



## **Accelops 4.3 BETA Installation Guide**

**Including Installation and Configuration for Accelops Visual Analytics**

**December 2014**

- 1. What's New in Release 4.3 ..... 3
- 2. Fresh Installation ..... 16
  - 2.1 Must Read Prior to Install ..... 17
  - 2.2 Configuring NFS Server ..... 20
  - 2.3 Installing in VMware ESX ..... 22
  - 2.4 Installing in KVM ..... 130
  - 2.5 Installing in HyperV ..... 134
  - 2.6 Installing in AWS ..... 143
- 3. Visual Analytics ..... 160
  - 3.1 AccelOps Visual Analytics: Overview ..... 161
  - 3.2 Installation and Configuration of AccelOps Visual Analytics ..... 165
    - 3.2.1 Requirements for AccelOps Visual Analytics ..... 166
    - 3.2.2 Installing and Configuring AccelOps Report Server ..... 168
    - 3.2.3 Installing and Configuring Visual Analytics Server ..... 198
    - 3.2.4 Installing and Configuring Visual Analytics Desktop ..... 199
  - 3.3 Creating and Managing Workbooks ..... 200
    - 3.3.1 Viewing Workbooks ..... 201
    - 3.3.2 Creating and Publishing Workbooks ..... 202
      - 3.3.2.1 Creating a Single Sheet Workbook ..... 203
      - 3.3.2.2 Creating a Multiple Sheet Workbook ..... 205
      - 3.3.2.3 Using AccelOps Workbooks with Tableau Visual Analytics Desktop and Server ..... 211
    - 3.3.3 Adding Users to Workbooks ..... 212
  - 3.4 Working with AccelOps Report Server ..... 213
    - 3.4.1 Report Server Architecture: phoenixdb and reportdb ..... 214
    - 3.4.2 Working with CMDB Data in AccelOps Report Server ..... 215
      - 3.4.2.1 Viewing phoenixdb Organization ..... 216
      - 3.4.2.2 Querying Incident Data in AccelOps Report Server ..... 217
        - 3.4.2.2.1 Reference: Attribute Columns in the ph\_incident\_view Table ..... 218
        - 3.4.2.2.2 Sample Incident Queries ..... 220
      - 3.4.2.3 Querying Other CMDB Tables in AccelOps Report Server ..... 221
        - 3.4.2.3.1 Querying Device Vendor and Model Distribution for Discovered Devices ..... 222
        - 3.4.2.3.2 Querying Discovered Devices ..... 223
    - 3.4.3 Working with Event Data in AccelOps Report Server ..... 224
      - 3.4.3.1 Viewing reportdb Organization ..... 225
      - 3.4.3.2 Syncing an AccelOps Report with Report Server ..... 226
      - 3.4.3.3 Deleting a Report from AccelOps Report Server ..... 227
      - 3.4.3.4 Modifying an Existing Report in AccelOps Report Server ..... 228
    - 3.4.4 Database Maintenance ..... 229
      - 3.4.4.1 Checking the Size of Tables ..... 230
      - 3.4.4.2 Creating Indices for Faster Queries ..... 231
      - 3.4.4.3 Deleting Old Tables ..... 232

# What's New in Release 4.3

This release adds features and functionality in several areas.

- **New Product**
  - AccelOps Visual Analytics
- **AccelOps Systems Features**
  - New chart types for search result visualization
  - Visualization for profiled metrics and drill down from incidents
- **Performance and Availability Monitoring**
  - Ability to specify per-device threshold in rules
  - Enhanced custom command output monitoring
  - Windows custom command output monitoring
- **Log Management and Security Incident Event Monitoring (SIEM)**
  - Amazon Web Services CloudTrail monitoring
  - Box.com file monitoring
  - Okta Single Sign-On (SSO) integration via SAML 2.0
  - Vendor default password usage detection
  - Detect malware via file hash and user agent match
  - Detect communication via Anonymity Networks (Open Proxies and TOR nodes)
- **Device Support**
- **Significant Enhancements**
  - Migrate SVN to local disk
  - Trigger event query optimization
  - Device location import
  - Event dropping rule enhancement
  - CIDR format for specifying discovery ranges
  - Launch discovery from CMDB
  - IP Address management enhancements
  - Critical interface selection usability enhancement
  - CMDB Report extensions
  - Dynamic EPS Adjustment algorithm
  - Incident table and Identity/Location table partitioning
  - Paged control support for Microsoft Active Directory LDAP discovery
  - Events when device performance monitoring status changes
  - Enhanced custom parser development graphical user interface
- **Fixed Issues and Minor Enhancements**
  - General GUI
  - Platform
  - Performance Monitoring / Event Pulling / Synthetic Transaction Monitoring (STM)
  - Rule / Query / Report Engine
  - Discovery
  - Device Support
  - Parsing area
  - Data: System Rules/Reports

## New Product

### AccelOps Visual Analytics

This release enables AccelOps data to be visualized using Tableau Visual Analytics and Visual Analytics Desktop in conjunction with the AccelOps Report Server. Two kinds of AccelOps data can be visualized:

- Data residing in Configuration Management Database (CMDB) e.g. Incidents, Device attributes
- Any event database report result as long as it contains an aggregation condition e.g. GROUP BY

You can find full information in the [Visual Analytics](#) section of this wiki.

## AccelOps Systems Features

### New chart types for search result visualization

This release enables users to visualize query results within AccelOps using scatter plots, bubble charts, tree maps and heat maps. This complements existing visualization via pie charts, bar charts, trend charts, and geo maps. Scatter plots enable users to see correlations between

any pair of calculated measures (e.g. CPU and memory utilization, Sent and received bytes etc). Bubble charts add a third dimension to scatter plots to reflect the size of the added dimension, e.g. in a scatter plot containing CPU and memory utilization as the two dimensions, the third dimension could be total sent and received bits/sec. The tree map is a hierarchical tree-structured visualization that is often used to analyze dominating components of multidimensional data e.g. IPS signatures, network traffic etc. Heat maps show the calculated measure for two dimensions using a color grade that helps users to understand severity. These charts are available both in Analytics and Dashboard areas.

This topic is discussed [here](#).

## Visualization for profiled metrics and drill down from incidents

AccelOps creates statistical baselines (profiles) for a large number of use cases. While earlier releases showed this information in tabular form, this data can now be visualized. For a specific dimension (such as host), up to four measures can be visualized on a trending hour-by-hour basis for weekdays and weekends as a multi-series column chart. In general, the profiles can be visualized as a scatter plot. From an incident indicating excessive deviation from statistical measures, it is now possible to drill down into the baseline with one click.

This topic is discussed [here](#).

## Performance and Availability Monitoring

### Ability to specify per-device threshold in rules

AccelOps has rules that trigger when certain thresholds are crossed. When the thresholds have to be fine-tuned on a per-device basis, rule conditions become complex and difficult to manage. This release solves this issue. Instead of explicit threshold values in rules, the thresholds are now defined as custom properties in which the user can choose to override the global thresholds by redefining the thresholds locally for a certain set of devices. Instead of hard-coding thresholds, rules are now written using a function that returns the appropriate values – local values if one is defined, global values otherwise. This approach keeps the number of rules the same, but allows users to set thresholds for any number of devices. Thresholds can be a simple number (e.g. CPU utilization) or a map (e.g. interface utilization for each interface, disk utilization for one or more disks).

This topic is covered [here](#) (see the section: "Thresholds as Custom Properties").

### Enhanced custom command output monitoring

This release enhances the way custom performance monitor command outputs are parsed into events. Prior to release 4.3.1, command outputs are parsed as one line into one event. This does not include, for example, "show version" commands for Cisco IOS routers that can span multiple lines. Release 4.3.1 improves this situation – multiple lines can be parsed into one event.

This topic is covered [here](#).

### Windows custom command output monitoring

A Windows custom performance monitor can be used to bring PowerShell command outputs into AccelOps. Prior to release 4.3.1, command output was obtained via Telnet/SSH, but that is not natural for Windows, and the user had to install Cygwin Telnet/SSH in Windows systems. This release enhances the situation by using winexe client on AccelOps - any Windows shell commands, such as PowerShell, can be remotely run on Windows servers using WMI credentials. There is no need to install any software on Windows. Additionally, multi-line command outputs can be parsed into one event in AccelOps. This enhancement now enables customers to run PowerShell commands in Windows servers.

This topic is covered [here](#).

## Log Management and Security Incident Event Monitoring (SIEM)

### Amazon Web Services CloudTrail monitoring

As more and more applications are deployed in the Cloud, monitoring user activity in the cloud is becoming increasingly important. For example, it is important to know when users are created, permissions are changed, virtual machines are spun up, network configurations are changed, or Virtual Private Clouds (VPCs) are created. This release enables AccelOps to efficiently collect, parse, report, and alert on Amazon Web Services activity via the AWS CloudTrail API.

This topic is discussed [here](#).



## Box.com file monitoring

Box.com is a cloud storage provider that is used by individuals as well as corporations to store and share files. This release enables AccelOps to monitor file activity within a Box.com account. AccelOps securely logs on to the Box.com account and monitors file creation, deletion, and modification activity within the account. More interestingly, for a specific file or all files in a folder, AccelOps can monitor file-sharing properties – is the file shared, is it password protected, is it preview/download enabled, and how many times was the file downloaded or viewed. If a particular file or directory contains confidential information, AccelOps can alert when any file in that directory was exposed to the outside or was viewed.

This topic is discussed [here](#).

## Okta Single Sign-On (SSO) integration via SAML 2.0

Okta is a cloud-based Single Sign-on (SSO) Service provider. This release enables AccelOps customers who are already authenticated in Okta to automatically log in to AccelOps without entering any credentials. AccelOps communicates via SAML 2.0 with Okta to verify user identity. In addition, AccelOps discovers all users defined in Okta (like Microsoft Active Directory) - the discovered users can be used in rule and report conditions and in notification policies. Finally, AccelOps collects Okta audit trails that can detect activity on the Okta web site such as account changes, logon activity, and other configuration changes.

This topic is discussed [here](#).

## Vendor default password usage detection

A common compliance requirement is to alert against the use of default vendor-defined credentials. This release enables AccelOps users to satisfy this requirement. AccelOps comes pre-built with a set of vendor and device specific default passwords. Users can add to this list. Whenever a device discovery succeeds with a credential from this list, an alert triggers.

This topic is discussed [here](#).

## Detect malware via file hash and user agent match

This release comes with a set of built-in well-known malware user agents and malware file hash signatures. Users can also import their own lists from outside sources. Since malware is known to use non-standard http user agents, AccelOps alerts when it sees a malware user agent, regular expression-based match from web server or web-proxy logs. If AccelOps is configured for file integrity monitoring, then it can alert when it detects a malicious file hash match in a monitored directory.

Malware hash is discussed [here](#).

User agent is discussed [here](#).

## Detect communication via Anonymity Networks (Open Proxies and TOR nodes)

A compromised host or a user with malicious intent uses various techniques to hide their identity, with two common examples being proxies and TOR networks. This release comes with a set of built in well-known proxies and TOR networks. Users can also import their own lists from outside sources. Whenever AccelOps sees an IP address match from firewall logs or Netflow, an alert is created.

This topic is discussed [here](#).

## Device Support

- Cisco VoIP infrastructure monitoring - see [here](#)
  - Cisco VoIP phone discovery from Cisco Call Manager via SNMP
  - Cisco Unity Connection - discovery and performance monitoring via SNMP
  - Cisco Presence Server - discovery and performance monitoring via SNMP
  - Cisco Contact Center - discovery and performance monitoring via SNMP
  - Cisco Tandberg VCS - discovery and performance monitoring via SNMP
  - Cisco Telepresence MCU - discovery and performance monitoring via SNMP
  - More detailed performance monitoring of Cisco Call Manager – SIP Trunk Status, Gateway Status, H323 Device Status, Voice mail Server Status, CTI Device Status, Media Device Status

- Parse 1000+ syslog messages from Cisco Call Manager and RTMT and create rules corresponding to RTMT Alerts
- Oracle ACME Packet Controller - discovery and performance monitoring via SNMP
- Brocade SAN Switch - discovery and performance monitoring via SNMP
- Dell Force10 Switch - discovery and performance monitoring via SNMP - see [here](#)
- Dell PowerConnect switches - discovery and performance monitoring via SNMP - see [here](#)
- Nimble Storage - discovery and performance monitoring via SNMP
- Cisco WAPX WLAN Controllers - discovery and performance monitoring via SNMP
- MS SQL Server 2014 - discovery, performance monitoring, audit log collection via SNMP, WMI, JDBC
- Oracle Audit log parsing via syslog
- Wireless LAN Controller "module" on Fortinet firewalls.
- TrippLite Environmental Monitors
- IBM WebSphere monitoring via HTTP(S) instead of JMX - see [here](#)
- Arista switches and routers - discovery and performance monitoring via SNMP - see [here](#)
- VMware vShield - log parsing via syslog

## Significant Enhancements

### Migrate SVN to local disk

AccelOps uses SVN to store device configuration data and installed software information. Over time, this repository can grow and contain a very large number of files. Earlier releases hosted SVN over NFS, and network performance could become an issue over time. Since all accesses to SVN are via the Supervisor node, this release moves SVN to a Supervisor local disk on a separate logical drive. Fresh AccelOps 4.3.1 install automatically create a separate partition for storing SVN files. During AccelOps 4.3.1 upgrade process, a special pre-upgrade step is invoked to copy the SVN files over NFS to local disk. Actual upgrade does not begin unless existing SVN data has been copied over to the new disk - so the system continues to work during the pre-upgrade process.

### Trigger event query optimization

Incidents are triggered by defined trigger events. When a user browses an incident in the graphical user interface, trigger events are shown, and incident notification emails can contain up to 10 trigger events. However, the AccelOps rule engine does not store raw events in memory, but only event identifiers, in an attempt to save memory. This means trigger events have to be retrieved from the event database by querying the event database. This query can be very expensive if the event is current, since the event may not have been indexed yet. This has been seen to create significant pressure on the AccelOps I/O system, especially if there is a sudden surge of incidents. This release addresses this issue by using an in-memory cache of raw messages for a short period of time.

### Device location import

CMDB devices typically belong to private address spaces and their location is only known to the system administrators. There is now an easy way to input this information into AccelOps. Users can define locations by IP range or sub-net, and the location in CMDB will be instantly updated. The locations can be input manually via the graphical user interface, or imported from a file. In addition, devices can be searched by location in both the summary dashboard and CMDB.

This topic is discussed [here](#).

### Event dropping rule enhancement

Devices are often chatty and send all kinds of uninteresting logs to AccelOps. Since online storage is expensive, it is often necessary to be able to efficiently drop events before they are processed or stored. This release enhances event dropping rule framework by:

- Including Source IP and Destination IP into the event dropping rule definition criteria
- Allowing two different actions: drop completely, or store but do not trigger rules
- Allowing the ability to automatically create drop rules from incidents in case the incident is a false positive, which is common in Network IPS event correlation scenarios

This topic is discussed [here](#).

### CIDR format for specifying discovery ranges

The test connectivity and discovery IP ranges can now be specified in CIDR notion as well.



## Launch discovery from CMDB

Rediscovery can be directly launched from the CMDB page.

CMDB > Devices			
<input type="button" value="New"/> <input type="button" value="Delete"/> <input type="button" value="Edit"/> <input type="text"/>			
Name	IP Address	Type	Ver
TwikiServer(Do_not_delete)		Generic Linux	Any
SJ-QA-F-Lnx-CHK	172.16.0.1	Checkpoint Firewall-1 SPLAT	R65
HQ-A-Pxy-blueCoat	172.16.0.141	Quick Info	
SJ-QA-A-Fdy-BigIron	172.16.0.2	Topology	
SJ-QA-Dmz-HPSW	172.16.0.254	Show Real-time events on this IP	
SJ-Dev-A-Fdy-FastIron	172.16.0.4	Show Events on this IP for last 5 mins	
PA-500_01_accelops	172.16.1.2	Add to WatchList	
QA-W2K3-X64	172.16.10.13	Re-Discover	
		Connect To...	

## IP Address management enhancements

When allocating new addresses to hosts, it is often important to know the hosts that currently are assigned to addresses in a specific network segment. In prior releases, AccelOps discovered the network segments and showed only CMDB devices in that network segment. These do not include user devices such as laptops, workstations, mobile devices etc., since these devices do not necessarily belong in CMDB. Starting with this release, the Network Segment page also shows the hosts in the Identity and Location page belonging to the same network segment. Since AccelOps accurately learns all the IP addresses in a network via DHCP and IP ARP cache scan, administrators can correctly see every active host belonging to a specific network segment.

This topic is discussed [here](#).

## Critical interface selection usability enhancement

AccelOps allows users to mark interfaces as critical, and such interfaces are always monitored for utilization and up/down status. A common example is switch trunk ports, since a trunk port going down can cause a widespread network outage. Currently there is no easy way in AccelOps to select the trunk ports of all switches. Instead, the user has to traverse every switch and select trunk ports within that switch, which can be very tedious for a large network containing a large number of switches. This release provides a flattened view of the network interfaces so that a user can quickly select a large number of interfaces matching some search criteria. This enables administrators to mark all critical interfaces for a large network with only a few clicks.

This topic is discussed [here](#).

## CMDB Report extensions

CMDB Reports are extended to include

- Successful Performance Monitor Reports
- Failed Performance Monitor Reports
- Identity and Location Report
- Scheduled Report
- Devices not updated in last N days

## Dynamic EPS Adjustment algorithm

AccelOps has an algorithm to re-distribute unused EPS at a collector to other collectors seeing an event spike. The algorithm is now adjusted to have the following property: A collector is now always guaranteed to have the events-per-second specified as "Guaranteed EPS." This EPS is never redistributed to other collectors. Only the excess EPS (defined as Overall EPS license minus the sum of all Guaranteed EPS) is redistributed on demand.

## Incident table and Identity/Location table partitioning

In AccelOps CMDB, there are two tables that grow with time:

- Incident table
- Identity/Location table

The incident table grows as new incidents are created, while the Identity/Location table grows as new computers and users appear in the system or change location. As these tables grow, eventually the database may become full and read performance may suffer with corresponding growth in the table indices. In this release the following enhancements are made:

### Incident Table Optimization:

- The incident table is partitioned by month, so recent queries access the current month and result in fast returns
- During migration to 4.3.1 release:
  - Data for the last three months is migrated to the new tables (based on Last Seen Time field)
  - All 'Active' incidents are migrated
  - Older incidents are archived. Scripts are provided for customer to migrate older incidents into 4.3.1 CMDB.
  - Scripts are provided to purge older incidents from 4.3.1 CMDB

### Identity/Location Table Optimization:

- The Identity/Location table is partitioned by month, so recent queries access the current month and result in fast returns
- During migration to 4.3.1 release:
  - Data for the last three months is migrated to the new tables (based on Last Seen Time field)
  - Older entries are archived. Scripts are provided for customer to migrate older identity/location entries into 4.3.1 CMDB.
  - Scripts are provided to purge older identity/location entries from 4.3.1 CMDB

## Paged control support for Microsoft Active Directory LDAP discovery

AccelOps discovers users in Microsoft Active Directory via LDAP protocol. By default, Microsoft LDAP search queries return up to 1000 entries per call (MaxPageSize limit - see [Microsoft KB article](#)). Earlier AccelOps releases required administrators to increase the MaxPageSize limit to a much higher number for user discoveries to work. This is generally inconvenient and may also cause resource issues on the server. This release enhances this situation. AccelOps LDAP discovery now uses the paged control version of the LDAP search API to fetch an arbitrarily large number of entries - 1000 at a time. Administrators are not required to increase the MaxPageSize limit beyond the default 1000.

## Events when device performance monitoring status changes

AccelOps now generates audit events when the performance monitoring status of a job changes.

1. User deleted a device or a collector:

```
<174>Nov 05 09:52:07
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[srcIpAddr]=192.168.20.164,[phCustId]=1,[sessionId]=11178d2aeae08e9c2babe2725fa1,[procName]=AppServer,[hostName]=HQ-A-Pxy-blueCoat,[hostIpAddr]=172.16.0.141,[eventSeverity]=PHL_INFO,[customer]=Super,[jobStatusType]=UserDeletedDevice,[user]=admin,[phLogDetail]=Monitors on device were deleted due to device being deleted
```

## 2. User disabled monitoring at a device level:

```
<174>Nov 05 09:53:58
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[srcIpAddr]=192.168.20.164,[phCustId]=1,[sessionId]=11178d2aeae08e9c2babe2725fa1,[procName]=AppServer,[hostName]=ACCELOPS-W2K3B4,[hostIpAddr]=192.168.64.124,[eventSeverity]=PHL_INFO,[custName]=Super,[jobStatusType]=UserDisabledDevice,[user]=admin,[phLogDetail]=Monitoring device, 192.168.64.124, is disabled by user
```

## 3. User enabled monitoring at a device level:

```
<174>Nov 05 09:54:38
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[srcIpAddr]=192.168.20.164,[phCustId]=1,[sessionId]=11178d2aeae08e9c2babe2725fa1,[procName]=AppServer,[hostName]=ACCELOPS-W2K3B4,[hostIpAddr]=192.168.64.124,[eventSeverity]=PHL_INFO,[custName]=Super,[jobStatusType]=UserEnabledDevice,[user]=admin,[phLogDetail]=Monitoring device, 192.168.64.124, is enabled by user
```

## 4. User disabled a specific job:

```
<174>Nov 05 09:55:17
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[phCustId]=1,[customer]=Super,[jobName]=System cpu usage,[srcIpAddr]=192.168.20.164,[appTransportProto]=SNMP,[sessionId]=11178d2aeae08e9c2babe2725fa1,[procName]=AppServer,[hostIpAddr]=192.168.64.124,[hostName]=ACCELOPS-W2K3B4,[eventSeverity]=PHL_INFO,[jobStatusType]=UserDisabledJob,[user]=admin,[pullInterval]=180,[phLogDetail]=Protocol to monitor is disabled
```

## 5. User enabled a specific job:

```
<174>Nov 05 09:55:59
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[phCustId]=1,[customer]=Super,[jobName]=System cpu usage,[srcIpAddr]=192.168.20.164,[appTransportProto]=SNMP,[sessionId]=11178d2aeae08e9c2babe2725fa1,[procName]=AppServer,[hostIpAddr]=192.168.64.124,[hostName]=ACCELOPS-W2K3B4,[eventSeverity]=PHL_INFO,[jobStatusType]=UserEnabledJob,[user]=admin,[pullInterval]=180,[phLogDetail]=Protocol to monitor is enabled
```

## 6. User changed job polling interval:

```
<174>Nov 05 09:57:21
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[phCustId]=1,
[customer]=Super,[jobName]=System real memory usage,
[srcIpAddr]=192.168.20.164,[appTransportProto]=SNMP,[sessionId]=11178d2aae08e9c2
babe2725fa1,[procName]=AppServer,[hostIpAddr]=192.168.64.124,[hostName]=ACCELOPS-
W2K3B4,[eventSeverity]=PHL_INFO,[jobStatusType]=UserChangedPollIntv,[user]=admin,
[pullInteval]=300,[phLogDetail]=Interval of protocol to monitor is changed
```

## 7. Discovery added a job for monitoring:

```
<174>Nov 05 10:05:14
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[phCustId]=1,
[customer]=Super,[jobName]=WMI Ping Status,
[appTransportProto]=WMI,[procName]=AppServer,[hostIpAddr]=192.168.64.124,
[hostName]=ACCELOPS-W2K3B4,
[eventSeverity]=PHL_INFO,[jobStatusType]=DiscoveryAdded,[user]=SYSTEM(su),[pullIn
teval]=120,[phLogDetail]=Monitor on device is added
```

## 8. Discovery removed a job for monitoring:

```
<174>Nov 05 10:12:08
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[phEventCategory]=2,[phCustId]=1,
[customer]=Super,[jobName]=Process Resource Usage via WMI,
[appTransportProto]=WMI,[procName]=AppServer,[hostIpAddr]=192.168.64.124,
[hostName]=ACCELOPS-W2K3B4,
[eventSeverity]=PHL_INFO,[jobStatusType]=DiscoveryRemoved,[user]=SYSTEM(su),[pull
Inteval]=180,[phLogDetail]=Status of protocol to monitor is changed
```

## 9. Performance Monitor module did not schedule a job:

```
<174>Nov 05 10:33:01
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[errReason]=Missing/Invalid WMI
credential for 192.168.20.207, PROC_RESOURCE,
[phEventCategory]=2,[phCustId]=1,[jobId]=1545818,[customer]=Super,[jobName]=Proce
ss Resource Usage via WMI,
[srcIpAddr]=192.168.64.153,[appTransportProto]=WMI,[sessionId]=13b348ad44270e0249
eafc9dfdc5,[procName]=AppServer,[hostIpAddr]=192.168.20.207,[hostName]=win-li5sip
p8s7s.accelops.net,[eventSeverity]=PHL_INFO,[jobStatusType]=DiscoveryNotScheduled
,[user]=1,[pullInteval]=180,[phLogDetail]=Monitor on device is not scheduled
```

## 10. Successful job:

```
<174>Nov 05 10:13:00
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[errReason]=,[phEventCategory]=2,
[phCustId]=1,[jobId]=1536112,[customer]=Super,[jobName]=Process Resource usage
via
SNMP,[srcIpAddr]=192.168.64.153,[appTransportProto]=SNMP,[sessionId]=128e159a35b9
f0f4cd71ca80222b,[procName]=AppServer,[hostIpAddr]=192.168.20.170,[hostName]=qa-w
in2008-217.accelops.net,[eventSeverity]=PHL_INFO,[jobStatusType]=ExecutionSuccess
,[user]=1,[pullInteval]=120,[phLogDetail]=Status of monitor is changed by Job
```

## 11. Failed job:

```
<174>Nov 05 10:15:00
[PH_AUDIT_DEV_MON_JOB_STATUS_CHANGE]:[custId]=1,[errReason]=Failed to get process
utilization in
executeGeneralProcResourceJobOpt,[phEventCategory]=2,[phCustId]=1,[jobId]=1536112
,[customer]=Super,[jobName]=Process Resource usage via SNMP,
[srcIpAddr]=192.168.64.153,[appTransportProto]=SNMP,[sessionId]=12ab68da5ca2ccea2
2a69cbda16e,[procName]=AppServer,[hostIpAddr]=192.168.20.170,[hostName]=qa-win200
8-217.accelops.net,[eventSeverity]=PHL_INFO,[jobStatusType]=ExecutionFailed,[user
]=1,[pullInterval]=120,[phLogDetail]=Monitoring device failed
```

## 12. Job stays in "Discovered Added" state for more than 15 minutes and is not scheduled:

```
<174>Nov 05 10:55:57
[PH_AUDIT_DEV_MON_JOB_NOT_STARTED]:[custId]=1,[phEventCategory]=2,[phCustId]=1,[c
ustomer]=Super,[jobName]=ICMP Ping
Status,[appTransportProto]=PING,[procName]=AppServer,[hostIpAddr]=172.16.10.110,[
hostName]=HOST-172.16.10.110,[eventSeverity]=PHL_INFO,[jobStatusType]=DiscoveryAd
ded,[user]=SYSTEM(phDiscovery),[pullInterval]=120,[phLogDetail]=Monitoring job did
not start yet
```

## Enhanced custom parser development graphical user interface

The custom parser development graphical user interface is enhanced to include the following:

1. Ability to search text within the XML file.
2. Add a line number in the XML file. When there is a Error in 'Validate' or 'Test', show the line number as a reference to help user fix the problem.
3. Allow user to reformat the text after block update for easy readability.
4. Allow an option to Clear XML in one shot to allow for bulk replace.
5. Color code the XML tags and text for easy readability.
6. Show the parsed fields in Test results in a nice tabular form.
7. Improve the scrolling/editing response for large XML files.
8. Show the XML in a tree form - allow cross-linking of the XML Tree and the text edit window.
9. Allow user to increase the size of the edit window.

