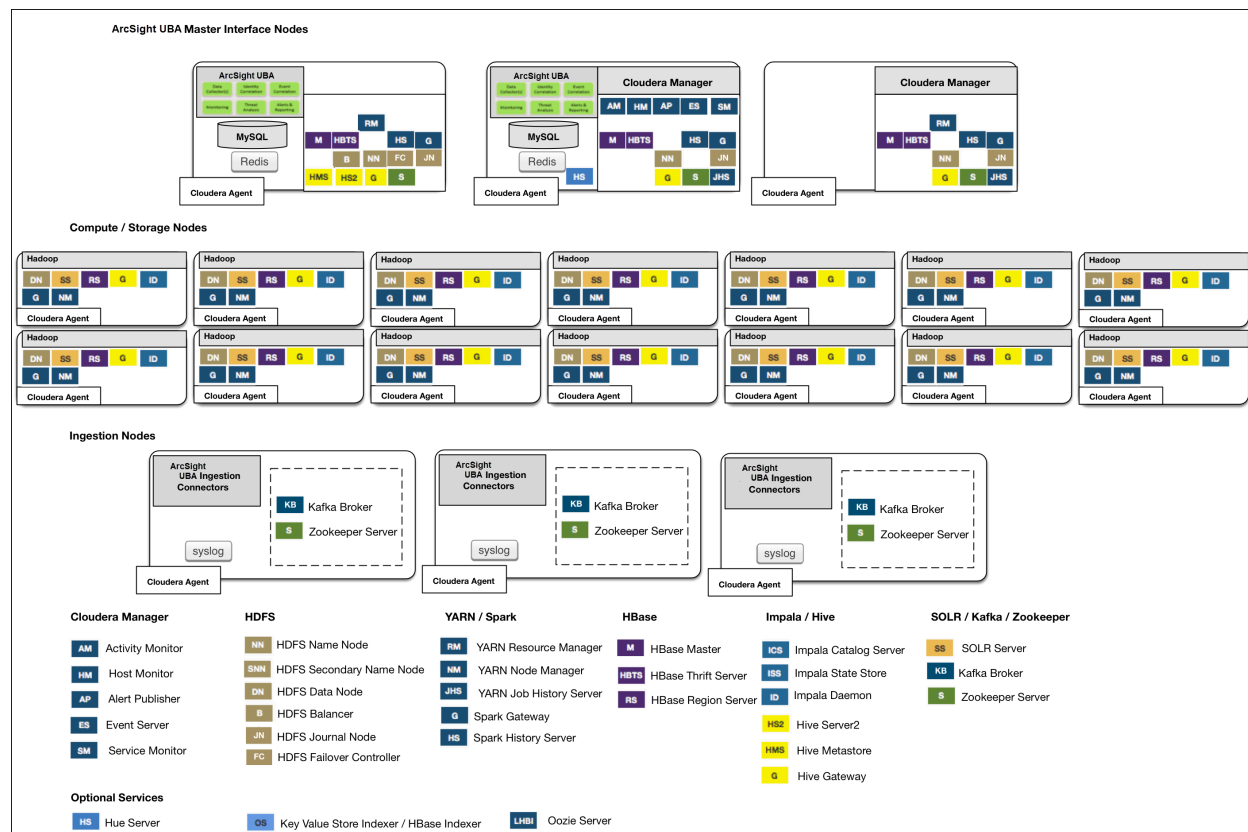
- ArcSight UBA Ingestion Nodes: These nodes are Edge nodes in a Hadoop cluster that are used to ingest security event log data into the environment with the ArcSight connectors.

Each ArcSight UBA Ingestion node performs the following tasks:

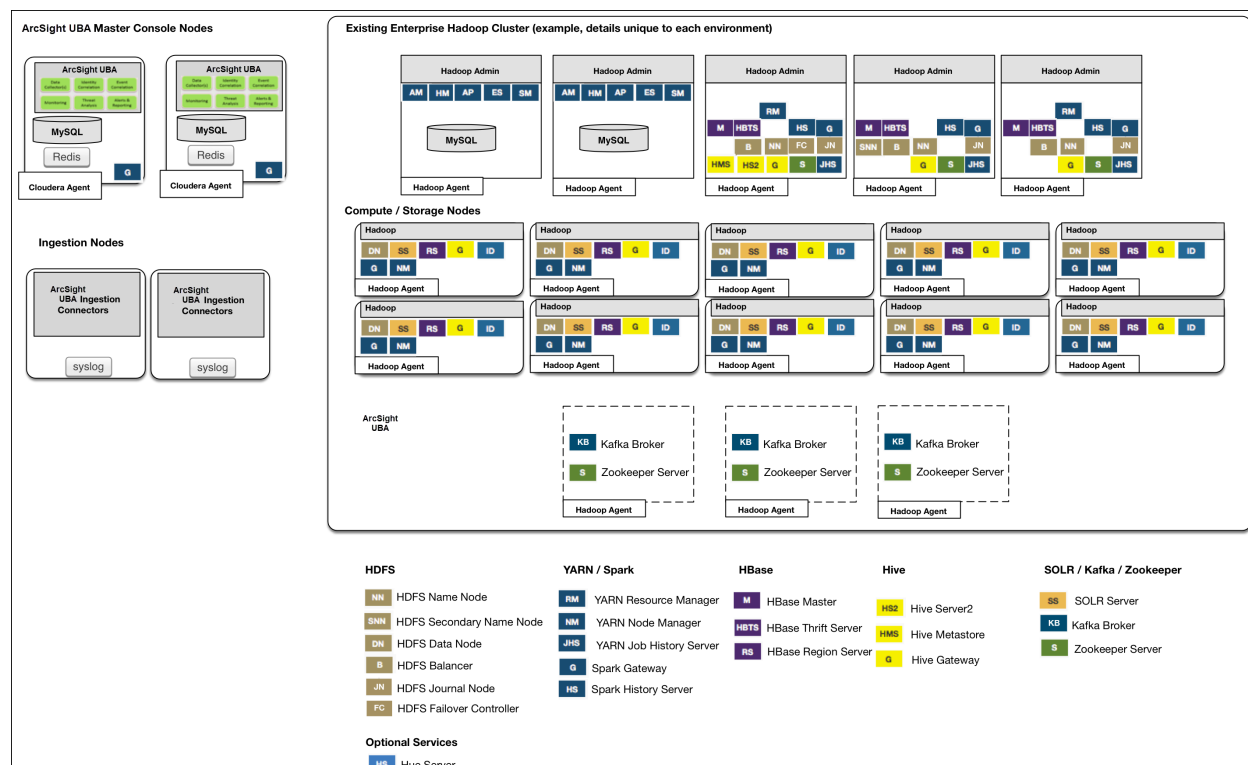
- Import Events from log sources
  - Publish events to Kafka Message Bus with batching, compression and encryption
  - Accept incoming log files on syslog
  - Cache In-transit messages
- ArcSight UBA Compute/Storage Nodes: These nodes are Edge nodes in a Hadoop cluster that are used to store compressed data and perform all the jobs associated with ArcSight UBA. Each ArcSight UBA Compute/Storage node performs the following tasks:
    - Fetch data from the ingestion nodes.
    - Perform all the jobs associated with ArcSight UBA based on the configuration stored in the Master node, including parsing, indexing, analytics, and storage.
    - Store data with 90% compression in structured JSON format.
    - Pass processed data to the ArcSight UBA console for review by the end user.

# ArcSight UBA Reference Architecture

## Option 1: Deployment with Dedicated Security Analytics Data



## Option 2: ArcSight UBA Deployment with Existing Hadoop Infrastructure



## Installation Components

The simplest deployment is installing the ArcSight UBA user interface and ingestion with the database and application running on the same server.

Data comes into Hadoop from the sources you've configured, and you log into the ArcSight UBA web interface on this same server to monitor and analyze data.

Depending on your needs, ArcSight UBA can be deployed with one or more ingestion nodes. This guide describes installing ArcSight UBA on a single node.

The following components are involved in the deployment:

- **Application Home Directory** – The application reads configuration data from files stored under the `securonix_home` folder. The `securonix_home` folder stores certain files required during application startup, configuration, and running of the application.
- **Relational database** – The application uses a relational database to store data and configuration for the ArcSight UBA application. ArcSight supports the MySQL database.
- **Ingestion node** – Ingestion nodes provide the capability to import and analyze activities and

security events.

- **Syslog forwarder** – Syslog forwarders are light-weight forwarders that have the capability to read log files incrementally and forward the logs as syslog to the ArcSight application or Real Time Analyzers.
- **Hadoop Cluster** – A Hadoop cluster with the services needed by ArcSight UBA.

Some of the core Hadoop components include the following in the ArcSight UBA architecture:

- **HDFS** - The Hadoop Distributed File System (HDFS) is used to store security events and violations. Data is stored in compressed parquet format
- **YARN** - Yet Another Resource Negotiator (YARN) provides resource management capabilities for jobs
- **Spark Streaming** - Processing framework for live streaming data
- **HBase** - Distributed no-SQL data store on HDFS to store the results of the analytics
- **Kafka** - Horizontally scalable message-bus used to manage the delivery of incoming security events
- **Impala(CDH) or Hive(HDP)** - Provides a SQL interface to the data stored in HDFS
- **Solr Cloud** - Provides distributed indexes of events for streaming fast, interactive access of security events and violations
- **ZooKeeper** - Cluster management software to maintain configurations and synchronization services across nodes within a cluster.
- **Redis**: In-memory database for fast in-memory look-ups. Redis is an open-source software project that implements data structure servers. It is networked and in-memory, and it stores keys with optional durability. In ArcSight UBA, Redis is used to store active lists of attribute values for specific durations of time that are used in Tier 2 Analytical Checks.

## Installation Flow

After installing and configuring the required Hadoop services, install ArcSight UBA using one of two available options:

- From the ArcSight UBA binary GUI installer on Linux
- From a Command Line installer

See [Installation Options](#) for information about how to install ArcSight UBA using the two available options.

## Installation Checklist

Before you begin the deployment, make sure that you have the following details ready:

- ArcSight User Behavior Analytics software.
- License files and license key.
- Deployment environment(s) – capacity planning, hardware, software, browser and port requirements.
- Database server (ArcSight UBA supports the MySQL 5.6 [and above] database).
- Source of identity data (human resource management system, LDAP source, database, others).
- Source of activity/event data (syslog server, audit tables, log management, SIEM, database monitoring, DLP, others).
- Email server configuration with an email account to send emails.
- Roles and associated privileges needed by users of the ArcSight User Behavior Analytics application.

## Prerequisites

- Oracle Java 8 (on all nodes, including YARN containers for Spark)
- MySQL 5.6.33 or above (on the ArcSight UBA Console Servers)

## Supported Operating Systems

The application supports the following operating systems:

- CentOS 6.8 and above, 7.1 and above
- RHEL 6.8 and above, 7.1 and above

## Supported Hadoop Deployment

ArcSight UBA can be deployed on an existing Hadoop cluster (Cloudera 5.x or Hortonworks 2.4, or 2.5).

The Hadoop distributions and the version of the services required are listed in the following table.

## Certification Matrix

<b>Service</b>	<b>CDH 5.11</b>	<b>CDH 5.12</b>	<b>CDH 5.13</b>	<b>CDH 5.10</b>	<b>Hortonworks 2.5</b>	<b>Hortonworks 2.6.x</b>
Hadoop	2.6.x	2.6.x	2.6.x	2.6.x	2.7.x	2.7.x
HBase	1.2.x	1.2.x	1.2.x	1.2.x	1.1.x	1.1.x
Hive	1.1.x	1.1.x	1.1.x	1.1.x	1.2.x	1.2.x
Impala	2.9.x	2.9.x	2.9.x	2.9.x	N/A	N/A
ZooKeeper	3.4.x	3.4.x	3.4.x	3.4.x	3.4.x	3.4.x
Kafka	0.10.x	0.10.x	0.10.x	0.10.x	0.10.x	0.10.x
Spark	1.6.x	1.6.x	1.6.x	1.6.x	1.6.x	1.6.x
	2.2	2.1x	2.1x	2.1x	2.1x	2.1x
Solr	4.10.x	4.10.x	4.10.x	4.10.x	5.5.x	5.5.x



## Required Communication Ports

1. Port for MySQL – Default: 3306
2. Tomcat Application Server Port – Default for HTTP: 8080. Default for HTTPS: 8443
3. Optional Ports:

Type	Description	Port
TCP Port	Optional	22
TCP Port	DNS host name lookup – DNS is used for name lookup and event enrichment.	UDP/TCP 53
TCP Port	DHCP/bootstrap protocol server is not needed when static IP addressing is used.	DHCP/port 67
TCP Port	Used for syslog server set up.	UDP 514
TCP Port	Only for server monitoring.	ICMP type 8
TCP Port	Get identity data from systems: connectivity varies by identity store.	Example: LDAP/389 LDAPS/636 to Active Directory
TCP Port	MSFT SMTP Gateway – SMTP notifications (email alerts from the application).	25/465
Master/Ingestion	communication uses:	3306/8443 (HTTPS)

4. Hadoop ports used by services in the cluster:

- SolrCloud
- ZooKeeper
- HBase
- Impala or Hive
- Kafka

## Supported Browsers

The application can be launched using any of the following browsers:

- Firefox 33 and above
- Internet Explorer 11
- Chrome (latest)

## Java

Oracle Java 1.8.0\_152

## Tomcat

Apache Tomcat 8.0.42

# Installation Prerequisites

Complete the following prerequisites before installing ArcSight UBA.

