



# HP User Behavior Analytics

HP UBA Version 1.1

Powered by  **SECURONIX**

## Installation Guide

August 31, 2015

# Table of Contents

Chapter 1: About This Guide.....	3
About HP User Behavior Analytics .....	3
Who Should Read this Guide.....	3
Chapter 2: Installation and Configuration.....	4
Installation Components.....	4
Checklist.....	4
Minimum Hardware Specifications.....	5
Supported OS .....	5
Supported Browsers.....	5
Required Communication Ports .....	5
Prerequisite – MySQL Installation.....	6
Prerequisite – XWindows .....	7
Prerequisite – Increase User Process Limit.....	7
Installation Steps.....	9
Update MySQL Server Settings .....	21
Chapter 3: Post Installation Activities.....	22
Start Using the Application.....	22
Chapter 4: Uninstall HP UBA .....	23
Chapter 5: Deployment Architecture.....	28
Deployment Considerations .....	28
Capacity planning .....	28
Deploy HP UBA in a Master/Child Architecture .....	30
Prerequisites.....	30
Add a Child Node.....	30
Appendix A: Tuning MySQL Configurations .....	32

## Chapter 1: About This Guide

Use the Install Guide to learn how to install HP User Behavior Analytics. In this manual, you can find:

- 1 What is HP User Behavior Analytics
- 2 Installation components
- 3 Installation procedure
- 4 Post-installation activities

### About HP User Behavior Analytics

HP User Behavior Analytics is an advanced security analytics solution that uses correlation, algorithms and visualizations to detect new threats, targeted attacks and privileged account misuse.

### Who Should Read this Guide

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

System administrators, responsible for ongoing operations and management should refer to the Administrators Guide. Users of the HP UBA application (security operations, information security professionals, security analysts, risk and compliance officers and IT specialists who need to use the functionalities within the product) should refer to the HP UBA User Guide. If you are responsible for integrating data sources refer to the HP UBA Administrator Guide/Importing Data.

## Chapter 2: Installation and Configuration

### Installation Components

The simplest deployment is the one you get by default when you install HP UBA: database and application running on the same server. Data comes in from the sources you've configured, and you log into the HP UBA web interface on this same server to monitor and analyze data.

Depending on your needs, HP UBA can be deployed in a Master – Child configuration or on a single node. This guide describes installing HP UBA on a single node.

Following components are involved in the deployment:

- Configuration folder – 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. HP UBA supports the MySQL database.
- Universal forwarder node\* – universal forwarders provide the capability to import and analyze activities and security events. They typically run on a separate computer than the computer running the application. They utilize the same database as the HP UBA application for storage of event data.
- Child node\* – Child nodes provide the capability to import and analyze activities and security events. They typically run on a separate computer from the computer running the application. They utilize a different database than the application.
- 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 HP UBA application or Real Time Analyzers.
- TPI aggregator – A text indexing engine used to index and store data aggregated from threat intelligence sources. Provides quick search and retrieval.

*\*These components are part of a master child deployment. They are not required for a single node deployment*

### Checklist

Before you get started with the deployment, make sure that you have the following details ready:

- HP User Behavior Analytics software
- License files and license key
- Deployment environment(s) – capacity planning, hardware, software, browser and port requirements as discussed above
- Database server (HP UBA supports the MySQL 5.6 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 HP User Behavior Analytics platform.

## Minimum Hardware Specifications

HP UBA is a high performance application and the system should meet the recommended hardware specification for an optimal experience

	Minimum (test environment)	Mid-Range (POC or small implementations)	Optimum	High Performance
<b>RAM</b>	16 GB	32 GB	64 GB	128GB
<b>Processors</b>	8 cores	16 cores	32 cores	64 cores
<b>Hard Disk</b>	500 GB	1TB	2TB	4TB
<b>Architecture</b>	64-bit architecture recommended.			

## Supported OS

The application supports Red Hat Enterprise Linux (RHEL) v6.5 and CentOS Linux v6.5.

## Supported Browsers

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

- Firefox 10.x and latest
- Internet Explorer 9, 10, 11
- Safari (latest)
- Chrome (latest)

**Note:** If you are using Internet Explorer 10, please turn off compatibility mode

## Required Communication Ports

- 1 Port for MySQL – Default port for MySQL is 3306
- 2 Tomcat Application Server Port – Default port for HTTP is 8080 and HTTPS is 8443
- 3 Optional Ports:
  - SSH port (Optional) – Port 22
    - ♦ UDP/TCP 53: DNS host name lookup – DNS is used for name lookup and event enrichment.
    - ♦ DHCP/port 67: DHCP/bootstrap protocol server is not needed when static IP addressing is used
    - ♦ UDP 514 used for syslog server set up.
    - ♦ ICMP type 8 only for server monitoring.
    - ♦ Get identity data from systems: connectivity varies by identity store, for example: LDAP/389 LDAPS/636 to Active Directory

## Prerequisite – MySQL Installation

Before running the HP UBA product installer (HPUBA11\_8\_17\_1.bin), please make sure MySQL is pre-installed and configured on the Linux server (RHEL and CentOS 6.5). If some of these items are not installed or misconfigured, the installer will let you know during installation.

### Install MySQL 5.6

- root> wget http://dev.mysql.com/get/mysql-community-release-el6-5.noarch.rpm
- root> rpm -ivh 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

### Start MySQL

- root> service mysqld start

### Enable MySQL Connections from Specific IP Addresses or Hostnames

**Note:** replace 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;".

### Change MySQL Data Directory

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

- 1 Shutdown MySQL
  - >> service mysqld stop
- 2 Create a backup of my.cnf
  - >> cp /etc/my.cnf /home/HPUBAUSER/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 Set enforce mode to permissive
  - 1 Run: \$ setenforce 0
  - 2 OR edit /etc/selinux/config and set SELINUX=permissive
- 8 Start MySQL

```
>> service mysqld start
```

## Recommended Best Practice

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.

## Prerequisite – XWindows

If XWindows has not been configured on the system, follow these steps prior to using the HPUBA installer:

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

If the UNIX machine is a VM and is accessed remotely, download a GUI client like ‘mobaxterm’ or ‘nomachine’ to access the remote server.

## Prerequisite – Increase 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
```

## Installation Guide

```
# 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 HPUBA user and run following command to verify output:

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



# Installation Steps

The HP 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 HPUBAUser` (for example – user whatever user name you want).
- 2 `passwd HPUBAUser` (give your non-root user a password).

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

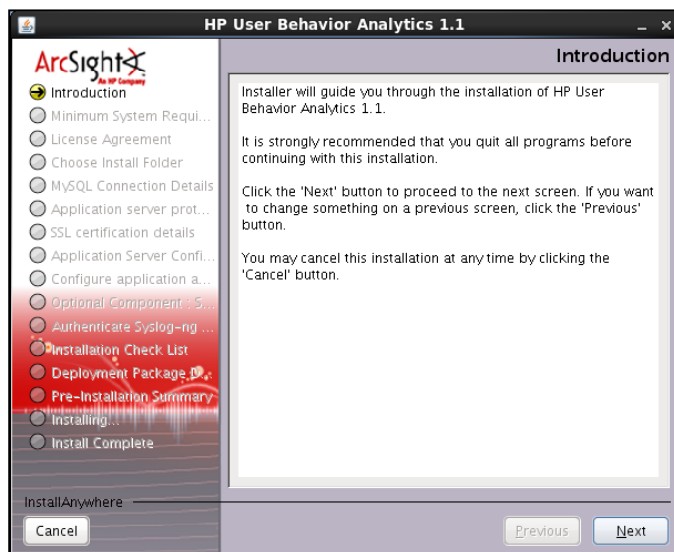
- 3 Login as root and go to `/etc/`  
`vi sudoers`
- 4 Scroll down through the sudoers file to the section below, add the user information for the non-root user that will start the installer (HPUBAUser from our example above) and save the file. Provide the non-root user password as sudo password on the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

```
## Allow root to run any commands anywhere
root ALL=(ALL) ALL
HP UBAUser ALL=(ALL) ALL
```