Clone Event Parser Definition

**Name \*** ArubaWLANParser-12:23:36

**Device Type \*** Aruba ArubaOS WLAN Controller

**Test Event \*** 2008-06-11 11:38:34 192.168.20.7 [192.168.20.7]:SNMPv2-MIB::sysUpTime.0 = Timeticks: (1355400) 3  
SMI::enterprises.14823.2.2.1.1.100.1003 SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.60 = Hex-STRIP  
SMI::enterprises.14823.2.2.1.1.2.1.1.2.192.168.180.1 = Hex-STRING: 00 1E 52 72 AF 4B

**Parser XML**  ⏪ Prev Next ⏩

```

115 STRING: "/>
116     </collectAndSetAttrByKeyValuePair>
117
118     <collectAndSetAttrByKeyValuePair sep="\t\\| SNMP\\|) Uptime" src="$_body">
119       <attrKeyMap attr="_eventType" key="SNMPv2-SMI::enterprises.14823.1.1.10 Enterprise Spe
120     </collectAndSetAttrByKeyValuePair>
121
122     <!-- set default event type, only event ID but no detail info -->
123     <when test="exist _eventType">
124       <setEventAttribute attr="eventType">combineMsgId("Aruba","-",$_eventType)</setEventA
125     </when>
126
127     <!-- get attributes on trap by trap basis -->
128
129     <choose>
130
131       <!-- system status related -->
132       <when test="$_eventType = "1000"">
133     <setEventAttribute attr="eventType">combineMsgId("Aruba","-",$_eventType,"-
134 ", "wlsxVlanLinkUp")</setEventAttribute>
135     <collectAndSetAttrByKeyValuePair sep="\t\\| SNMP" src="$_body">
136       <attrKeyMap attr="hostVLAN" key="SNMPv2-SMI::enterprises.14823.2.3.1.11.1.1.23.0 =
137     </collectAndSetAttrByKeyValuePair>
138     </when>
139
140     <when test="$_eventType = "1001"">
141     <setEventAttribute attr="eventType">combineMsgId("Aruba","-",$_eventType,"-
142 ", "wlsxVlanLinkDown")</setEventAttribute>
143     <collectAndSetAttrByKeyValuePair sep="\t\\| SNMP" src="$_body">

```



Enabled

Description

## Fixed Issues and Minor Enhancements

### General GUI

- Bug 7489: Created a CMDB report named "Active Dependent Rules" that tracks which rules depend on other rules. This helps users to tweak/enable/disable chained system rules
- Bug 8021: Added indices in ph\_task and ph\_alter tables in PostgreSQL - this improves the GUI experience when user visits Alert and Task tabs
- Bug 8054: Allow an option to search on 'Origin' field in every tab in Admin > Device Support area. This allows users to quickly see user defined Device/Application types, Event Types, Event Attribute Types, Parsers and Dashboard columns
- Bug 8165: Show VLAN as a column in Analytics > Identity and Location Report
- Bug 8181: Need to get result of scheduled report even if the report has no data
- Bug 8291: Allow user to unlock an AccelOps account
- Bug 8896: Allow scheduled reports to skip charts and only contain tables
- Bug 9266: Added "errReason" attribute to system event PH\_REPORT\_ACTION\_STATUS - the attribute states why notification failed
- Bug 9670: CMDB Device shows under Scheduled Maintenance even after device removed from Schedule Maintenance Calendar
- Bug 9900: Expose Last Updated Time and Discover Method fields of a device for use in CMDB reports
- Bug 10083: Display a warning when user disables or deletes a rule that is referenced in other rules
- Bug 10172: Change the AccelOps GUI CMDB > Users so that all locally created users cannot edit the "domain" field
- Bug 10198: Removal of devices or organizations from CMDB sometimes display foreign key violation errors
- Bug 10250: Remove "Show Password" check box for credential
- Bug 10382: enhancement: allow HTML tags in custom e-mail templates
- Bug 10394: Allow bulk disabling for blocked IP in the CMDB through the GUI
- Bug 10450: Add 'Apply To' option to facilitate applying multiple authentication profiles to one or more users
- Bug 10563: Add an Export button for Related Incidents screen
- Bug 10830: Add locations view in summary dashboard
- Bug 11343: Long device names truncated on Widget dashboard
- Bug 11371: Allow import/ export of user defined watch list
- Bug 11498: A rule with CLEAR conditions becomes invalid after clone process - constraints between main and clear rules are not properly copied over
- Bug 11508: Ability to set locations for a large number of devices in one shot
- Bug 11596: Ability to add Notes to Rule exceptions. One should also be report on Rule exceptions.
- Bug 11597: Add Remediation section to Rule definition. Add this to default email template. Make this part of CMDB report. Add this to custom notification template.
- Bug 12583: User can not manually add important processes that have the same name but different process parameters
- Bug 12613: Columns on Amazon EC2 performance view should be same as EC2 dashboard
- Bug 12694: Provide an option to not have charts in exported PDF reports
- Bug 12698: Enable search on "Monitor Errors" and "Error Description" on Admin > Setup wizard > Monitor Change/Performance > Monitor Errors popup
- Bug 12760: Edit a Report Schedule and the Report automatically Runs
- Bug 12786: Make error message clearer for event dropping rule creation on grouped incidents

### Platform

- Bug 9518: Glassfish log rotation is now configured for saving space - only keep 20 files and each of them max-sized 2000000 bytes.
- Bug 9828: EPS Pulling functionality has limitations that lead to dropping of events by collectors
- Bug 9938: Allow modular 'yum upgrades' for non-base-CentOS packages like JVM, Chrome, PostgreSQL, Glassfish
- Bug 10144: Do not overwrite customer's ssl.conf during upgrade
- Bug 11926: DNS caching code has performance issues
- Bug 12130: AccelOps uses rsyslog to receive our internally generated events. There is a throttle defined there (200 messages in 5 minutes interval). This will result message loss in high throughput situation like VoIP phone discovery, Layer 2 port mapping discovery etc. This throttle is removed since this is intra-computer communication and can handle much higher message rates.
- Bug 12538: Detailed events, rules and reports for performance monitoring status changes
- Bug 12584: Collectors sometimes fail to negotiate HTTP(S) connection to Super/Worker if ever they choose SSLV3 (because of poodle

- vulnerability - possibly because of a man-in-the-middle device like a IPS or a firewall disallowing all SSLV3 negotiations.
- Bug 12585: The phMonitor module crashes when it sees a 3.7.6's rest\_cache\_api list entry in phoenix\_config.txt
- Bug 12586: The configuration file phoenix\_config.txt needs to be upgraded properly by maintaining user's changes from previous versions
- Bug 12644: Run script notification may fail if the raw message contains special XML characters
- Bug 12649: Updating Dynamic Watch List by incidents causes Application Server to run out of memory when there are many many-to-many relationships between incidents and dynamic watch lists
- Bug 12650: Full VM build does not 'yum update' packages as previously designed

## Performance Monitoring / Event Pulling / Synthetic Transaction Monitoring (STM)

- Bug 9848: Packet transmission timeouts for SNMP v1 and v2 phoenix\_config needs to be extended from 1 minute to 5 minutes
- Bug 11423: Add Custom command output monitoring via winexe for windows environments
- Bug 12066: Parse CVSS\_BASE score for vulnerabilities into (vuln CvssBaseScore attribute)
- Bug 12213: Cisco IOS CPU can not be monitored in some cases with multiple CPUs - performance monitoring has to identify the control plane cpus
- Bug 12214: PerfMonitor module will stop sending PH\_DEV\_MON\_PING\_STAT events for a gateway if its immediate down steam device are down
- Bug 12387: NexPose vulnerability report XML parsing takes a long time
- Bug 12458: Checkpoint needs a resume event handler
- Bug 12561: Discovery never removes a PING job even if the device is not reachable by PING during discovery. This is done since a PING is fundamental for measuring up time. Only a user can manually disable PING jobs
- Bug 12601: Admin > Setup wizard > Monitor Change/Performance tab does not reflect the status of successful discovery after correcting device credentials.
- Bug 12604: Events are not picked up by parser module if Supervisor node is down for an extended period of time
- Bug 12625: For Qualys Vulnerability Scanner, Test connectivity succeeds, but the discover method and event pull methods are not set in discovery, resulting in no job creation for Java agent.
- Bug 12661: Don't trigger config change while getting config error
- Bug 12730: Allow pulling interval to be less than 1 minute to pull windows logs at a faster rate. Added a phoenix config entry of "wmi\_pull\_interval\_scale" that can have a range from 1 (default) to 10. If users want to have a shorter interval for WMI event pulling, they can change it to 6 to make it 6 times faster; if the pulling interval in GUI is 1 minute, then events are pulled every 10 seconds.
- Bug 12754: Enhance custom command output monitoring to generate an event to indicate no matching lines for regular expression. This can be used to detect e.g. a process is down from running the top command. If the regular custom command output monitoring command is PH\_DEV\_MON\_CUST\_DF then AccelOps would generate PH\_DEV\_MON\_CUST\_DF\_NOT\_FOUND when there are no matching lines in the command output
- Bug 12787: Enhance performance monitoring status job upload to keep uploading if failed last time
- Bug 12801: Custom command output monitoring - deleted items are still monitored
- Bug 12803: Custom SNMP job monitoring sometimes fails to distinguish when one key is a prefix of the other; e.g. SNMPv2-SMI::enterprises.9.9.48.1.1.1.5.1 from SNMPv2-SMI::enterprises.9.9.48.1.1.1.5.16. So the obtained value for SNMPv2-SMI::enterprises.9.9.48.1.1.1.5.1 may be the value for SNMPv2-SMI::enterprises.9.9.48.1.1.1.5.16
- Bug 12804: Custom Performance Monitoring - when custom transforms are nested, the order may not be preserved resulting in wrong calculations. For example, if custom transforms are defined as "used = transform(used/1024)", "total = transform(total/1024)" and "memUtil = transform(used/total)", then the transforms must be calculated in the order used -> total -> memUtil. Earlier releases did not do this.

## Rule / Query / Report Engine

- Bug 8512: Turn off inaccurate system rule 'Windows Disk controller problem'
- Bug 9847: Add Identity and location to CMDDB Report
- Bug 10996: Worker down rule did not trigger after license expires
- Bug 11008: RFE: when retrieving Triggered Events for an Incident in GUI or e-mail notification, search 60 minutes before and after Incident Time. This is fixed by the 'Trigger Event Query Optimization' feature
- Bug 12322: Increase per rule GROUP BY thresholds - per cust and over all customers
- Bug 12558: Improved cache miss handling for profile anomaly rule handling: The profiles for anomaly detection are stored in a SQLite database. When rule engine attempts to look up the average and standard deviation values for a particular key (e.g. IP address, port number etc), and an exact match is not found, then earlier releases picked up the lowest values for that profile in that database. This heuristic often causes unnecessary rule triggers. This release makes tightens the cache miss handling case - profile anomaly rules do not trigger for a key value if the average and standard deviation values are not found for that key value for that time period.
- Bug 12640: When performing analytical searches when individual countries from the country groups are referenced in filter conditions as objects, then no results are returned

## Discovery

- Bug 10363: Disallow discovery via Virtual IP
- Bug 10533: Add ability to define IP subnets in the 172.16.16.0/22 type format
- Bug 11308: Foundry router becomes generic-generic
- Bug 11972: Telnet discovery of HP Procurve switches fails due to error in expect script

- Bug 12713: If host name contains unprintable characters like backspace(x08) and enquiry(x05), then parsing this XML causes app server to throw exceptions and run out of memory

## Device Support

- Bug 9942: RFE: add performance monitoring for Cisco WAPX (lightweight) devices
- Bug 10006: RFE: add device support for Tripp lite UPS devices
- Bug 10307: Microsoft cluster services incorrectly assigned to Microsoft Exchange Application Group
- Bug 10362: Add support for (Oracle) Acme Packet Border Controller
- Bug 10366: Support for Dell Force 10 Networking devices

## Parsing area

- Bug 8894: Cisco ASA parser: trailing white space in User attribute causing searches to miss events with condition User EQUALS 'string'
- Bug 10418: Parse username in Windows MSSQL Event 18453
- Bug 11351: Parse username in Win-Security-5145
- Bug 12485: Parse jitter field in Cisco VoIP CDR/CMR record
- Bug 12626: Snort events collected via database have wrong severity
- Bug 12641: Added more Symantec Anti-virus events
- Bug 12643: A null pointer exception can happen during pulling performance monitor config for discovery
- Bug 12656: Some FortiGate traffic logs do not parse if "status=" is missing in the logs
- Bug 12668: More Foundry switch logs to be parsed
- Bug 12714: Enhance McAfee EPO parser to parse more logs
- Bug 12716: Enhance Cisco IOS and NX-OS parser to parse more logs
- Bug 12757: Put Fortigate firewall DHCP messages into the identity and location section - so the IP to user and host name mapping shows up when FortiGate is acting as the DHCP server
- Bug 12765, 12767: Add a "Total Bit Rate" attribute for interface utilization and Netflow events - so user can quickly run 95th percentile on the total bandwidth for an interface

## Data: System Rules/Reports

- Bug 8512: Turn off inaccurate system rule 'Windows Disk controller problem'
- Bug 10113: Added description for windows security events: 5142, 5143, 5144, 5168, 4985, 5145
- Bug 12660: Enhance "Heavy TCP Port Scan" rule to exclude Windows Security Firewall logs (Win-Security-5156), ASA/PIX Teardown events (ASA-302014, PIX-302014, FWSM-302014)
- Bug 12676: The event type group for JUNOS\_KMD\_VPN\_DOWN\_ALARM\_USER event is incorrect

# Fresh Installation

## Must Read Prior to Install

- [Enterprise Deployment Scenarios](#)
- [Service Provides Deployment Scenarios](#)
- [Hardware Requirements](#)
  - [Hardware Requirements for Supervisor and Worker nodes](#)
  - [Hardware Requirements for Collector nodes](#)
  - [Hardware Requirements for Report Server nodes](#)
- [Generic Steps during installation](#)

## Enterprise Deployment Scenarios

Scenario	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
AO-VA 1 node	x						This is the most basic single site enterprise deployment
AO-VA 1 node with Collectors	x		x				This is also an enterprise deployment covering multiple sites. Data collection is simplified by deploying a collector for the satellite sites.
AO-VA Cluster	x	x		x			This is the scalable enterprise deployment. NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.
AO-VA Cluster with Collectors	x	x	x	x			This deployment adds collectors to the mix and is the most comprehensive enterprise deployment
AO-VA 1 node with additional Tableau Reporting	x				x	x	This is the most basic single site enterprise deployment with added capability for Reporting via Tableau or other Business Intelligence tools
AO-VA 1 node with Collectors and additional Tableau Reporting	x		x		x	x	This is also an enterprise deployment covering multiple sites with added capability for Reporting via Tableau or other Business Intelligence tools. Data collection is simplified by deploying a collector for the satellite sites.
AO-VA Cluster with additional Tableau Reporting	x	x		x	x	x	This is the scalable enterprise deployment with added capability for Reporting via Tableau or other Business Intelligence tools. NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.
AO-VA Cluster with Collectors and additional Tableau Reporting	x	x	x	x	x	x	This deployment adds collectors to the mix and is the most comprehensive enterprise deployment with added capability for Reporting via Tableau or other Business Intelligence tools.

## Service Provides Deployment Scenarios

Scenario	Supervisor Node	Worker Node	Collector Node	NFS Server	Report Server	Visual Analytics Server	Description
AO-SP 1 node	x						This is the most basic single site service provider deployment mostly suitable for Hosting providers. Organizations can be created by splitting up the IP address space.
AO-SP 1 node with Collectors	x		x				This is also a service provider deployment covering multiple sites. Data collection is simplified by deploying a collector for the satellite sites. Organizations can be added by assigning a collector to an organization and/or by splitting up the IP address space.
AO-SP Cluster	x	x		x			This is the scalable service provider deployment suitable for deployments with large compute and storage needs. NFS Server is required in the data sharing architecture between Supervisor and Worker nodes. Organizations can be created by splitting up the IP address space.
AO-SP Cluster with Collectors	x	x	x	x			This deployment adds collectors to the mix and is the most comprehensive service provider deployment. Organizations can be added by assigning a collector to an organization and/or by splitting up the IP address space.
AO-SP 1 node with additional Tableau Reporting	x				x	x	This is the most basic single site service provider deployment with added capability for Reporting via Tableau or other Business Intelligence tools.
AO-SP 1 node with Collectors and additional Tableau Reporting	x		x		x	x	This is also an service provider deployment covering multiple sites with added capability for Reporting via Tableau or other Business Intelligence tools. Data collection is simplified by deploying a collector for the satellite sites.

AO-SP Cluster with additional Tableau Reporting	x	x		x	x	x	This is the scalable service provider deployment with added capability for Reporting via Tableau or other Business Intelligence tools. NFS Server is required in the data sharing architecture between Supervisor and Worker nodes.
AO-SP Cluster with Collectors and additional Tableau Reporting	x	x	x	x	x	x	This deployment adds collectors to the mix and is the most comprehensive service provider deployment with added capability for Reporting via Tableau or other Business Intelligence tools.

## Hardware Requirements

### Hardware Requirements for Supervisor and Worker nodes

The AccelOps All-In-One Virtual Appliance can be installed using either storage configured within the ESX server or NFS storage.

#### Event Data Storage Requirements

The storage requirement shown in the Event Data Storage column is only for the `eventdb` data, but the `/data` partition also includes CMDB backups and queries. You should set the `/data` partition to a larger amount of storage to accommodate for this.

**Note: Starting Release 3.6, the minimum memory requirements for Super and Worker Virtual Appliances have gone up from 8GB to 16GB**

Overall EPS	Quantity	Host SW	Processor	Memory	OS/App and CMDB Storage	Event Data Storage (1 year)
1,500	1	ESXi (4.0 or later preferred)	4 Core 3 GHz, 64 bit	16 GB	200GB (80GB OS/App; 60GB CMDB; 60GB B SVN)	3 TB
4,500	1	ESXi (4.0 or later preferred)	4 Core 3 GHz, 64 bit	16 GB	200GB (80GB OS/App; 60GB CMDB; 60GB B SVN)	8 TB
7,500	1 Super 1 Worker	ESXi (4.0 or later preferred)	Super: 8 Core 3 GHz, 64 bit Worker: 4 Core 3 GHz, 64 bit	Super: 24 GB Worker: 16 GB	Super: 200GB (80GB OS/App; 60GB CMDB; 60GB SVN) Worker: 200GB (80GB OS/App)	12 TB
10,000	1 Super 1 Worker	ESXi (4.0 or later preferred)	Super: 8 Core 3 GHz, 64 bit Worker: 4 Core 3 GHz, 64 bit	Super: 24 GB Worker: 16 GB	Super: 200GB (80GB OS/App; 60GB CMDB; 60GB SVN) Worker: 200GB (80GB OS/App)	17 TB
20,000	1 Super 3 Workers	ESXi (4.0 or later preferred)	Super: 8 Core 3 GHz, 64 bit Worker: 4 Core 3 GHz, 64 bit	Super: 24 GB Worker: 16 GB	Super: 200GB (80GB OS/App; 60GB CMDB; 60GB SVN) Worker: 200GB (80GB OS/App)	34 TB
30,000	1 Super 5 Workers	ESXi (4.0 or later preferred)	Super: 8 Core 3 GHz, 64 bit Worker: 4 Core 3 GHz, 64 bit	Super: 24 GB Worker: 16 GB	Super: 200GB (80GB OS/App; 60GB CMDB; 60GB SVN) Worker: 200GB (80GB OS/App)	50 TB
Higher than 30,000	Consult AccelOps					

### Hardware Requirements for Collector nodes

Model	Quantity	Host SW	Processor	Memory	OS/App Storage
Collector	1	ESX	2 Core 2 GHz, 64 bit	4 GB	40 GB

### Hardware Requirements for Report Server nodes

Model	Quantity	Host SW	Processor	Memory	OS/App Storage	Reports Data Storage (1 year)
Report Server	1	ESX	8 Core 3 GHz, 64 bit	16 GB	200GB (80GB OS/App; 60GB CMDB )	?

## Generic Steps during installation

- Import a Virtual Appliance into a VMware ESX Server
- Edit the VA hardware settings
- Start and Configure the AccelOps VA from the VMware Console
- Configuring the Timezone
- Register the Virtual Appliance and download license

# Configuring NFS Server

- [Step 1: Download and Install the Required Software packages](#)
- [Step 2: Export the Shared Directory](#)

This section details the steps for installing NFS server on CentOS Linux. For other operating systems, follow similar steps.

## Step 1: Download and Install the Required Software packages

1. Log as root into CentOS 6.x
2. Download and Install NFS packages

### Install NFS packages

```
yum install nfs-utils nfs-utils-lib
```

3. Run following startup scripts for the NFS server to start

### Start NFS scripts

```
chkconfig nfs on  
service rpcbind start  
service nfs start
```

4. Check NFS service status, make sure 'nfsd' service is running.

### NFS Status

```
service nfs status
```

## Step 2: Export the Shared Directory

1. Create new Directory in the large Volume to share with the AccelOps clients (Supervisor and Workers) and change the access permissions to everyone.

### Create Shared Directory

```
mkdir /accelops  
chmod -R 777 /accelops
```

2. Edit /etc/exports file to share /accelops directory to AccelOps clients (Supervisor and Workers).



### Export Shared Directory

```
vi /etc/exports

/accelops <Supervisor_IP_Address>(rw, sync, no_root_squash)
/accelops <Worker1_IP_Address>(rw, sync, no_root_squash)
/accelops <Worker2_IP_Address>(rw, sync, no_root_squash)
```

3. Save /etc/exports changes and restart NFS server

### Restart NFS Server

```
service nfs restart
```

4. Check Shared directories

### Check Shared directories

```
showmount -e localhost
```

Example:

Export list for localhost:

```
/accelops <Supervisor_IP_Address>, <Worker1_IP_Address>, <Worker2_IP_Address>
```

# Installing in VMware ESX

- Prerequisites for Installing AccelOps on ESX
- Setting Network Time Protocol (NTP) for ESX
- Installing a Supervisor node
  - Section 1: Import a Supervisor image into VMware ESX server
  - Section 2: Edit the Supervisor hardware settings
  - Section 3: Select mount point to Store Events data
  - Section 4: Configure the Supervisor from the VMware Console
  - Section 5: Register the Supervisor
  - Section 6: Troubleshoot Supervisor installation
- Installing a Worker node
  - Section 1: Import a Worker image into VMware ESX server
  - Section 2: Edit the Worker hardware settings
  - Section 3: Select mount point to Store Events data
  - Section 4: Configure the Worker from the VMware Console
  - Section 5: Register the Worker to Supervisor
- Installing a Collector node
  - Section 1: Import a Collector image into VMware ESX server
  - Section 2: Edit the Collector hardware settings
  - Section 3: Configure the Collector from the VMware Console
  - Section 4: Register the Collector to Supervisor
- Installing a Report Server node
  - Section 1: Import a Report Server image into VMware ESX server
  - Section 2: Edit the Report Server hardware settings
  - Section 3: Select mount point to Store Reports data
  - Section 4: Configure the Report Server from the VMware Console
  - Section 5: Register the Report Server to Supervisor

## Prerequisites for Installing AccelOps on ESX

The following are required prior to importing the AccelOps Virtual Appliance:

- Static IP address
  - Determine the static IP address and subnet mask for your virtual appliance.
- Enter the AccelOps host name within your local DNS server
- Proxy server information (IP address and port number)
  - The authenticated proxy server is not supported in this version of AccelOps. Turn off authentication on proxy server for this host or completely disable the proxy for this host.
- Determine the VMWare ESX Datastore location where the Virtual Appliance image will be stored
- Determine the IP address of NFS mount point and NFS share name

## Setting Network Time Protocol (NTP) for ESX

It is important that your Virtual Appliance has the accurate time in order to correlate events from multiple devices within the environment.

1. Log in to your VMWare ESX server.
2. Select your ESX host server.
3. Click the **Configuration** tab.
4. Under **Software**, select **Time Configuration**.
5. Click **Properties**.
6. Select **NTP Client Enabled**.
7. Click **Options**.
8. Under General, select **Start automatically**.
9. Under **NTP Setting**, click **Add...**
10. Enter the IP address of the NTP servers to use.

If you don't have an internal NTP server, you can access a publicly available one at <http://tf.nist.gov/tf-cgi/servers.cgi>

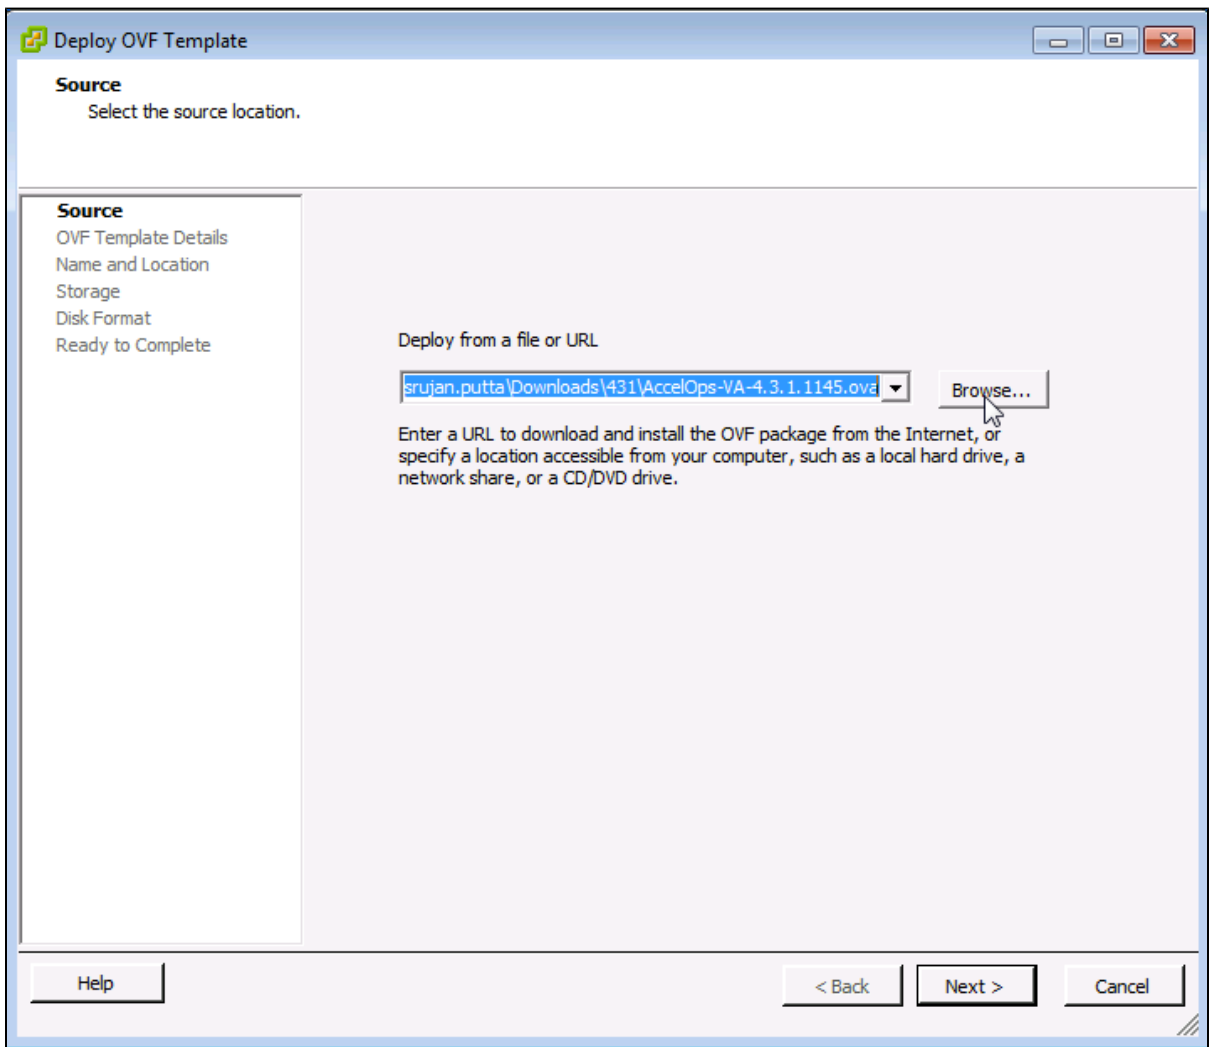
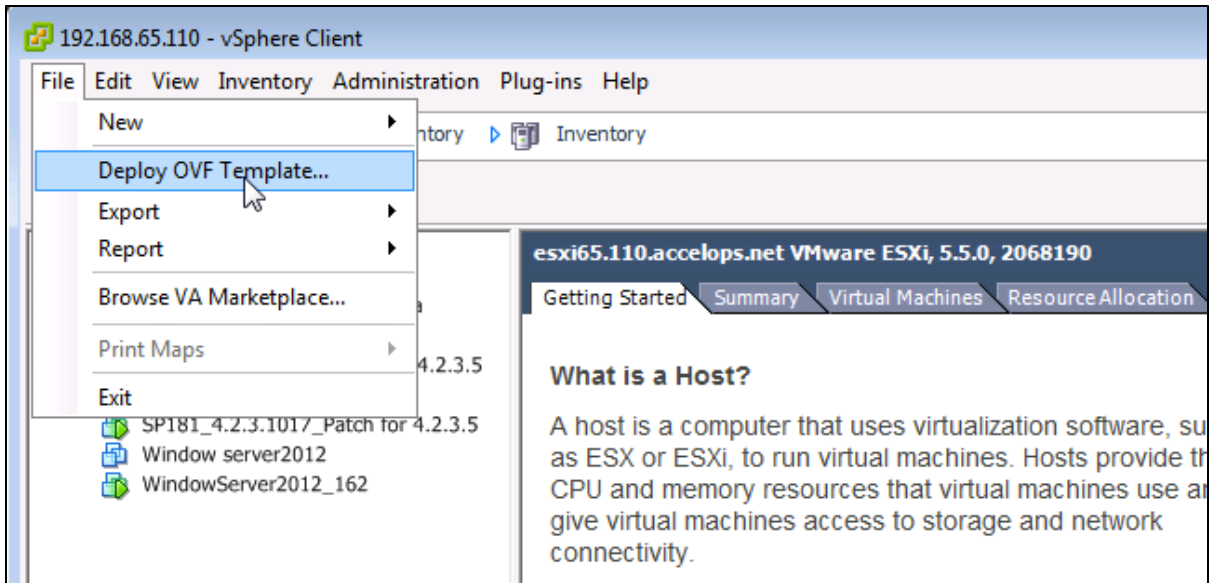
11. Click **Restart NTP service**.
12. Click **OK** to apply the changes.

## **Installing a Supervisor node**

### **Section 1: Import a Supervisor image into VMware ESX server**

#### ***Step 1: Begin importing Supervisor***

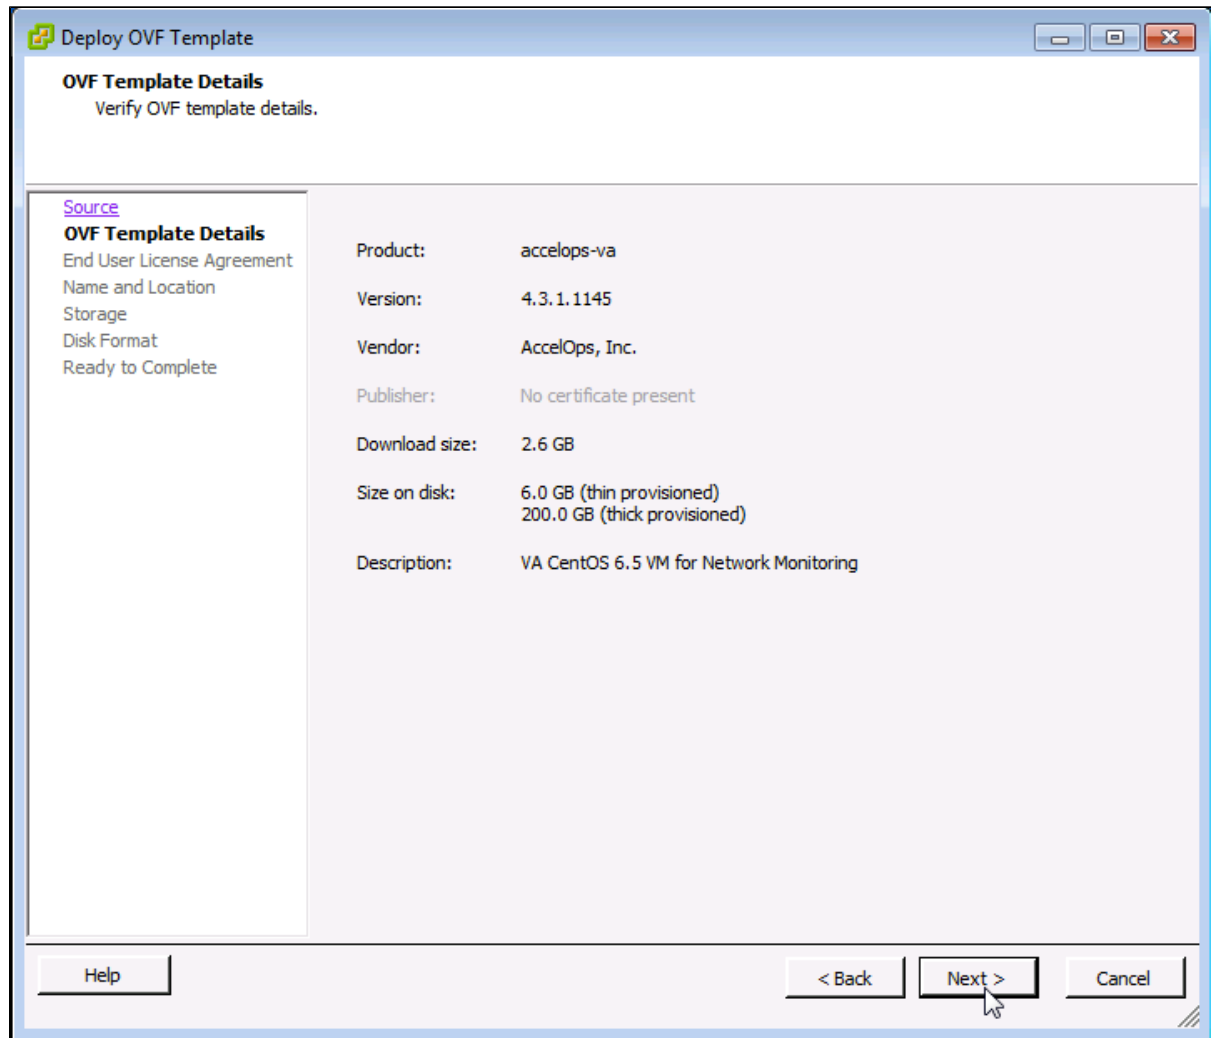
1. Download Supervisor OVA package from AccelOps image server
2. Log into VMware vSphere Client
  - a. Click on File -> Select Deploy OVF Template
  - b. Click on Browse button to choose OVA file
  - c. Locate Supervisor .ova file (Example: AccelOps-VA-4.3.1.1145.ova)



**Step 2: Review the Virtual Appliance details**

OVF Template details will tell you how big the file is to import and how much disk space it will take up in your selected ESX server.