## MySQL Installation

Before running the ArcSight UBA product installer (ArcSight\_UBA\_6.10\_Application\_XXXXXX.bin), ensure MySQL is pre-installed and configured on the Linux server (RHEL6.8+ or CentOS 6.8+). If some of these items are not installed or are improperly configured, the installer will notify you during installation. After installation, you must install and integrate the Hadoop cluster to complete the configuration of the solution.

### Install MySQL

MySQL 5.6+ is required. This is used for configuration information. The ArcSight UBA data is stored in Hadoop.

- `root> wget http://dev.mysql.com/get/mysql-community-release-el6-5.noarch.rpm`
- `root> rpm -Uvh mysql-community-release-el6-5.noarch.rpm`
- `root> yum install mysql-server`

### Set up MySQL to Start at Boot Time

- `root> chkconfig mysqld on`

### Modify my.cnf file

- `root> vi /etc/my.cnf`

Add the following lines in the [mysqld] section:

- `lower_case_table_names=1`  
If set to 0, table names are stored as specified and comparisons are case sensitive. If set to 1, table names are stored in lowercase on disk and comparisons are not case sensitive. If set to 2, table names are stored as given but compared in lowercase.
- `innodb_file_per_table = 1`



**Note:** You must update the my.cnf file if you are using MySQL version 5.7, and you receive the following error:

Caused by: com.mysql.jdbc.exceptions.jdbc4.MySQLSyntaxErrorException:  
Expression #3 of SELECT list is not in GROUP BY clause and contains  
nonaggregated column 'securonix50beta.riskscorec0.userid' which is  
not functionally dependent on columns in GROUP BY clause; this is  
incompatible with sql\_mode=only\_full\_group\_by...



**Note:** To resolve this error, you must make the following change to the my.cnf file:

```
sql_mode="STRICT_TRANS_TABLES,NO_ZERO_IN_DATE,NO_ZERO_DATE,ERROR_
FOR_DIVISION_BY_ZERO,
NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION"
```

- max\_allowed\_packet=6000000

This increases the size of the packet that is transferred to MySQL during installation.

## Start MySQL

- root> service mysqld start

## Enable MySQL Connections from Specific IP Addresses or Hostnames



**Note:** Replace the following terms with those specific to your configuration:

- grant all on \*.\* to 'root'@'192.168.1.1' identified by 'password';

**OR**

- grant all on \*.\* to 'root'@'myhostname' identified by 'password';

After either grant statement above, run "flush privileges;".



**Note:** The ArcSight UBA Spark jobs connect to MySQL to obtain information.

## Change MySQL Data Directory

If MySQL has previously been installed on the hardware, or if you need to allocate more disk space for MySQL, use the following steps (commands are shown running as root):

## 1. Shutdown MySQL:

```
>> service mysqld stop
```

## 2. Create a backup of my.cnf:

```
>> cp /etc/my.cnf /home/securonix/my.cnf.bak
```

## 3. Create a folder called "data" in /storage:

```
>> chmod 755 /storage
>> mkdir /storage/data
```

## 4. Move contents of MySQL to the new data storage:

```
>> mv /var/lib/mysql /storage/data/
```

## 5. Change ownership of the "data" folder to mysql (recursively):

```
>> chown -R mysql:mysql /storage/data
```

## 6. Change the data directory and socket to the new location in "/etc/my.cnf".

## 7. Start MySQL:

```
>> service mysqld start
```

## Secure MySQL

After installing MySQL, run the secure installation:

```
/usr/bin/mysql_secure_installation
```

Choose the following options:

Enter the root password

```
Change Root password: Y
Remove anonymous user : Y
Disallow root login remotely: N
Remove test database: Y
Reload privilege tables now? Y
```

Enter the password that was randomly generated during installation, change the **<password>** above with a strong password for this installation.

```
service mysql restart
```

## MySQL System Privileges

Ensure that the MySQL database files are owned by the user "mysql" and the group "mysql".

```
ls -l /var/lib/mysql
```

Confirm that only the user "mysql" and "root" have access to the directory /var/lib/mysql.

The mysql binaries, which reside under the /usr/bin/ directory, should be owned by "root" or the specific system "mysql" user. Other users should not have write access to these files.

```
ls -l /usr/bin/my*
```

Grant privileges for remote access (change the password in the command below):

```
mysql -u root -p
mysql> show grants;
mysql> GRANT ALL PRIVILEGES ON *.* TO 'root'@'%' IDENTIFIED BY
PASSWORD '*DD10536B12A3478A0E2529893F6FB2EDBEA60310' WITH GRANT
OPTION;
Query OK, 0 rows affected (0.00 sec)
mysql> flush privileges;
mysql> quit
```

## Recommended Best Practice

The following best practices are recommended.

### Ensure Hostnames Resolve in DNS

Configure the host name and ensure it resolves in DNS. HTTPS/SSL certificates are recommended for secure access, and must exactly match the host name during connectivity. Using the hostname will allow for IP address changes later without re-configuration of SSL certificates.

### Increase OS User Process Limit

The operating system's default user-process limit may not be sufficient. Increase this limit to ensure the system has adequate processing capacity by following these steps:

1. Edit the `sysctl.conf` file as root (or sudo):

```
vi /etc/sysctl.conf
fs.file-max = 65536
Increase the maximum receive socket buffer size
net.core.rmem_max = 524280
Increase the default receive socket buffer size
net.core.rmem_default = 524280
Increase the maximum send socket buffer size
net.core.wmem_max = 524280
Increase the default send socket buffer size
net.core.wmem_default = 524280
Increase the maximum amount of option memory buffers
net.core.optmem_max = 57344
```

2. Edit the `90-nproc.conf` file (or sudo):

```
vi /etc/security/limits.d/90-nproc.conf
* soft nproc 65535
* hard nproc 65535
* soft nofile 65535
* hard nofile 65535
```

3. Edit the `limits.conf` file:

```
vi /etc/security/limits.conf
* soft nproc 65535
* hard nproc 65535
* soft nofile 65535
* hard nofile 65535
```

4. Reboot the machine.
5. Log in as ArcSight UBA user and run following command to verify output:

```
ulimit -a
<check commandline>
```

### Increase Open Files Limit for MySQL

Increasing the `open_files_limit` parameter in `my.cnf` works differently in CentOS7. Follow the steps below to increase the limit:

1. Edit `/usr/lib/systemd/system/mysqld.service`
2. Add the below two lines under the `[Service]` section

```
LimitNOFILE=65535
LimitNPROC=65535
```
3. Increase the `table_open_cache` and `open_files_limit` in `my.cnf`
4. Perform `sudo systemctl daemon-reload`
5. Restart MySQL

# Installation Options

You can install ArcSight UBA using two different options:

- [From the ArcSight UBA binary GUI installer on Linux](#)
- [From a Command Line installer](#)

The installation methods above can be used to install the following components:

- ArcSight UBA Console on the Master Nodes
- ArcSight UBA Spark Gateway on the Master Nodes or wherever the YARN manager is deployed

The following sections describe how to install ArcSight UBA using these above options.



**Note:** Redis is deployed along with the Console.



**Note:** Syslog is deployed upon selecting the Console or Ingester.

## Linux GUI Installer

This section describes how to install ArcSight UBA using the GUI Installer on Linux.

### Running the Linux GUI Installer

#### Prerequisite – X Window

If you want to install the application on Linux with GUI mode, and if X Window has not been configured, follow these steps prior to using the ArcSight UBA installer:

```
Run: yum -y groupinstall "Desktop" "Desktop Platform" "X Window
System" "Fonts"
Run: yum install xorg-x11-xauth xterm
```

Verify X Forwarding parameters in both SSH files and restart SSH:

```
/etc/ssh/ssh_config
ForwardX11Trusted yes
/etc/ssh/sshd_config32-
X11Forwarding yes
Run: service sshd restart
Run: export DISPLAY=localhost:10.0
```

User XAUTH to create the following file:

```
/home/Securonix/.Xauthority
serverhostname/unix:10 MIT-MAGIC-COOKIE-1 33d3b333e33f333e-
f33333a3c333e3bd
```

The ArcSight UBA software should be installed by a non-root user. To create a non-root user, open a terminal session and run the following commands:



1. `useradd securonix` (For example, choose any user name.)
2. `passwd securonix` (Provide your non-root user a password.)

In order for Syslog-ng installation to proceed, the user must provide the SUDO password in installer screen. Since the installation binary is started as non-root user, this information must be provided in the sudoers file.

3. Log in as root and go to `/etc/`

```
vi sudoers
```

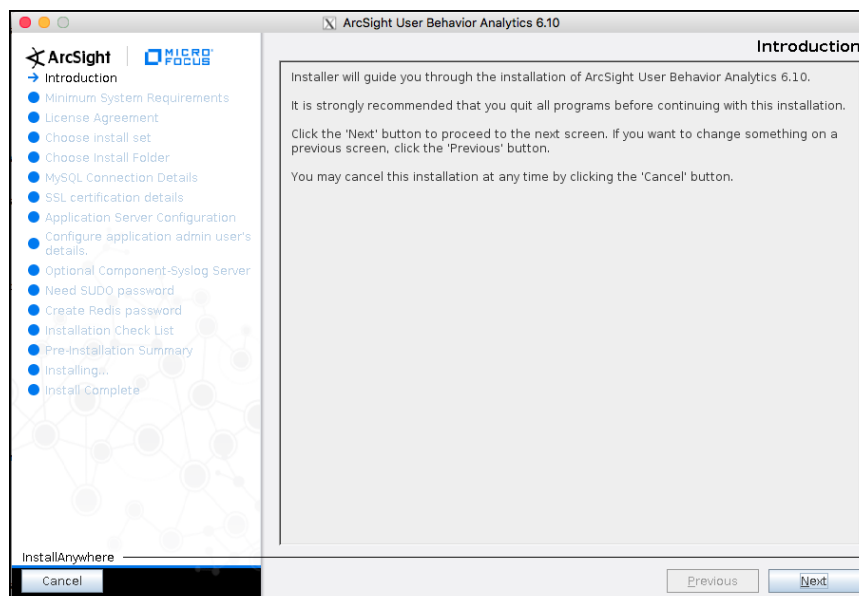
4. Scroll down through the sudoers file to the section below, add the user information for the non-root user that will start the installer (`securonix` from the previous example ), and save the file. Provide the non-root user password as sudo password on the installer screen.

```
## Allow root to run any commands anywhere
root ALL=(ALL) ALL
securonix ALL=(ALL) ALL
## Allows members of the 'sys' group to run networking, soft-
ware,
## service management apps and more.
# %sys ALL = NETWORKING, SOFTWARE, SERVICES, STORAGE,
DELEGATING, PROCESSES, LOCATE, DRIVERS
## Allows people in group wheel to run all commands
```

5. Open a terminal and execute the following command to begin the installation of the ArcSight UBA application:

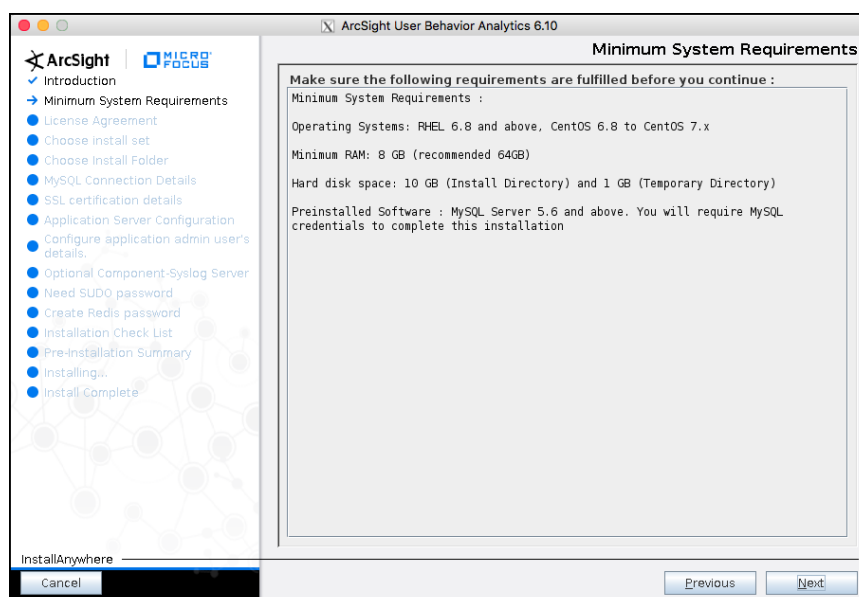
```
ArcSight_UBA_6.10_Application_xxxxxxx.bin
```

The installer will open. Click **Next**.