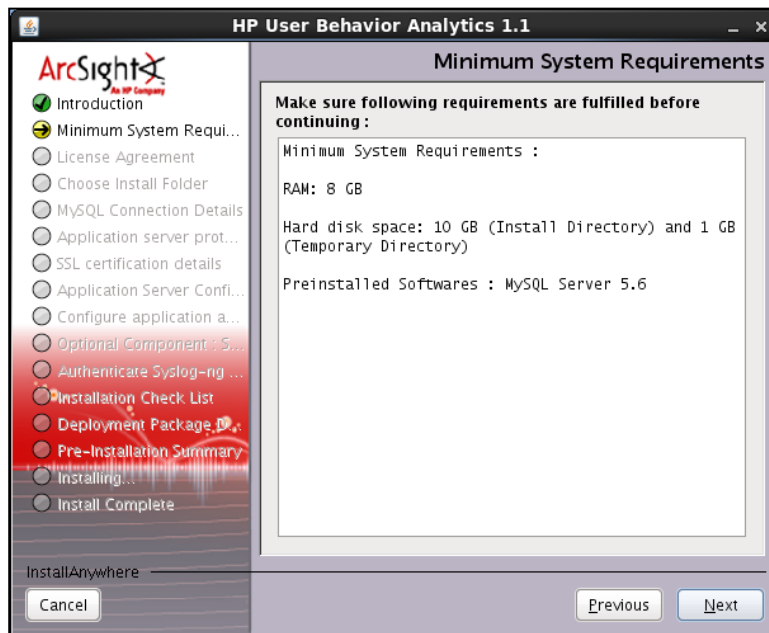
```
## Allows members of the 'sys' group to run networking, software,
## service management apps and more.
# %sys ALL = NETWORKING, SOFTWARE, SERVICES, STORAGE, DELEGATING, PROCESSES, LOCATE,
DRIVERS
```

```
## Allows people in group wheel to run all commands
```

- 5 Please make sure the installer has the right permissions to be executed:  
`chmod 755 HPUBA11_8_17_1.bin`
- 6 In order to begin the installation of HP UBA application, open a terminal and execute the following command:  
`./ HPUBA11_8_17_1.bin`



- 7 Review system requirements, and continue. If your system does not meet minimum requirements, you may accept the risk, but the installation may not be supported.

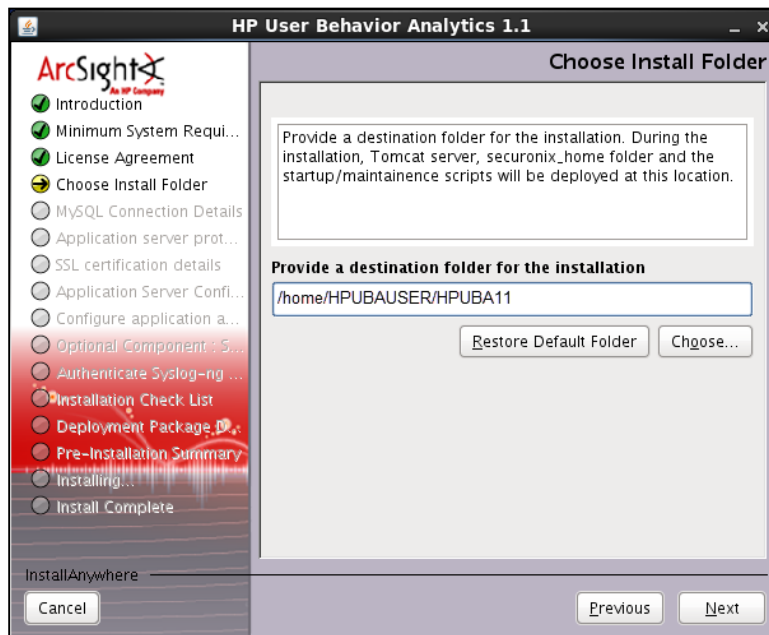


- 8 On the license screen, read and scroll to the bottom of the license, to enable and check the “I accept...” button.

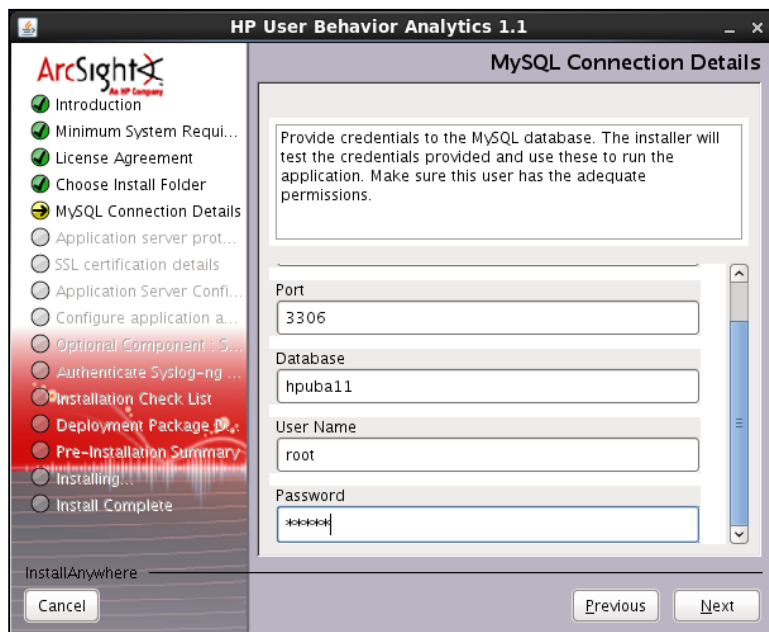


## Installation Guide

- 9 Select the location to install the files to and click **Next**. The default will install files to a HPUBA11 subdirectory in the home directory of the currently logged in user.

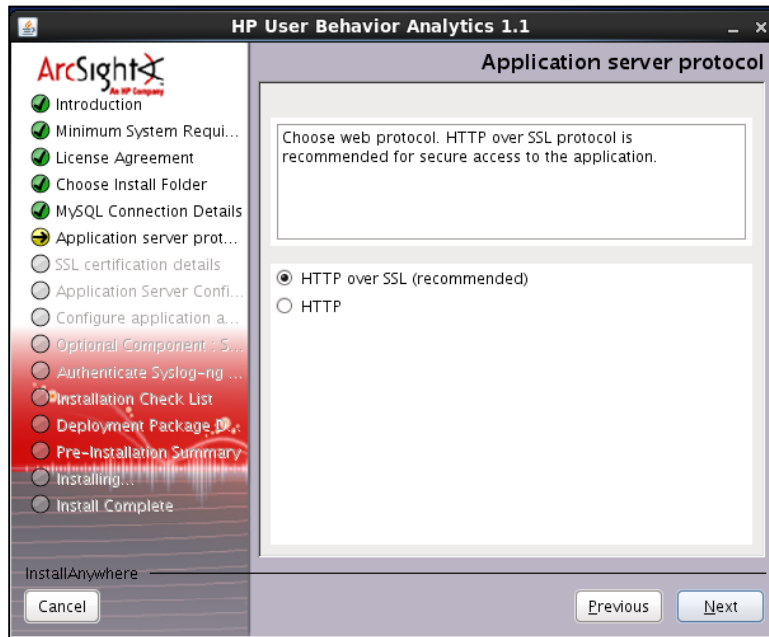


- 10 Modify the MySQL connection details to include the system DBA user and password and port modifications as necessary, and click **Next**. The database will be created on the MySQL instance specified.



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

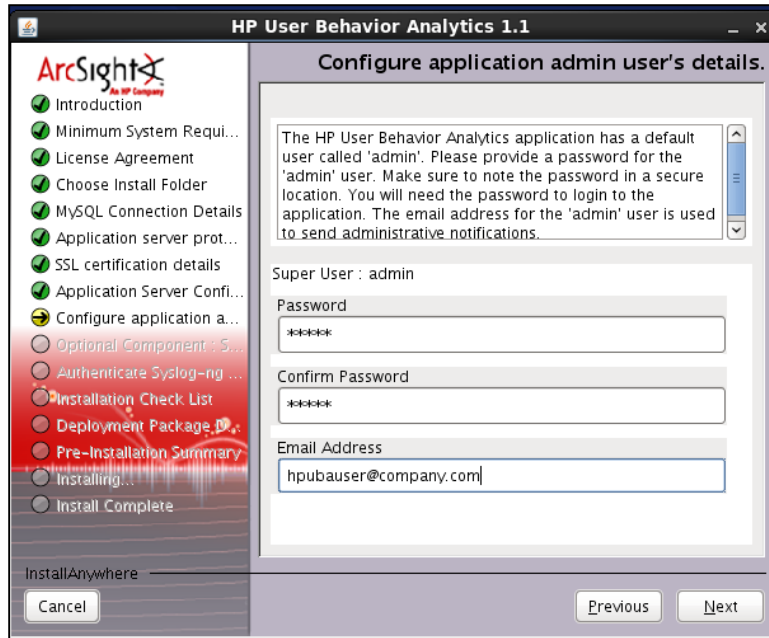
- 11 Select the connection protocol to be used. HTTPS (port 8443) is recommended for secure communications.



- 12 Enter the SSL certification details as shown below.