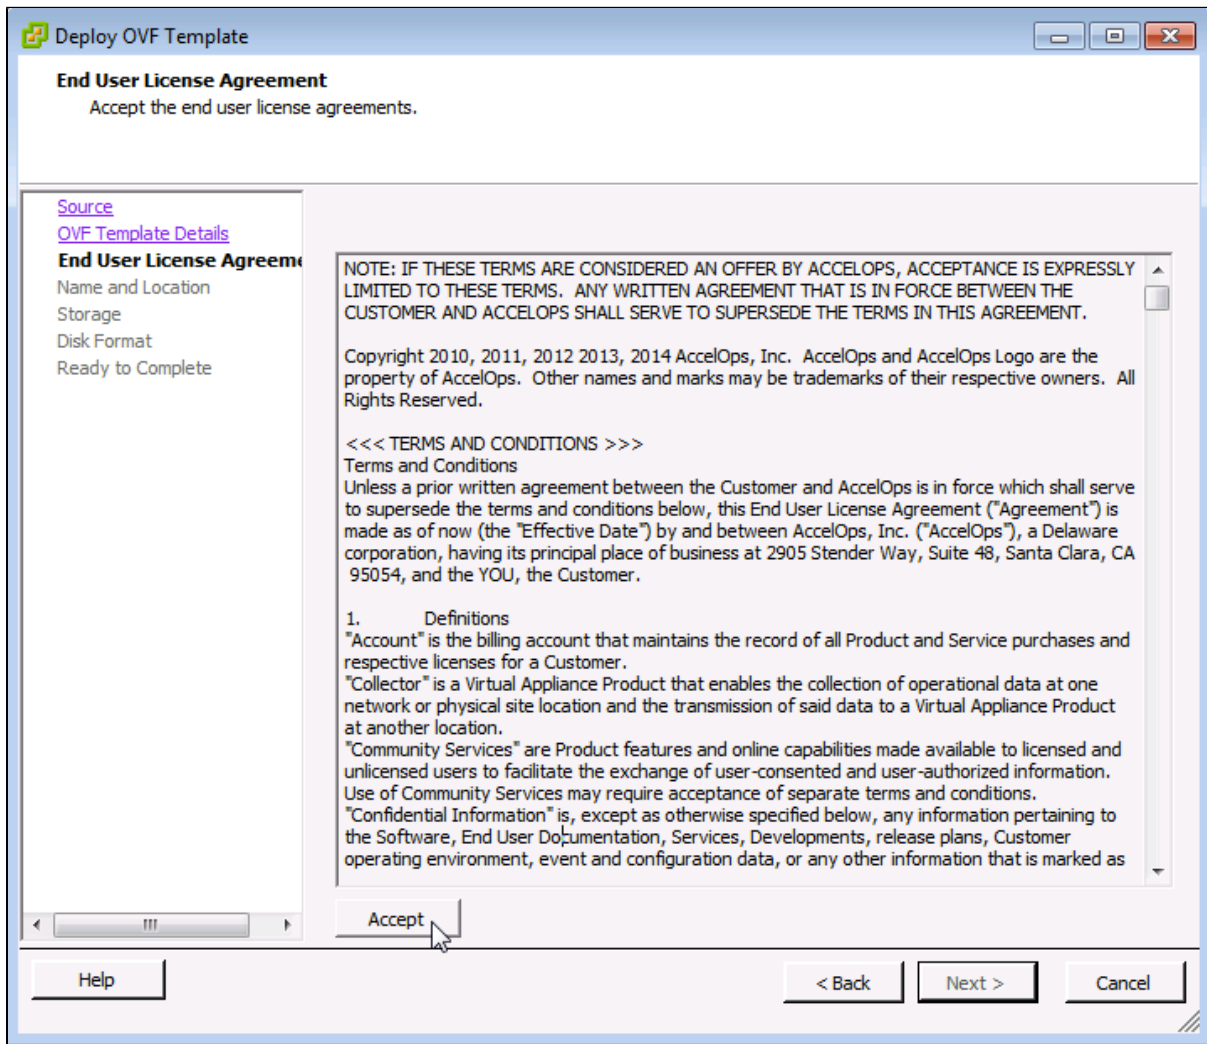
1. Click the 'Next' button



### Step 3: End User License Agreement

To continue you must accept the End User License Agreement.

1. Click on 'Accept' button on License Agreement
2. Click on 'Next' button



**Step 4: Name and Location**

Enter a name for the Virtual Appliance that is being imported. This name is independent of the host name and is only visible within the ESX environment.

1. Enter a Supervisor name and Click on 'Next' button

**Deploy OVF Template**

**Name and Location**  
Specify a name and location for the deployed template

[Source](#)  
[OVF Template Details](#)  
[End User License Agreement](#)  
**Name and Location**  
Storage  
Disk Format  
Ready to Complete

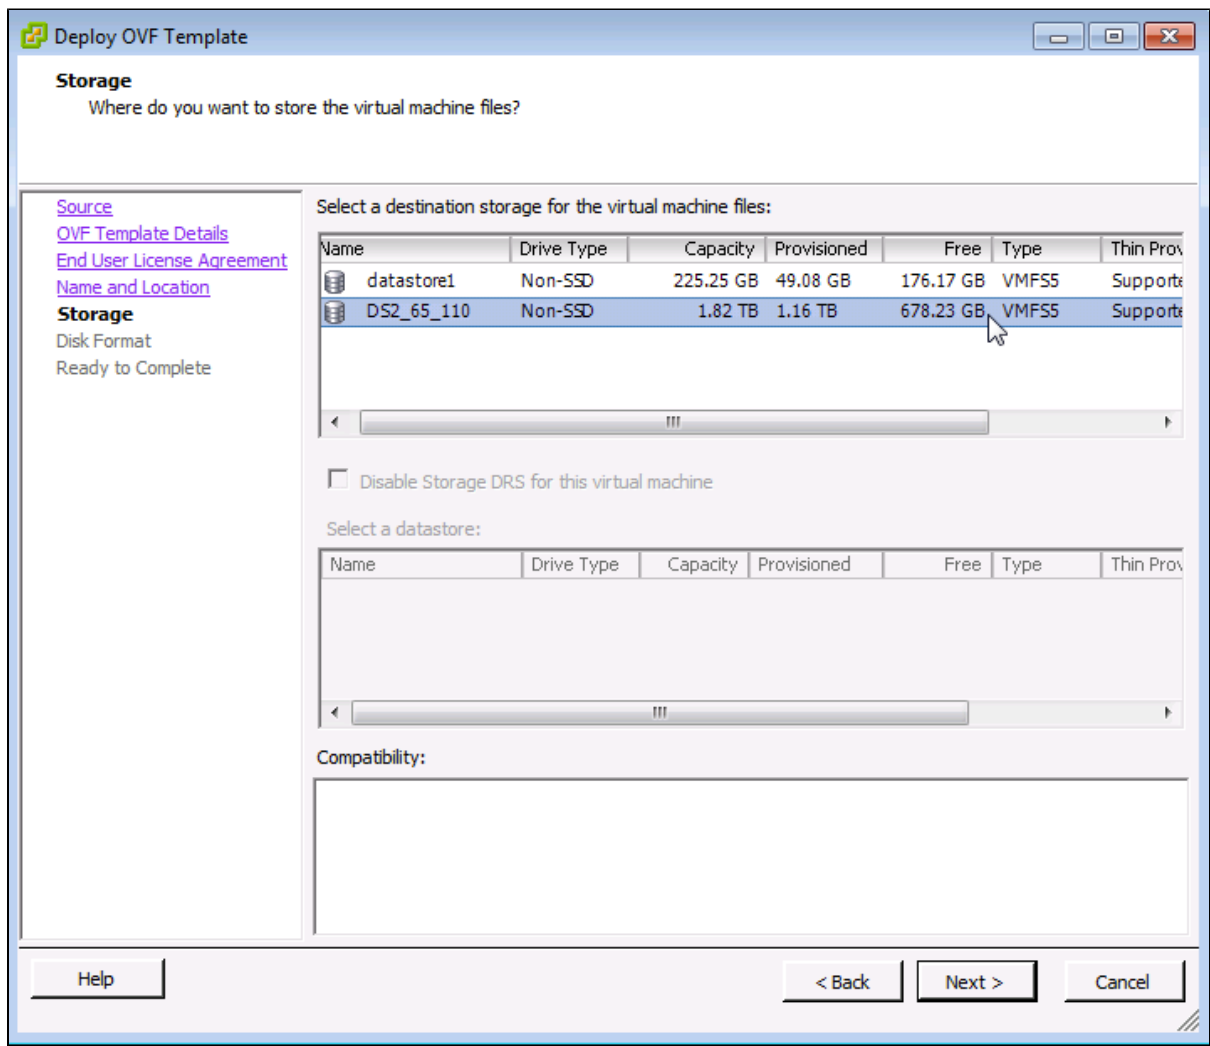
Name:  
AccelOps-Supervisor

The name can contain up to 80 characters and it must be unique within the inventory folder.

Help      < Back      Next >      Cancel

### Step 5: Storage

1. Select the desired storage location(Data store) and Click on 'Next' button



**Step 6: Disk Format**

1. Select the Disk format as 'Thick Provision Lazy Zeroed' (recommended) and Click on 'Next' button



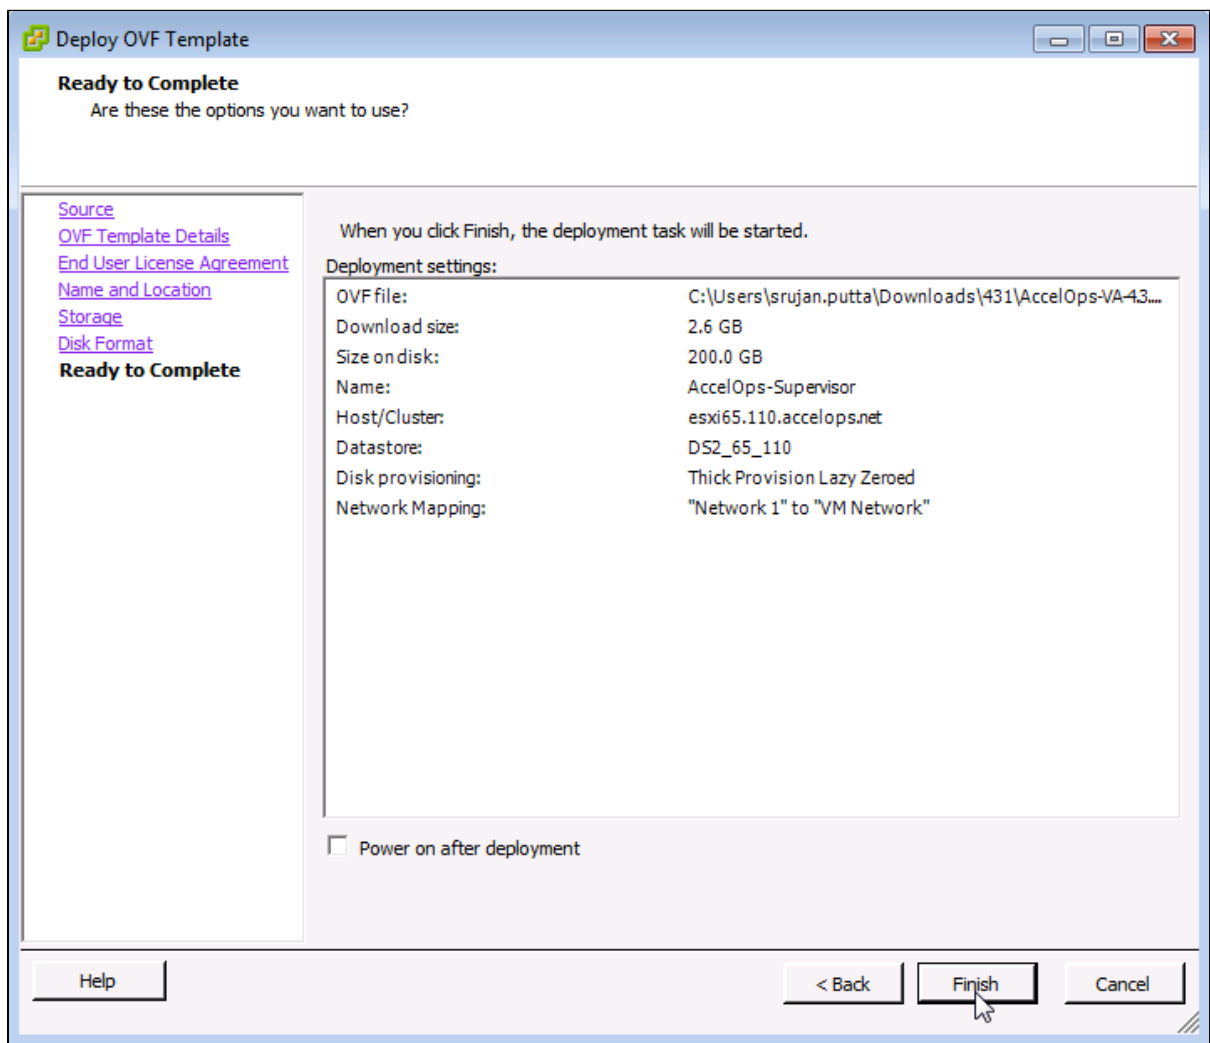
The screenshot shows a window titled "Deploy OVF Template" with a sub-header "Disk Format". Below the sub-header is the question "In which format do you want to store the virtual disks?". On the left side, there is a navigation pane with links for "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Storage", and "Disk Format". The "Disk Format" section is currently selected and shows "Ready to Complete". The main area contains the following fields and options:

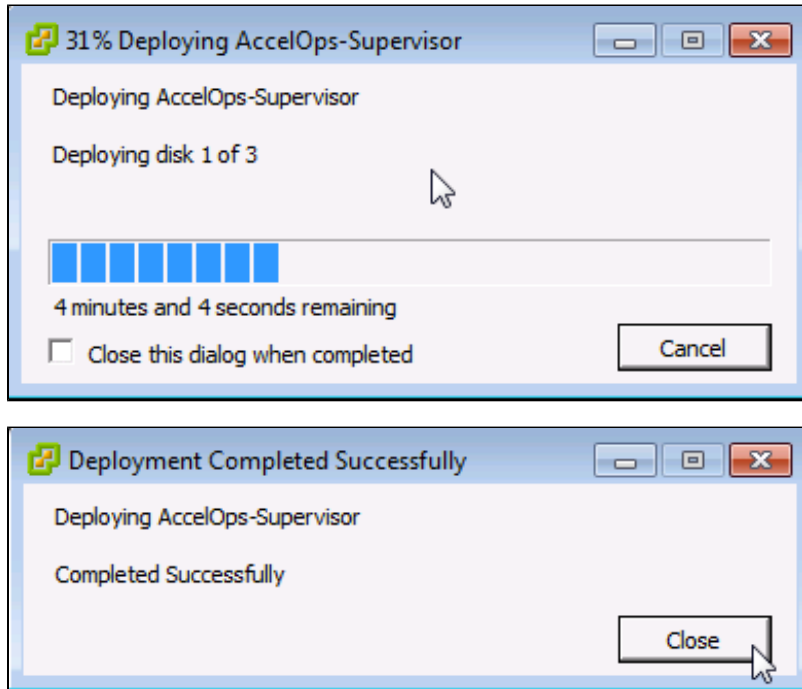
- Datastore:** A text box containing "DS2\_65\_110".
- Available space (GB):** A text box containing "678.2".
- Provisioning Options:** Three radio buttons are present:
  - Thick Provision Lazy Zeroed
  - Thick Provision Eager Zeroed
  - Thin Provision

At the bottom of the window, there are three buttons: "Help", "< Back", and "Next >", with a mouse cursor pointing at the "Next >" button. A "Cancel" button is also present to the right of the "Next >" button.

### Step 7: Ready to Complete

1. Review the Supervisor resources and Click on 'Finish' button to begin importing Virtual Appliance
2. Virtual Appliance import will takes 7 to 10mins to complete. Do not turn off or reboot the system during this time.



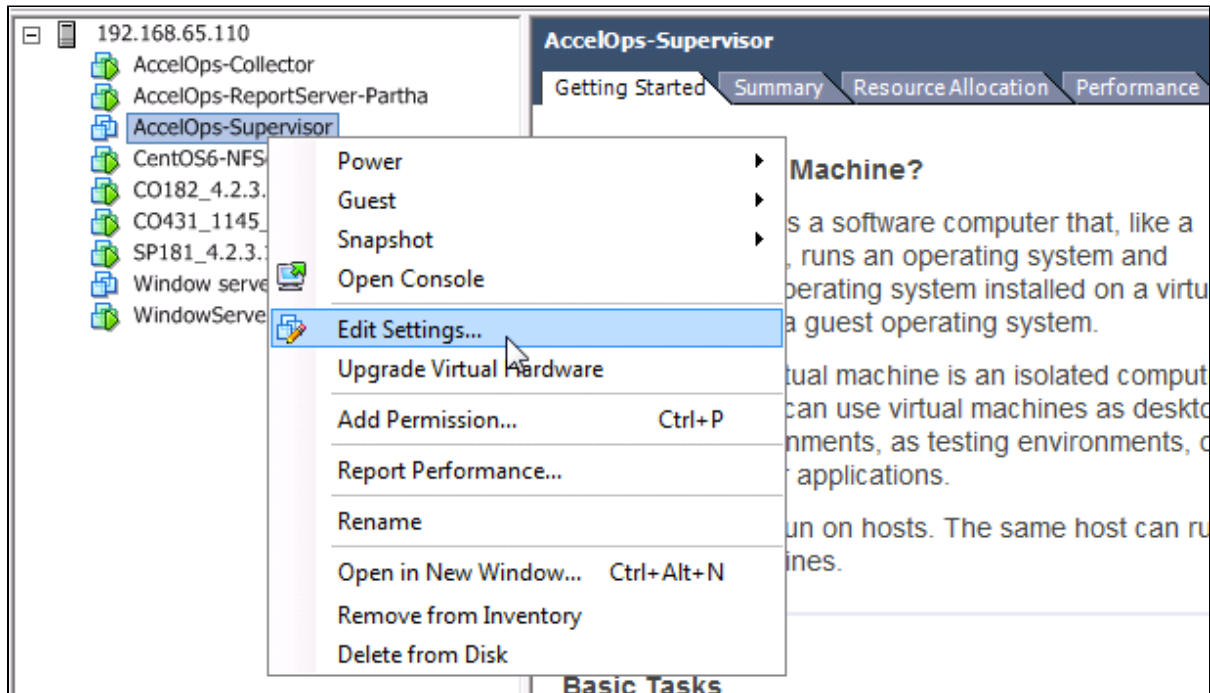


## Section 2: Edit the Supervisor hardware settings

Prior to starting the Supervisor for the first time it is required to make some modifications to the hardware setting for the Virtual Appliance.

### Step 1: Edit Settings

1. From the VMware vSphere client, Select the imported Supervisor Virtual appliance
2. Right mouse click and select Edit Settings



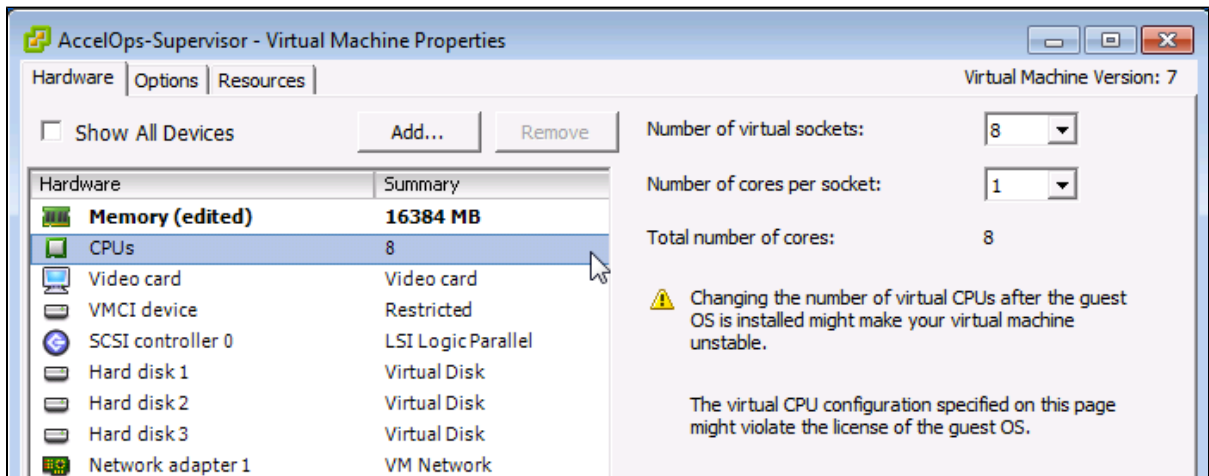
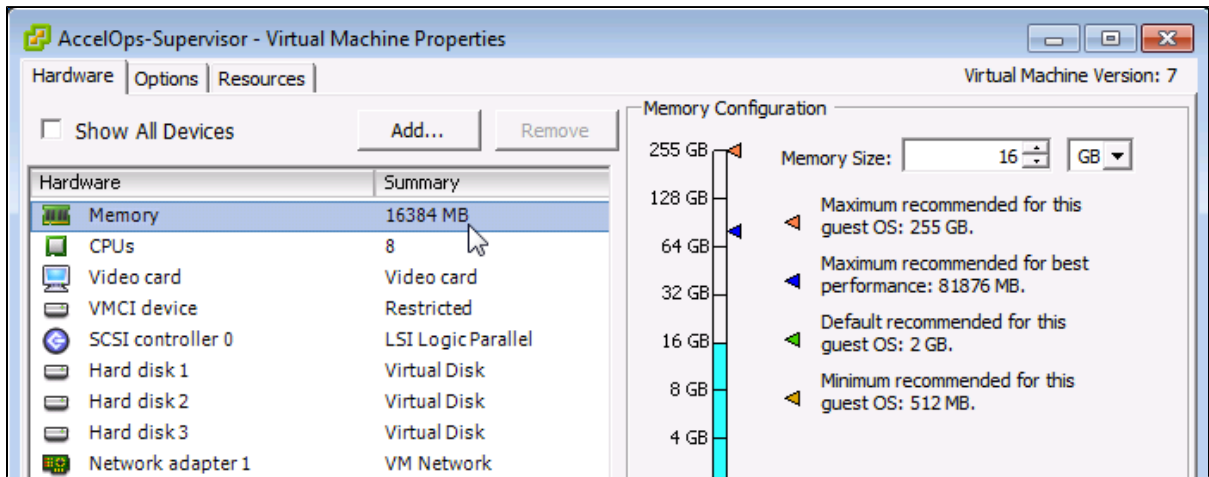
**Step 2: Verify Allocated Memory**

1. On Hardware tab, Click on Memory and verify the memory to at least 16 GB

**Allocate more memory for large deployments**

For large deployments you should allocate at least 24GB of memory. See the section on Supervisor Hardware requirements in the topic [Must Read Prior to Install](#) for more information.

2. Click on CPUs, verify the CPU cores to at least 8 cores



**Section 3: Select mount point to Store Events data**

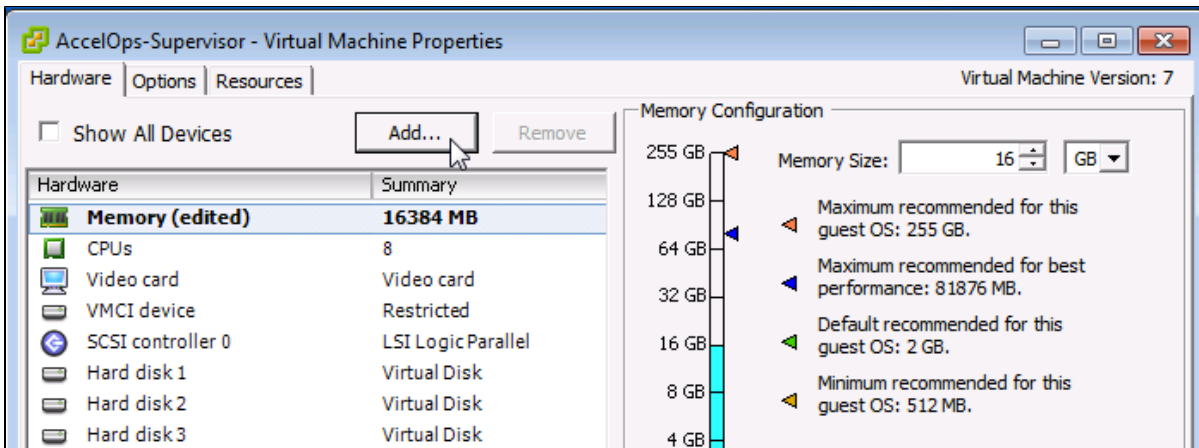
The AccelOps Supervisor can be installed using either storage configured within the ESX server or NFS storage.

**Note**

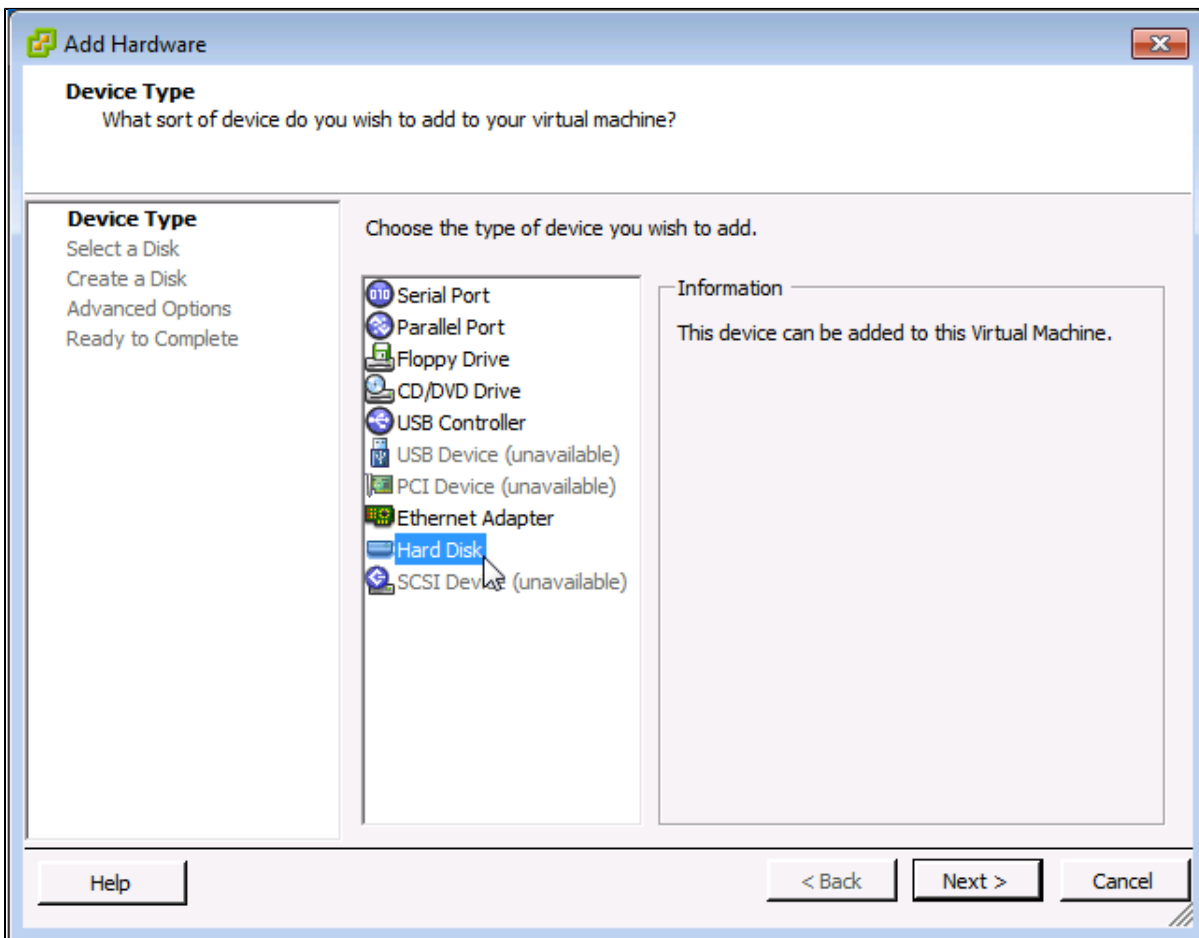
Skip Step 1 if you decided to use NFS storage to store the events data. Else, continue Step 1 to create local disk to store the events data.