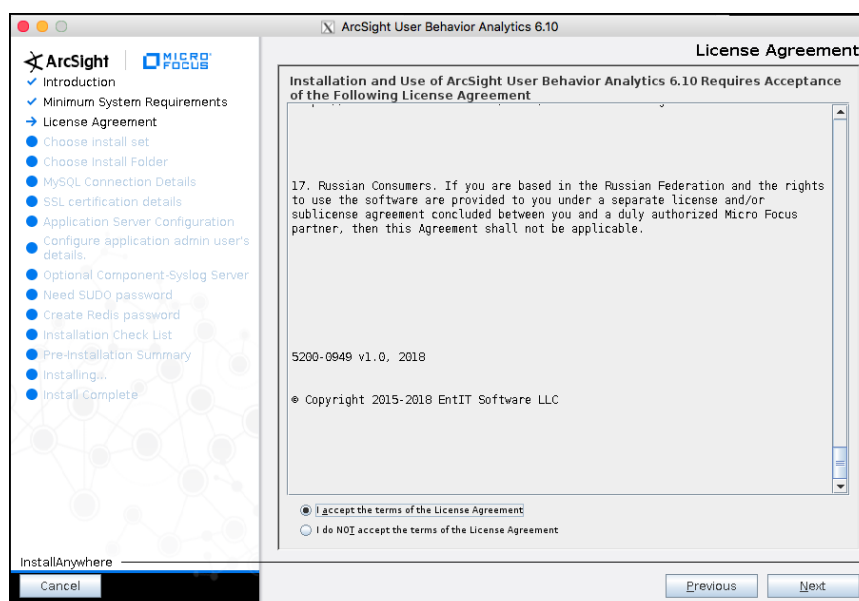


6. **Minimum System Requirements:** Review the system requirements and click **Next**. If your

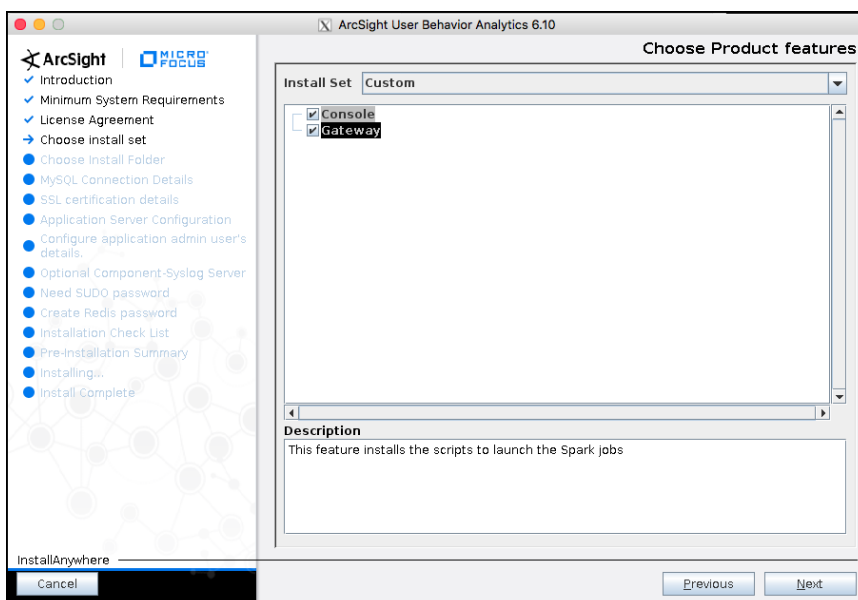
system does not meet minimum requirements, you may accept the risk, but the installation may not be supported.



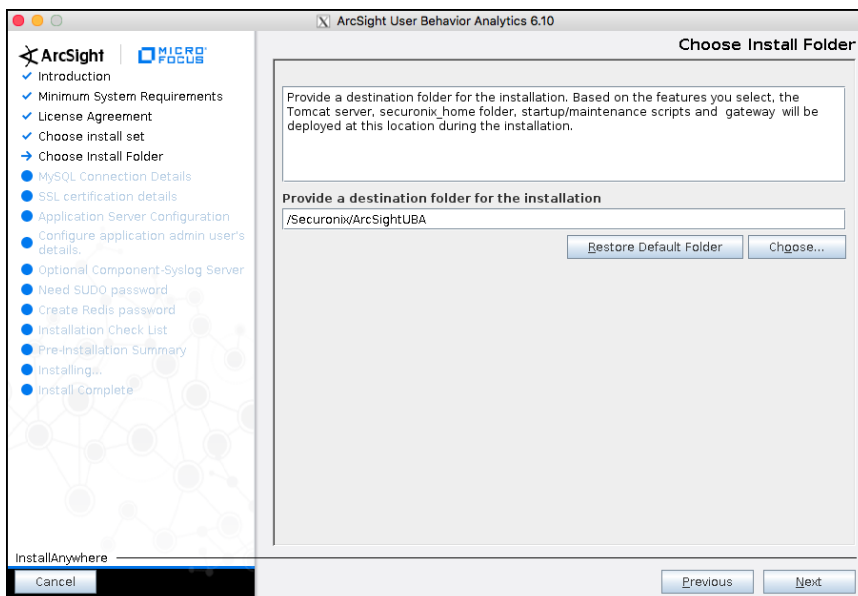
- License Agreement:** On the license screen, read and scroll to the bottom of the license, to enable and check the “I accept...” button.



- Choose Install Set:** Choose the product features to install or select an Install Set from the drop-down.



9. **Choose Install Folder:** Select the location to install the ArcSight UBA application and click **Next**. The default will install the application to a subdirectory in the home directory for the current user.



10. **MySQL Connection Details:** Modify the MySQL connection details to include the system DBA user and password and port modifications as necessary, and click **Next**. The database schema will be created on the MySQL instance specified.

**ArcSight User Behavior Analytics 6.10**

**MySQL Connection Details**

Provide credentials to the MySQL database. The installer will test the credentials provided and use these to run the application. Make sure this user has the adequate permissions.

Host Name / IP Address: 10.0.51.193

Port: 3306

Database: arcsightuba

User Name: root

Password: .....

Previous Next

The credentials supplied will be tested for valid connectivity and you will receive an error screen if the connection to MySQL fails to proceed, correct the host/user/password as needed, and re-enter the credentials.

11. **SSL certification details:** Enter the SSL certification details as shown in the following image (use information appropriate for your organization).

**ArcSight User Behavior Analytics 6.10**

**SSL certification details**

The information provided on this screen will be used to generate the SSL keystore and certificate. It is required to establish the connection using HTTP over SSL protocol.

IP Address / Host Name: 10.0.51.193

Port: 8443

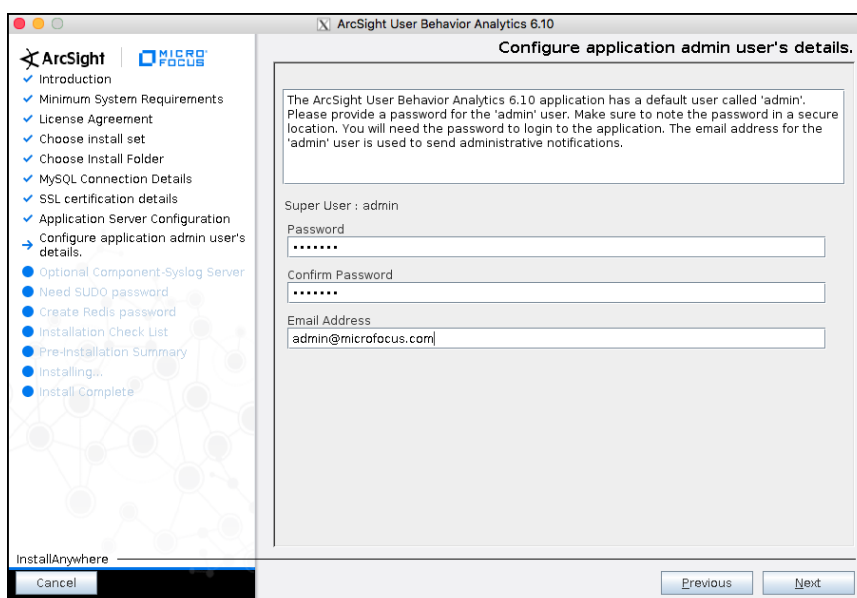
Organization Name: MicroFocus

Keystore Password: .....

Confirm Password: .....

Previous Next

12. **Configure application admin user's details:** Enter the credentials for the ArcSight UBA administrative user (this account will be used for login using a browser). The password must be a minimum of 6 characters. You will get an error screen if you do not meet this requirement.

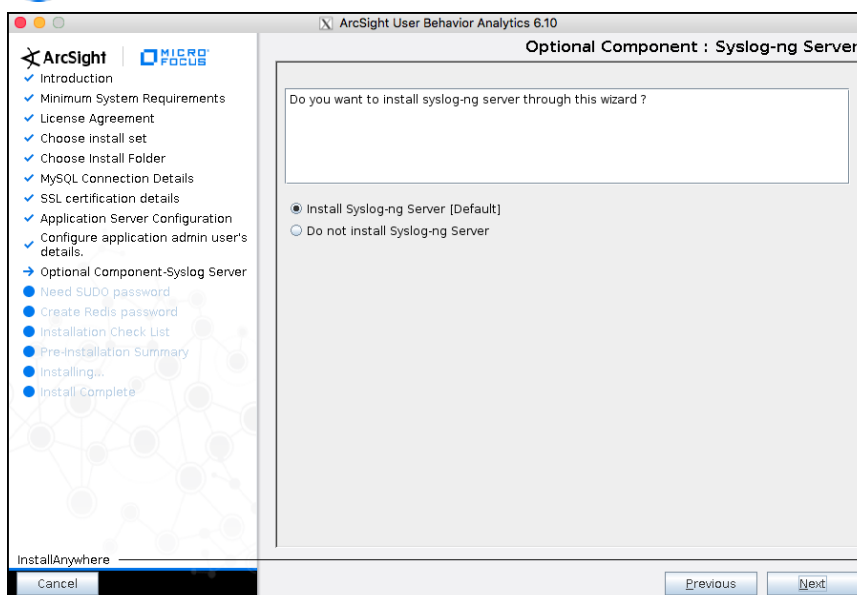


13. **Optional Component: Syslog-ng Server:** Select one of the following to indicate if you want to install syslog-ng server through the install wizard:

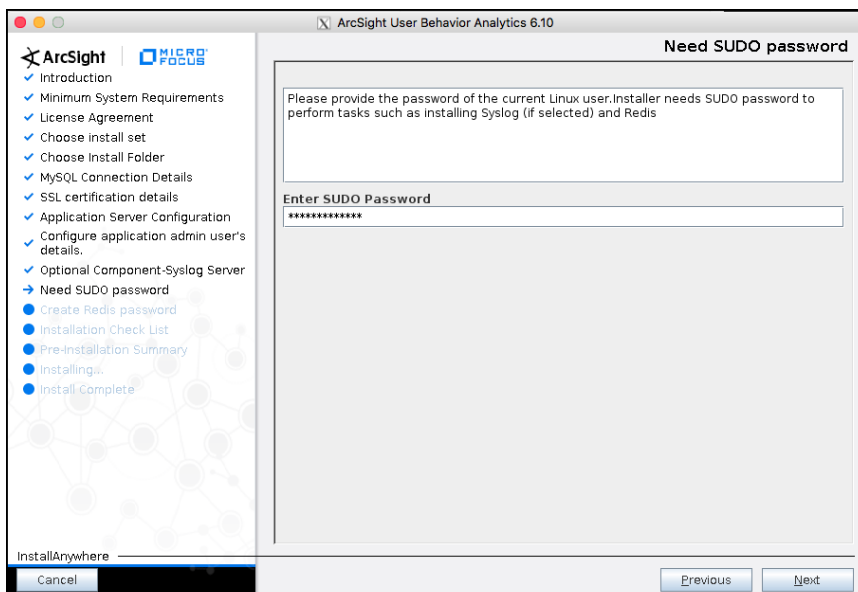
- **Install Syslog-ng Server [Default]**
- **Do not install Syslog-ng Server**



**Note:** ArcSight UBA requires Syslog-ng server to be running.

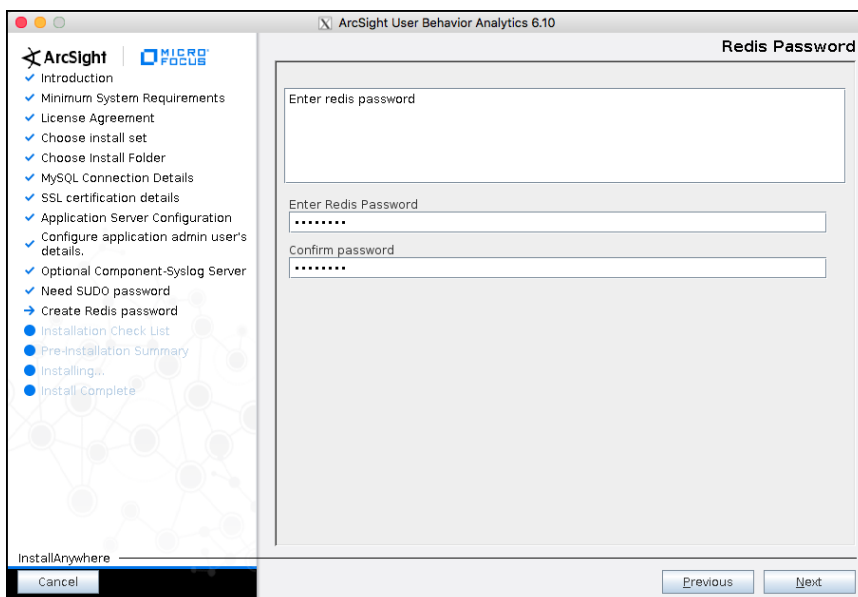


14. **Need SUDO password:** Enter the SUDO password (required to install Syslog and Redis).

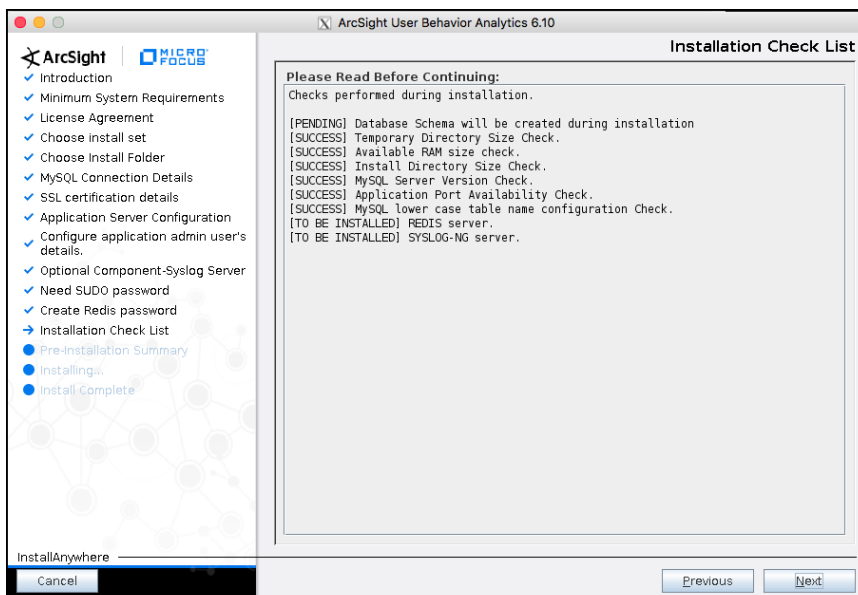


The installer checks the SUDO password before continuing.

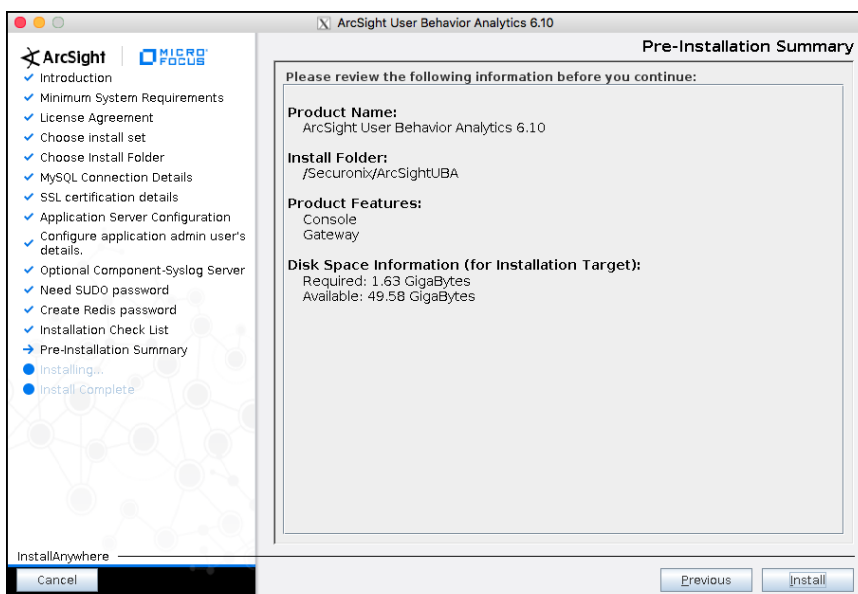
15. **Create Redis Password:**



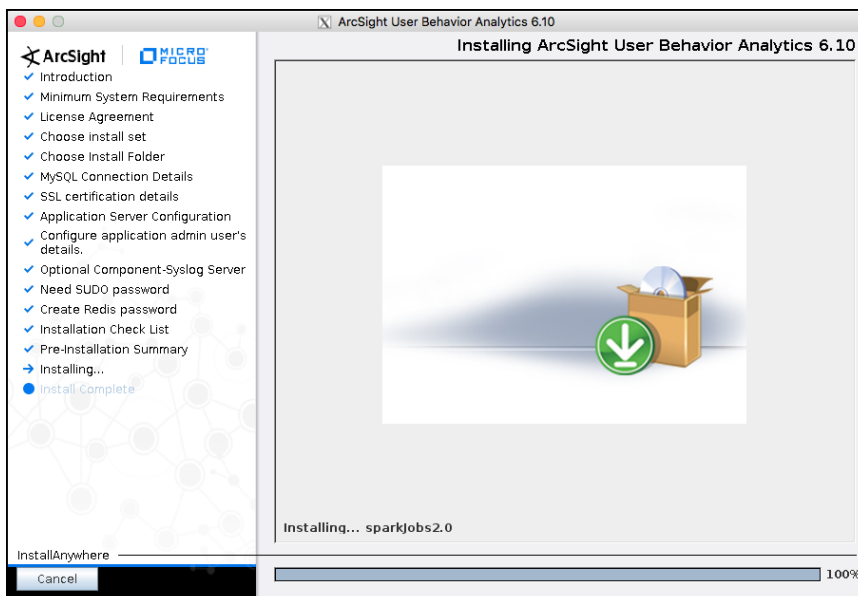
16. **Installation Checklist:** Review the installation checklist and click **Next** to continue.



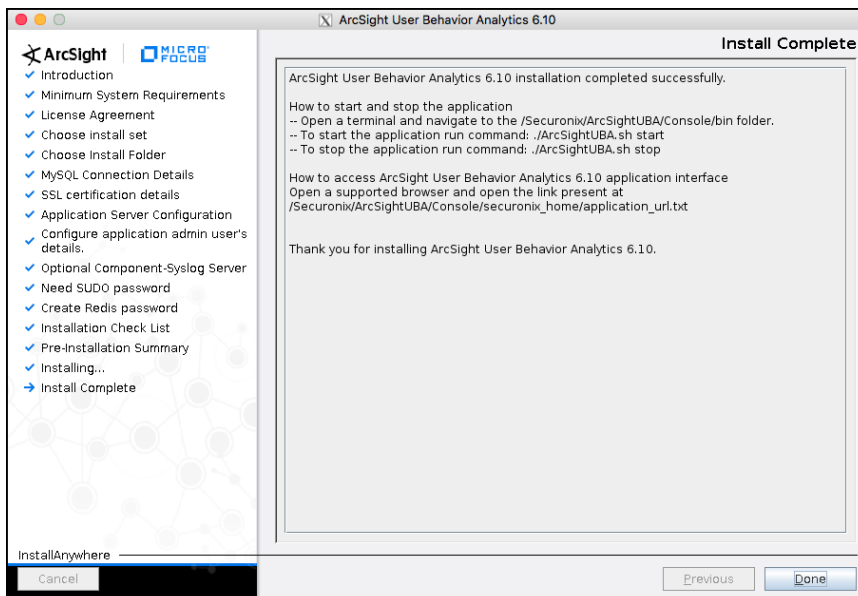
17. **Pre-Installation Summary:** Review the Pre-Installation Summary and click **Install**.



Additional components will be installed. After the status bar shows 100% you should see an installation complete screen. The time required will vary based on hardware performance.



33. **Install Complete:** Click **Done** to continue from the Successful installation – Summary Screen.



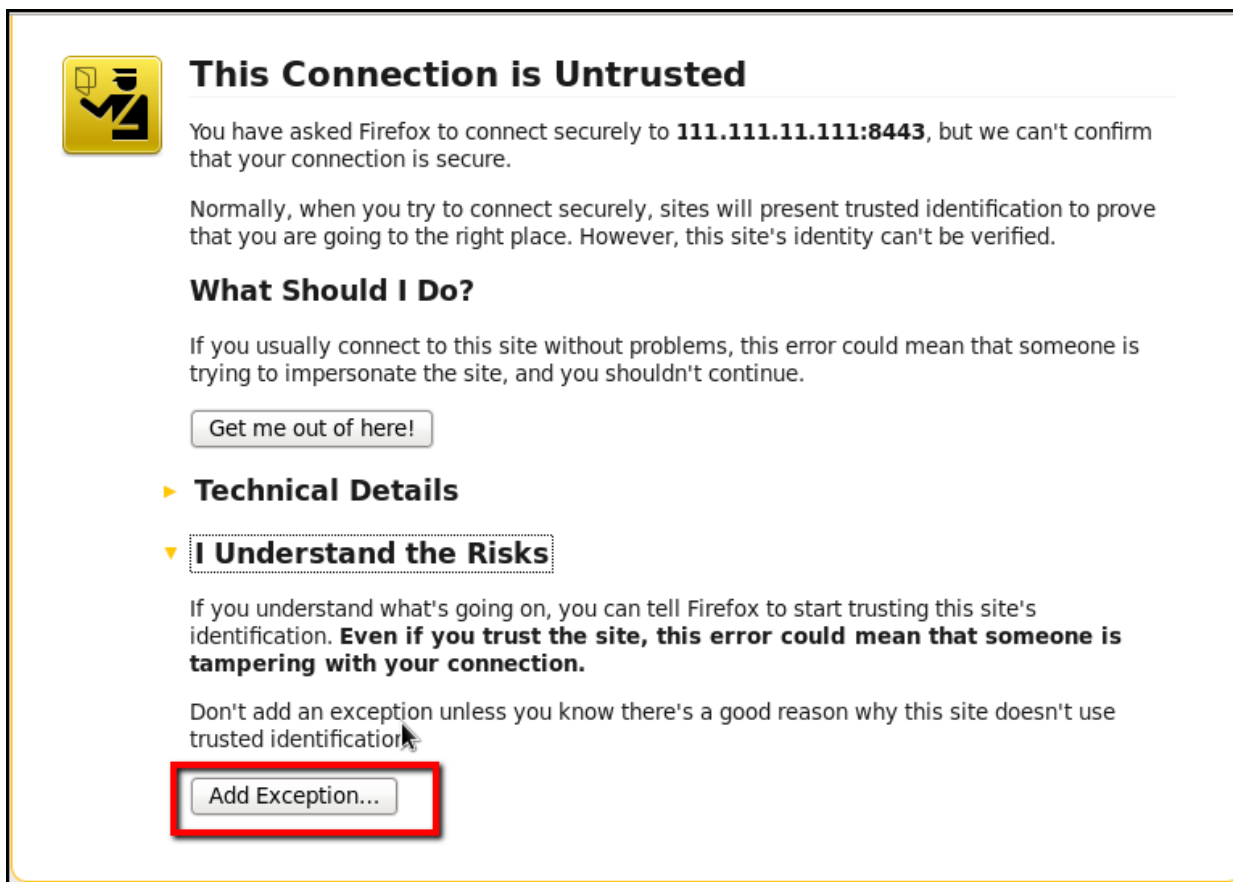
34. Start the ArcSight UBA application from a command line, by running the following command:  
`../ArcSightUBA.sh start` (Stop or restart the application by substituting stop or restart in place of start as shown in the next image.)