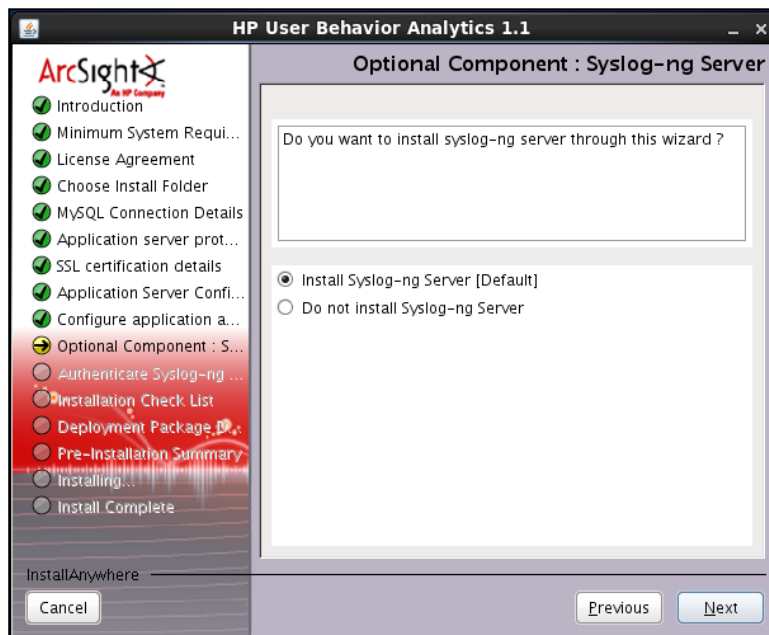


- 13 Enter the credentials for the HP UBA administrative user (this account will be used for login using a browser). The password must be a minimum of 6 characters. You will get an error screen if you do not meet this requirement.

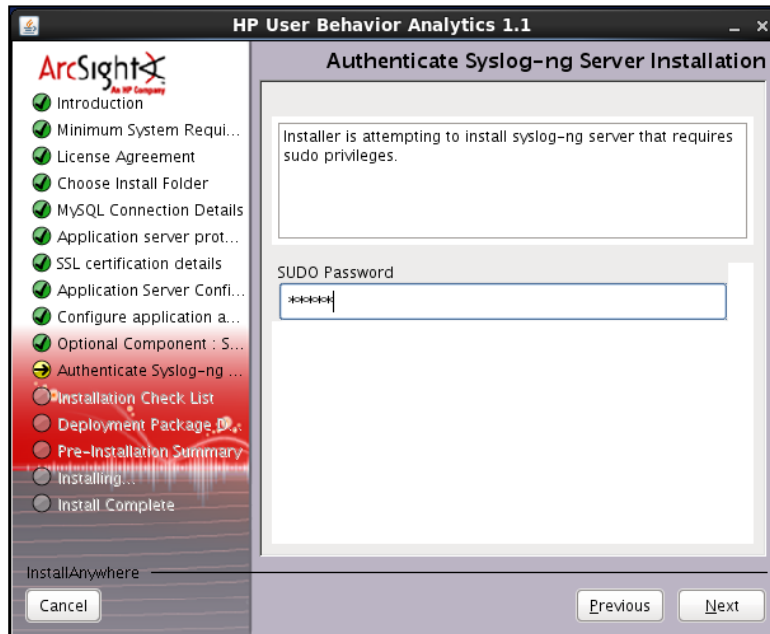


## 14 Configure Syslog-NG


HP UBA requires Syslog-ng server to be running. If you do not already have it installed, the installer will install it by default as part of the process.



15 Enter the sudo password (required to configure a service).



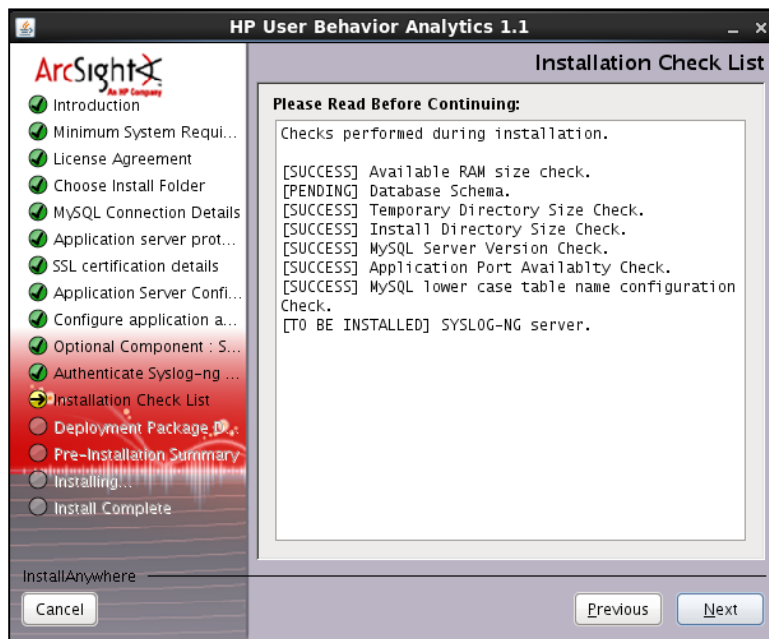
16 The installer will check the sudo password before continuing:



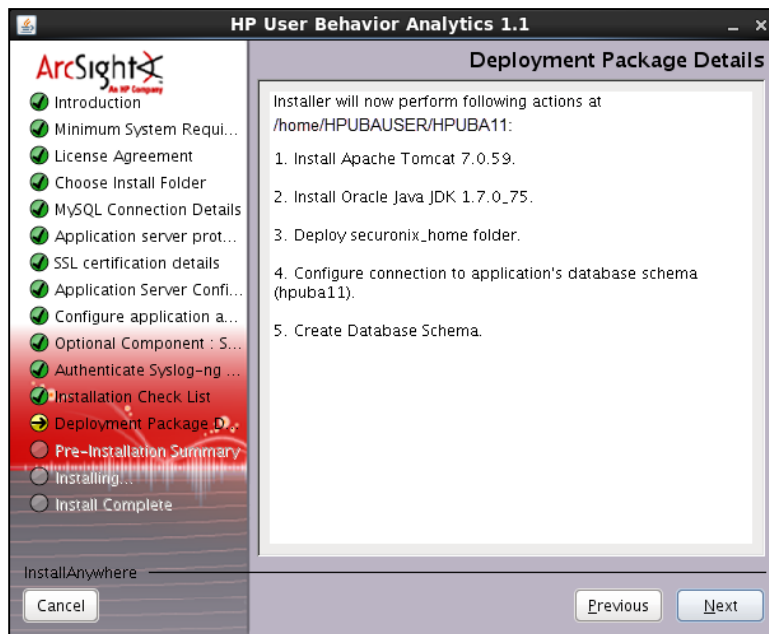
17 If syslog has already been installed, the system will show the following screen, and this step can be skipped.



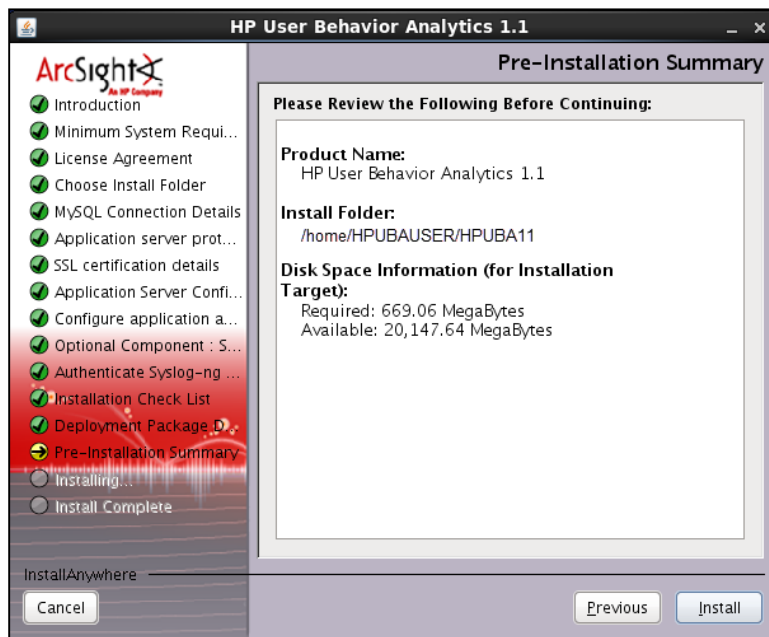
18 Installation checklist. Select **Next** to continue.