**Step 1: Adding an Additional Hard Disk to store the Events data**

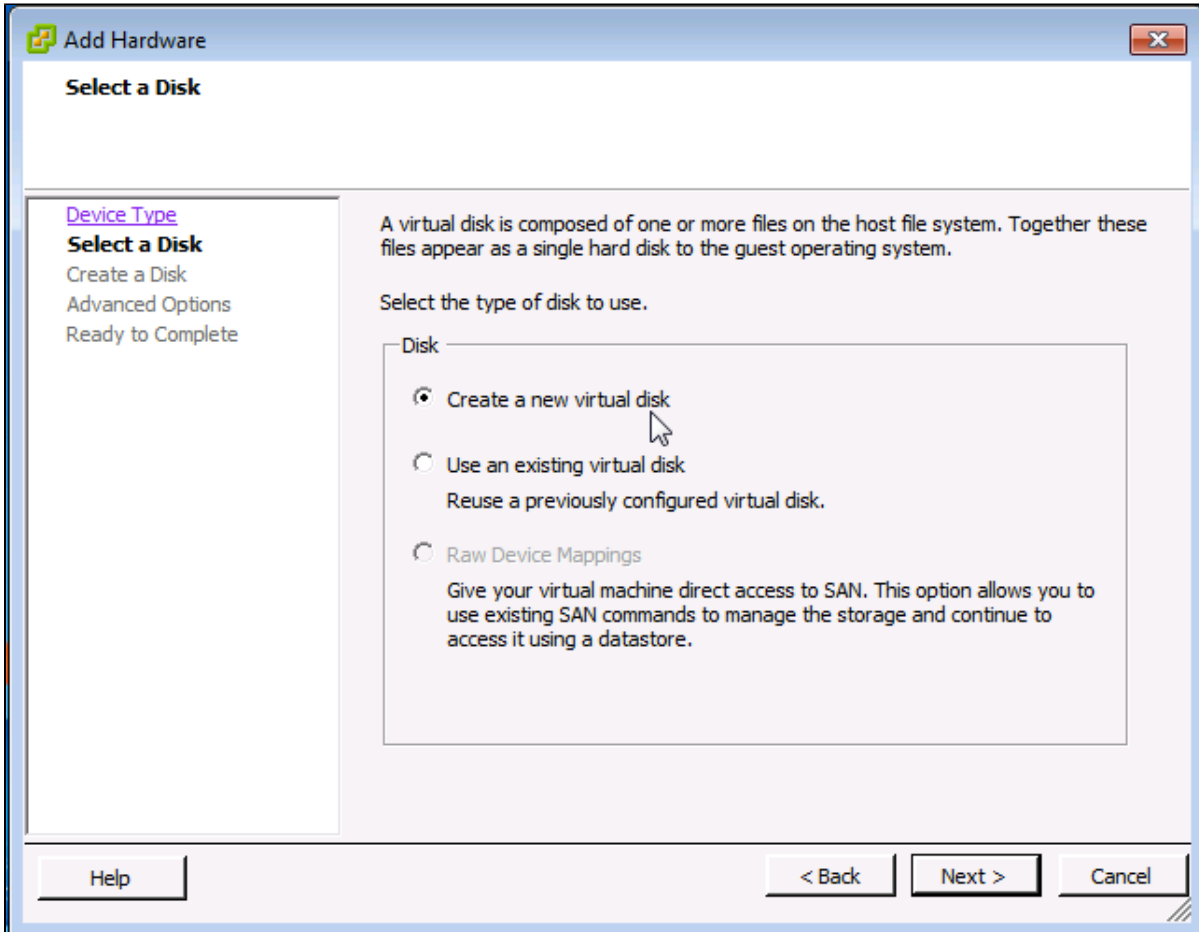
1. On Hardware tab, Click on the 'Add' button



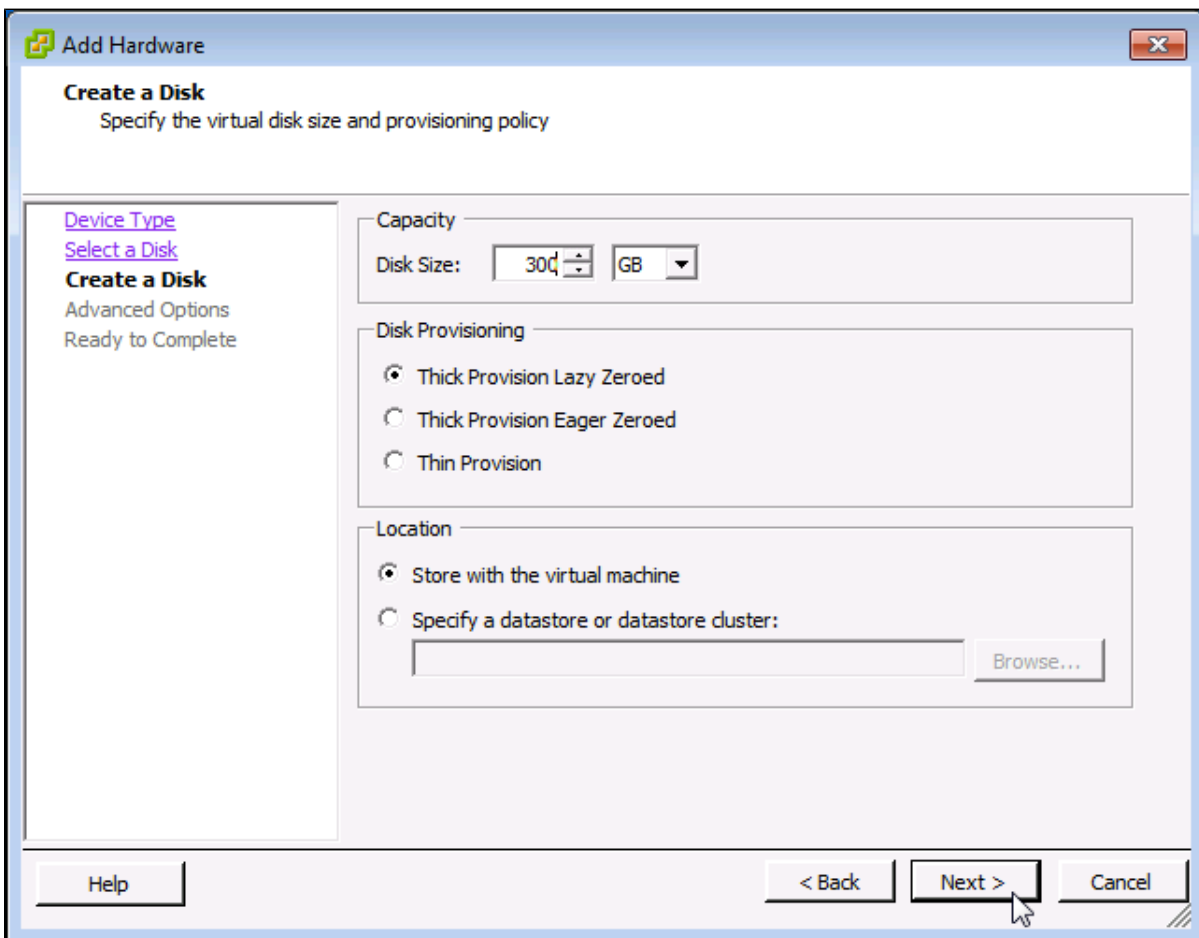
2. On Hardware pop-up, Select 'Hard Disk' and click on 'Next' button



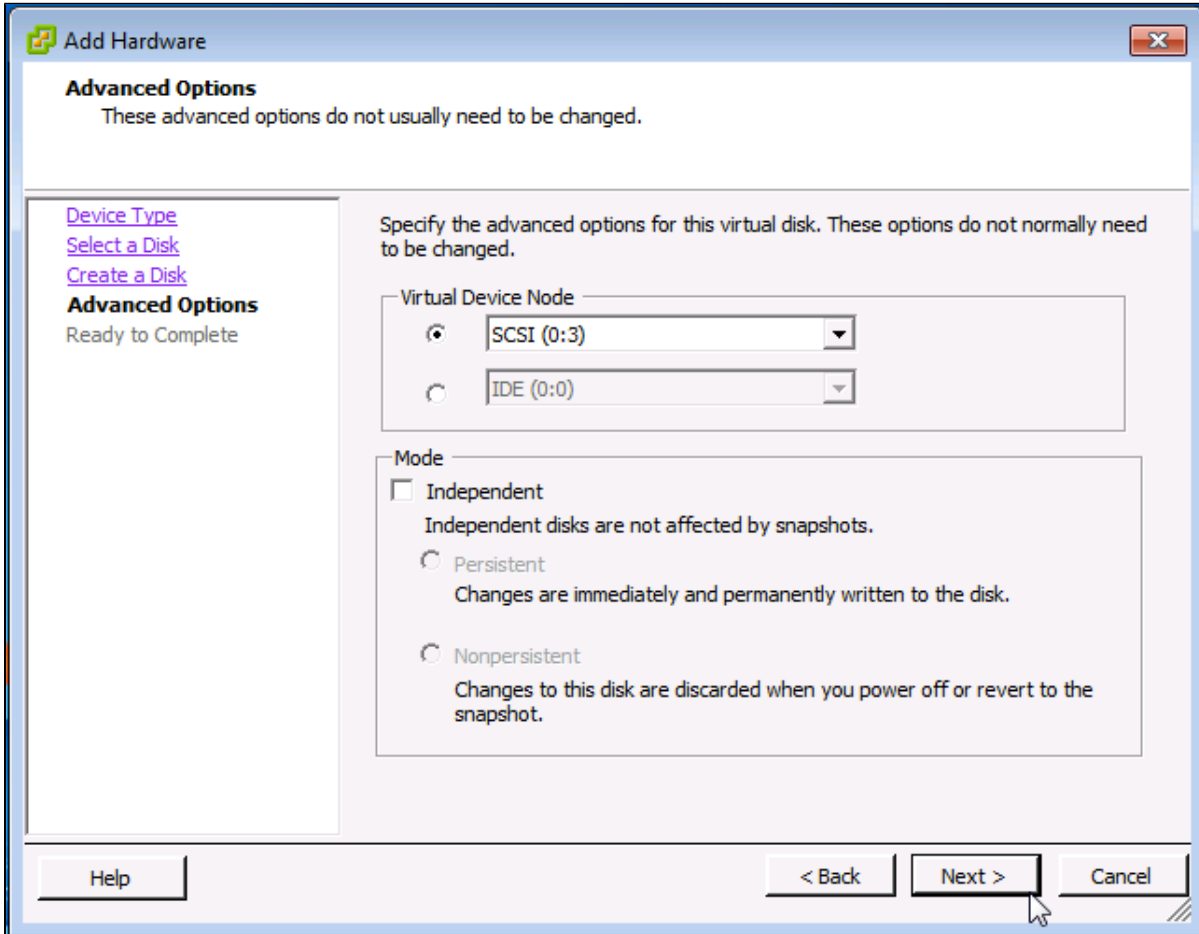
3. Select 'Create a new virtual disk' and click on 'Next' button



4. Enter the Disk size and click on 'Next' button  
Refer to Hardware Requirements for Report Server nodes

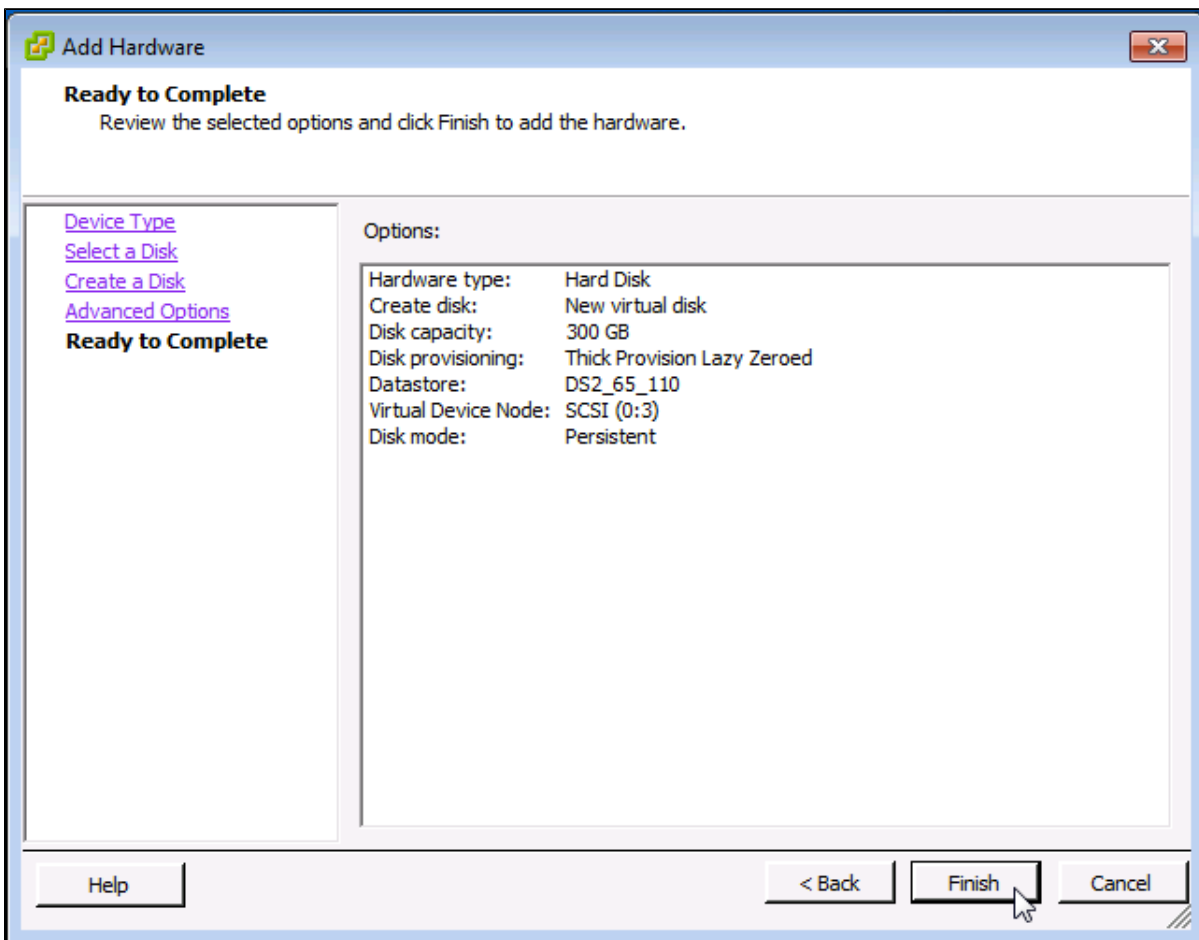


5. Ensure that 'Independent' is not checked and click on 'Next' button

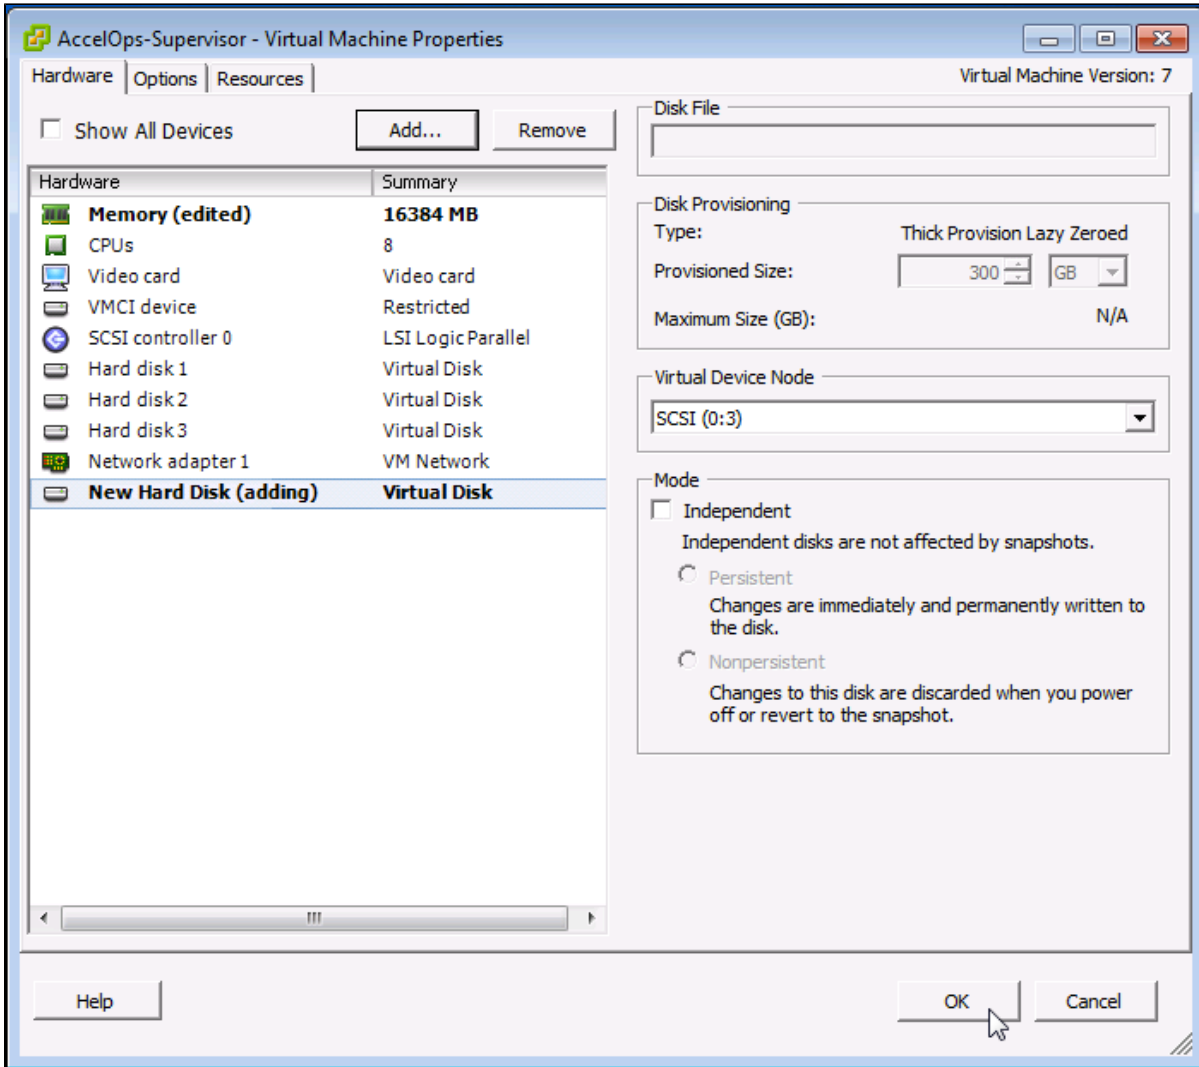


6. Confirm the Options are correct and click the 'Finish' button






7. Click 'OK' on Virtual Machine properties pop-up



8. Wait for until Adding Disk task complete on vSphere client



Note

Skip Step 2 if you choose Step 1 to store the AccelOps event data in local Hard disk. Else, continue Step 2 to create NFS mount point to store the events data.

**Step 2: Create new NFS mount point to store the Events data**

1. Log into NFS server, Create new directory to store Events data
2. Refer to [NFS documentation](#) to check NFS directory permission

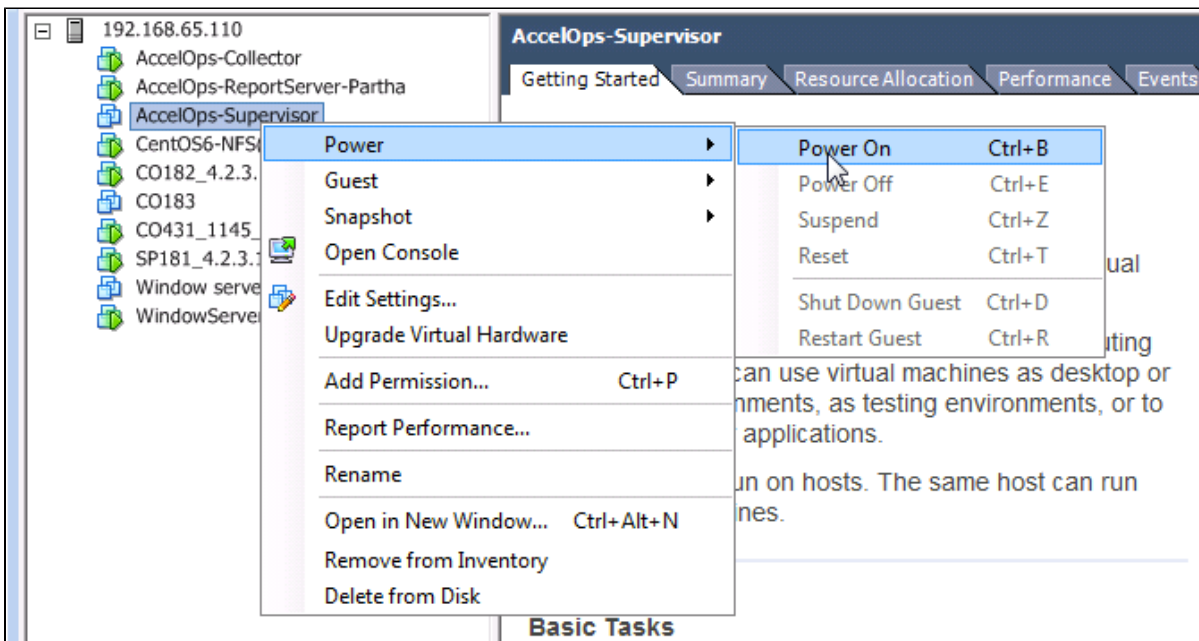
**Section 4: Configure the Supervisor from the VMware Console**

**Note**

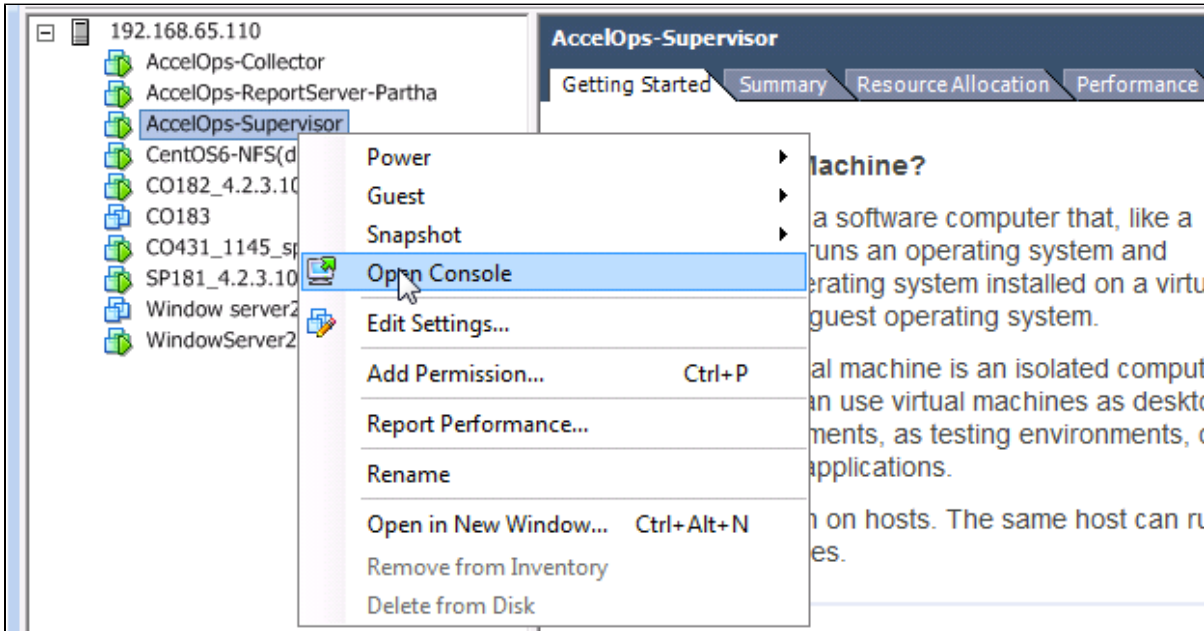
While in the VM console window do not press any control keys during the installation. Doing so might cause the installation to stop and not to precede any further. For example Ctrl – C or Ctrl – Z. If this occurs you must erase the Supervisor and start over from the beginning.

**Step 1: Start Supervisor**

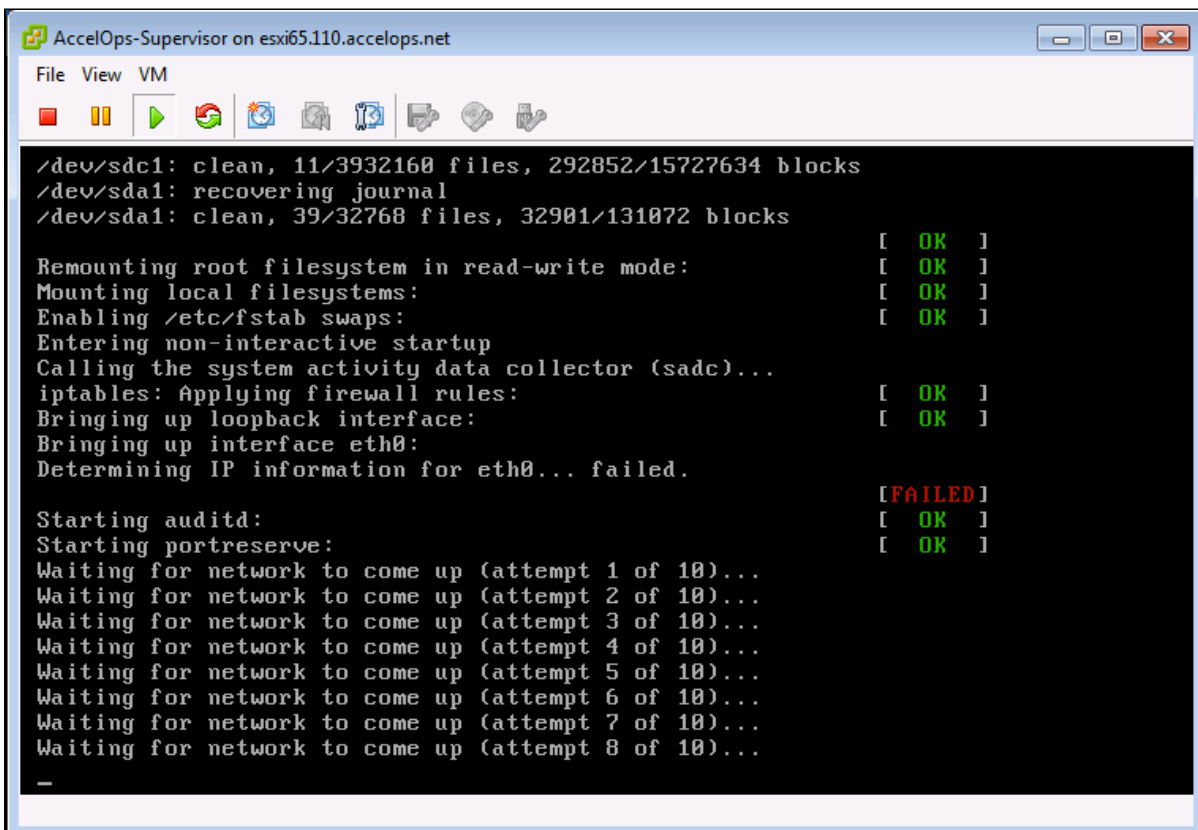
1. From the VMware vSphere client, Select recently imported Supervisor Virtual appliance and Power On (Right click > Power > Power On)



2. Open Supervisor VM Console (Right click > Open Console)



3. Network eth0 Failed message may appear on the screen during the loading, but this is to be expected the first time the AccelOps Supervisor is started.




```
AccelOps-Supervisor on esxi65.110.accelops.net
File View VM
/dev/sdc1: clean, 11/3932160 files, 292852/15727634 blocks
/dev/sda1: recovering journal
/dev/sda1: clean, 39/32768 files, 32901/131072 blocks
[ OK ]
Remounting root filesystem in read-write mode: [ OK ]
Mounting local filesystems: [ OK ]
Enabling /etc/fstab swaps: [ OK ]
Entering non-interactive startup
Calling the system activity data collector (sadc)...
iptables: Applying firewall rules: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0:
Determining IP information for eth0... failed.
[FAILED]
Starting auditd: [ OK ]
Starting portreserve: [ OK ]
Waiting for network to come up (attempt 1 of 10)...
Waiting for network to come up (attempt 2 of 10)...
Waiting for network to come up (attempt 3 of 10)...
Waiting for network to come up (attempt 4 of 10)...
Waiting for network to come up (attempt 5 of 10)...
Waiting for network to come up (attempt 6 of 10)...
Waiting for network to come up (attempt 7 of 10)...
Waiting for network to come up (attempt 8 of 10)...
-
```

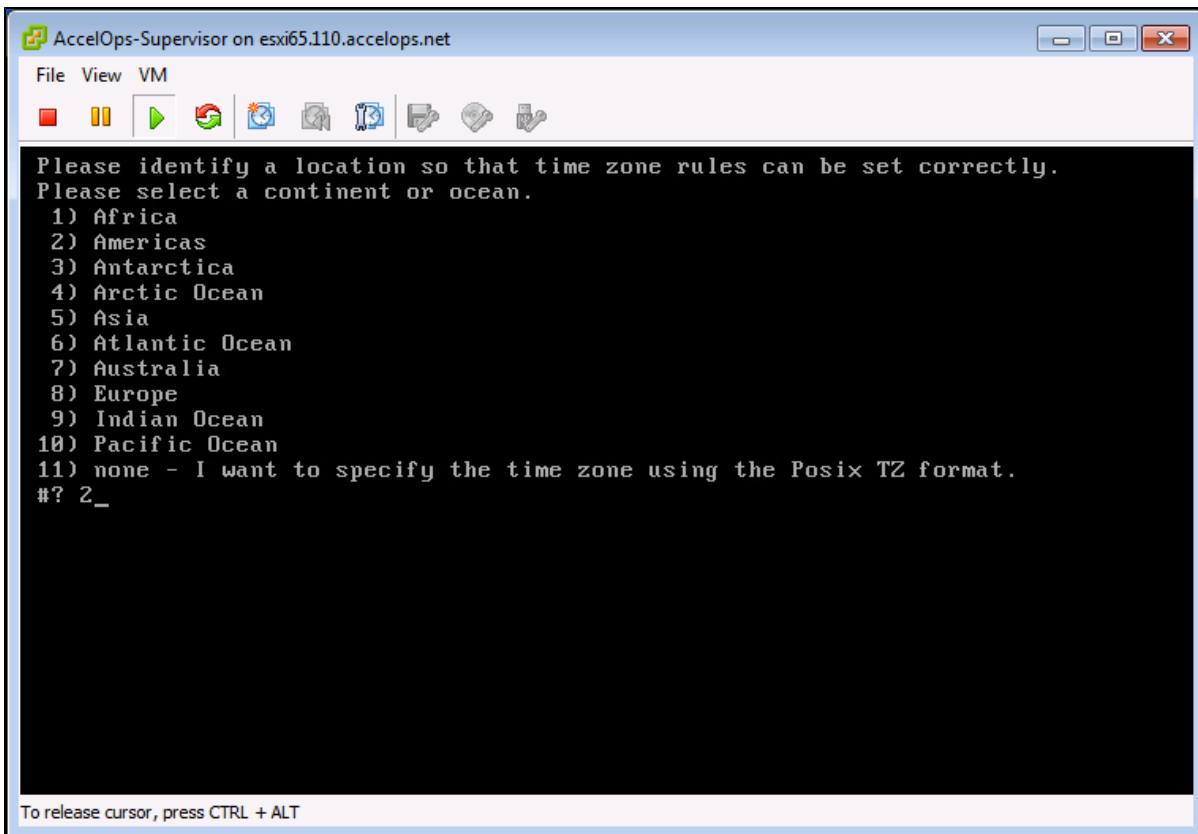
## Step 2: Configure the Timezone

It is important to select the correct time zone that you are located in.