When startup has completed, you should see the following message:

The ArcSightUBA Web Interface is at ...



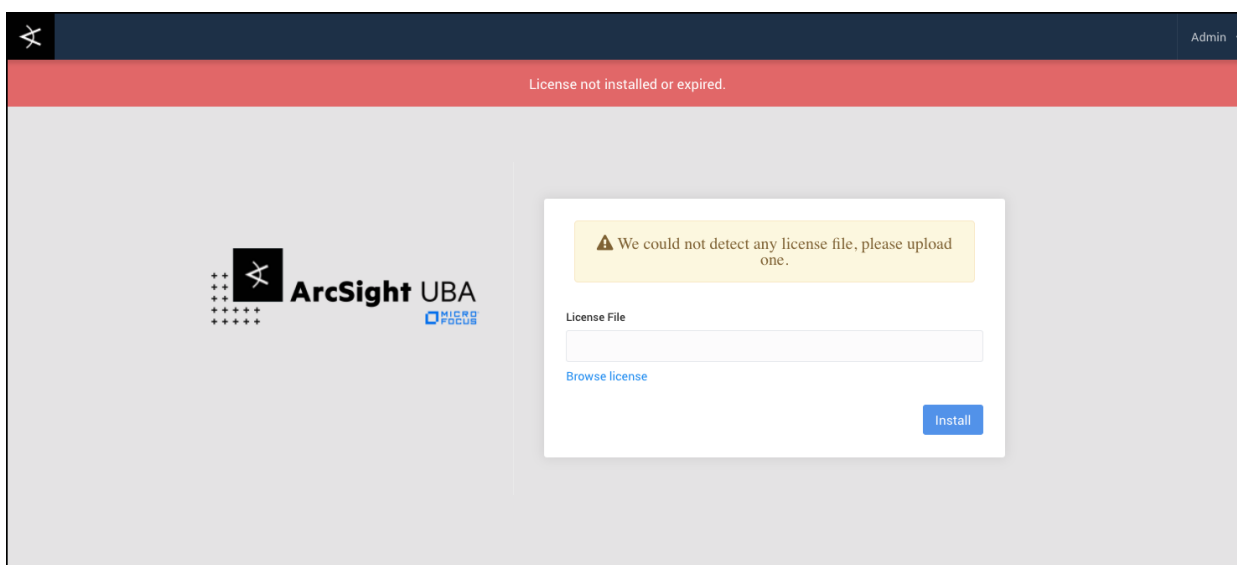
35. Validate the installation by connecting to the URL shown at the end of the startup script.
36. Accept and add an exception for the self-signed certificate of the host you have just configured.



37. Log in to the system using the admin user and password previously configured during the installation.



38. During the first installation, you will see the following screen. You must install the license provided to continue. Select the **Install License** option.



39. Select the \*.lic file provided with your licensing agreement.

After installing the license, you are returned to the login screen.

## Command Line Installer

This section describes how to install ArcSight UBA using the Command Line Installer.

## Running the Command Line Installer

The ArcSight UBA software should be installed by a non-root user. To create a non-root user, open a terminal session and run the following commands:

1. `useradd securonix` (For example, choose any user name.)

2. `passwd securonix` (Provide your non-root user a password.)

In order for Syslog-ng installation to proceed, the user must provide the SUDO password in the installer screen. Since the installation binary is started as non-root user, this information must be provided in the sudoers file.

3. Log in as root and go to `/etc/`

```
vi sudoers
```

4. Scroll down through the sudoers file to the section below, add the user information for the non-root user that will start the installer (`securonix` from the example above), and save the file. Provide the non-root user password as sudo password on the installer screen.

```
## Allow root to run any commands anywhere
root ALL=(ALL) ALL
securonix ALL=(ALL) ALL
## Allows members of the 'sys' group to run networking, soft-
ware,
## service management apps and more.
# %sys ALL = NETWORKING, SOFTWARE, SERVICES, STORAGE,
DELEGATING, PROCESSES, LOCATE, DRIVERS
## Allows people in group wheel to run all commands
```

Complete the following steps to install the ArcSight UBA application:

1. Open a terminal and execute the following command:  
`./ArcSight_UBA_6.10_Application_xxxxxxx.bin -i console`
2. **Introduction:** The introductory text will explain how to quit the installation or go back to a previous step. Respond to each prompt to proceed to the next step in the installation.

```
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...

Launching installer...

=====
ArcSight User Behavior Analytics 6.10          (created with InstallAnywhere)
=====

Preparing CONSOLE Mode Installation...

=====
Introduction
=====

InstallAnywhere will guide you through the installation of ArcSight User
Behavior Analytics 6.10.

It is strongly recommended that you quit all programs before continuing with
this installation.

Respond to each prompt to proceed to the next step in the installation.  If
you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE:
```

3. **Minimum System Requirements:** The minimum system requirements are displayed. If the system does not meet the minimum requirements, you will have an option to quit (default) or continue with the installation. Click **Enter** to continue.

```
=====
Minimum System Requirements
=====

Make sure the following requirements are fulfilled before you continue :

Minimum System Requirements :

Operating Systems: RHEL 6.8 and above, CentOS 6.8 to CentOS 7.x

Minimum RAM: 8 GB (recommended 64GB)

Hard disk space: 10 GB (Install Directory) and 1 GB (Temporary Directory)

Preinstalled Software : MySQL Server 5.6 and above. You will require MySQL
credentials to complete this installation
```

4. **License Agreement:** Click **Enter** when prompted to move through the License Agreement. Click **Y** at the end to accept the license and continue.

```
=====
License Agreement
=====
```

Installation and Use of ArcSight User Behavior Analytics 6.10 Requires Acceptance of the Following License Agreement:

Micro Focus End User License Agreement – Enterprise Version

1. Applicability. This end user license agreement (the "Agreement") governs the use of accompanying software, unless it is subject to a separate agreement between you and Micro Focus International plc and its subsidiaries (Micro Focus). By downloading, copying, or using the software you agree to this Agreement. Micro Focus provides translations of this Agreement in certain languages other than English, which may be found at: <https://software.microfocus.com/about/software-licensing>.

2. Terms. This Agreement includes supporting material accompanying the software or referenced by Micro Focus, which may be software license information, additional license authorizations, software specifications, published warranties, supplier terms, open source software licenses and similar content ("Supporting Material"). Additional license authorizations are

[PRESS <ENTER> TO CONTINUE:

at: <https://software.microfocus.com/about/software-licensing>.

3. Authorization. If you agree to this Agreement on behalf of another person or entity, you warrant you have authority to do so.

4. Consumer Rights. If you obtained software as a consumer, nothing in this Agreement affects your statutory rights.

5. Electronic Delivery. Micro Focus may elect to deliver software and related software product or license information by electronic transmission or download.

5. **Choose Install Set:** Choose the product features to be installed by the installer.



**Note:** Feature 1: Console is selected by default, indicated by an **X**. To add features, type the number of the feature (example, type 2 to add Gateway). Typing 1 will deselect Console.

```
=====
Choose Install Set
=====

=====
Choose Product Features (By default Console is selected. Please type 2 to
select both Console and Gateway features)
=====

ENTER A COMMA_SEPARATED LIST OF NUMBERS REPRESENTING THE FEATURES YOU WOULD
LIKE TO SELECT, OR DESELECT. TO VIEW A FEATURE'S DESCRIPTION, ENTER
'?<NUMBER>'. PRESS <RETURN> WHEN YOU ARE DONE:

    1- [X] Console
    2- [ ] Gateway

[Please choose the Features to be installed by this installer.: 2

=====
Chosen Install features
=====

Features chosen for Installation :
Console,Gateway

->1- Continue
    2- Re-select features

ENTER THE NUMBER OF THE DESIRED CHOICE, OR PRESS <ENTER> TO ACCEPT THE
[   DEFAULT:
```

6. **Choose the Install Folder:** Choose the destination folder for the installation. Accept the default or enter the desired absolute path for the installation. Default: /Securonix/ArcSightUBA.

```
=====
Choose Install Folder
-----

Provide a destination folder for the installation. Based on the features you
select, the Tomcat server, securonix_home folder, startup/maintenance scripts
and gateway will be deployed at this location during the installation.

Provide a destination folder for the installation

Default Install Folder: /Securonix/ArcSightUBA

ENTER AN ABSOLUTE PATH, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
: █
```

7. **MySQL Connection details:** Provide credentials to the MySQL database. Ensure the user has the appropriate permissions.



```
ENTER THE NUMBER OF THE DESIRED CHOICE, OR PRESS <ENTER> TO ACCEPT THE
[  DEFAULT: 2

=====
MySQL Connection Details
=====

Provide credentials to the MySQL database. The installer will test the
credentials provided and use these to run the application. Make sure this user
has the adequate permissions.

[Host Name / IP Address (Default: ): 10.0.51.193

=====
MySQL Connection Details
=====

[Port (Default: 3306):

=====
MySQL Connection Details
=====

[Database (Default: arcsightuba):

=====
MySQL Connection Details
=====

[User Name (Default: ): root

=====
MySQL Connection Details
=====

Password: █
```

The credentials supplied will be tested for valid connectivity and you will receive an error message if the connection to MySQL fails. To proceed, correct the host/user/password as needed, and re-enter the credentials.

8. **SSL certification details:** Enter the SSL certification details (if applicable).

```
=====
SSL certification details
-----

The information provided on this screen will be used to generate the SSL
keystore and certificate. It is required to establish the connection using
HTTP over SSL protocol.

[IPAddress / Host Name (Default: )]: 10.0.51.193

=====
SSL certification details
-----

[Port (Default: 8443):

=====
SSL certification details
-----

[Organization Name (Default: )]: MicroFocus

=====
SSL certification details
-----

[Keystore Password:
```

14. **Configure application admin user's details:** Enter the following credentials for the ArcSight UBA administrator. The Super User name is admin.

```
=====
Configure application admin user's details.
=====

The ArcSight User Behavior Analytics 6.10 application has a default user
called 'admin'. Please provide a password for the 'admin' user. Make sure to
note the password in a secure location. You will need the password to login to
the application. The email address for the 'admin' user is used to send
administrative notifications.

Super User : admin

[Email Address (Default: )]: admin@microfocus.com


=====
Configure application admin user's details.
=====

[Password:
```

15. **Optional Component: Syslog Service:** ArcSight UBA requires Syslog-ng server to be running. Select whether to install Syslog-ng Server (default) through the install wizard. If Syslog-ng is already installed, installation of the Syslog-ng is skipped.

```
=====
Optional Component : Syslog-ng Server
=====

Do you want to install syslog-ng server through this wizard ?

->1- Install Syslog-ng Server [Default]
   2- Do not install Syslog-ng Server

[ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT::
```

16. **Need SUDO password for Redis/Syslog-ng Installation:** Enter the SUDO password (password of the current Linux user) to install the Redis server/Syslog-ng. The system checks the password and whether syslog has been installed.

```
=====
Need SUDO password
=====

Please provide the password of the current Linux user. Installer needs SUDO
password to perform tasks such as installing Syslog (if selected) and Redis

SUDO Password: █
```

17. **Password for Redis Server:** Set a password for the Redis server.