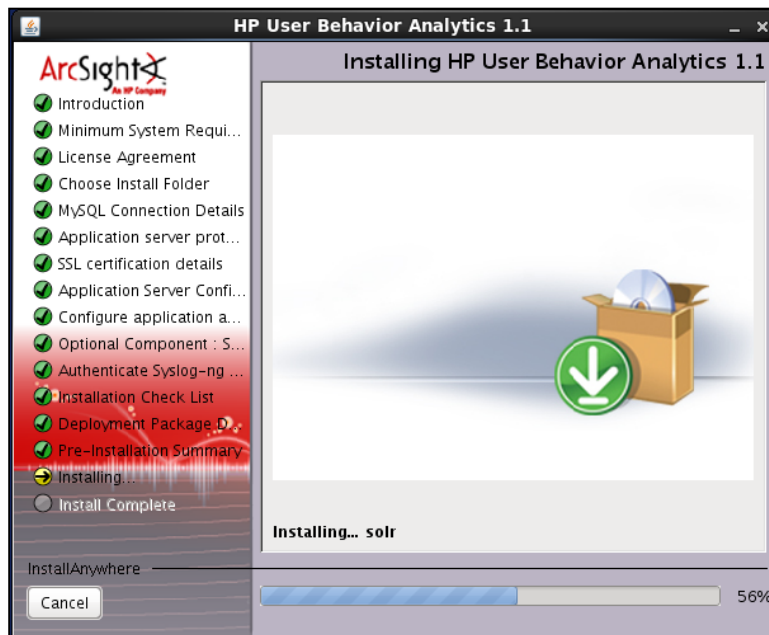
19 Deployment package details. click **Next** to continue.



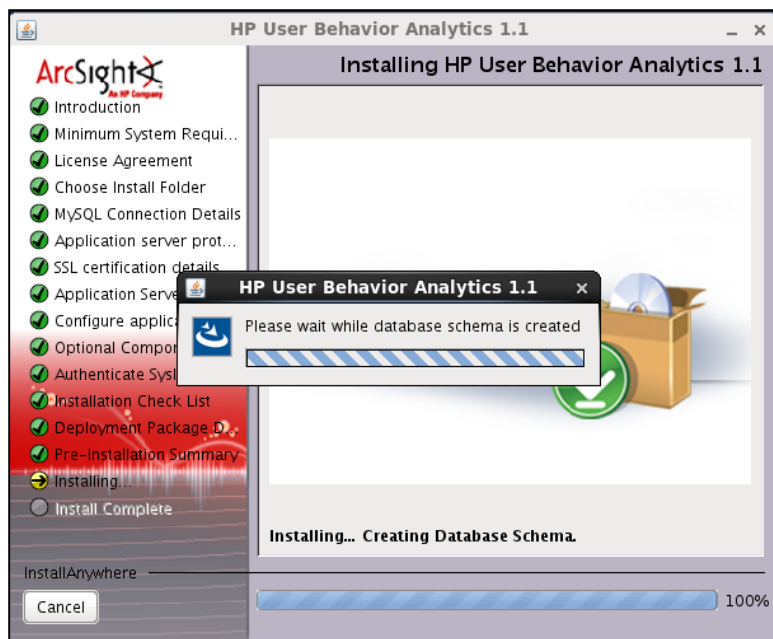
20 Click **Install** on the Pre-Installation Summary screen.



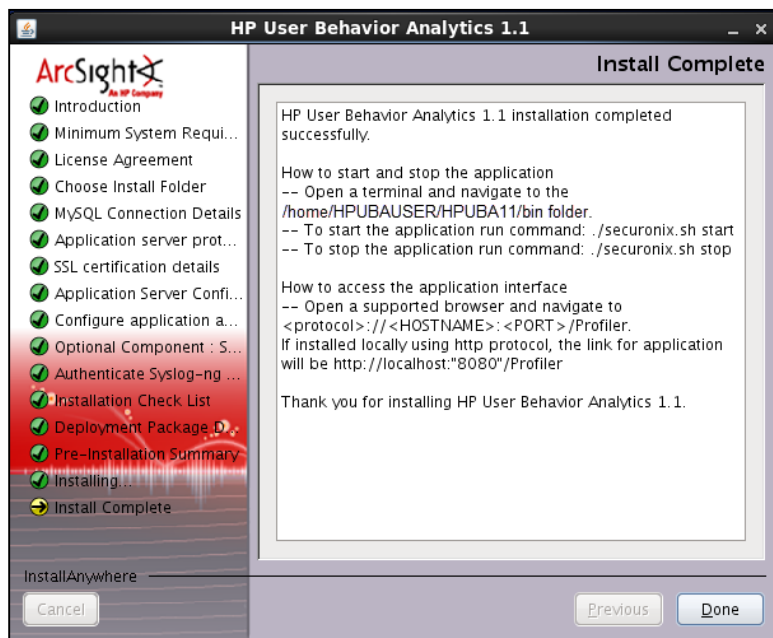
21 Additional components will be installed. After the status bar shows 100% you should see an installation complete screen. The HP UBA database and schema will be created. The time required will vary based on hardware performance.







22 Successful installation – Summary Screen. Click **Done** to continue.



23 Start the HP UBA application from a command line, by running:

```
./securonix.sh start
```

Start, stop or restart the application with the commands shown below:

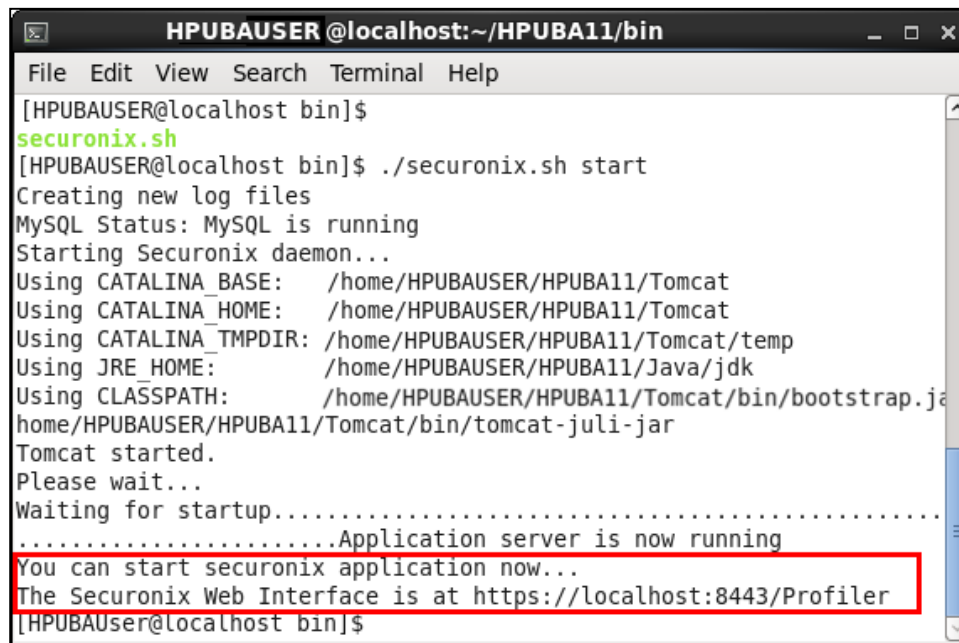
```
[HPUBAUSER@localhost bin]$ ./securonix.sh
Usage : To start Securonix Application - ./securonix.sh start
        To stop Securonix Application - ./securonix.sh stop
        To restart Securonix Application - ./securonix.sh restart
[HPUBAUSER@localhost bin]$
```

## Installation Guide

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

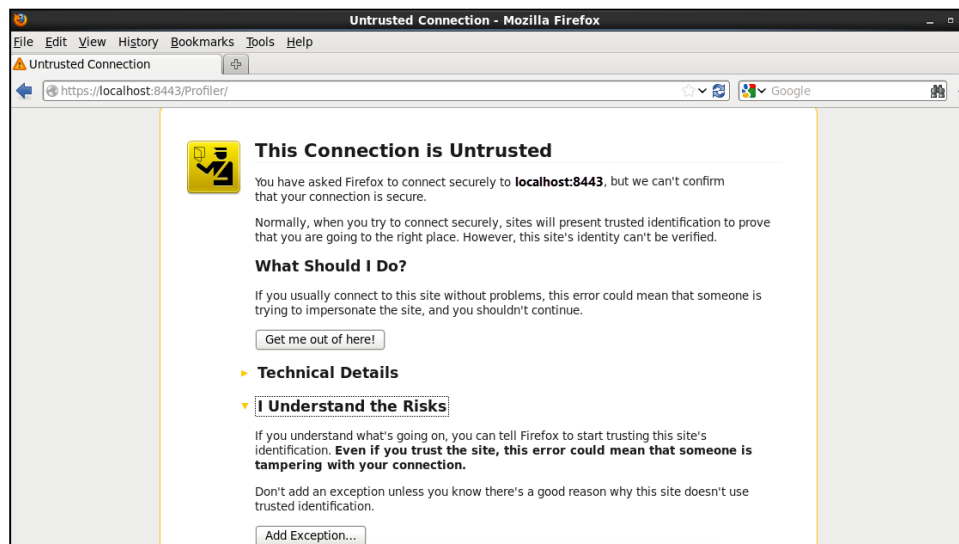
*You can start securonix application now...*