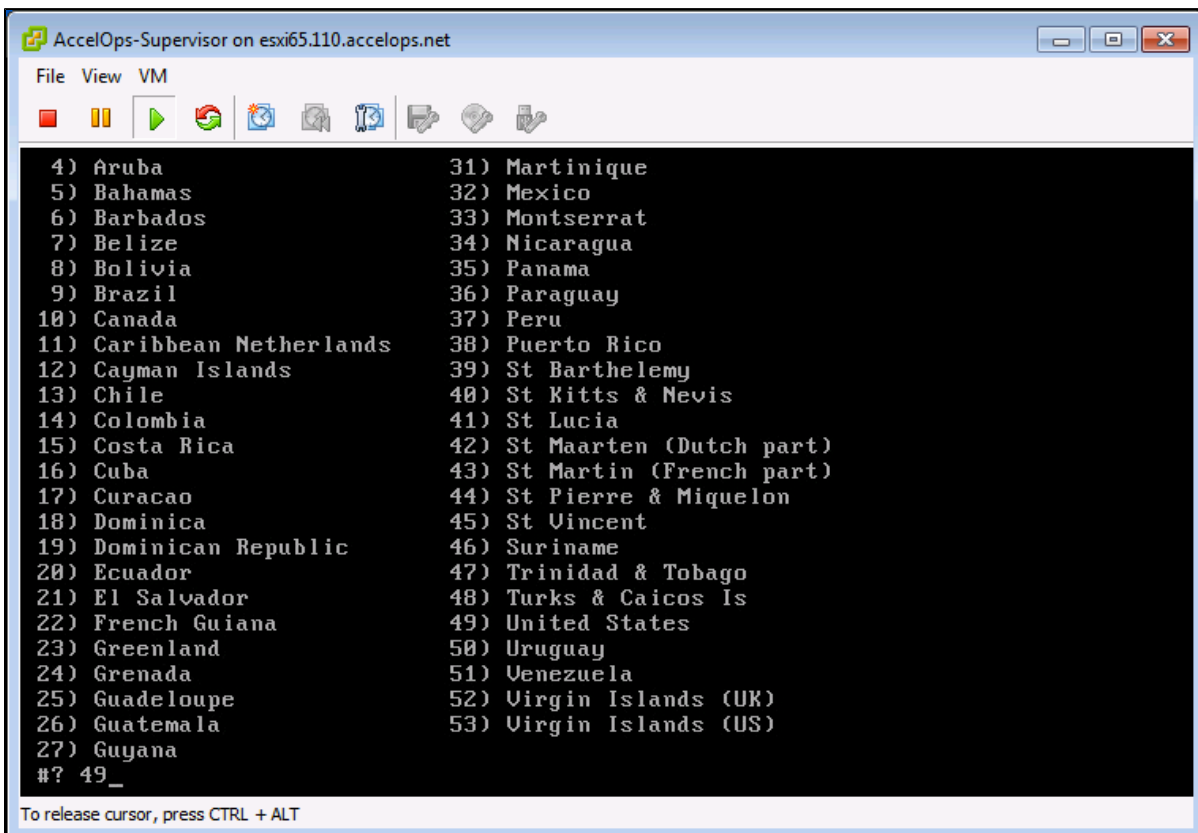
1. On VM console, Select 'Set Timezone' and press the 'Enter' key



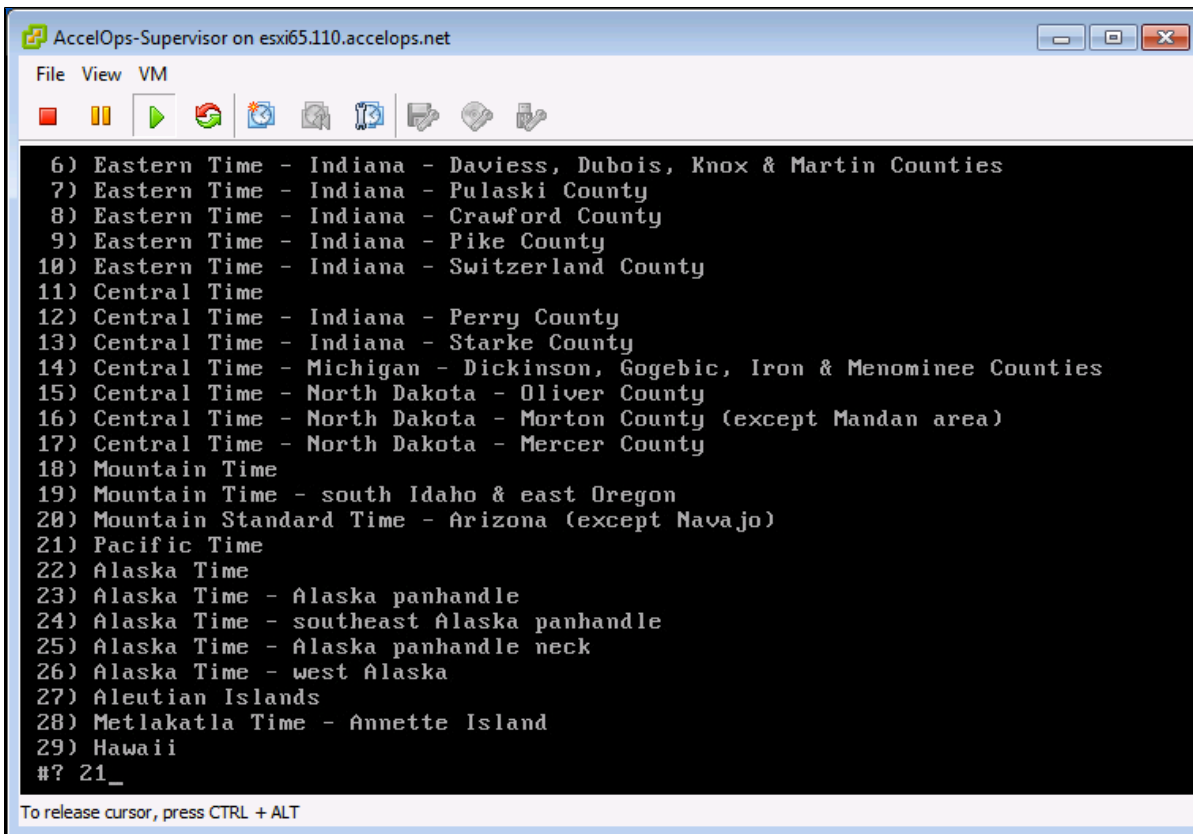
2. Select your Location and press the 'Enter' key



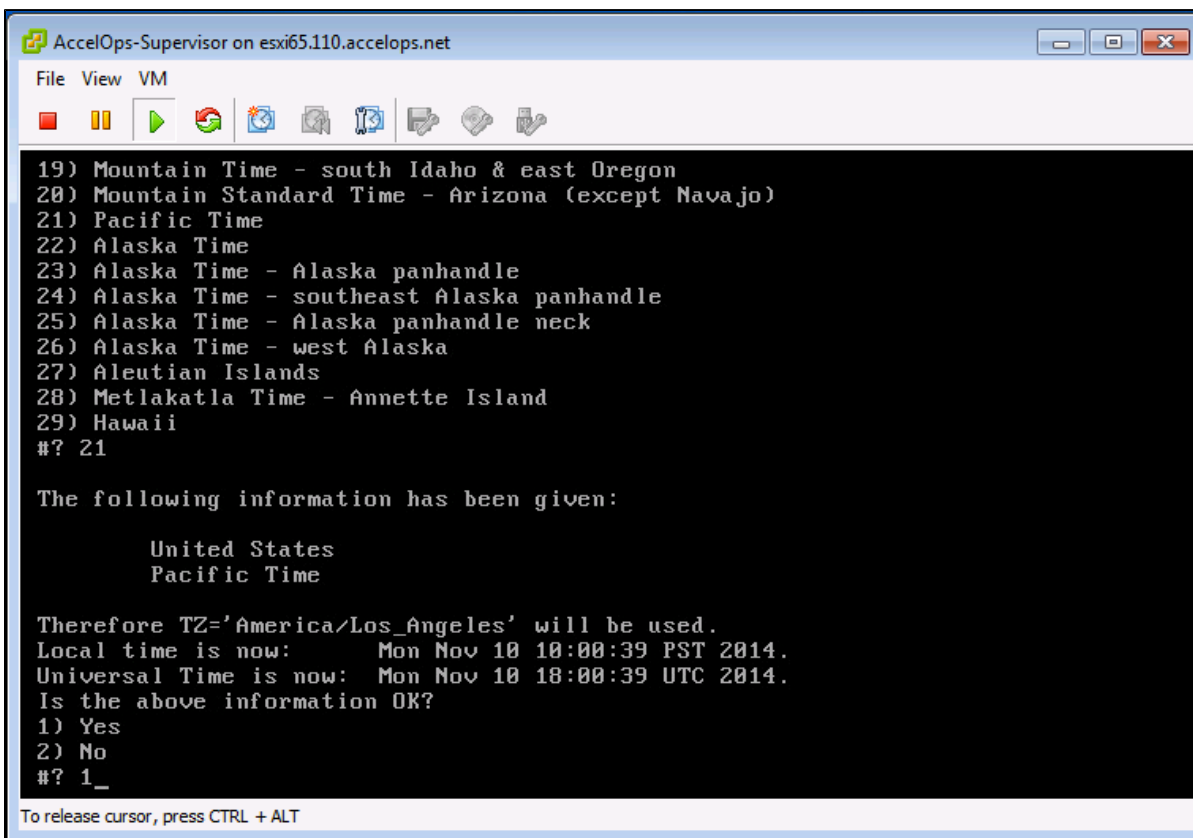
3. Select your Country and press the 'Enter' key



4. Select your Timezone and press the 'Enter' key




5. Review and Confirm your Timezone selection






**Step 3: Configure the Network**

1. On VM console, Select 'Login' and press the 'Enter' key



2. Enter default log in information.

Login: **root**  
 Password: **ProspectHills**



3. Run vami\_config\_net script to configure network and installing AccelOps Supervisor

```

Configure the Network
/opt/vmware/share/vami/vami_config_net
    
```

```
localhost login: root
Password:
[root@localhost ~]# /opt/vmware/share/vami/vami_config_net

CHANGE NETWORK CONFIGURATION for eth0

Please enter the desired network parameters.
To exit, type q at any prompt.

IP Address []: _
```

4. Configure Network with Static IP address
  - a. The following information is required to configure Static IP address:
    - i. IP Address
    - ii. Netmask
    - iii. Gateway
    - iv. DNS Server(s)

#### Proxy Server

Do not configure the proxy server setting. The authenticated proxy server is not supported in this version of AccelOps. Turn off authentication on proxy server for this host or completely disable the proxy for this host.

- b. Press the 'y' key to accept the changes
- c. Enter 'Host name' and press 'Enter' key

```
Gateway []: 192.168.64.1
DNS Server 1 []: 192.168.0.10
DNS Server 2 [192.168.0.10]: 192.168.0.40
Is a proxy server necessary to reach the Internet? y/n [n]:

IP Address:      192.168.65.148
Netmask:         255.255.252.0
Gateway:         192.168.64.1
Proxy Server:
DNS Servers:     192.168.0.10, 192.168.0.40

Is this correct? y/n [y]: y

Reconfiguring the network...
DNS server settings updated
Determining if ip address 192.168.65.148 is already in use for device eth0...
vami_login: no process killed
Network parameters successfully changed to requested values
Input hostname: super.accelops.net
```

#### Step 4: Select Mount Point to store your data

The AccelOps Supervisor can be installed using either storage configured within the ESX server or NFS storage.

Choose only following Step 4.1 or 4.2 based on Section 3

1. Enter your Local hard disk mount point

#### Note

Skip Step 4.1 if you decided use NFS storage to store the Events data. Else, continue below steps.

**Local Storage Mount point**

---

/dev/sdd

```
Mount Point For Your Data: local storage (/dev/sdd) or NFS (ip:/dir) with right
permission to proceed or use Ctrl+C to go back to console.
Mount Point: /dev/sdd
mke2fs 1.41.12 (17-May-2010)
/dev/sdd is entire device, not just one partition!
Proceed anyway? (y,n) y_
```

2. Enter your NFS mount point

**Note**  
Skip Step 4.2 if you choose Step 4.1 to store Events data in local Hard disk. Else, continue below steps.

**NFS Mount point**

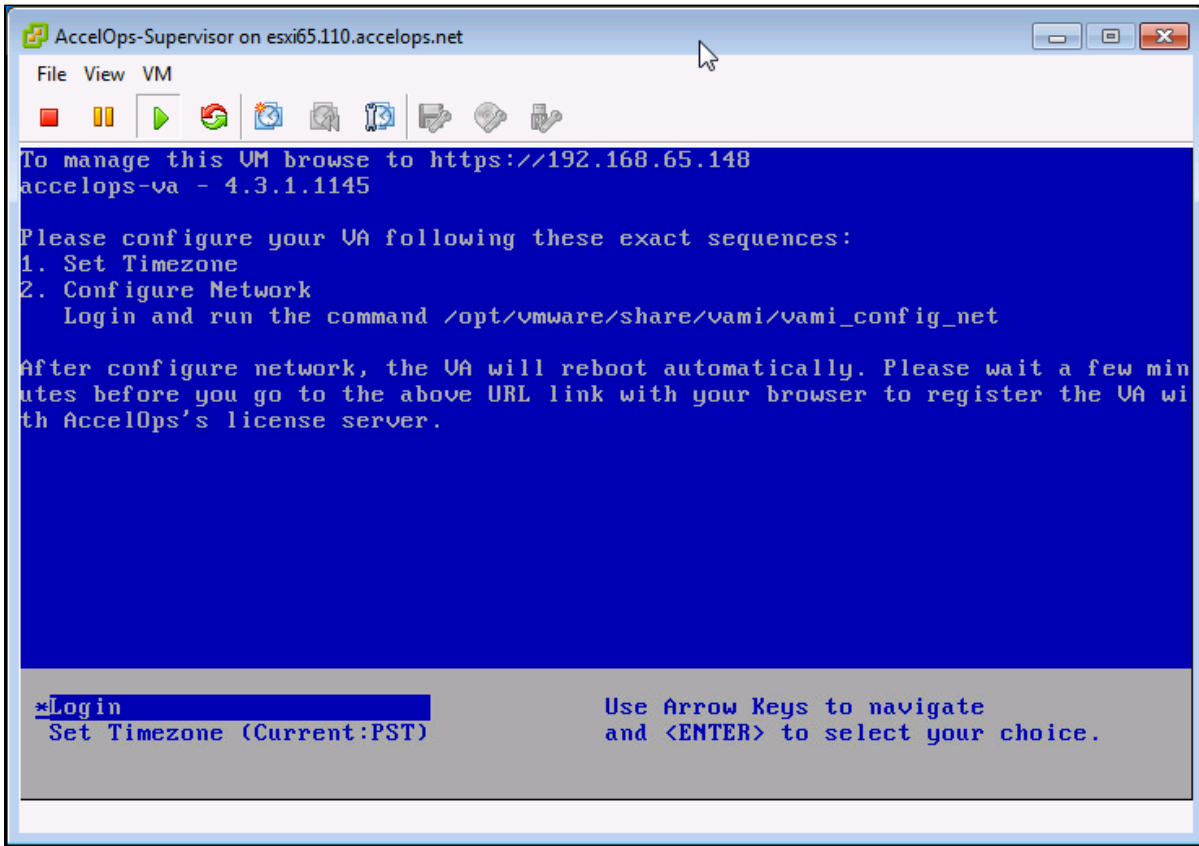
---

<NFS\_Server\_IP\_Address> : /<Directory\_Path>

```
Mount Point For Your Data: local storage (/dev/sdd) or NFS (ip:/dir) with right
permission to proceed or use Ctrl+C to go back to console.
Mount Point: 192.168.67.168:/accelops/SP431
192.168.67.168:/accelops/SP431 mounted
Running initialization of Super host.
Stopping crond: [ OK ]
Running initsys

Committed revision 1.
Install AccelOps db .....
Initialize AccelOps DB .....
Populate AccelOps DB .....
_
```


3. After mount point, Supervisor will reboot automatically  
In 15 to 25mins, Supervisor will be configure successful.



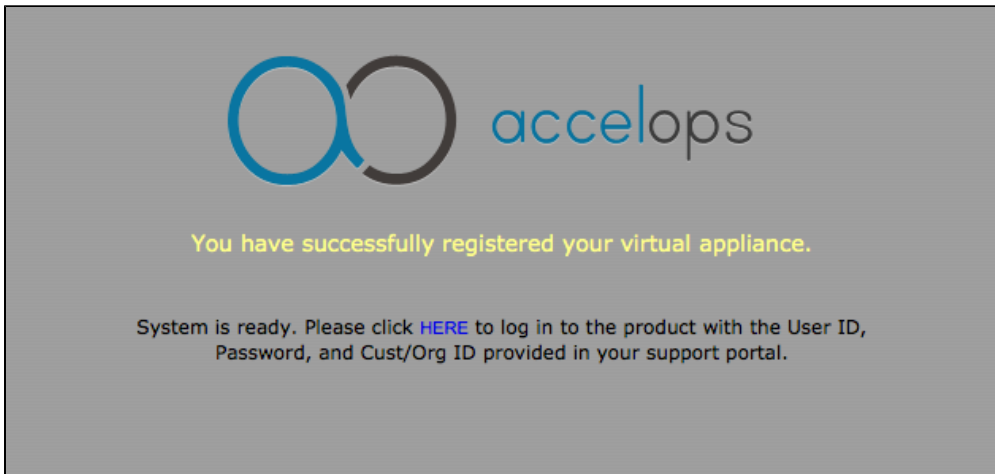
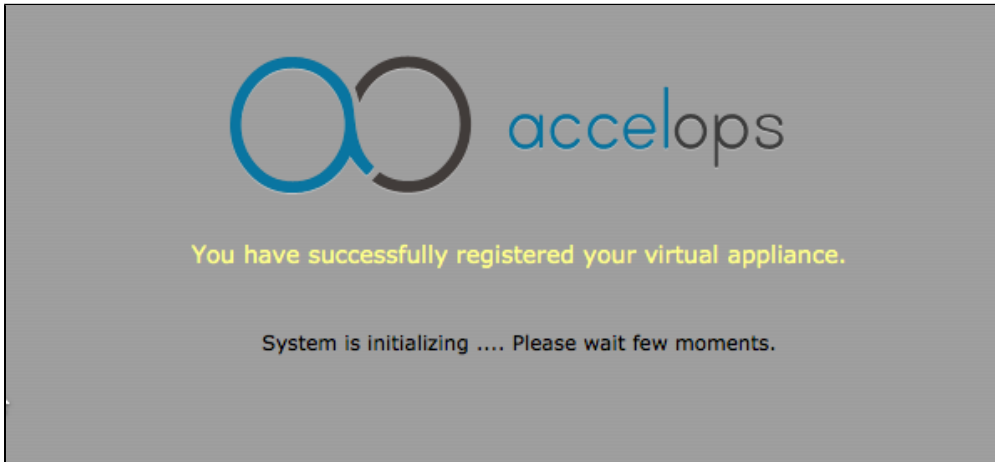
## Section 5: Register the Supervisor

### Step 1: Register AccelOps Supervisor

1. In browser windows, enter 'https://<Supervisor IP>'



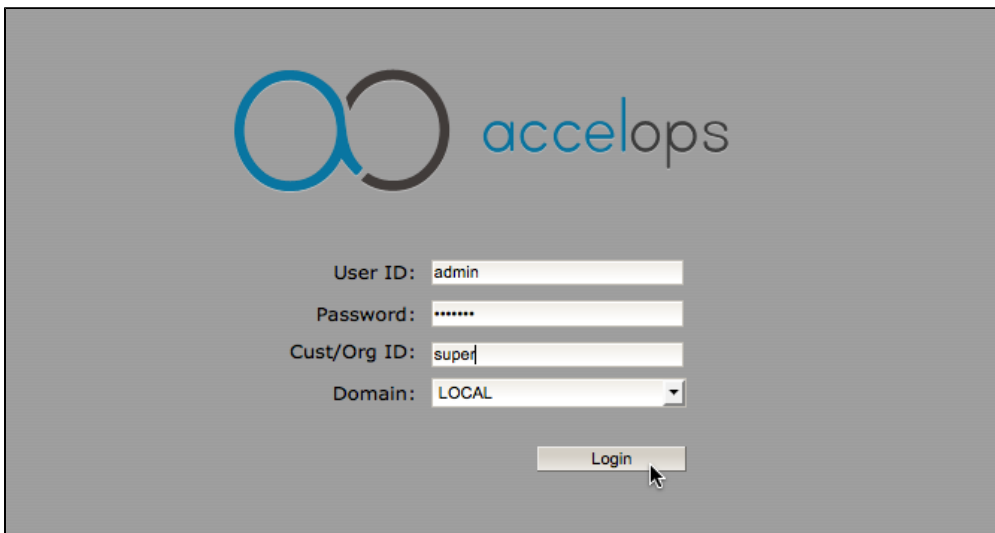
2. Enter license user name and password and click on Register



#### **Offline License Registration**


Contact AccelOps Support if you need to register Supervisor offline.

3. After Registration, Click blue link '[HERE](#)' on AccelOps UI to login
4. Enter the default log in credentials  
User ID: **admin**  
Password: **admin\*1**  
Cust/Org ID: **super**



The image shows the AccelOps login interface. At the top left is the AccelOps logo, consisting of two interlocking circles (one blue, one grey) followed by the text "accelops" in a blue, lowercase sans-serif font. Below the logo are four input fields: "User ID:" with the text "admin", "Password:" with seven dots, "Cust/Org ID:" with the text "super", and "Domain:" with a dropdown menu showing "LOCAL". A "Login" button is positioned below the domain field, with a mouse cursor hovering over it.

5. Go to Admin > Cloud Health and check Supervisor health.

Admin > Cloud Health									
 Refresh    Last Updated: 00:39:16									
Name	IP Address	Health	Module Role	Version	Build Date	Cores	Load Average	CPU	Swap Size
super.accelops.net	192.168.65.148	Normal	supervisor	4.3.1.1145	10:04:55 AM Oct 15 ...	8	1.17,0.71,0.25	2.13	6 GB
super.accelops.net ( 192.168.65.148 ) - supervisor: Process Details									
Process Name	Status	Up Time	CPU	Event Rate	Physical Mem	Virtual Mem	Shared		
phMonitorSupervisor	Up	2m 53s	0%	0/s	73 MB	1177 MB	0		
phParser	Up	1m 51s	0%	0/s	263 MB	1788 MB	99		
phQueryMaster	Up	1m 51s	0%	0/s	58 MB	874 MB	0		
phRuleMaster	Up	1m 51s	0%	0/s	39 MB	551 MB	0		
phCheckpoint	Up	1m 51s	0%	0/s	21 MB	98 MB	0		
phReportLoader	Up	1m 51s	0%	0/s	233 MB	630 MB	0		
phReportWorker	Up	1m 51s	0%	0/s	232 MB	1322 MB	3		
phDataManager	Up	1m 51s	0%	0/s	49 MB	1050 MB	1		
phIdentityMaster	Up	1m 51s	0%	0/s	26 MB	373 MB	0		
phQueryWorker	Up	1m 51s	0%	0/s	234 MB	1263 MB	0		
phAgentManager	Up	1m 51s	0%	0/s	200 MB	881 MB	0		
phRuleWorker	Up	1m 51s	0%	0/s	239 MB	1256 MB	2		
phPerfMonitor	Up	1m 51s	0%	0/s	40 MB	700 MB	0		
phDiscover	Up	1m 51s	0%	0/s	32 MB	373 MB	0		
phIdentityWorker	Up	1m 51s	0%	0/s	36 MB	907 MB	4		
phReportMaster	Up	1m 51s	0%	0/s	36 MB	415 MB	0		
rsyslogd	Up	5h 25m 6s	0%	0/s	3 MB	243 MB	0		
httpd	Up	5h 24m 40s	0%	0/s	9 MB	260 MB	0		
pgsql DB	Up	5h 24m 51s	0%	0/s	17 MB	448 MB	0		
glassfish	Up	5h 24m 33s	3%	0/s	1299 MB	8183 MB	0		

## Section 6: Troubleshoot Supervisor installation

### Step 1: Verify Supervisor system level permissions

1. SSH to Supervisor and Check 'cmdb, data, query, querywkr and svn' permissions

```

System level permissions
[root@super ~]# ls -l /
dr-xr-xr-x.  2 root    root    4096 Oct 15 11:09 bin
dr-xr-xr-x.  5 root    root    1024 Oct 15 14:50 boot
drwxr-xr-x   4 postgres postgres 4096 Nov 10 18:59 cmdb
drwxr-xr-x   9 admin    admin    4096 Nov 11 11:32 data
drwxr-xr-x  15 root    root    3560 Nov 10 11:11 dev
-rw-r--r--   1 root    root      34 Nov 11 12:09 dump.rdb
drwxr-xr-x. 93 root    root   12288 Nov 11 12:12 etc
drwxr-xr-x.  4 root    root    4096 Nov 10 11:08 home
dr-xr-xr-x. 11 root    root    4096 Oct 15 11:13 lib
dr-xr-xr-x.  9 root    root   12288 Nov 10 19:13 lib64
drwx-----. 2 root    root   16384 Oct 15 14:46 lost+found
drwxr-xr-x.  2 root    root    4096 Sep 23 2011 media
drwxr-xr-x.  2 root    root    4096 Sep 23 2011 mnt
drwxr-xr-x. 10 root    root    4096 Nov 10 09:37 opt
drwxr-xr-x.  2 root    root    4096 Nov 10 11:10 pbin
dr-xr-xr-x 289 root    root      0 Nov 10 11:13 proc
drwxr-xr-x.  8 admin    admin    4096 Nov 11 00:37 query
drwxr-xr-x.  8 admin    admin    4096 Nov 10 18:58 querywkr
dr-xr-x---.  7 root    root    4096 Nov 10 19:13 root
dr-xr-xr-x.  2 root    root   12288 Oct 15 11:08 sbin
drwxr-xr-x.  2 root    root    4096 Oct 15 14:47 selinux
drwxr-xr-x.  2 root    root    4096 Sep 23 2011 srv
drwxr-xr-x.  4 apache  apache  4096 Nov 10 18:58 svn
drwxr-xr-x. 13 root    root      0 Nov 10 11:13 sys
drwxrwxrwt.  9 root    root    4096 Nov 11 12:12 tmp
drwxr-xr-x. 15 root    root    4096 Oct 15 14:58 usr
drwxr-xr-x. 21 root    root    4096 Oct 15 11:01 var

```

2. Check '/data , /cmdb and /svn' permissions

```

EventDB, CMDDB and SVN permssions
[root@super ~]# ls -l /data
drwxr-xr-x 3 root    root    4096 Nov 11 02:52 archive
drwxr-xr-x 3 admin    admin    4096 Nov 11 12:01 cache
drwxr-xr-x 2 postgres postgres 4096 Nov 10 18:46 cmdb
drwxr-xr-x 2 admin    admin    4096 Nov 10 19:04 custParser
drwxr-xr-x 5 admin    admin    4096 Nov 11 00:29 eventdb
drwxr-xr-x 2 admin    admin    4096 Nov 10 19:04 jmxXml
drwxr-xr-x 2 admin    admin    4096 Nov 11 11:33 mibXml

[root@super ~]# ls -l /cmdb
drwx----- 14 postgres postgres 4096 Nov 10 11:08 data

[root@super ~]# ls -l /svn
drwxr-xr-x 6 apache  apache  4096 Nov 10 18:58 repos

```



**Step 2: Check Backend system health**

1. SSH to Supervisor, run 'phstatus' command

```

System Status
[root@super ~]# phstatus

Every 1.0s: /opt/phoenix/bin/phstatus.py

System uptime: 12:37:58 up 17:24, 1 user, load average: 0.06, 0.01, 0.00
Tasks: 20 total, 0 running, 20 sleeping, 0 stopped, 0 zombie
Cpu(s): 8 cores, 0.6%us, 0.7%sy, 0.0%ni, 98.6%id, 0.0%wa, 0.0%hi, 0.1%si, 0.0%st
Mem: 16333720k total, 5466488k used, 10867232k free, 139660k buffers
Swap: 6291448k total, 0k used, 6291448k free, 1528488k cached

PROCESS                UPTIME                CPU%                VIRT_MEM            RES_MEM
phParser                12:00:34             0                   1788m               280m
phQueryMaster           12:00:34             0                   944m                63m
phRuleMaster            12:00:34             0                   596m                85m
phRuleWorker            12:00:34             0                   1256m               252m
phQueryWorker           12:00:34             0                   1273m               246m
phDataManager           12:00:34             0                   1505m               303m
phDiscover              12:00:34             0                   383m                32m
phReportWorker          12:00:34             0                   1322m               88m
phReportMaster          12:00:34             0                   435m                38m
phIpIdentityWorker      12:00:34             0                   907m                47m
phIpIdentityMaster      12:00:34             0                   373m                26m
phAgentManager          12:00:34             0                   881m                200m
phCheckpoint            12:00:34             0                   98m                 23m
phPerfMonitor           12:00:34             0                   700m                40m
phReportLoader          12:00:34             0                   630m                233m
phMonitor               31:21                0                   1120m               25m
Apache                  17:23:23             0                   260m                11m
Node.js                 17:20:54             0                   656m                35m
AppSvr                  17:23:16             0                   8183m               1344m
DBSvr                   17:23:34             0                   448m                17m