Please Enter the Password:

```
=====
Create a password for Redis server
=====

Set a new password for Redis server

[Enter REDIS Password:
```

18. **Installation Checklist:** Read before continuing. Click **Enter** to continue when finished.

```
=====
Installation Check List
-----

Please read before continuing:

Checks performed during installation.

[PENDING] Database Schema will be created during installation
[SUCCESS] Temporary Directory Size Check.
[SUCCESS] Available RAM size check.
[SUCCESS] Install Directory Size Check.
[SUCCESS] MySQL Server Version Check.
[SUCCESS] Application Port Availability Check.
[SUCCESS] MySQL lower case table name configuration Check.
[TO BE INSTALLED] REDIS server.
[TO BE INSTALLED] SYSLOG-NG server.

[IMPORTANT INFORMATION COMPLETE. PRESS <ENTER> TO CONTINUE:
```

19. **Pre-Installation Summary:** Review before continuing. Click **Enter** to continue when finished.

```
=====
Pre-Installation Summary
=====

Please review the following information before you continue:

Product Name:
  ArcSight User Behavior Analytics 6.10

Install Folder:
  /Securonix/ArcSightUBA

Product Features:
  Console,
  Gateway

Disk Space Information (for Installation Target):
  Required:  1.63 GigaBytes
  Available: 48.21 GigaBytes

PRESS <ENTER> TO CONTINUE: █
```

20. **Installation Complete:** Review instructions for the application and click **Enter** to exit the installer.

```

-----|-----|-----|-----]
Certificate stored in file </Securonix/ArcSightUBA/Console/Tomcat/conf/Securonixhost.cer>
Certificate was added to keystore

=====
Installation Complete
=====

ArcSight User Behavior Analytics 6.10 installation completed successfully.

How to start and stop the application
-- Open a terminal and navigate to the /Securonix/ArcSightUBA/Console/bin
folder.
-- To start the application run command: ./ArcSightUBA.sh start
-- To stop the application run command: ./ArcSightUBA.sh stop

How to access ArcSight User Behavior Analytics 6.10 application interface
Open a supported browser and open the link present at
/Securonix/ArcSightUBA/Console/securonix_home/application_url.txt

Thank you for installing ArcSight User Behavior Analytics 6.10.
PRESS <ENTER> TO EXIT THE INSTALLER: █

```

21. Start the ArcSight UBA application from a command line by running the following command:
 

```
./ArcSightUBA.sh start
```

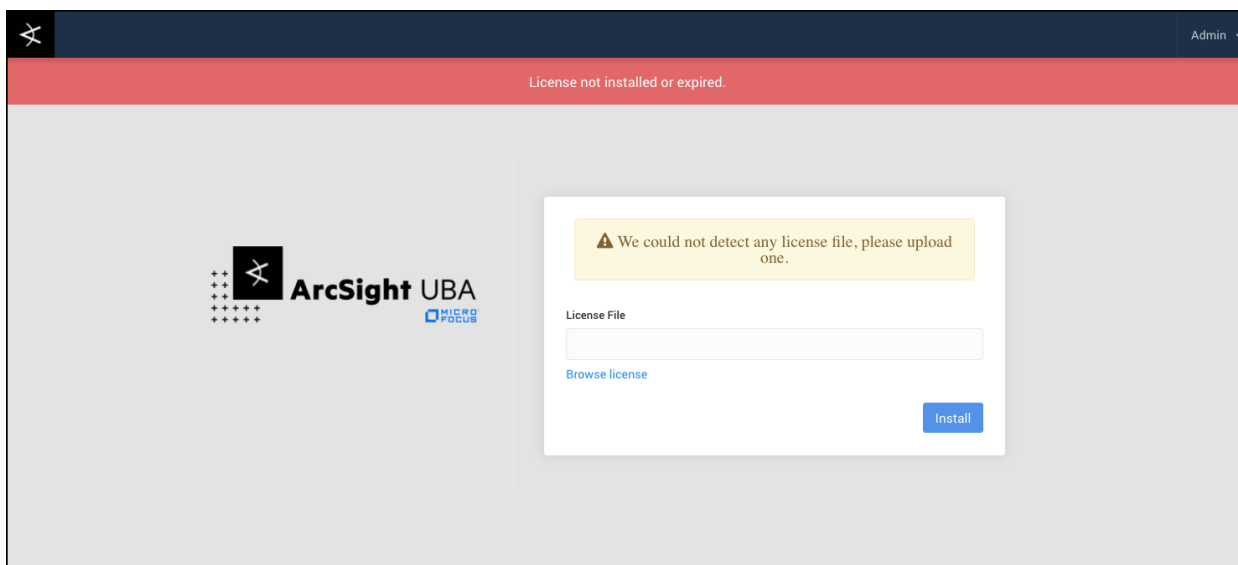
 Stop or restart the application by adding `stop` or `restart` in place of `start`.  
 When startup has completed, you should see the following message:
 

```
The ArcSightUBA Web Interface is at ...
```
22. Validate the installation by opening your browser and connecting to the URL shown at the end of the startup script.
23. Log in to the system using the admin user and password previously configured during the install-

ation.



24. During the first installation, you will see the following screen. You must install the license provided to continue. Select the **Install License** option.



25. After installing the license, you are returned to the login screen.



# Post-Installation Procedures

When ArcSight UBA is successfully installed, complete the following Post-Installation Procedures:

- [Run the Spark Application Utility Script](#)
- [Configure Redis](#)
- [Configure Hadoop](#)

## Running the Spark Application Utility Script

When you have successfully installed ArcSight UBA, complete the following steps to run the Spark jobs:



**Note:** For complete details to run the scripts with different options, see the PFA readme file in **<Installation\_folder>/Spark\_jobs/sparkJobs/readme\_snypr\_apps.txt**.

1. Navigate to the installation folder designated during Installation. Example: <Install\_folder>/Spark\_jobs/sparkJobs/.
2. (Optional) Edit Spark jobs parameters in **snypr\_apps.properties** in **<Install\_folder>/Spark\_jobs/sparkJobs/conf/snypr\_apps.properties**.

ArcSight\_apps.properties

- The scripts to run each spark job have been converted to accept the configurations as parameters
- Configurable parameters:
  - Name (name)
  - Driver memory (driver-memory)
  - Number of executors (num-executors)
  - Executor Memory (executor-memory)
  - Drive cores (driver-cores)
  - Executor cores (executor-cores)
  - Consumer group (cg)
  - Max rate per partition (mrpp)
  - Duration (d)
  - HDFS user having permission to run the spark applications (hdfs.user)
  - Unique identifier to distinguish your jobs, attached as prefix to the name of each spark job initiated (tenant.id)  
Eg if tenant.id =CS, the jobs will be run as "CS\_Event\_Enrichment"
  - queue name

- These configurations are maintained in a file, "snypr\_apps.properties". This file can be edited to run your custom configurations for each job.



**Note:** For a complete list of the default property values for each Spark Job, see [Appendix A: Spark Job Properties](#).

### 3. Execute the following command to run all Spark jobs at once:

```
sh snypr_apps.sh -a all
```

```
snypr_apps.sh
```

- Utility script to start the spark application by reading from the properties file, ArcSight\_apps.properties
- The script will start each application, wait for 20s and check the status of the application. If it is in "Running" state, then it generates a <jobname.pid> file with the applicationId of the application. Example: jobname=CS\_Event\_Enrichment, then CS\_Event\_Enrichment.pid will have the applicationId for CS\_Event\_Enrichment
- If it doesn't go to "Running" state, then the <Job-name=applicationId> is stored in a file < tenant.id\_pending\_apps.txt> Eg.CS\_pending\_apps.txt
- This file is once again checked after 30s to check if the application has moved to "Running" state or "Failed" state. The logs for status check for pending applications are maintained in < tenant.id\_pending\_apps.log> Eg.CS\_pending\_apps.log
- Following are the ways the script can start the spark application:
  - o -r (or) range : Series of applications at once i.e to start jobs from enrichment through behavior analytics, we can use:  
 sh snypr\_apps.sh -r <from\_number-to\_number>  
 Eg:  
 sh snypr\_apps.sh -r 1-4  
 sh snypr\_apps.sh -range 1-4
  - o -i (or) initiate : Start individual jobs or a list of discontinuous jobs:  
 sh snypr\_apps.sh -i <job\_number>  
 Eg.  
 sh snypr\_apps.sh -i 1  
 sh snypr\_apps.sh -i 1,5,9  
 sh snypr\_apps.sh -initiate 1,5,9

- o -a (or) alias: By alias
 

```
sh snypr_apps.sh -a <alias_name>
```

 Eg:
 

```
sh snypr_apps.sh -a enrichment
sh snypr_apps.sh -alias enrichment
```
- To start all applications
 

```
sh snypr_apps.sh -a all
```

### Example

```
[root@10-0-0-90 sparkjobs]# sh snypr_apps.sh -a all
17/05/18 12:11:12 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Event_Enrichment running successfully, CS_Event_Enrich-
ment.pid generated in /Securonix/ten-
ants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:11:34 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Event_Ingestion running successfully, CS_Event_Ingestion.pid
generated in /Securonix/tenants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:11:55 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Event_Indexer running successfully, CS_Event_Indexer.pid gen-
erated in /Securonix/tenants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:12:16 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Warning: CS_Behaviour_Analytics not running, appended to
/Securonix/tenants/ArcSight/Gateway/sparkJobs/logs/CS_pending_
app.txt
17/05/18 12:12:37 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Policy_Engine_IEE running successfully, CS_Policy_Engine_
IEE.pid generated in /Securonix/ten-
ants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:12:58 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Policy_Engine_AEE running successfully, CS_Policy_Engine_
AEE.pid generated in /Securonix/ten-
ants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:13:20 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_ThreatModel_RiskScoring_App running successfully, CS_
ThreatModel_RiskScoring_App.pid generated in /Securonix/ten-
ants/ArcSight/Gateway/sparkJobs/logs
17/05/18 12:13:41 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Warning: CS_Analytics_App not running, appended to
/Securonix/tenants/ArcSight/Gateway/sparkJobs/logs/CS_pending_
app.txt
17/05/18 12:14:02 INFO client.RMPProxy: Connecting to
```

```
ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Behaviour_Profile running successfully, CS_Behaviour_Profile.pid generated in /Securonix/tenants/ArcSight/Gateway/sparkJobs/logs
Checking all pending app status, check /Securonix/tenants/ArcSight/Sparkjobs/logs/CS_pending_app.log for any further errors
```

#### 4. Execute the following command to terminate all Spark jobs at once:

```
sh snypr_apps.sh -t all CS
```

- -t (or) -terminate : To terminate a particular application, you can pass by range, one job number / list of jobs, alias or kill all jobs for a particular tenant

```
sh snypr_apps.sh -t <option> <tenantId>
```

- Options:

- Option1 : to terminate a continuous range of jobs

```
sh snypr_apps.sh -t <start_job_num>-<end_job_num> <tenantId>
```

Eg:

```
sh snypr_apps.sh -t 1-4 CS
```

- Option2: to kill one job

```
sh snypr_apps.sh -t <job_num> <tenantId>
```

```
sh snypr_apps.sh -t 4 CS
```

- Option 3: to kill a list of discontinuous jobs

```
sh snypr_apps.sh -t <job_num1>,<job_num2>,<job_num3>.. <tenantId>
```

Eg:

```
sh snypr_apps.sh -t 1,5,8 CS
```

- Option4: to kill using an alias

```
sh snypr_apps.sh -t <job_alias_name> <tenantId>
```

Eg:

```
sh snypr_apps.sh -t enrichment CS
```

- Option5: to kill all applications for a tenant

```
sh snypr_apps.sh -t all <tenantId>
```

Eg:

```
sh snypr_apps.sh -t all CS
```

```
sh snypr_apps.sh -terminate all CS
```