*The Securonix Web Interface is at ...*

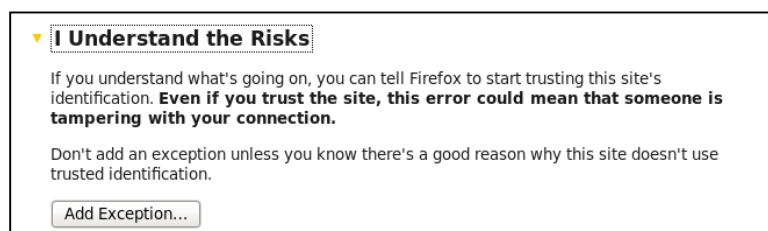


```
HPUBAUSER @localhost:~/HPUBA11/bin
File Edit View Search Terminal Help
[HPUBAUSER@localhost bin]$
securonix.sh
[HPUBAUSER@localhost bin]$ ./securonix.sh start
Creating new log files
MySQL Status: MySQL is running
Starting Securonix daemon...
Using CATALINA_BASE: /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_HOME: /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_TMPDIR: /home/HPUBAUSER/HPUBA11/Tomcat/temp
Using JRE_HOME: /home/HPUBAUSER/HPUBA11/Java/jdk
Using CLASSPATH: /home/HPUBAUSER/HPUBA11/Tomcat/bin/bootstrap.jar
/home/HPUBAUSER/HPUBA11/Tomcat/bin/tomcat-juli.jar
Tomcat started.
Please wait...
Waiting for startup.....
.....Application server is now running
You can start securonix application now...
The Securonix Web Interface is at https://localhost:8443/Profiler
[HPUBAUSER@localhost bin]$
```

24 Validate the installation by connecting to the URL shown at the end of the startup script.



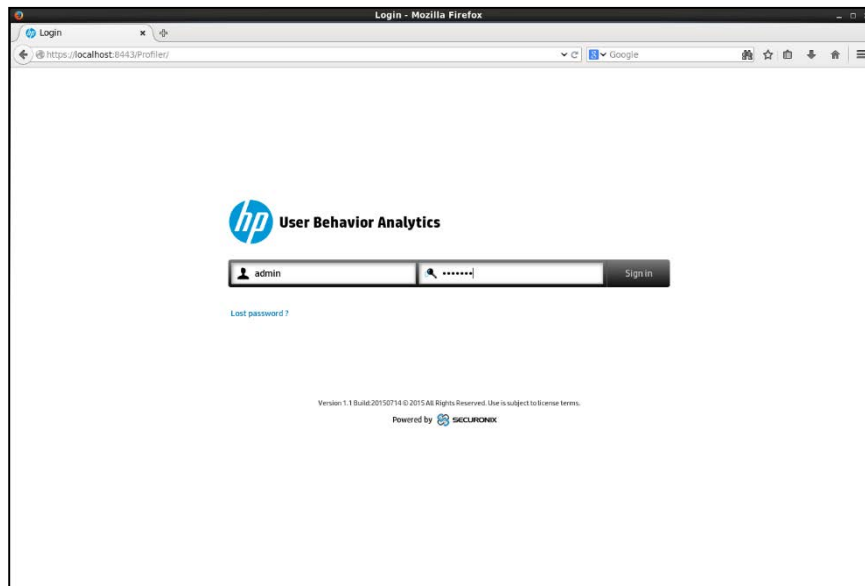
25 Accept and add an exception for the self-signed certificate of the host you have just configured.




## Installation Guide



26 Login to the system using the admin user and password configured earlier during the installation.



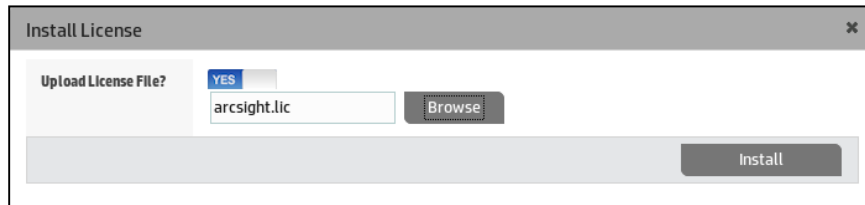
27 On the first installation you should see the following screen.



You will need to install the license provided to continue. Select the **Install License** option.

## Installation Guide

28 Upload the \*.lic file provided with your licensing agreement



The screenshot shows a web-based dialog box titled "Install License" with a close button (X) in the top right corner. Inside the dialog, there is a section labeled "Upload License File?". To the right of this label is a blue "YES" button. Below the "Upload License File?" label is a text input field containing the filename "arcsight.lic". To the right of the input field is a "Browse" button. At the bottom right of the dialog is a large "Install" button.

29 After installing the license, you'll be returned to the login screen.

# Update MySQL Server Settings

To update the MySQL server settings, follow these steps:

- 1 Use a text editor to open the file **my.cnf**.
- 2 Add following entries in the [mysqld] section as shown below (The actual settings may vary depending on the server size and volume of data you plan to bring into the HPUBA application):

```
[mysqld]
lower_case_table_names=1
skip-external-locking
key_buffer_size = 256M
max_allowed_packet = 1G
bulk_insert_buffer_size = 100M
table_open_cache = 256
sort_buffer_size = 512M
read_buffer_size = 1G
read_rnd_buffer_size = 512M
myisam_sort_buffer_size = 512M
federated
thread_cache_size = 8
query_cache_size = 512M
thread_concurrency = 0
innodb_buffer_pool_size = 2G
innodb_additional_mem_pool_size = 256M

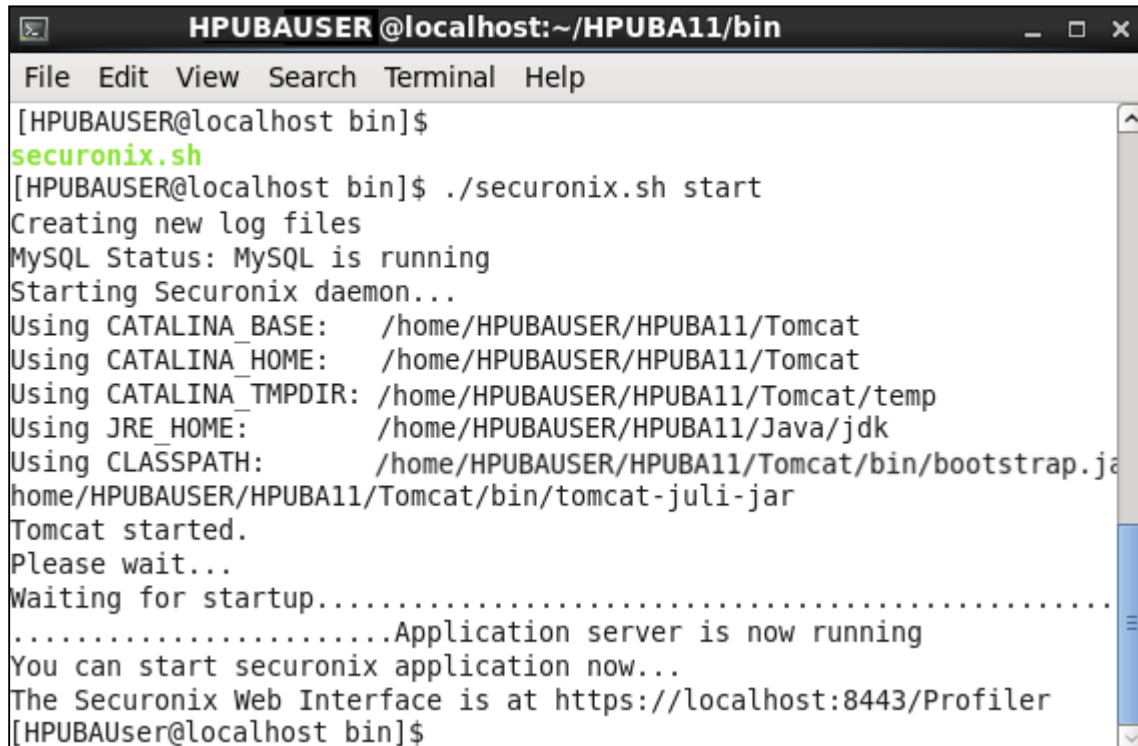
# Set ..log_file_size to 25 % of buffer pool size
innodb_log_file_size = 256M
innodb_log_buffer_size = 8M
innodb_flush_log_at_trx_commit = 0
innodb_doublewrite = 0
innodb_support_xa = 0
#innodb_lock_wait_timeout = 50
innodb_file_per_table=1
tmp_table_size = 512M
innodb_lock_wait_timeout = 180
connect_timeout = 100000
innodb_flush_method = O_DIRECT
```

## Chapter 3: Post Installation Activities

### Start Using the Application

Open a terminal, and navigate to the `/bin/` folder of your UP UBA installation. (In our installation sample, that folder is `/home/HPUBAUSER/HPUBA11/bin.`)