```

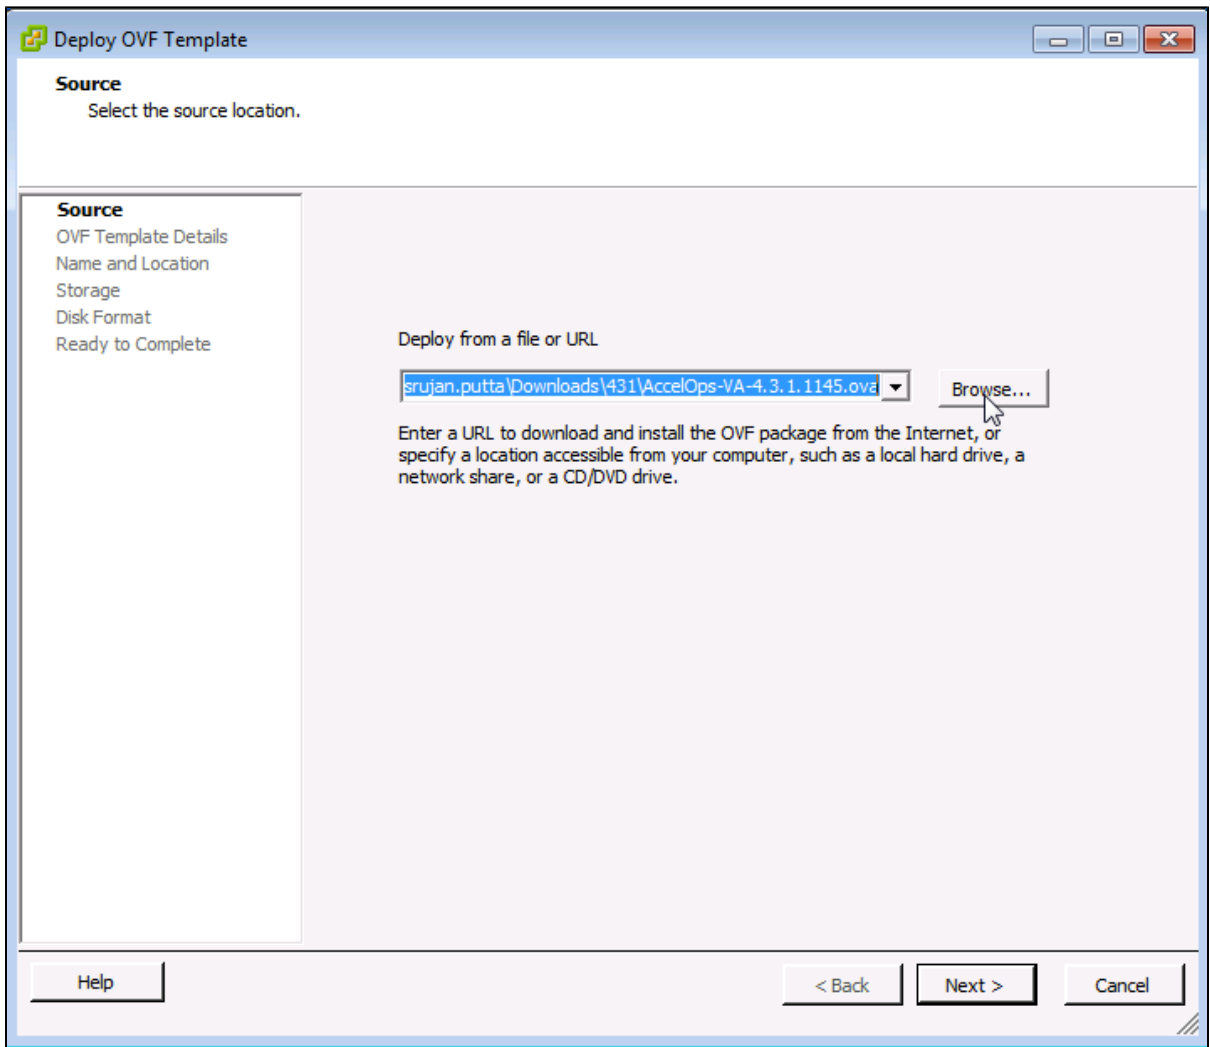
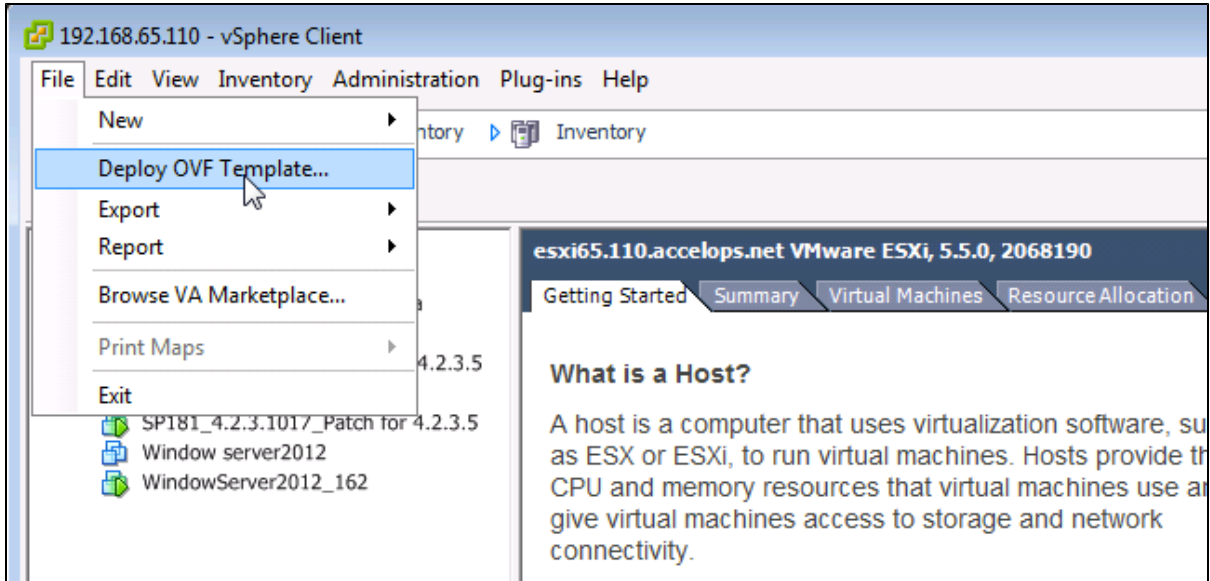
**Installing a Worker node****Section 1: Import a Worker image into VMware ESX server****NOTE**

Supervisor and Worker installation packages are same. Please use Supervisor's NFS mount point location to install Worker.

**Step 1: Begin importing Worker**

1. Download Supervisor OVA package from AccelOps image server
2. Log into VMware vSphere Client
  - a. Click on File -> Select Deploy OVF Template
  - b. Click on Browse button to choose OVA file

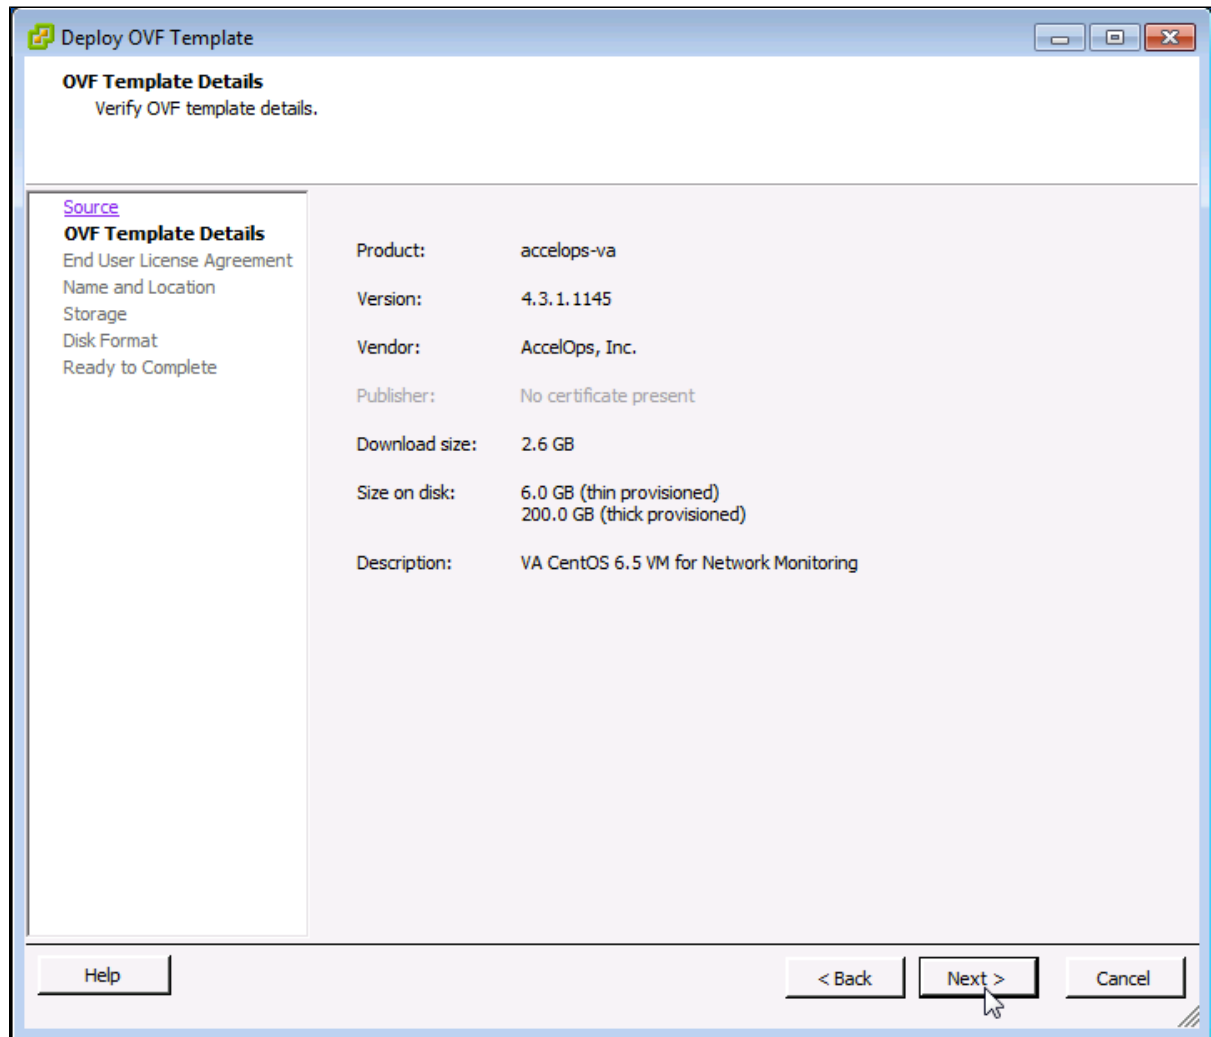
c. Locate Supervisor .ova file (Example: AccelOps-VA-4.3.1.1145.ova)



**Step 2: Review the Virtual Appliance details**

OVF Template details will tell you how big the file is to import and how much disk space it will take up in your selected ESX server.

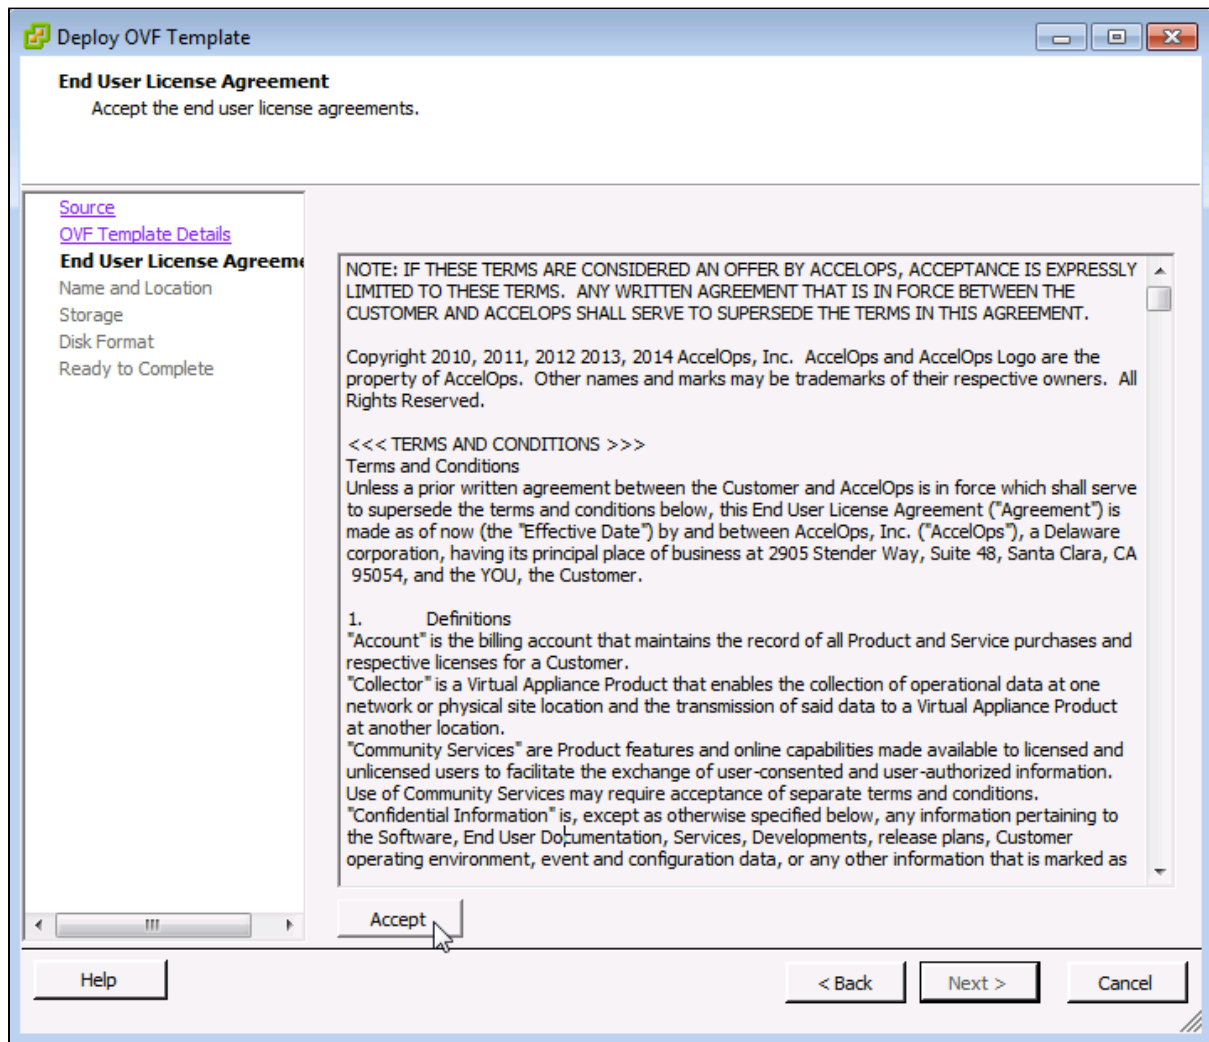
1. Click the 'Next' button



### Step 3: End User License Agreement

To continue you must accept the End User License Agreement.

1. Click on 'Accept' button on License Agreement
2. Click on 'Next' button



#### Step 4: Name and Location

Enter a name for the Virtual Appliance that is being imported. This name is independent of the host name and is only visible within the ESX environment.

1. Enter a Worker name and Click on 'Next' button

**Deploy OVF Template**

**Name and Location**  
Specify a name and location for the deployed template

[Source](#)  
[OVF Template Details](#)  
[End User License Agreement](#)

**Name and Location**  
Storage  
Disk Format  
Ready to Complete

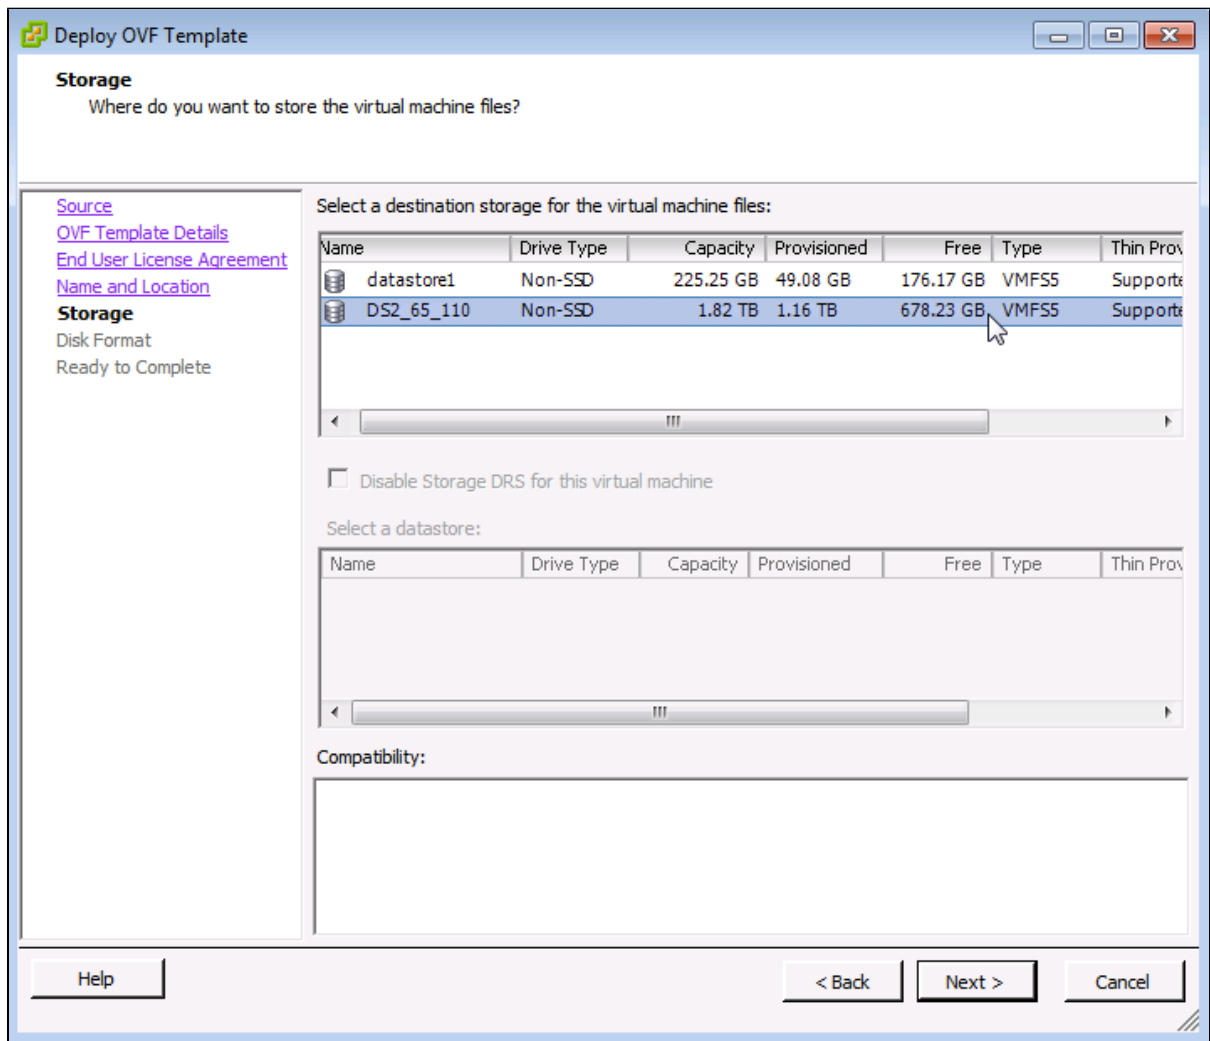
Name:  
AccelOps-Worker

The name can contain up to 80 characters and it must be unique within the inventory folder.

Help      < Back      Next >      Cancel

### Step 5: Storage

1. Select the desired storage location(Data store) and Click on 'Next' button



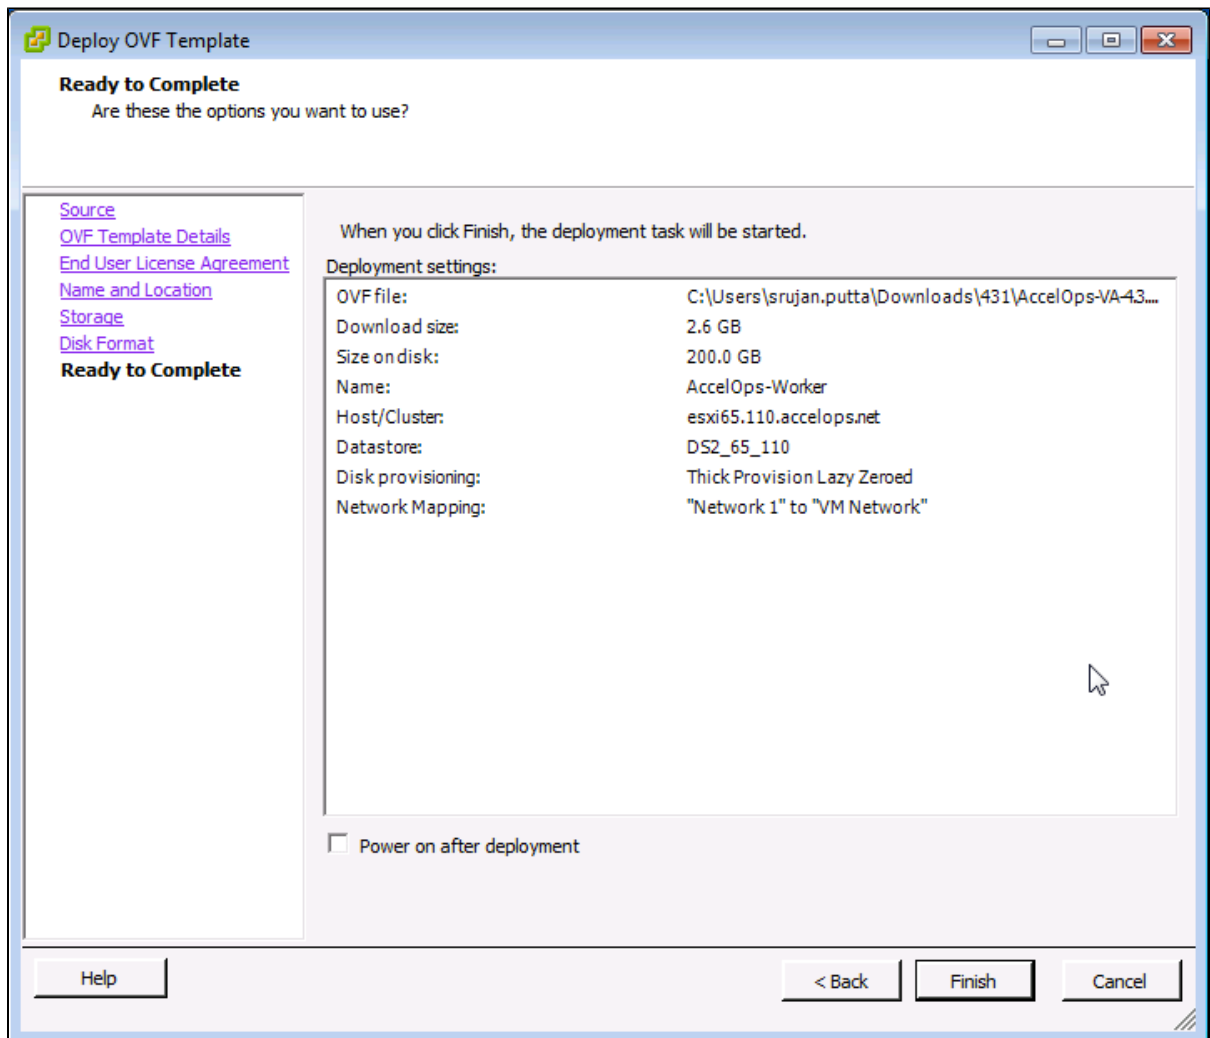
**Step 6: Disk Format**

1. Select the Disk format as 'Thick Provision' (recommended) and Click on 'Next' button

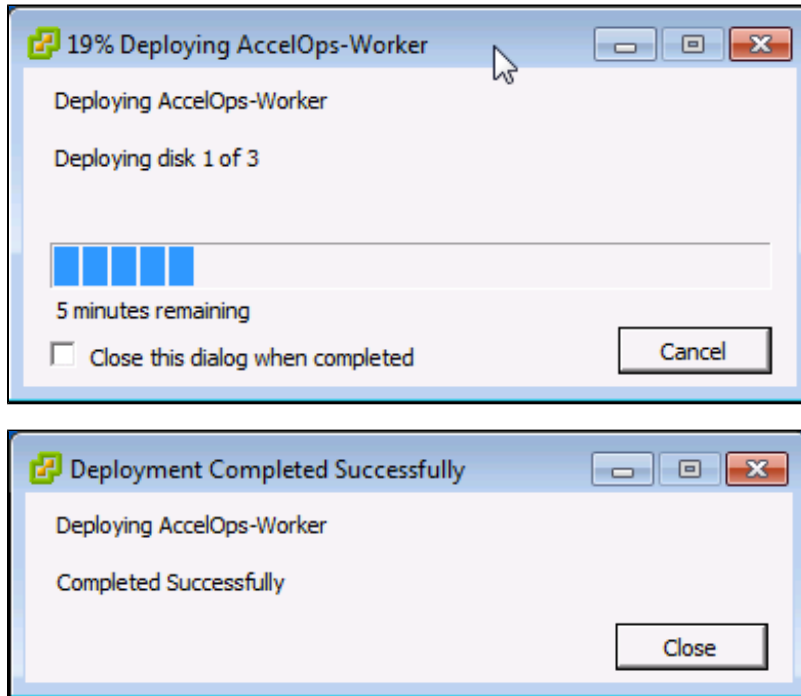
The screenshot shows a window titled "Deploy OVF Template" with a sub-header "Disk Format". Below the sub-header is the question "In which format do you want to store the virtual disks?". On the left side, there is a navigation pane with links: "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Storage", and "Disk Format" (which is highlighted). Below "Disk Format" is the text "Ready to Complete". The main area contains two input fields: "Datastore:" with the value "DS2\_65\_110" and "Available space (GB):" with the value "678.2". Below these fields are three radio button options: "Thick Provision Lazy Zeroed" (selected), "Thick Provision Eager Zeroed", and "Thin Provision". At the bottom of the window, there are three buttons: "Help", "< Back", and "Next >" (with a mouse cursor over it), and a "Cancel" button.

**Step 7: Ready to Complete**

1. Review the Worker resources and Click on 'Finish' button to begin importing Virtual Appliance
2. Virtual Appliance import will takes 7 to 10mins to complete. Do not turn off or reboot the system during this time.





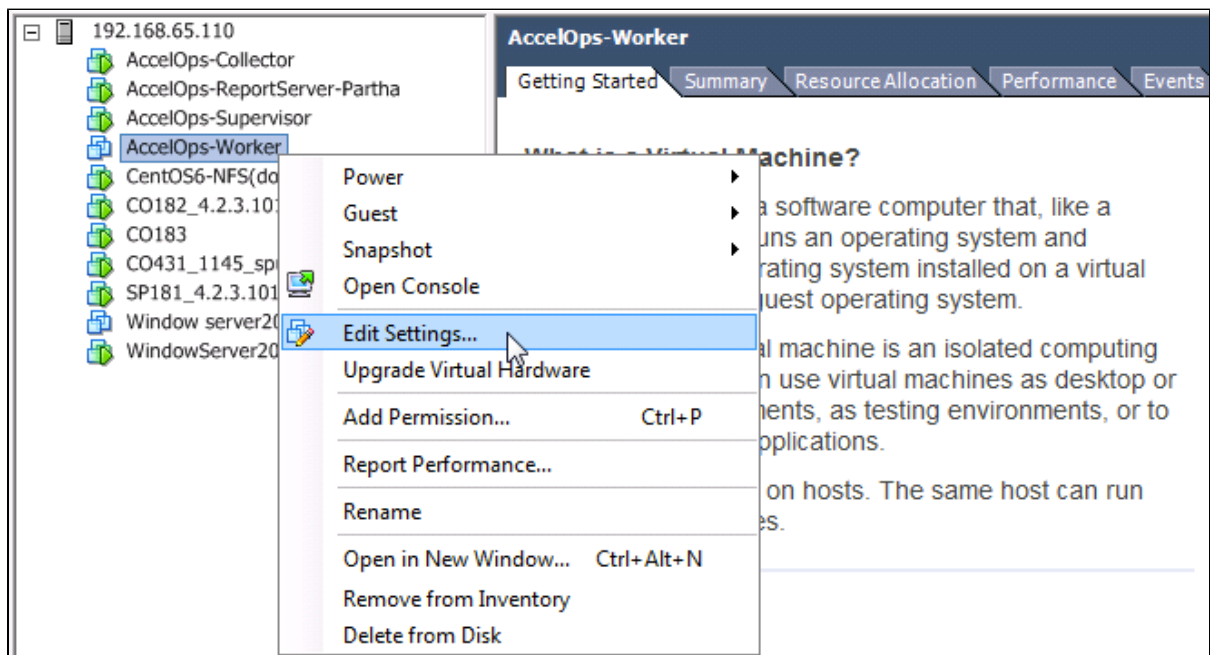


## Section 2: Edit the Worker hardware settings

Prior to starting the Worker for the first time it is required to make some modifications to the hardware setting for the Virtual Appliance.

### Step 1: Edit Settings

1. From the VMware vSphere client, Select the imported Worker Virtual appliance
2. Right mouse click and select Edit Settings



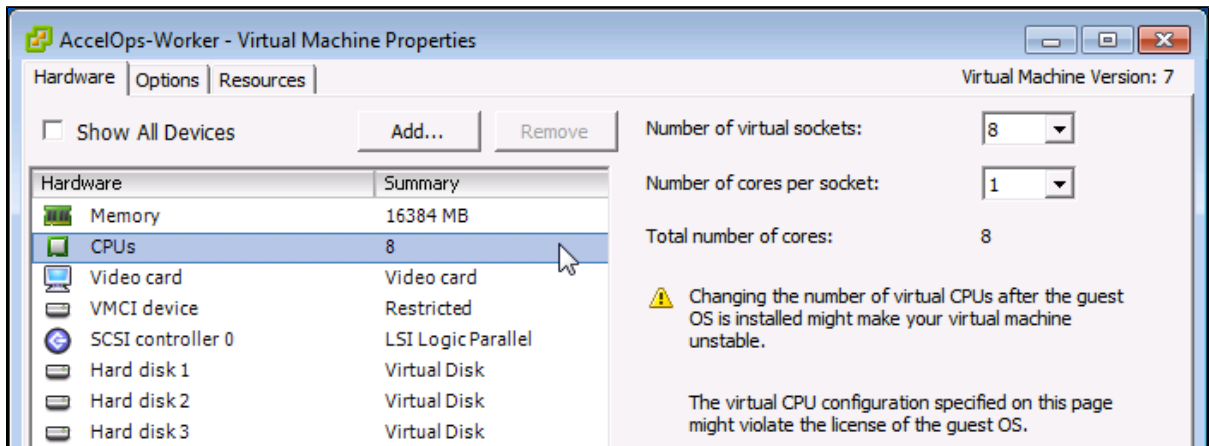
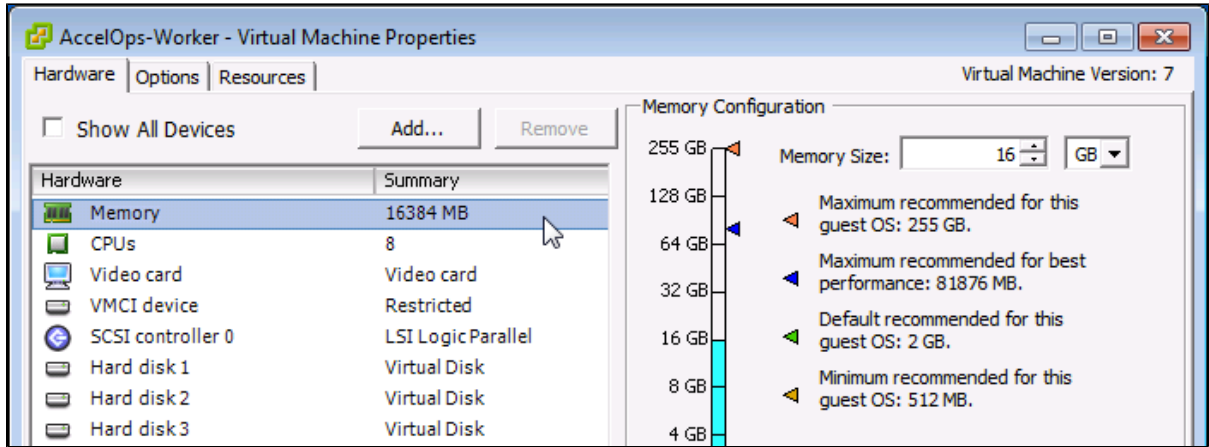
**Step 2: Verify Allocated Memory**

1. On Hardware tab, Click on Memory and verify the memory to at least 16 GB

**Allocate more memory for large deployments**

For large deployments you should allocate at least 24GB of memory. See the section on [Worker Hardware requirements](#) in the topic [Must Read Prior to Install](#) for more information.