#### Example

```
[root@10-0-0-90 sparkjobs]# sh snypr_apps.sh -t all CS
17/05/18 12:41:20 INFO client.RMProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
No such yarn application with name = CS_Analytics_App present
17/05/18 12:41:21 INFO client.RMProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
No such yarn application with name = CS_Behaviour_Profile present
```

```
17/05/18 12:41:22 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_Event_Enrichment
17/05/18 12:41:23 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:24 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0199
17/05/18 12:41:25 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0199
17/05/18 12:41:26 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_Event_Indexer
17/05/18 12:41:27 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:28 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0201
17/05/18 12:41:29 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0201
17/05/18 12:41:29 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_Event_Ingestion
17/05/18 12:41:30 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:32 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0200
17/05/18 12:41:32 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0200
17/05/18 12:41:33 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_Policy_Engine_AEE
17/05/18 12:41:34 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:35 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0204
17/05/18 12:41:36 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0204
17/05/18 12:41:37 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_Policy_Engine_IEE
17/05/18 12:41:38 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:39 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0203
17/05/18 12:41:39 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0203
17/05/18 12:41:40 INFO client.RMPProxy: Connecting to
```

```

ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
killing CS_ThreatModel_RiskScoring_App
17/05/18 12:41:41 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:41:42 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Killing application application_1494916983246_0205
17/05/18 12:41:43 INFO impl.YarnClientImpl: Killed application
application_1494916983246_0205

```

5. Execute the following command to check the status of all pending Spark jobs:

```
snypr_check-pending-apps.sh
```

- Checks if the applications started by the snypr\_apps.sh has moved to "running" state and if in running state, generates <job-name.pid> file containing the applicationId of the application
- Argument required:
- File containing the jobname & applicationId <tenant.id\_pending\_apps.txt>

```
Eg.CS_pending_apps.txt
```

```
Usage:
```

```
sh snypr_check-pending-apps.sh CS_pending_apps.txt
```

### Example

```

[root@10-0-0-90 sparkjobs]# cat logs/CS_pending_app.log
17/05/18 12:14:54 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
CS_Behaviour_Analytics failed, try re-initiating

```

6. Execute the following command to check the status of all Spark jobs: sh snypr\_apps.sh -s 1 CS

- -s (or) -status : To get status a particular application by specifying the jobnum or alias name

```
sh snypr_apps.sh -s <job_num> <tenantId>
```

```
Eg.
```

```
sh snypr_apps.sh -s 1 CS
```

```
sh snypr_apps.sh -s <alias_name> <tenantId>
```

```
Eg.
```

```
sh snypr_apps.sh -s enrichment <tenantId>
```

### Example

```

[root@10-0-0-90 sparkjobs]# sh snypr_apps.sh -s 1 CS
17/05/18 12:21:32 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:21:33 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:21:34 INFO client.RMPProxy: Connecting to ResourceMan-
ager at 10-0-0-90.securonix.com/10.0.0.90:8032
Application Report :

```

```

Application-Id : application_1494916983246_0199
Application-Name : CS_Event_Enrichment
Application-Type : SPARK
User : securonix
Queue : root.root
Start-Time : 1495127477673
Finish-Time : 0
Progress : 10%
State : RUNNING
Final-State : UNDEFINED
Tracking-URL : http://10.0.0.90:43713
RPC Port : 0
AM Host : 10.0.0.90
Aggregate Resource Allocation : 2500214 MB-seconds, 1220 vcore-seconds
Log Aggregation Status : NOT_START
Diagnostics :
[root@10-0-0-90 sparkjobs]# sh snypr_apps.sh -s 2 CS
17/05/18 12:21:54 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:21:56 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:21:57 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
Application Report :
Application-Id : application_1494916983246_0200
Application-Name : CS_Event_Ingestion
Application-Type : SPARK
User : securonix
Queue : root.root
Start-Time : 1495127498602
Finish-Time : 0
Progress : 10%
State : RUNNING
Final-State : UNDEFINED
Tracking-URL : http://10.0.0.94:34406
RPC Port : 0
AM Host : 10.0.0.94
Aggregate Resource Allocation : 2510521 MB-seconds, 1225 vcore-seconds
Log Aggregation Status : NOT_START
Diagnostics :
[root@10-0-0-90 sparkjobs]# sh snypr_apps.sh -s 3 CS
17/05/18 12:22:11 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:22:12 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
17/05/18 12:22:14 INFO client.RMPProxy: Connecting to ResourceManager at 10-0-0-90.securonix.com/10.0.0.90:8032
Application Report :
Application-Id : application_1494916983246_0201

```

```

Application-Name : CS_Event_Indexer
Application-Type : SPARK
User : securonix
Queue : root.root
Start-Time : 1495127520070
Finish-Time : 0
Progress : 10%
State : RUNNING
Final-State : UNDEFINED
Tracking-URL : http://10.0.0.91:43255
RPC Port : 0
AM Host : 10.0.0.91
Aggregate Resource Allocation : 2489124 MB-seconds, 1214 vcore-seconds
Log Aggregation Status : NOT_START
Diagnostics :

```

7. Execute the following script to get logs for a particular application:

```
sh snypr_apps.sh -l <application_name>
```

- -l (or) -logs : To get logs a particular application

```
sh snypr_apps.sh -l <application_name>
```

Eg.

```
sh snypr_apps.sh -l CS_Event_Enrichment
```

```
sh snypr_apps.sh -logs CS_Event_Enrichment
```

8. Execute the following script to run the Spark applications in either multitenant mode or single mode:
- ```
sh snypr_apps.sh <script intializing parameters> -m <mode>
```

- -m (or) -mode : To run the spark applications in either "multitenant" mode to pass the dynamic parameters or "single" mode

```
sh snypr_apps.sh <script intializing parameters> -m <mode>
```

Eg.

```
sh snypr_apps.sh -a all -m multitenant
```

```
sh snypr_apps.sh -i 1 -m single
```

```
sh snypr_apps.sh -r 1-9 -m multitenant
```



**Note:** If the mode is not explicitly mentioned in the parameters, the applications will be initiated with mode specified in the properties file.



**Note:** Files starting with "mt" are used to run the scripts in multi-tenant mode.

## Configuring Redis

The redis.conf file contains a number of directives that have a very simple format: keyword argument1 argument2 ... argumentN.

Configure the following settings related to IP Address, Password, and Port for Redis from the redis.conf file. The path to the redis.conf file is determined during the installation process.



1. Bind the IP address to the server where Redis is installed. Example: bind 10.0.0.60
2. Accept connections on the specified port (default is 6379). Example: port 6379
3. Specify the password for the Redis. Example: requirepass open24X7

# Hadoop Integration

This section covers the configuration of Hadoop services for ArcSight UBA.

## Configure Hadoop Initial Settings

Configure these settings before configuring Hadoop Services for ArcSight UBA.

### Memory Tuning for Services

1. Impala: Impala Daemon Memory Limit - mem\_limit - 12 GB
2. YARN: Container Memory - yarn.nodemanager.resource.memory-mb - 100 GB
3. HDFS: Java Heap Name Node: 8 GB
4. HDFS: Java Heap Secondary Name Node: 8 GB
5. HDFS: Java Heap Size of DataNode in Bytes: 1 GB
6. Zookeeper: Maximum Client Connections - maxClientCnxns - 2000
7. Zookeeper-Kafka (if a second zookeeper is configured): Maximum Client Connections - maxClientCnxns - 2000
8. KAFKA: Java Heap Size of Broker - broker\_max\_heap\_size - 1 GB
9. KAFKA: Maximum Message Size - message\_max\_bytes - 10 MiB
10. KAFKA: Replica Maximum Fetch Size - replica.fetch.max.bytes - 10 MiB
11. HBase: Java Heap Size Master in Bytes: 1 GB
12. HBase: Java Heap Size Thrift in Bytes: 1 GB
13. HBase: Java Heap Size Region Server in Bytes: 31 GB
14. Solr: Java Heap Size Solr Server in Bytes: 31 GB
15. Solr: Java Direct Memory Size of Solr Server in Bytes: 38 GB

### Cloudera Manager Service Configuration

1. Java Heap Size of Service Monitor in Bytes - increase - 1 GB
2. Java Heap Size of Host Monitor in Bytes - increase - 1 GB
3. Maximum Non-Java Memory of Host Monitor - firehose\_non\_java\_memory\_bytes - increase - 4 GB
4. Maximum Non-Java Memory of Service Monitor - firehose\_non\_java\_memory\_bytes - increase - 4 GB

### HBase Logging

Add the value below to the HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml:

```
name: hbase.ipc.warn.response.time
value: 500
```

## High Availability (HA) HDFS Name Node (NN)

1. Add an HA YARN Resource Manager role with the actions wizard for YARN.



**Note:** Ensure the Secondary Name Node (SNN) is on a different server. If the installer created the SNN and NN on the same server, move the SNN to a different server.

2. Ensure Zookeeper is started.
3. Enable Name Node HA on the HDFS wizard to set up HA:
  1. Select Enable High Availability.
  2. Provide Nameservice: nameservice1.
  3. Add Journal Node directories:  
/dfs/jn

## HA YARN Resource Manager Configuration

1. Add resources manager to another node:
  1. Select **Yarn**.
  2. Select **Instance**.
  3. Select **Actions -> Enable High Availability**.
  4. Select the HA Resource Manager host to use (<server>.<domain>.<com>).
  5. Click **Continue**.
2. Increase the HDFS Handler counts:

**NameNode Handler Count:** dfs.namenode.handler.count 60

**NameNode Service Handler Count:** dfs.namenode.service.handler.count 60

## Configure Hadoop Services for ArcSight UBA

The configuration details specified in this section for Hadoop services impact the changes you make in the ArcSight UBA user interface [Configure ArcSight UBA Hadoop Settings](#) section.

### Kafka

#### Create Kafka Topics



**Note:** Kafka uses ZooKeeper, so you must first start a ZooKeeper server if you don't already have one. After you start the Zookeeper server, you can start Kafka.

The following command format is used to create a Kafka topic with the required replication factor and number of partitions. Specify the IP address of the nodes you want to create Kafka topics:

```
kafka-topics --create --zookeeper <IP address1>,<IP address2>,<IP address3> --replication-factor 1 --partitions 5 --topic <custom topic name>
```

### Example :

```
kafka-topics --create --zookeeper 10.0.0.63:2181,10.0.0.64:2181,10.0.0.65:2181 --replication-factor 1 --partitions 5 --topic SecuronixResource-Enriched-Preview
```

You can create the following kafka topics using the command shown:

- **Preview topic:** Used while configuring the Ingestor data source.
- **Access topic:** Used for ingesting access data.
- **Users topic:** Used for ingesting user data.
- **Enriched Topic:** Once the raw feed is ingested, it is converted into enriched format using spark streaming job and is stored in enriched format.
- **Raw Topic:** Used to store the raw incoming feed into Kafka.
- **Configuration Messages Topic (Control):** Sent when a particular policy or anything in the application is created or updated that needs to be sent to Spark Streaming jobs using the Control topic in Kafka.
- **Log Message Topic (Ops):** The logging for the Spark application are published to this topic in CEF Format.
- **Violations Topic:** The flagged Violations are published to Violation topic for Risk Score calculation. Example: ArcSight6-Violations.
- **Tier2 Topic:** Summary and the Behavior analytics are performed on the events that get published to the Tier2 topic.
- **AEE Tier2 Topic:** Provides a topic where events from IEE is channeled to AEE Spark processor.
- **Job Tracker Topic (Count):** Used to store the statistics like how many events are imported, re-imported etc.
- **IndexerCountTopic:** Used to store the statistics of the events which are indexed. Stores how many events are indexed in SolrCloud and how many Violations are indexed in SolrCloud.