Start the application with the command `./securonix.sh start`



```
HPUBAUSER @localhost:~/HPUBA11/bin
File Edit View Search Terminal Help
[HPUBAUSER@localhost bin]$
securonix.sh
[HPUBAUSER@localhost bin]$ ./securonix.sh start
Creating new log files
MySQL Status: MySQL is running
Starting Securonix daemon...
Using CATALINA_BASE:  /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_HOME:  /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_TMPDIR: /home/HPUBAUSER/HPUBA11/Tomcat/temp
Using JRE_HOME:       /home/HPUBAUSER/HPUBA11/Java/jdk
Using CLASSPATH:      /home/HPUBAUSER/HPUBA11/Tomcat/bin/bootstrap.jar
/home/HPUBAUSER/HPUBA11/Tomcat/bin/tomcat-juli-jar
Tomcat started.
Please wait...
Waiting for startup.....
.....Application server is now running
You can start securonix application now...
The Securonix Web Interface is at https://localhost:8443/Profiler
[HPUBAUSER@localhost bin]$
```

Once the application starts, you can open your favorite browser and navigate to `http://ipaddress:portnumber/Profiler`

## Chapter 4: Uninstall HP UBA

Should you need to uninstall the application, follow these steps:

- 1 To remove the application, first be certain that you have stopped the application.

(**Note:** if you would like the uninstaller to delete the database schema as well, please make sure to keep MySQL running.)

- 2 To stop the HP UBA application, log out of any open web sessions, open a terminal session and type:

```
./securonix.sh stop
```

- 3 Next, navigate to the folder where the application is installed, then to the \_HPUBA\_installation (if you accepted the defaults during installation) folder. Example:

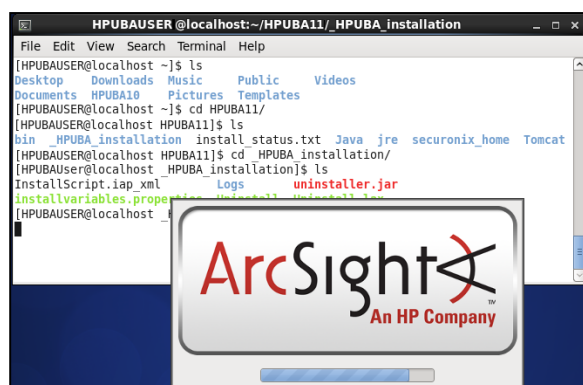
```
cd HPUBA11
cd _HPUBA_installation
```

Issue the following command:

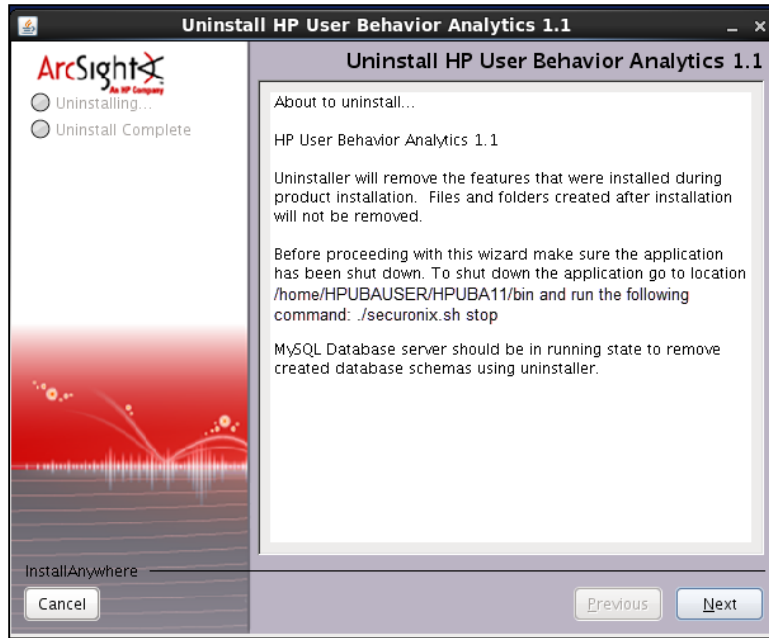
```
./Uninstall
```

```
[HPUBAUSER@localhost bin]$ ./securonix.sh stop
Stopping Securonix daemon
Stopping Application Server
Using CATALINA_BASE:   /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_HOME:   /home/HPUBAUSER/HPUBA11/Tomcat
Using CATALINA_TMPDIR: /home/HPUBAUSER/HPUBA11/Tomcat/temp
Using JRE_HOME:        /home/HPUBAUSER/HPUBA11/Java/jdk
Using CLASSPATH:        /home/HPUBAUSER/HPUBA11/Tomcat/bin/bootstrap.jar:/home/HPUBAUSER
/HPUBAUSER/HPUBA11/Tomcat/bin/tomcat-juli.jar
Waiting for application server to stop
Waiting for application server to stop
Waiting for application server to stop
Application server has been stopped successfully
NOTE: MySQL is running
[HPUBAUSER@localhost bin]$ cd ..
[HPUBAUSER@localhost HPUBA11]$ ls
bin  HPUBA_installation  install  status.txt  Java  jre  securonix_home  Tomcat
[HPUBAUSER@localhost HPUBA11]$ cd HPUBA_installation
[HPUBAUSER@localhost HPUBA_installation]$ ls
InstallScript.iap.xml  Logs  uninstaller.jar
installvariables.properties  Uninstall  Uninstall.lax
[HPUBAUSER@localhost HPUBA_installation]$ ./Uninstall
```

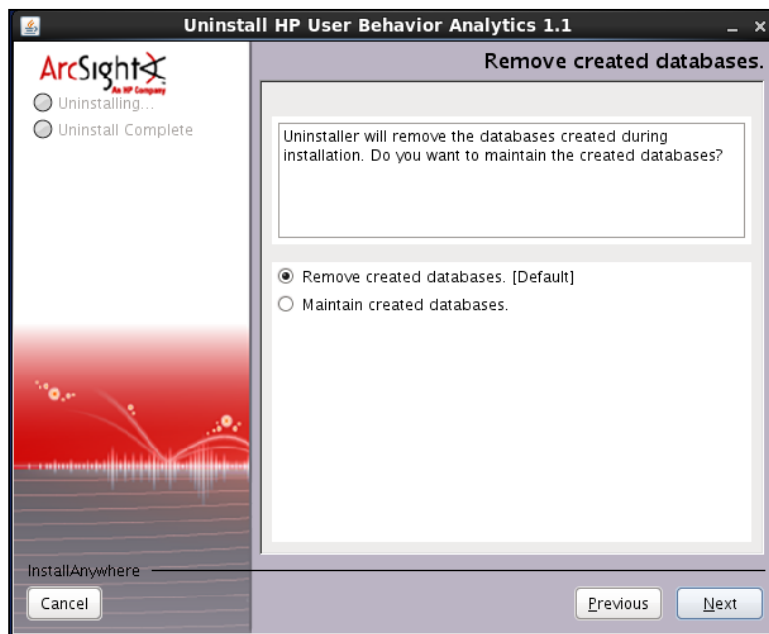
This will start the uninstallation wizard:



- 4 Click **Next** to start the uninstall process.

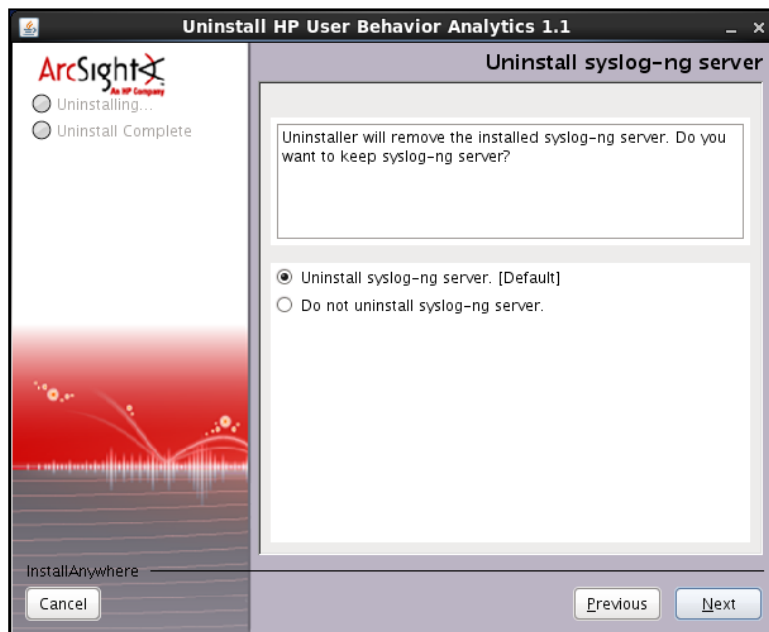


- 5 Select the option to remove or keep databases as part of the uninstall process.