2. Click on CPUs, verify the CPU cores to at least 8 cores



**Section 3: Select mount point to Store Events data**

The AccelOps Worker can be installed using only NFS storage. Please use Supervisor's NFS mount point location to install Worker.

1. SSH to Supervisor, copy NFS mount point location to store Events data

### Supervisor's Mount point location

```
[root@super ~]# mount
/dev/sda3 on / type ext3 (rw)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
devpts on /dev/pts type devpts (rw,gid=5,mode=620)
tmpfs on /dev/shm type tmpfs (rw)
/dev/sda1 on /boot type ext3 (rw)
/dev/sdb1 on /cmdb type ext3 (rw)
/dev/sdc1 on /svn type ext3 (rw)
none on /proc/sys/fs/binfmt_misc type binfmt_misc (rw)
sunrpc on /var/lib/nfs/rpc_pipefs type rpc_pipefs (rw)
192.168.67.168:/accelops/SP431 on /data type nfs
(rw,noatime,nfsvers=3,timeo=14,nolock,addr=192.168.67.168)
```

Example Supervisor Mount point location: 192.168.67.168:/accelops/SP431

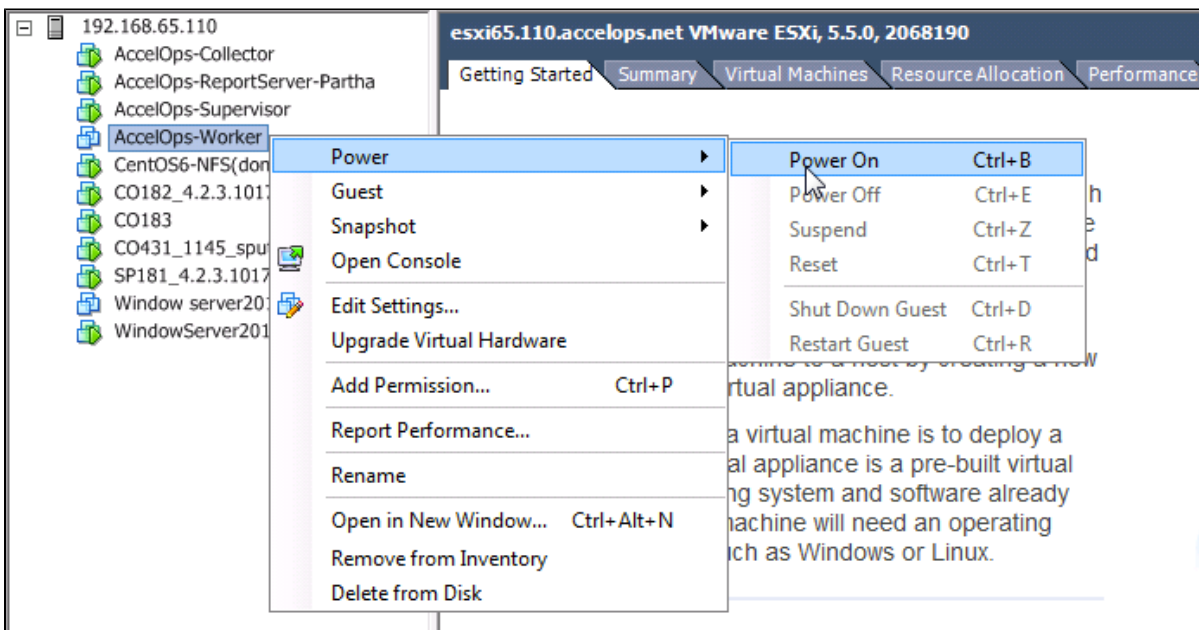
## Section 4: Configure the Worker from the VMware Console

### Note

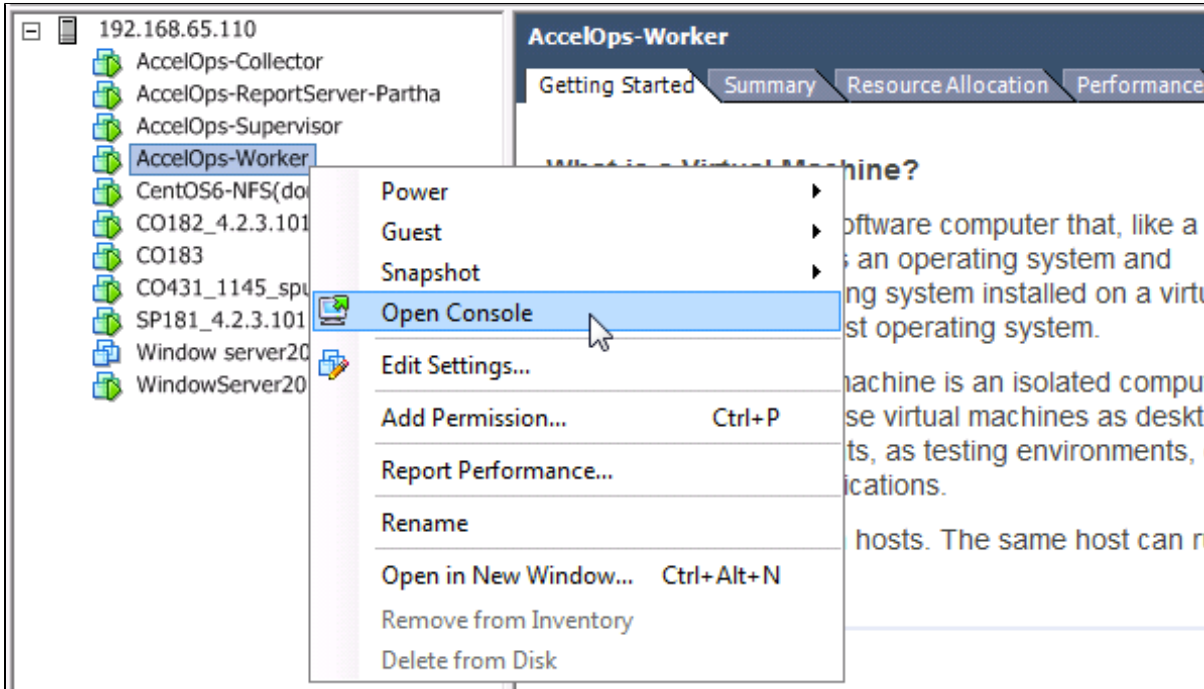
While in the VM console window do not press any control keys during the installation. Doing so might cause the installation to stop and not to precede any further. For example Ctrl – C or Ctrl – Z. If this occurs you must erase the Supervisor and start over from the beginning.

### Step 1: Start Worker

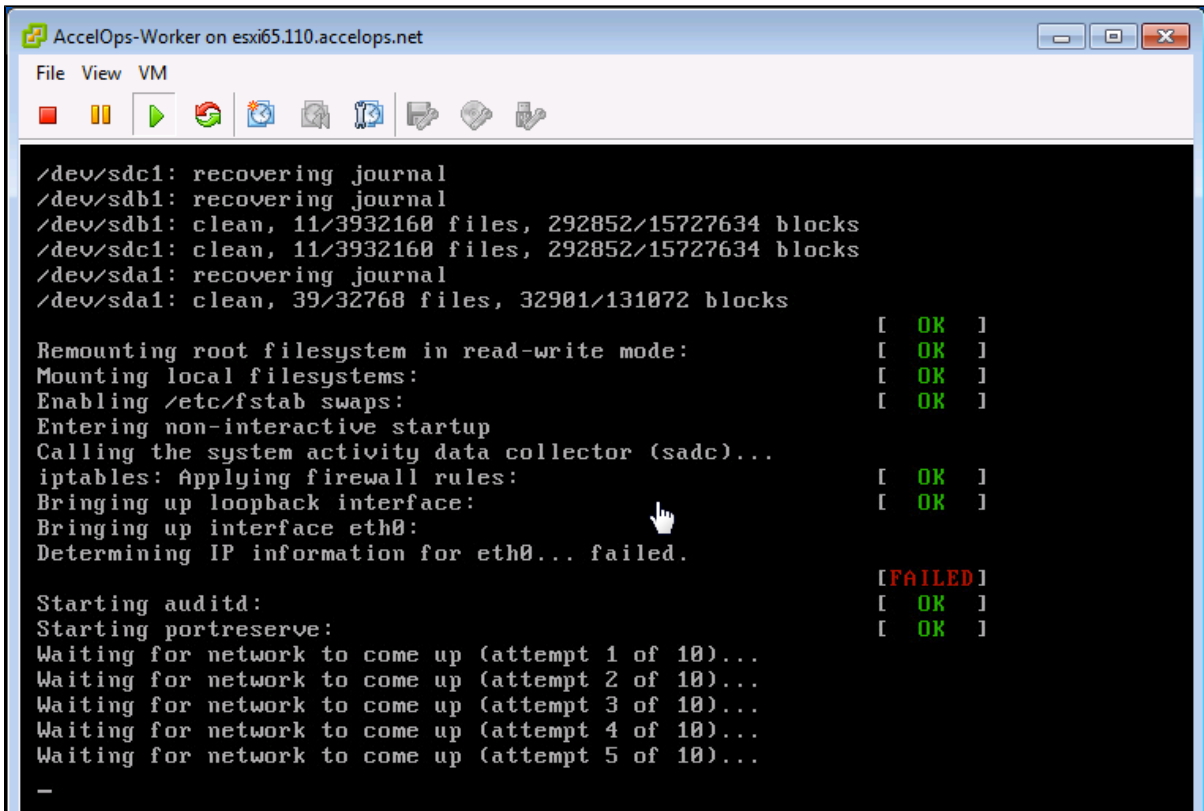
1. From the VMware vSphere client, Select recently imported Worker Virtual appliance and Power On (Right click > Power > Power On)



2. Open Worker VM Console (Right click > Open Console)




3. Network eth0 Failed message may appear on the screen during the loading, but this is to be expected the first time the AccelOps Worker is started.




**Step 2: Configure the Timezone**

It is important to select the correct time zone that you are located in.

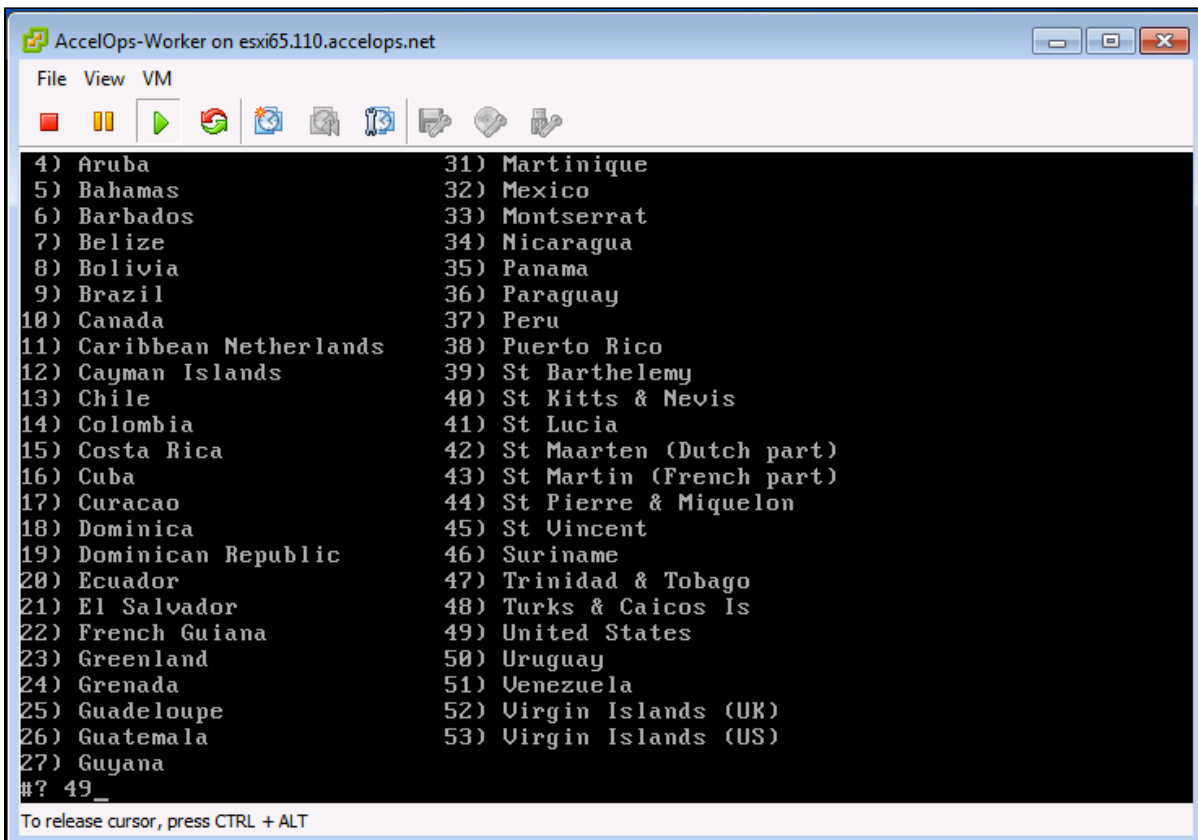
1. On VM console, Select 'Set Timezone' and press the 'Enter' key



2. Select your Location and press the 'Enter' key



3. Select your Country and press the 'Enter' key



4. Select your Timezone and press the 'Enter' key

```

AccelOps-Worker on esxi65.110.accelops.net
File View VM
6) Eastern Time - Indiana - Daviess, Dubois, Knox & Martin Counties
7) Eastern Time - Indiana - Pulaski County
8) Eastern Time - Indiana - Crawford County
9) Eastern Time - Indiana - Pike County
10) Eastern Time - Indiana - Switzerland County
11) Central Time
12) Central Time - Indiana - Perry County
13) Central Time - Indiana - Starke County
14) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties
15) Central Time - North Dakota - Oliver County
16) Central Time - North Dakota - Morton County (except Mandan area)
17) Central Time - North Dakota - Mercer County
18) Mountain Time
19) Mountain Time - south Idaho & east Oregon
20) Mountain Standard Time - Arizona (except Navajo)
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - southeast Alaska panhandle
25) Alaska Time - Alaska panhandle neck
26) Alaska Time - west Alaska
27) Aleutian Islands
28) Metlakatla Time - Annette Island
29) Hawaii
#? 21_
To release cursor, press CTRL + ALT

```

5. Review and Confirm your Timezone selection

```

AccelOps-Worker on esxi65.110.accelops.net
File View VM
19) Mountain Time - south Idaho & east Oregon
20) Mountain Standard Time - Arizona (except Navajo)
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - southeast Alaska panhandle
25) Alaska Time - Alaska panhandle neck
26) Alaska Time - west Alaska
27) Aleutian Islands
28) Metlakatla Time - Annette Island
29) Hawaii
#? 21

The following information has been given:

    United States
    Pacific Time

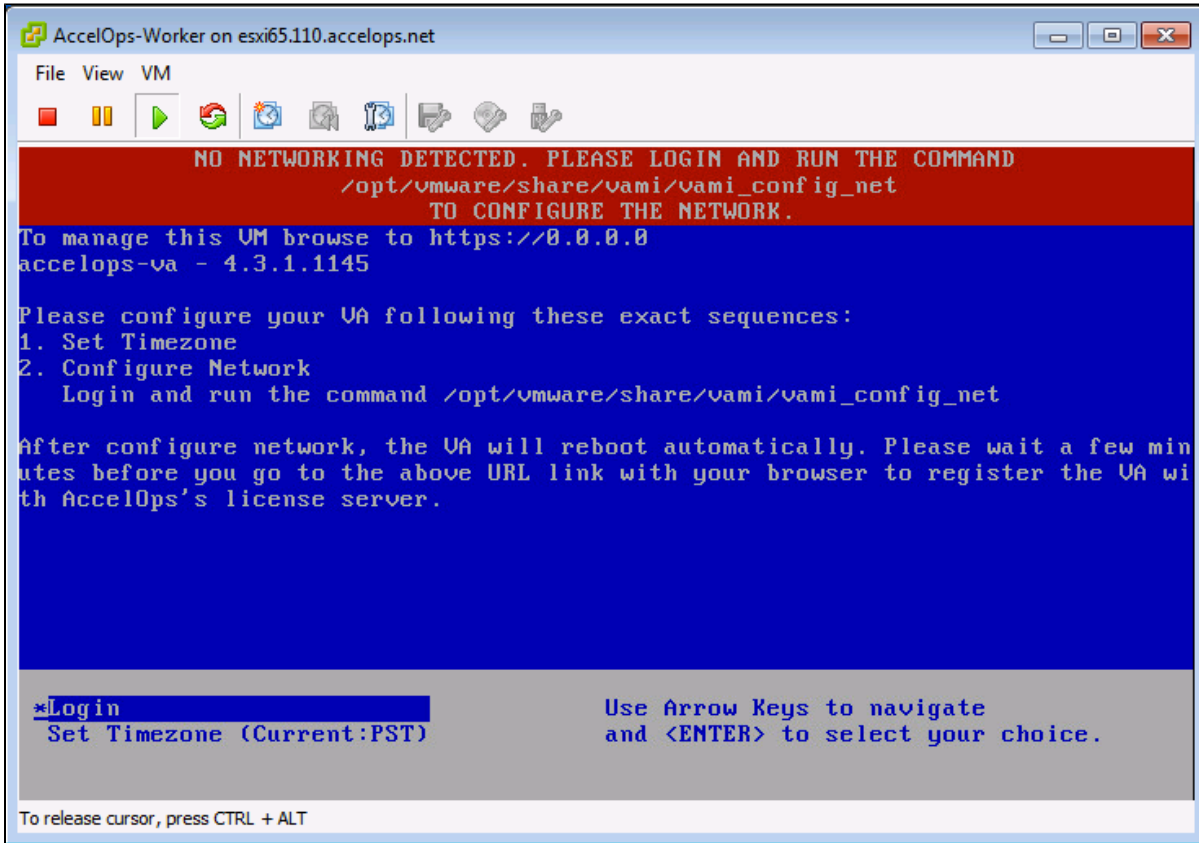
Therefore TZ='America/Los_Angeles' will be used.
Local time is now:   Mon Nov 10 18:38:50 PST 2014.
Universal Time is now: Tue Nov 11 02:38:50 UTC 2014.
Is the above information OK?
1) Yes
2) No
#? 1_
To release cursor, press CTRL + ALT

```




**Step 3: Configure the Network**

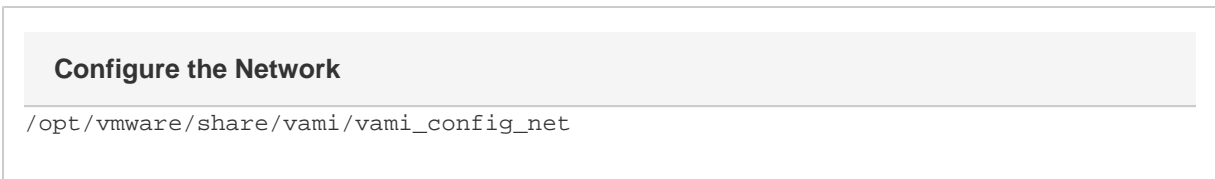
1. On VM console, Select 'Login' and press the 'Enter' key



2. Enter default log in information.  
Login: **root**  
Password: **ProspectHills**



3. Run vami\_config\_net script to configure network and installing AccelOps Worker





```

AccelOps-Worker on esxi65.110.accelops.net
File View VM
localhost login: root
Password:
[root@localhost ~]# /opt/vmware/share/vami/vami_config_net

CHANGE NETWORK CONFIGURATION for eth0

Please enter the desired network parameters.
To exit, type q at any prompt.

IP Address []: _

```

4. Configure Network with Static IP address
  - a. The following information is required to configure Static IP address:
    - i. IP Address
    - ii. Netmask
    - iii. Gateway
    - iv. DNS Server(s)

#### Proxy Server

Do not configure the proxy server setting.

- b. Press the 'y' key to accept the changes
- c. Enter 'Host name' and press 'Enter' key

```

Gateway []: 192.168.64.1
DNS Server 1 []: 192.168.0.10
DNS Server 2 [192.168.0.10]: 192.168.0.40
Is a proxy server necessary to reach the Internet? y/n [n]:
IP Address: 192.168.65.149
Netmask: 255.255.252.0
Gateway: 192.168.64.1
Proxy Server:
DNS Servers: 192.168.0.10, 192.168.0.40
Is this correct? y/n [y]: y
Reconfiguring the network...
DNS server settings updated
Determining if ip address 192.168.65.149 is already in use for device eth0...
vami_login: no process killed
Network parameters successfully changed to requested values
Input hostname: worker1.accelops.net

```

**Step 4: Select Mount Point to store your data**

The AccelOps Worker can be installed using NFS storage.

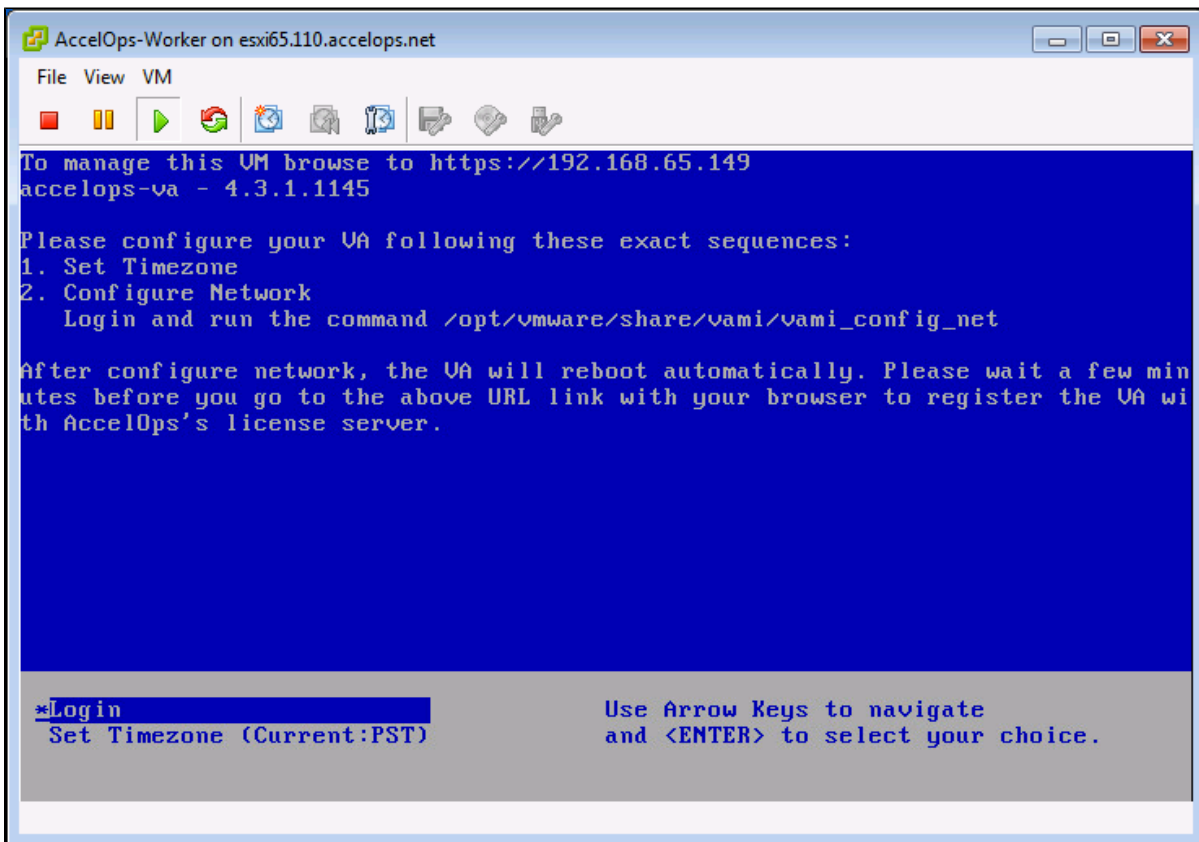
Choose only following Step 4.1 or 4.2 based on Section 3

1. Enter Supervisor's NFS mount point location

```
NFS Mount point
<NFS_Server_IP_Address>:/<Supervisor_Mount_Directory_Path>
```

```
Mount Point For Your Data: local storage (/dev/sdd) or NFS (ip:/dir) with right
permission to proceed or use Ctrl+C to go back to console.
Mount Point: 192.168.67.168:/accelops/SP431
192.168.67.168:/accelops/SP431 mounted
Running initialization of Worker host
```

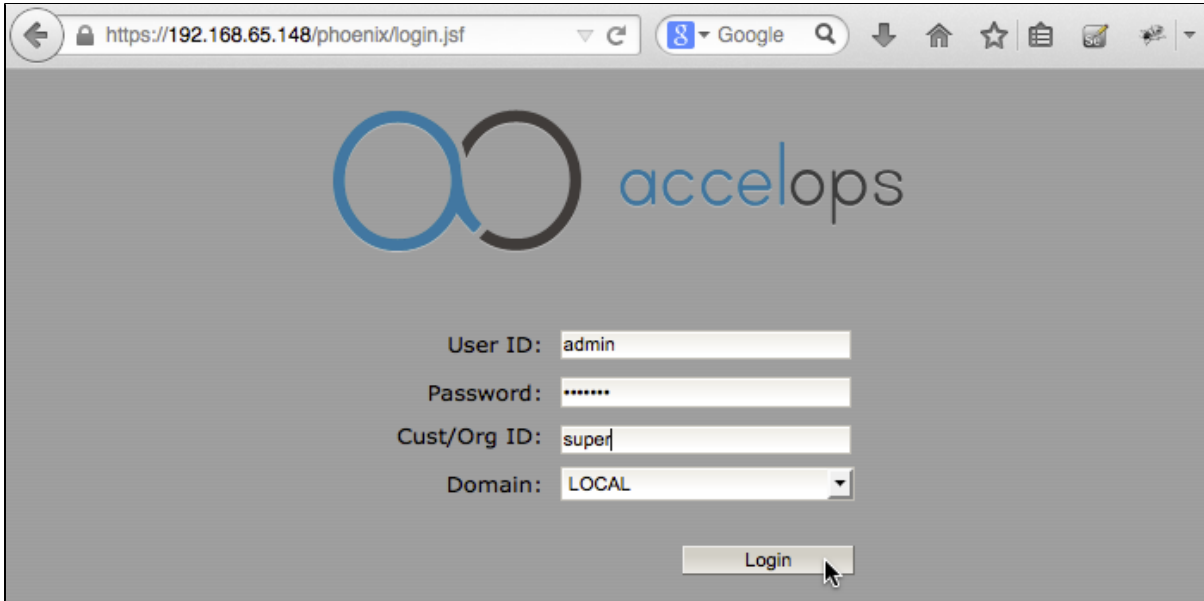
2. After mount point, Worker will reboot automatically  
In 3mins, Worker will be configure successful.



## Section 5: Register the Worker to Supervisor

### Step 1: Register AccelOps Worker to Supervisor

1. In browser windows, enter 'https://<Supervisor IP>'
2. Enter the default log in credentials  
User ID: **admin**  
Password: **admin\*1**  
Cust/Org ID: **super**



3. Go to Admin > License Management > VA Information, Click on 'Add' button and Enter new Worker IP address.

**Admin > License Management**

**License Information**  Refresh

Your current license details and current usages are shown below.  
You may request additional workers and/or EPS by upgrading your license.