To check that the topics were created successfully, use the following command format:

```
kafka-topics -list --zookeeper <IP address of the zookeeper server> | grep <kafka topic name>
```

**Example :**

```
kafka-topics -list --zookeeper 10.0.0.10:2181 | grep ArcSight611082017
```

```
[securonix@sc-10-0-0-10 sparkJobs]$ kafka-topics -list --zookeeper 10.0.0.10:2181
Securonix-Control
SecuronixResource-Enriched
SecuronixResource-Enriched1
SecuronixResource-Raw
SecuronixResource-Raw1
SecuronixResource-Unprocessed
SecuronixResource-Unprocessed1
ash
ash-Control
ashControl
ashResource-Enriched
ashResource-Raw
ashResource-Unprocessed
indexerCountTopic
omkar-Control
omkar-Resource-Enriched
omkar-Resource-Raw
omkar-Resource-Unprocessed
[securonix@sc-10-0-0-10 sparkJobs]$
```

**Update Kafka Topic Time**

Data is stored for a specific number of days in a kafka topic. To update the kafka topic time, use this command:

```
kafka-topics --zookeeper <IP address1>, <IP address2> --alter --topic
<kafka topic> --config retention.ms=<time in milliseconds>
```

**Example :**

```
kafka-topics --zookeeper 10.0.0.63:2181,10.0.0.64:2181,10.0.0.65:2181 --alter --topic
SecuronixResource-Access --config retention.ms=7200000
```

**HBase**

Create HBase Namespace

1. Log in to HBase shell.

```
hbase shell
```

2. Create the namespace using the following commands:

```
hbase(main):002:0> create_namespace 'securonix'
```

3. View the list of namespaces available:

```
hbase(main):002:0> list_namespace 'securonix'
```

#### 4. Exit HBase shell:

```
hbase(main):002:0> exit
```

## Hive

Create the Hive Schema (on Prod Master and UAT Master):

#### 1. Log in to Hive:

```
hive shell
```

#### 2. Create the database schema:

```
create database securonix;
```

#### 3. Confirm that the database schema has been created successfully:

```
show databases;
```

#### 4. Quit Hive shell:

```
quit;
```

## Impala

Verify database availability with Impala shell:

```
impala-shell
Starting Impala Shell without Kerberos authentication
Connected to <FQDN>:21000
Server version: impalad version 2.3.0-cdh5.5.4 RELEASE (build
e65ded24350974ae6b4e475557b358b718fad29e)
*****_
*****
Welcome to the Impala shell. Copyright (c) 2015 Cloudera, Inc. All
rights reserved.
(Impala Shell v2.3.0-cdh5.5.4 (e65ded2) built on Mon Apr 25
10:55:11 PDT 2016)
After running a query, type SUMMARY to see a summary of where time
was spent.
*****_
*****
[<FQDN>:21000] > create database securonix;
show databases;
quit;
Query: invalidate metadata
Fetched 0 row(s) in 2.80s
[sc-10-0-0-10.securonix.com:21000] > show databases;
Query: show databases
+-----+
| name |
+-----+
```

```
| _impala_builtins |
| default |
| securonix |
+-----+
Fetched 3 row(s) in 0.00s
exit
```

## HDFS

1. Create the HDFS Service Account Folder for ArcSight UBA:

```
su hdfs
hdfs dfs -mkdir /user/securonix
hdfs dfs -chown securonix:securonix /user/securonix
hdfs dfs -ls /user
exit
```

2. Log in as Securonix user and do the following:

```
su securonix
hdfs dfs -mkdir /user/securonix/ArcSightUBA
hdfs dfs -mkdir /user/securonix/ArcSightUBA/resources-enriched
hdfs dfs -mkdir /user/securonix/ArcSightUBA/resources-enriched-
incoming
hdfs dfs -mkdir /user/securonix/ArcSightUBA/resources-summary
```

3. Optional: Check that the directory was created successfully:

```
hdfs dfs -ls /user/securonix/ArcSightUBA
```

# Uninstall ArcSight UBA

To manually uninstall the ArcSight UBA application, complete the following steps:

1. Stop Tomcat, Redis, and Syslog-ng services.
2. Run the uninstaller by running the following command: `/<Installation_folder>/_ArcSightUBA_installation/Uninstall`
3. Run the following command:  

```
sudo rm -rf /etc/init.d/syslog-ng /etc/init.d/redis  
rm -rf <Installation_folder>
```



# Appendix A: Spark Job Properties

This section includes all default property values for each Spark Job.

## Job 1: Event Enrichment

| Property                            | Default Value    |
|-------------------------------------|------------------|
| Name                                | Event_Enrichment |
| Driver-Memory                       | 1g               |
| Number of Executors                 | 5                |
| Executor Memory                     | 1g               |
| Driver Cores                        | 1                |
| Executor Cores                      | 1                |
| Consumer Group Name                 | RG-enrichment    |
| Max rate per Partition (mrpp)       | 5                |
| Job Duration                        | 1s               |
| Driver Memory Overhead              | 1024             |
| Executor Memory Overhead            | 1024             |
| Minimum Executors                   | 1                |
| Maximum Executors                   | 20               |
| Initial Executors                   | 1                |
| Sustained Scheduler Backlog Timeout | 60s              |
| Executor Idle Timeout               | 300s             |

## Job 2: Event Ingestion

| Property                            | Default Value   |
|-------------------------------------|-----------------|
| Job Name                            | Event_Ingestion |
| Driver-Memory                       | 1g              |
| Number of Executors                 | 5               |
| Executor Memory                     | 1g              |
| Driver Cores                        | 1               |
| Executor Cores                      | 1               |
| Consumer Group Name                 | RG-ingestion    |
| Max rate per Partition (mrpp)       | 5               |
| Job Duration                        | 1s              |
| Driver Memory Overhead              | 1024            |
| Executor Memory Overhead            | 1024            |
| Minimum Executors                   | 1               |
| Maximum Executors                   | 20              |
| Initial Executors                   | 1               |
| Sustained Scheduler Backlog Timeout | 60s             |
| Executor Idle Timeout               | 300s            |

## Job 3: Event Indexer

| Property                            | Default Value |
|-------------------------------------|---------------|
| Job Name                            | Event_Indexer |
| Driver-Memory                       | 1g            |
| Number of Executors                 | 5             |
| Executor Memory                     | 1g            |
| Driver Cores                        | 1             |
| Executor Cores                      | 1             |
| Consumer Group Name                 | RG-indexer    |
| Max rate per Partition (mrpp)       | 5             |
| Job Duration                        | 1s            |
| Driver Memory Overhead              | 1024          |
| Executor Memory Overhead            | 1024          |
| Minimum Executors                   | 1             |
| Maximum Executors                   | 20            |
| Initial Executors                   | 1             |
| Sustained Scheduler Backlog Timeout | 60s           |
| Executor Idle Timeout               | 300s          |

## Job 4: Behavior Analytics

| Property                            | Default Value      |
|-------------------------------------|--------------------|
| Job Name                            | Behavior_Analytics |
| Driver-Memory                       | 1g                 |
| Number of Executors                 | 5                  |
| Executor Memory                     | 1g                 |
| Driver Cores                        | 1                  |
| Executor Cores                      | 1                  |
| Consumer Group Name                 | RG-behaviorsummary |
| Max rate per Partition (mrpp)       | 100                |
| Job Duration                        | 30                 |
| Driver Memory Overhead              | 1024               |
| Executor Memory Overhead            | 1024               |
| Minimum Executors                   | 1                  |
| Maximum Executors                   | 20                 |
| Initial Executors                   | 1                  |
| Sustained Scheduler Backlog Timeout | 60s                |
| Executor Idle Timeout               | 300s               |

## Job 5: Policy Engine IEE (Individual Event Evaluator)

| Property                            | Default Value     |
|-------------------------------------|-------------------|
| Job Name                            | Policy_Engine_IEE |
| Driver-Memory                       | 1g                |
| Number of Executors                 | 5                 |
| Executor Memory                     | 1g                |
| Driver Cores                        | 1                 |
| Executor Cores                      | 1                 |
| Consumer Group Name                 | RG-policy-iee     |
| Max rate per Partition (mrpp)       | 5                 |
| Job Duration                        | 2s                |
| Driver Memory Overhead              | 1024              |
| Executor Memory Overhead            | 1024              |
| Minimum Executors                   | 1                 |
| Maximum Executors                   | 20                |
| Initial Executors                   | 1                 |
| Sustained Scheduler Backlog Timeout | 60s               |
| Executor Idle Timeout               | 300s              |

## Job 6: Policy Engine AEE (Aggregated Events Evaluator)

| Property                            | Default Value     |
|-------------------------------------|-------------------|
| Job Name                            | Policy_Engine_AEE |
| Driver-Memory                       | 1g                |
| Number of Executors                 | 5                 |
| Executor Memory                     | 1g                |
| Driver Cores                        | 1                 |
| Executor Cores                      | 1                 |
| Job Duration                        | 60s               |
| Driver Memory Overhead              | 1024              |
| Executor Memory Overhead            | 1024              |
| Minimum Executors                   | 1                 |
| Maximum Executors                   | 20                |
| Initial Executors                   | 1                 |
| Sustained Scheduler Backlog Timeout | 60s               |
| Executor Idle Timeout               | 300s              |

## Job 7: Risk Generation

| Property                            | Default Value               |
|-------------------------------------|-----------------------------|
| Job Name                            | ThreatModel_RiskScoring_App |
| Driver-Memory                       | 1g                          |
| Number of Executors                 | 5                           |
| Executor Memory                     | 1g                          |
| Driver Cores                        | 1                           |
| Executor Cores                      | 1                           |
| Consumer Group Name                 | RG-riskgeneration           |
| Max rate per Partition (mrpp)       | 20                          |
| Job Duration                        | 10s                         |
| Driver Memory Overhead              | 1024                        |
| Executor Memory Overhead            | 1024                        |
| Minimum Executors                   | 1                           |
| Maximum Executors                   | 20                          |
| Initial Executors                   | 1                           |
| Sustained Scheduler Backlog Timeout | 60s                         |
| Executor Idle Timeout               | 300s                        |

## Job 8: Analytics

| Property                            | Default Value |
|-------------------------------------|---------------|
| Job Name                            | Analytics_App |
| Driver-Memory                       | 1g            |
| Number of Executors                 | 5             |
| Executor Memory                     | 1g            |
| Driver Cores                        | 1             |
| Executor Cores                      | 1             |
| Consumer Group Name                 | RG-analytics  |
| Max rate per Partition (mrpp)       | 20            |
| Job Duration                        | 10s           |
| Driver Memory Overhead              | 1024          |
| Executor Memory Overhead            | 1024          |
| Minimum Executors                   | 1             |
| Maximum Executors                   | 20            |
| Initial Executors                   | 1             |
| Sustained Scheduler Backlog Timeout | 60s           |
| Executor Idle Timeout               | 300s          |



## Job 9: Behavior Profile

| Property                            | Default Value     |
|-------------------------------------|-------------------|
| Job Name                            | Behaviour_Profile |
| Driver-Memory                       | 1g                |
| Number of Executors                 | 5                 |
| Executor Memory                     | 1g                |
| Driver Cores                        | 1                 |
| Executor Cores                      | 1                 |
| Minimum Executors                   | 1                 |
| Maximum Executors                   | 20                |
| Initial Executors                   | 1                 |
| Sustained Scheduler Backlog Timeout | 60s               |
| Executor Idle Timeout               | 300s              |

## Job 10: Event Archiver

| Property                            | Default Value  |
|-------------------------------------|----------------|
| Job Name                            | Event Archiver |
| Driver-Memory                       | 1g             |
| Number of Executors                 | 5              |
| Executor Memory                     | 1g             |
| Driver Cores                        | 1              |
| Executor Cores                      | 1              |
| Consumer Group Name                 | RG-archiver    |
| Max rate per Partition (mrpp)       | 10             |
| Job Duration                        | 5s             |
| Driver Memory Overhead              | 1024           |
| Executor Memory Overhead            | 1024           |
| Minimum Executors                   | 1              |
| Maximum Executors                   | 20             |
| Initial Executors                   | 1              |
| Sustained Scheduler Backlog Timeout | 60s            |
| Executor Idle Timeout               | 300s           |

## Job 11: HDFS Indexer

| Property                  | Default Value    |
|---------------------------|------------------|
| Job Name                  | HDFS Indexer Job |
| Driver-Memory             | 3g               |
| Number of Executors       | 10               |
| Executor Memory           | 5g               |
| Driver Cores              | 2                |
| Executor Cores            | 4                |
| Driver Memory Overhead    | 1024             |
| Executor Memory Overhead  | 1024             |
| Streaming Concurrent Jobs | 1                |
| Max App Attempts          | 20               |

# 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 system, click the link above and an email window opens with the following information in the subject line:

**Feedback on Documentation (Micro Focus ArcSight User Behavior Analytics 6.10)**

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 [arst-techpubs@hpe.com](mailto:arst-techpubs@hpe.com).

We appreciate your feedback!