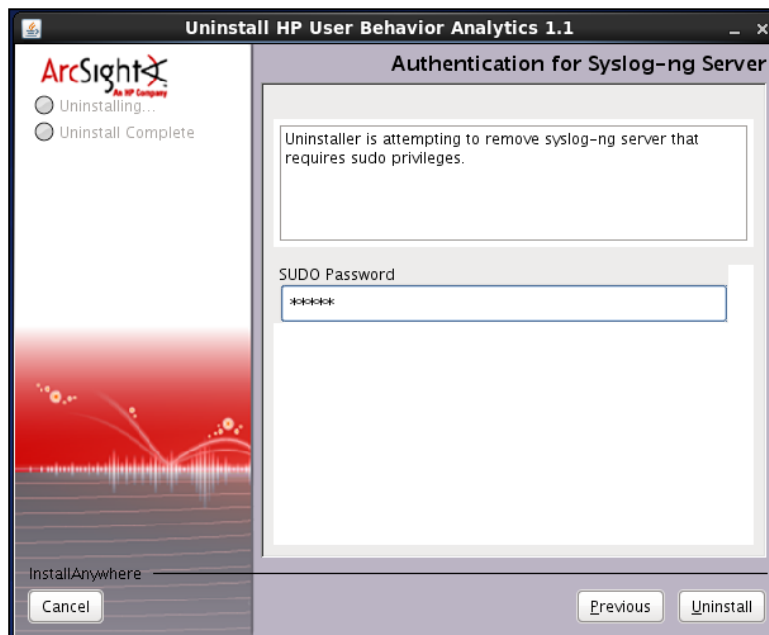


- 6 Select the option to remove Syslog or leave it in place during the uninstall process.

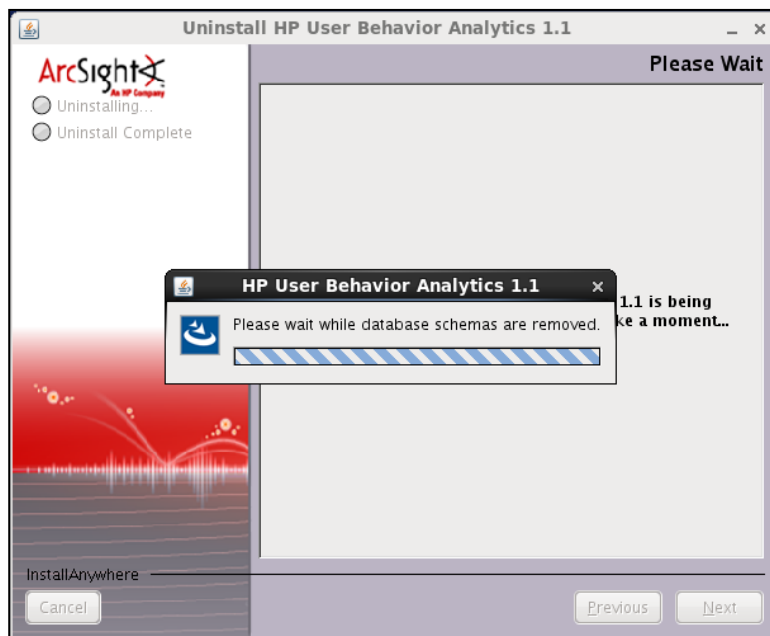
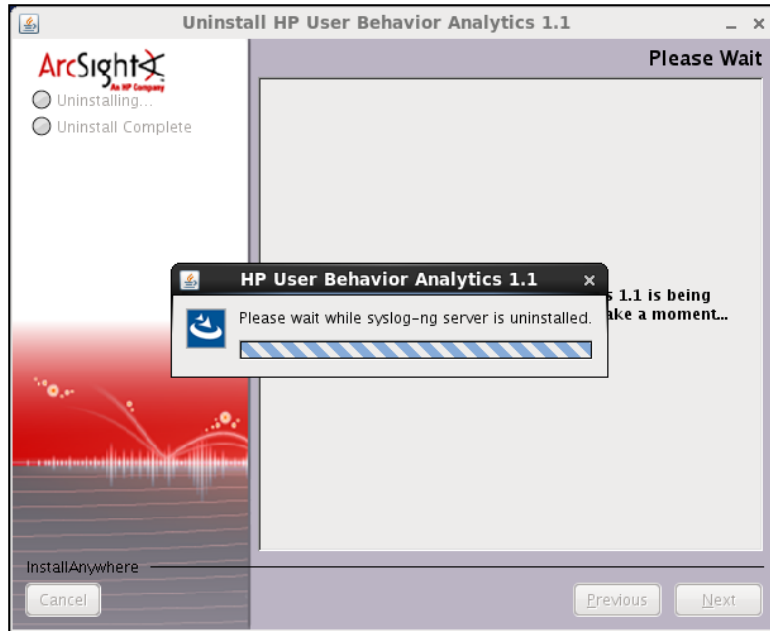


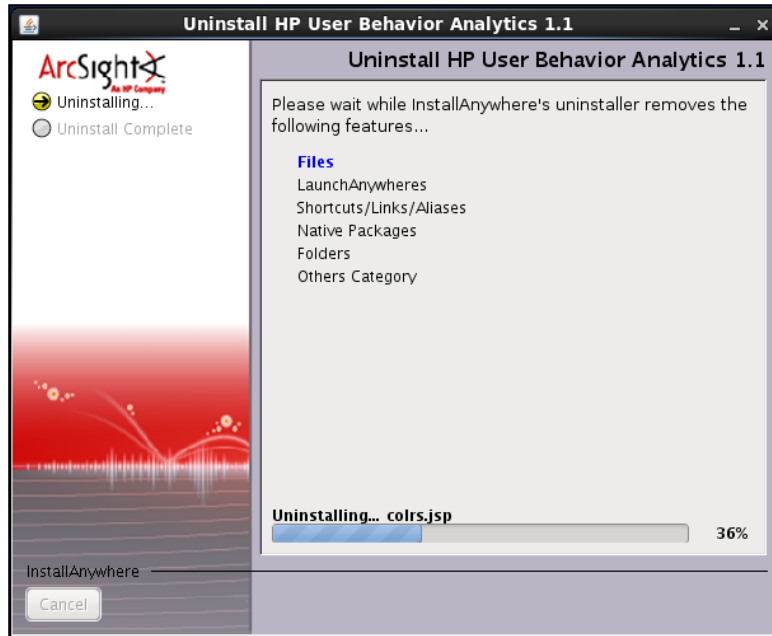
- 7 Click **Uninstall** to execute.

- 8 The uninstaller will ask for and verify the sudo password. Enter it and click **Uninstall**.



## Installation Guide





- 9 When the process is complete, you can remove the **HPUBA11** and any underlying folders manually.

## Chapter 5: Deployment Architecture

### Deployment Considerations

The appropriate architecture for each deployment must consider the following environmental requirements:

- High availability and failover requirements
- Disaster recovery plans
- Hardware requirements for each environment
- Network design
- Capacity planning
- Data integration considerations
- Implementation considerations and tasks
- Number of environments

Additionally, the architecture for the deployment of HP User Behavior Analytics platform solution depends on a number of throughput requirements including the function fulfilled by the solution and the amount of data to be analyzed. The throughput characteristics of the solution include:

- How many users will interact with HP UBA?
- Does the solution require high availability?
- How much data is being ingested to the solution?
  - Number of user identities
  - Number of access privileges
  - Number of devices (resources)
  - Number of activities
  - Type of analytics
  - Amount of live data (number of days \* amount of data per day)
  - Amount of warm data (number of days \* amount of data per day)

Based on the evaluation of the environmental and throughput variables above, the appropriate deployment architecture is selected. This includes the number of nodes required to analyze the data, the amount of storage required on the servers, the number of CPUs, RAM and processor speed required for each node.

### Capacity planning

When HP UBA imports your data, it creates two main types of files: the “rawdata” file that contains the original data in compressed form and the database files that store this data in the MySQL database. HP UBA also creates log files and additional configuration files that do not require a significant amount of storage. Based on the throughput requirements, you can estimate the disk space will be needed for based on the amount of incoming data.

Typically, the compressed raw data file is approximately 10% the size of the incoming, pre-imported raw data. The database files range in size from approximately 10% to 50% of the raw data file. This value varies based on the number of unique transactions in the data.

The best way to anticipate the storage requirements is to test importing a representative sample of your data, and then reviewing the sizes of the resulting data directories

When planning a deployment, it is important to decide which nodes are going to be responsible for the import of what data.

Single Node Deployment: All the events will be going to a single node.

Single Node – Multiple Forwarders: All the events ultimately go into the same database server.

Multiple Nodes – Multiple Forwarders: With this configuration, the nodes responsible for importing events must anticipate the storage requirements and allocate the appropriate resources to each node.

The following considerations are used to make this decision:

- How many events are generated by each device per day?
- How many days' worth of event data need to be keep online? How many events does this constitute?

Example of devices being monitored with the Events/Day anticipated:

Device Name	Device Type	Number of Events per day
AD-DC-001	Active Directory	10,000,000
AD-DC-002	Active Directory	30,000,000
ORA-DB-001	Oracle Database	20,000
ORA-DB-002	Oracle Database	40,000
HIDS Events	Host Intrusion Detection	20,000,000
SMTP DLP Events	Data loss prevention	80,000
Network DLP	Data loss prevention	1,280,000
Web Proxy Server	Web proxy	90,000,000
SMF	Mainframe	800,000
<b>Total Events Per Day</b>		<b>152,220,000</b>

The table above assumes that the requirements are to maintain 30 days' worth of live data for investigation and monitoring, and archive that data after the 30-day period. Over a 30-day period, the number of events would be 4,528,200,000 events (~ 4.520 billion events) for the entire solution. In this example deployment, three nodes, a single master and 2 child nodes can be used. The child nodes will divide the processing of the incoming event data.

Master Node: User Interface operations, and accepting events from child nodes

Child Node #1: Event data from AD-DC-001, AD-DC-002, ORA-DB-001, ORA-DB-002 and HIDS Events

Child Node #2: Event data from SMTP DLP Events, Network DLP, Web Proxy Server, SMF

Note: In a Multiple Application Node architecture, child nodes that are processing event sources should be in close proximity. The network bandwidth requirements should also be considered.

# Deploy HP UBA in a Master/Child Architecture

## Prerequisites

A child node is another instance of the HP UBA application that is running on another server. The Child node has its own database and application server. Before you begin, make sure that the following prerequisites are met:

- 1 The Master and child node must be running the same version (same build number) of the HP UBA application.
- 2 The Master and child node must be able to communicate with each other via web services (usually this is port 8080 or 8443 depending on the port on which the HPUBA application is running).
- 3 The child node must be able to connect to the database of the Master node.
- 4 The master node must be able to connect to the database of the Child node.
- 5 Be certain that the SecuronixDB connection type has been configured on the Master node. If not already configured, navigate to **Configure>Connection Types**, then **Actions>Add New Connection**.

## Add a Child Node

The configuration of clustering for HPUBA is done with in the web interface on the master node.

- 1 Navigate to **Configure>Clustering**.
- 2 Set **Enable Clustering** switch to **Yes**. (A pop up alert screen will appear. Click **Close**.)
- 3 Select to **Actions>Register New Node** to add new nodes to the application.
  - **Name:** Name of the node.
  - **Node Application URL:** Internet address for the node (for example: <https://192.168.1.155:8443/Profiler>).
  - **Child Node Management Server URL:** Management Server URL (for example: <https://192.168.1.154:8443/manager>).
  - **Enabled:** Select **Yes** to enable the node.

- **Application User Name:** Admin credentials for the HP UBA application running on the node. This is used to call the web services.
  - **Application Password:** Admin credentials for the HP UBA application running on the node. This is used to call the web services.
- 4 Enter credentials to connect to the HP UBA application running on the child server.
- **Health Check Interval:** Specify a value in seconds for the child node to provide a heartbeat to the master.
  - **Specify database timezone for node:** Specify the time zone of the child database.
  - **Advanced Settings:** Provide information for synchronizing data between the master and child nodes.
    - **Accounts:** synchronize the activityaccount table from master to child.
    - **Users:** synchronize the users table from the master to child.
    - **Organizations:** synchronize the organizations table from master to child.
    - **Lookuptables:** synchronize the lookup tables from master to child.
    - **Peer:** synchronize the peer tables.
    - **Resources:** synchronize the resources table.
    - **Watchlist:** synchronize the watchlist table.

You can choose to receive notifications when the child node fails to provide the master with a heartbeat.

- 5 Set **Enable notifications when a node goes down?** to **Yes**, and configure template to be used and period (in minutes) for notifications to be sent.
- 6 Click **Test Node URL** to test the created node.
- 7 Click **Next**.
- **URL:** The JDBC URL (for example: jdbc://mysql://192.168.1.155:3306/HPUBA) to the database.
  - **Username:** Database Username.
  - **Password:** Database password.

- 8 Click **Register**.

Once clustering is configured, the application configuration file (securonix\_home/conf/application-context.xml) can be reviewed to verify the details of the clustering.

Example Master application-context.xml:

```
<clustering enabled="true" master="true" nodeName="master" healthCheckInterval="30"
syncNodesInterval="5" resourceGroupsPerNode="150" />
```

Example Child Node 1 application-context.xml:

```
<clustering enabled="true" master="false" nodeName="ChildNode1" healthCheckInterval="30"
syncNodesInterval="5" resourceGroupsPerNode="150" />
```

## Appendix A: Tuning MySQL Configurations

Consider the following parameters when tuning your MySQL configuration:

### **Key\_buffer**

Change this parameter based on the amount of RAM in the system. Ideally, it should set at 20% of free memory available.

### **Key\_buffer\_size**

Change this parameter based on the amount of RAM in the system. Ideally, it should set at 20% of free memory available.

### **Sort\_buffer\_size**

This is a memory buffer used when ordering is required (Group By, Order By). Increasing the `sort_buffer_size` means allowing more memory to be used for the sorting process. However, increasing the `sort_buffer_size` can be detrimental to performance because the full size of the sort buffer is allocated for each thread that needs to do a sort, even if a large sort buffer is not required.

### **Read\_buffer\_size**

This parameter is used for caching the indexes in a temp file when sorting rows, bulk insert into partitions or caching results of nested queries. Set to a value in multiples of 4. If you do many sequential scans, you might want to increase this value from the default value which is 131072.

### **Read\_rnd\_buffer\_size**

Setting the variable to a large value can significantly improve ORDER BY performance. Change the variable only from within those clients that need to run large queries.

### **Join\_buffer\_size**

This is the minimum size of the buffer that is used for plain index scans, range index scans and joins that do not use indexes and thus perform full table scans. Increase the value of `join_buffer_size` to get a faster full join when adding indexes is not possible.

### **max\_heap\_table\_size:**

This parameter should be set at the available RAM divided by the maximum number of connections (`max_connections`).

### **Tmp\_table\_size:**

This parameter should be set at the available RAM divided by the maximum number of connections (`max_connections`).

### **Myisam\_sort\_buffer\_size**

This is the size of the buffer allocated when sorting indexes during a REPAIR TABLE or when creating indexes with CREATE INDEX or ALTER TABLE. The default size is 8 MB and the max size is platform dependent.

### **Query\_cache\_size**

This is the amount of memory allocated for caching query results. The default value is 0, which disables the query cache. The permissible values are multiples of 1024.



### **Thread\_concurrency**

This function enables application to give the threads system a hint about the desired number threads that should be run at the same time.

### **Innodb\_buffer\_pool\_size**

The size in bytes of the memory buffer InnoDB uses to cache data and indexes of its tables. The default value is 8MB. The larger you set this value, the less disk I/O is needed to access data in tables. On a dedicated database server, you may set this to up to 80% of the machine's physical memory size. However, do not set it too large because competition for physical memory might cause paging in the operating system. In addition, the time to initialize the buffer pool is roughly proportional to its size.

### **innodb\_additional\_mem\_pool\_size**

The size in bytes of a memory pool InnoDB uses to store data dictionary information and other internal data structures. The more tables you have in your application, the more memory you need to allocate here. The default value is 1MB.

### **innodb\_flush\_log\_at\_trx\_commit**

- If the value of `innodb_flush_log_at_trx_commit` is 0, the log buffer is written out to the log file once per second and the flush to disk operation is performed on the log file, but nothing is done at a transaction commit.
- When the value is 1 (the default), the log buffer is written out to the log file at each transaction commit and the flush to disk operation is performed on the log file.
- When the value is 2, the log buffer is written out to the file at each commit, but the flush to disk operation is not performed on it. However, the flushing on the log file takes place once per second also when the value is 2.
- The default value is set to 1.

### **innodb\_file\_per\_table**

If `innodb_file_per_table` is disabled (the default), InnoDB creates tables in the shared tablespace. If `innodb_file_per_table` is enabled, InnoDB creates each new table using its own .ibd file for storing data and indexes, rather than in the shared tablespace.

### **Max\_allowed\_packet**

It is safe to increase the value of this variable because the extra memory is allocated only when needed. For example, `mysqld` allocates more memory only when you issue a long query or when `mysqld` must return a large result row. The small default value of the variable is a precaution to catch incorrect packets between the client and server and to ensure that you do not run out of memory by using large packets accidentally. The default value is 16MB and can go up to 1GB.