**Customer Name** Super

License Attribute	Allowed (per license)	Current Usage
# of VAs	10	1
EPS	100000	See table
Number of Devices	30000	1
Valid Time	12:00:00 AM Dec 15 2012 - 12:00:00 AM Dec 24 2016	774 days left
Number of Organizations	500	0
License Type	AO-SP	

**Add VA**

VA IP Address:

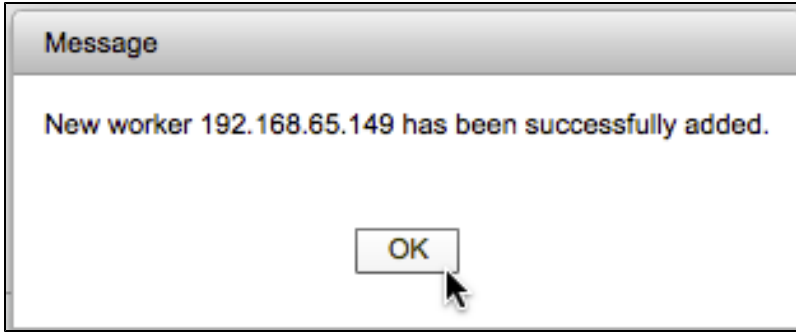
Events Per Second	MAX(Event Rate (/ sec))	MIN(Event Rate (/ sec))
Super	42	0.38

**VA Information**


Mode	IP Address
VA(Super)	192.168.65.148

**Report Server Information**

IP Address	Expires	Status



4. New Worker will configure successful in 2mins and Check new Worker IP address in VA information tab.



5. Go to Admin > Cloud Health and check Worker health.

**Admin > Cloud Health**

Refresh    Last Updated: 12:08:22

Name	IP Address	Health	Module R	Version	Build Date	Cores	Load Ave	CPU	Sw
super.accelops.net	192.168.65.148	Normal	supervis	4.3.1.1145	10:04:55 AM Oct 15 ...	8	0.04,0.04	1.51	6 G
worker1.accelops.net	192.168.65.149	Normal	worker	4.3.1.1145	10:04:55 AM Oct 15 ...	8	0.21,0.25	1.26	6 G

**worker1.accelops.net ( 192.168.65.149 ) - worker: Process Details**

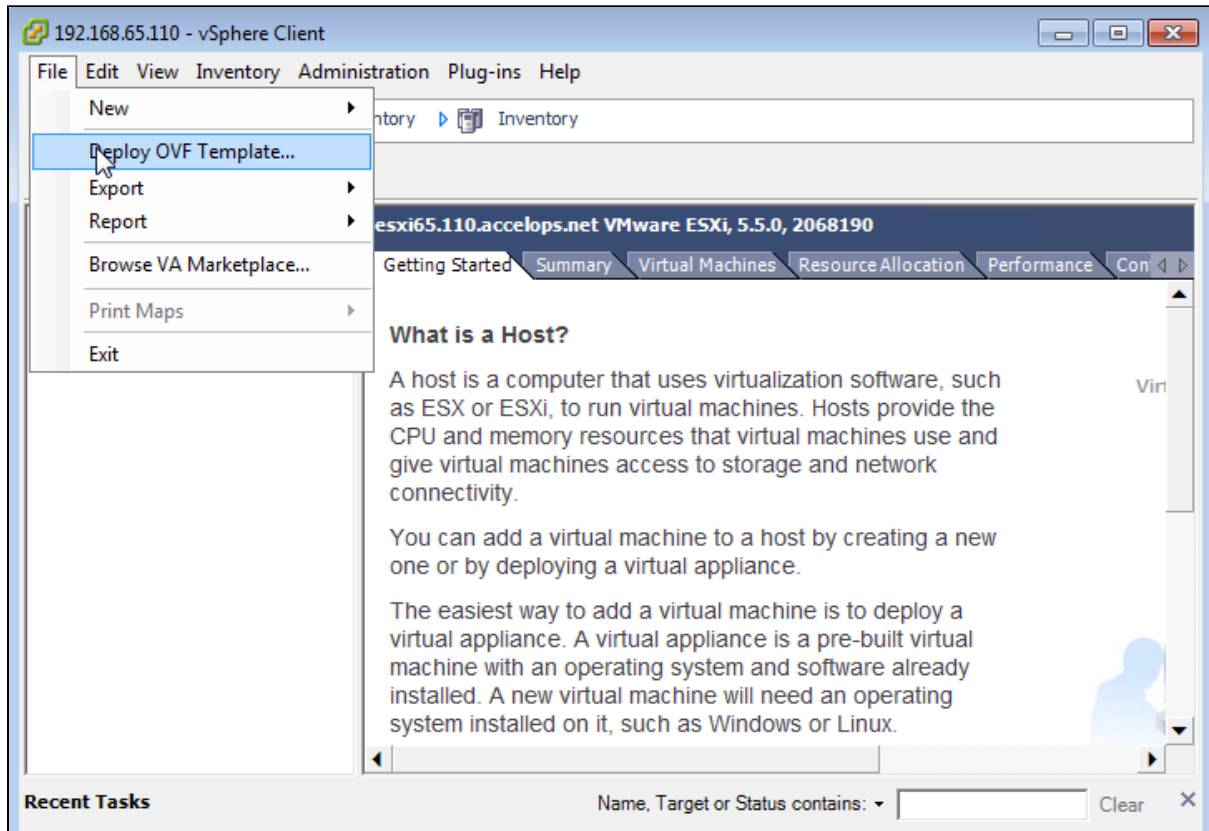
Process Name	Status	Up Time	CPU	Event Rate	Physical Mem	Virtual Mem
phMonitorWorker	Up	2m 18s	0%	0/s	111 MB	1131 MB
phParser	Up	2m 17s	0%	0/s	261 MB	1768 MB
phCheckpoint	Up	2m 17s	0%	0/s	19 MB	98 MB
phReportWorker	Up	2m 17s	0%	0/s	231 MB	1322 MB
phDataManager	Up	2m 17s	0%	0/s	40 MB	992 MB
phQueryWorker	Up	2m 17s	0%	0/s	234 MB	1253 MB
phAgentManager	Up	2m 17s	0%	0/s	203 MB	881 MB
phRuleWorker	Up	2m 17s	0%	0/s	234 MB	1256 MB
phPerfMonitor	Up	2m 17s	0%	0/s	37 MB	700 MB
phIdentityWorker	Up	2m 17s	0%	0/s	35 MB	907 MB
rsyslogd	Up	28m 23s	0%	0/s	1 MB	243 MB
httpd	Up	2m 17s	0%	0/s	9 MB	260 MB

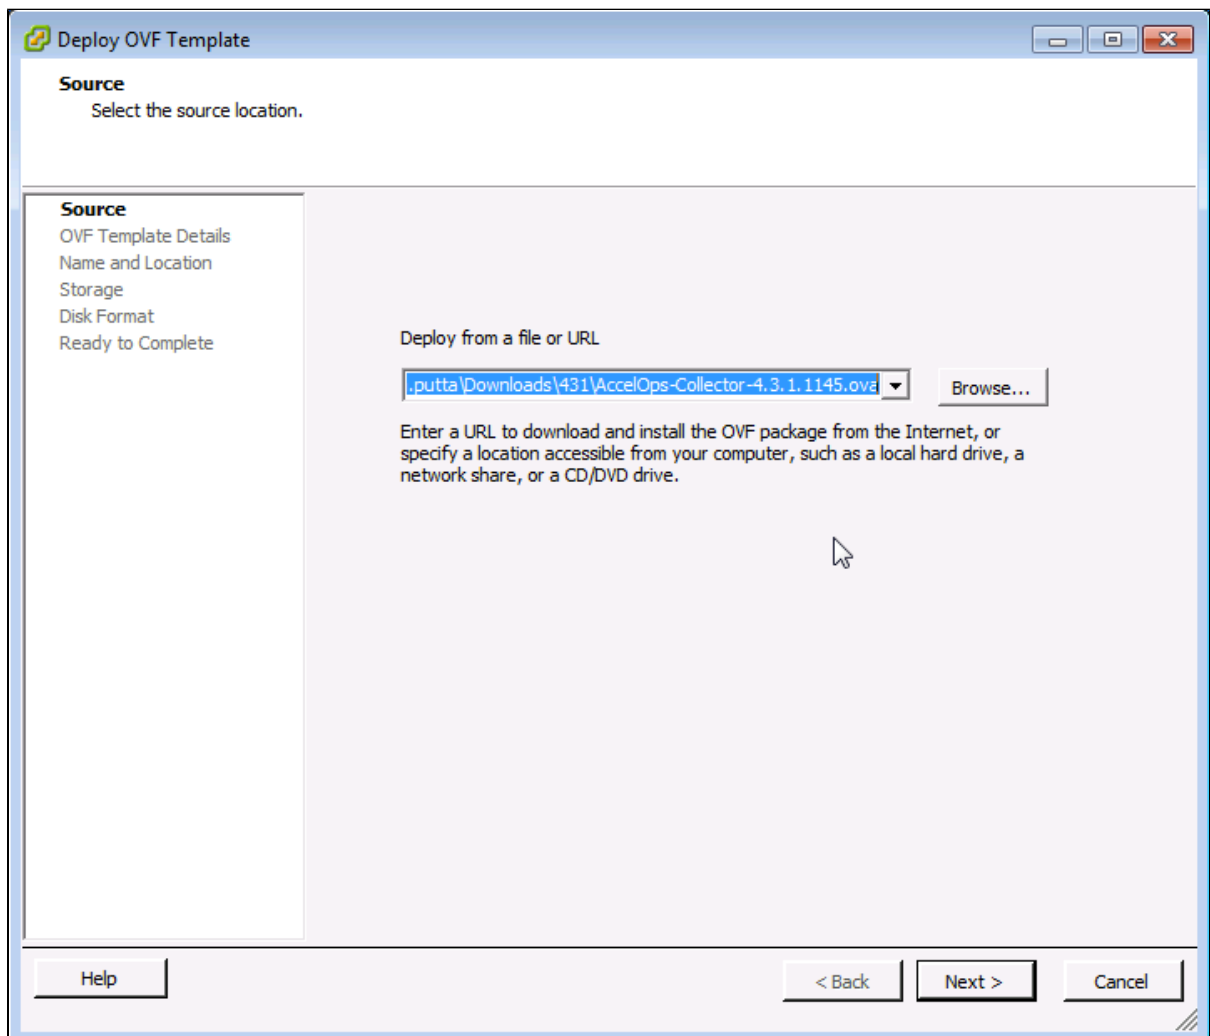
## Installing a Collector node

### Section 1: Import a Collector image into VMware ESX server

**Step 1: Begin importing Collector**

1. Download Collector OVA package from AccelOps image server
2. Log into VMware vSphere Client
  - a. Click on File -> Select Deploy OVF Template
  - b. Click on Browse button to choose OVA file
  - c. Locate Collector .ova file (Example: AccelOps-Collector-4.3.1.1145.ova)

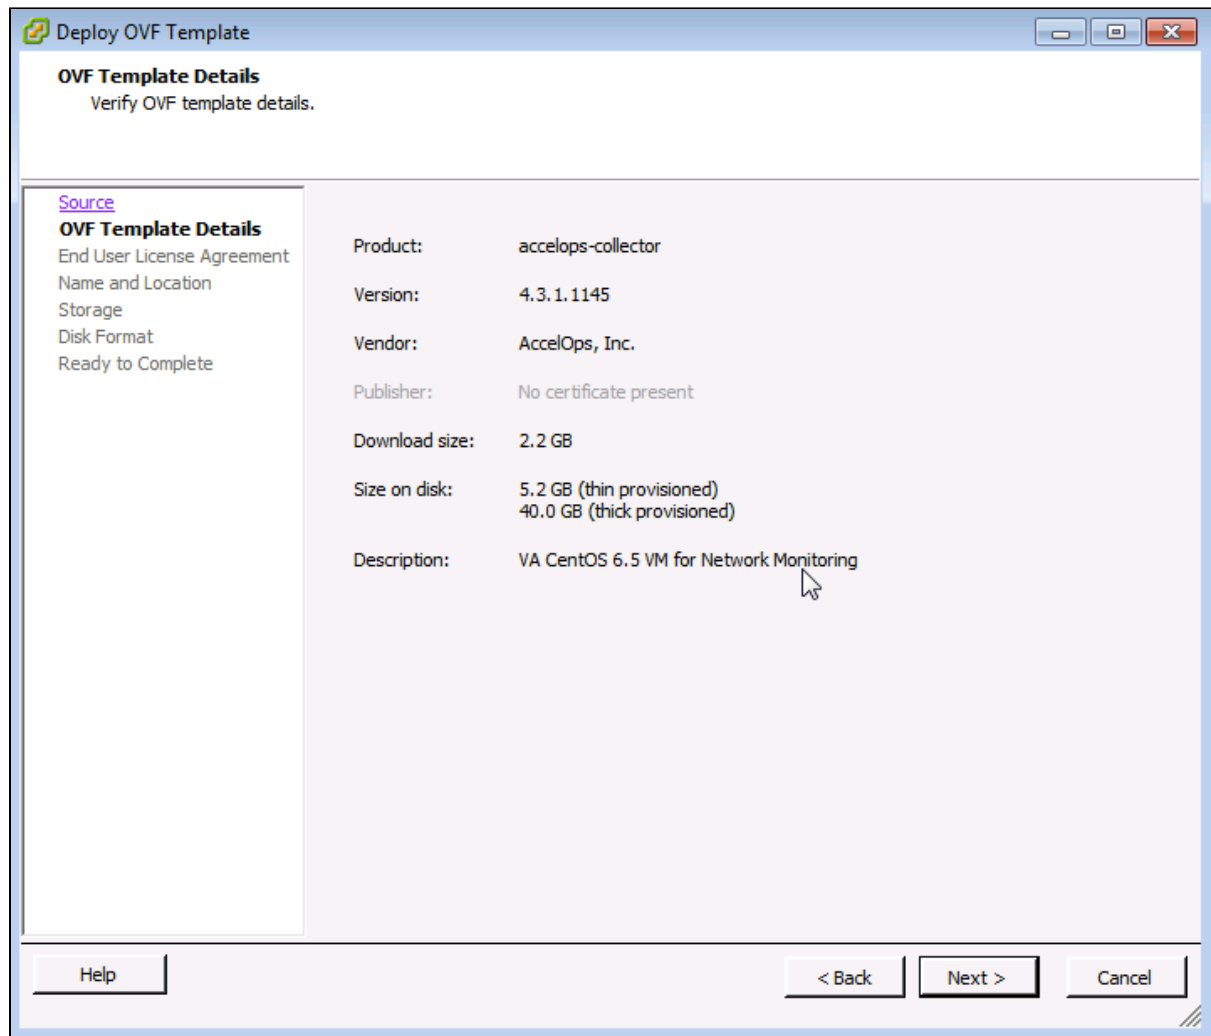




**Step 2: Review the Virtual Appliance details**

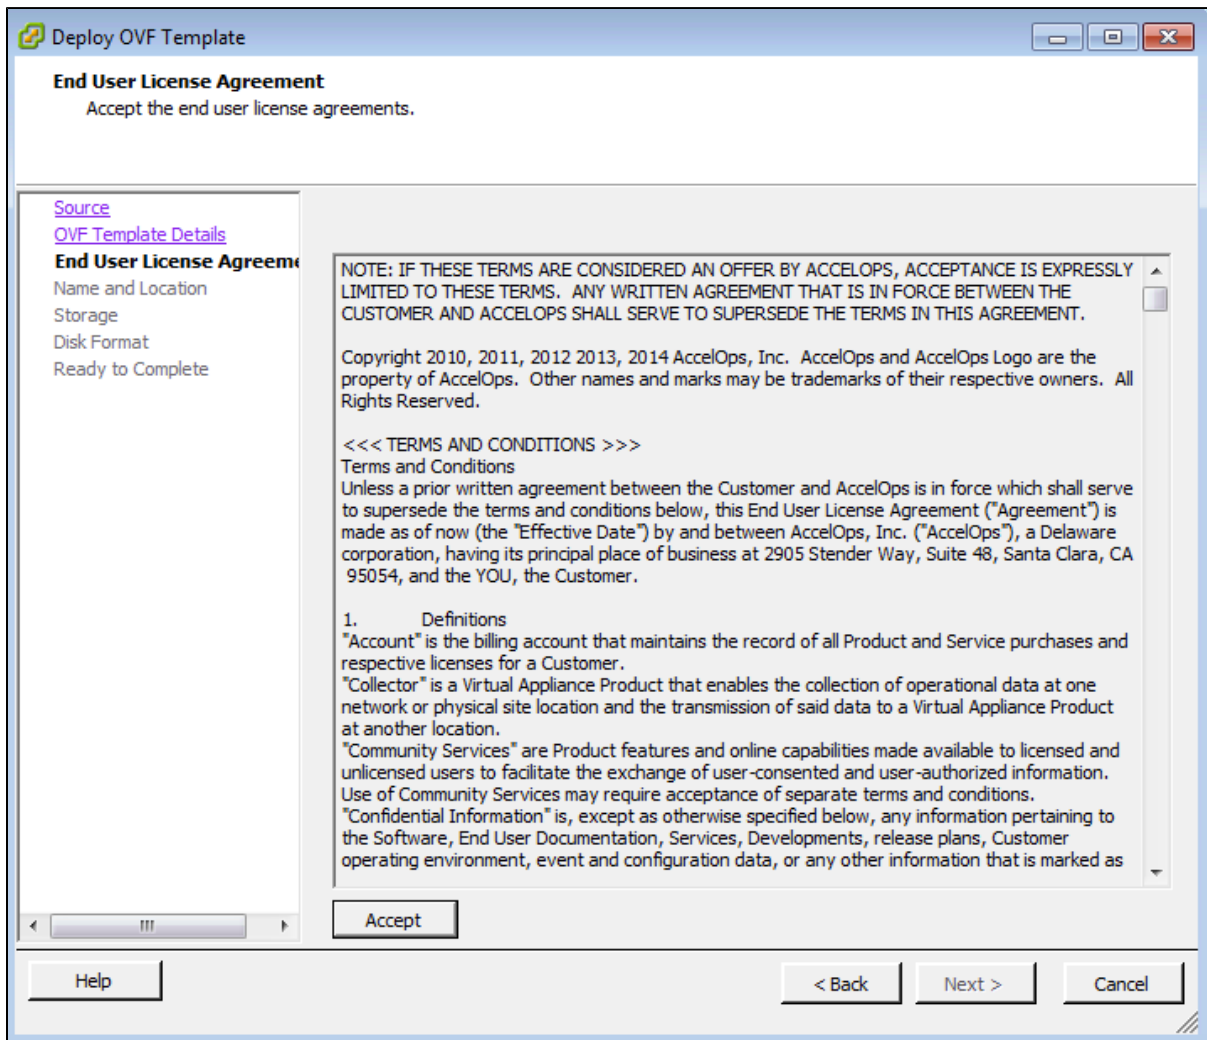
OVF Template details will tell you how big the file is to import and how much disk space it will take up in your selected ESX server.

1. Click the 'Next' button

**Step 3: End User License Agreement**

To continue you must accept the End User License Agreement.

1. Click on 'Accept' button on License Agreement
2. Click on 'Next' button



**Step 4: Name and Location**

Enter a name for the Virtual Appliance that is being imported. This name is independent of the host name and is only visible within the ESX environment.

1. Enter a Collector name and Click on 'Next' button



**Deploy OVF Template**

**Name and Location**  
Specify a name and location for the deployed template

[Source](#)  
[OVF Template Details](#)  
[End User License Agreement](#)  
**Name and Location**  
Storage  
Disk Format  
Ready to Complete

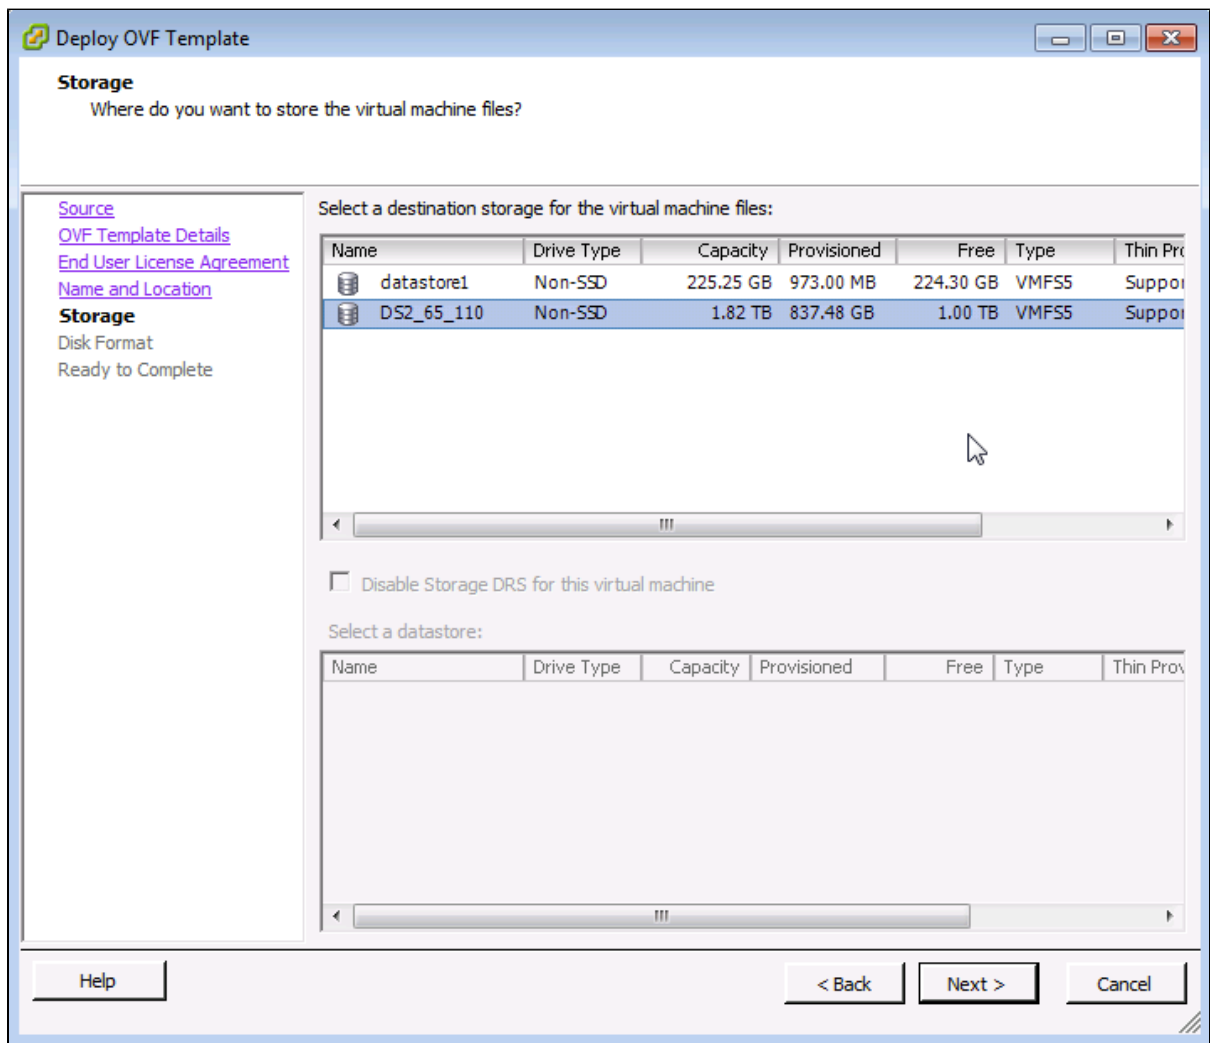
Name:  
AccelOps-Collector

The name can contain up to 80 characters and it must be unique within the inventory folder.

Help < Back Next > Cancel

**Step 5: Storage**

1. Select the desired storage location(Data store) and Click on 'Next' button



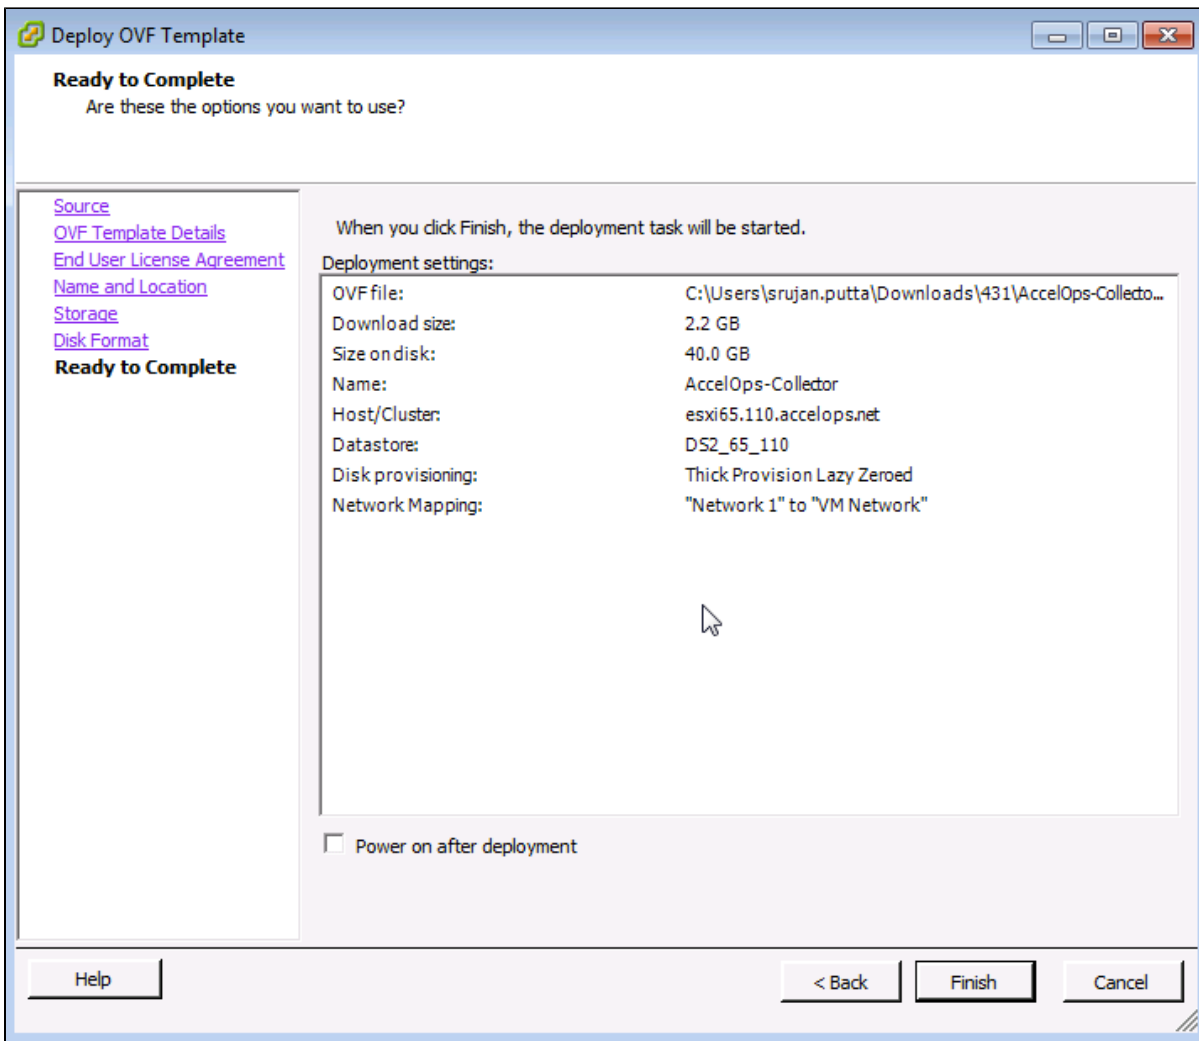
**Step 6: Disk Format**

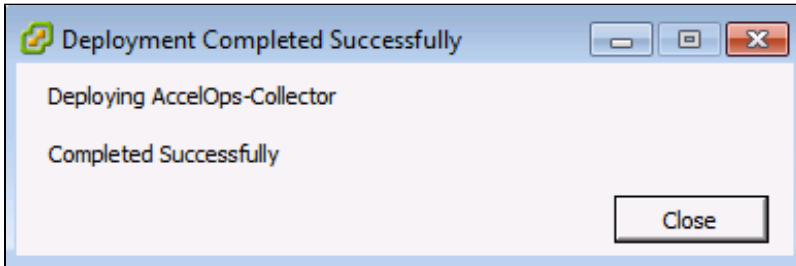
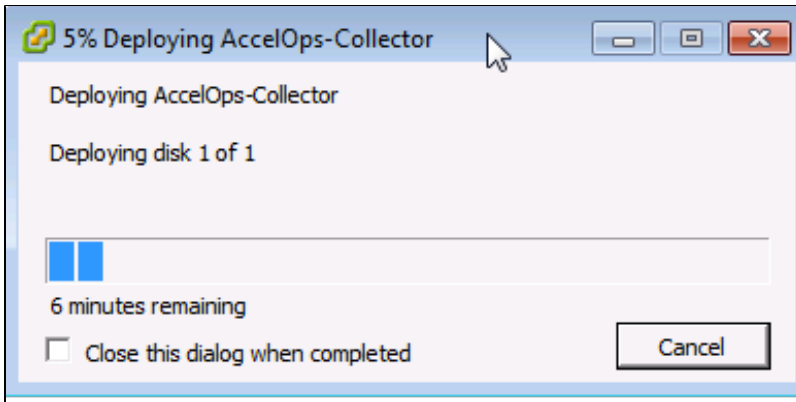
1. Select the Disk format as 'Thick Provision' (recommended) and Click on 'Next' button

The screenshot shows a window titled "Deploy OVF Template" with a mouse cursor over the title bar. The window contains a "Disk Format" section with the question "In which format do you want to store the virtual disks?". On the left, there is a sidebar with links: "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Storage", and "Disk Format" (which is highlighted). Below "Disk Format" is the text "Ready to Complete". The main area shows "Datastore:" with a text box containing "DS2\_65\_110" and "Available space (GB):" with a text box containing "1025.3". Below these are three radio button options: "Thick Provision Lazy Zeroed" (selected), "Thick Provision Eager Zeroed", and "Thin Provision". At the bottom of the window are three buttons: "Help", "< Back", and "Next >", and a "Cancel" button on the far right.

**Step 7: Ready to Complete**

1. Review the Collector resources and Click on 'Finish' button to begin importing Virtual Appliance
2. Virtual Appliance import will takes 5 to 7mins to complete. Do not turn off or reboot the system during this time



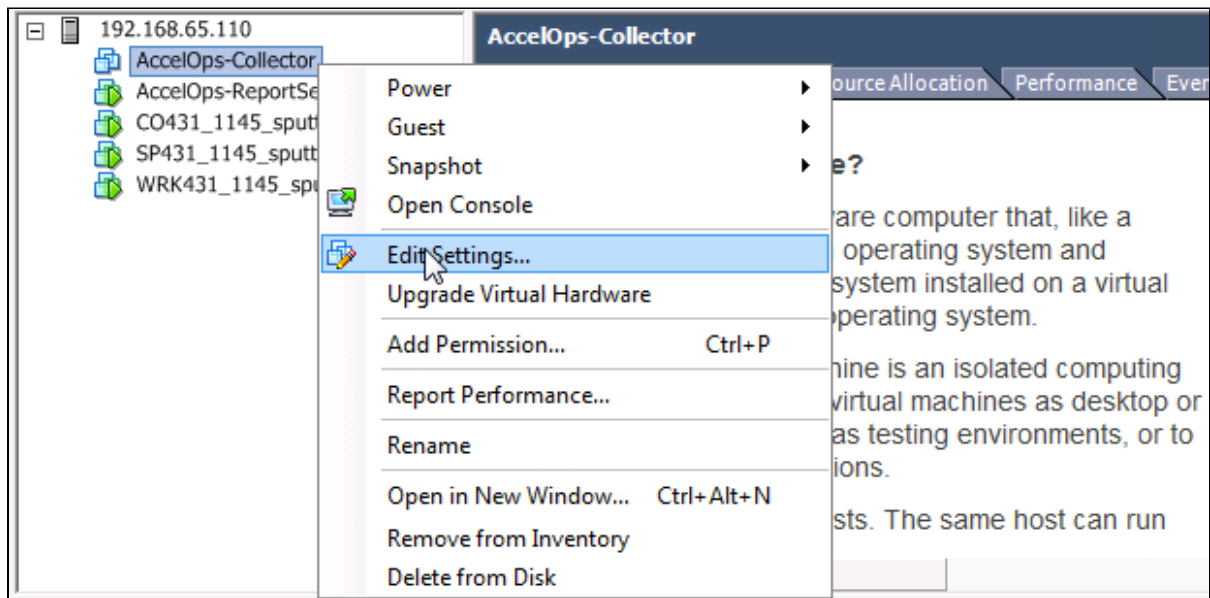


## Section 2: Edit the Collector hardware settings

Prior to starting the Collector for the first time it is required to make some modifications to the hardware setting for the Virtual Appliance.


### Step 1: Edit Settings

1. From the VMware vSphere client, Select the imported Collector Virtual appliance
2. Right mouse click and select Edit Settings



### Step 2: Verify Allocated Memory

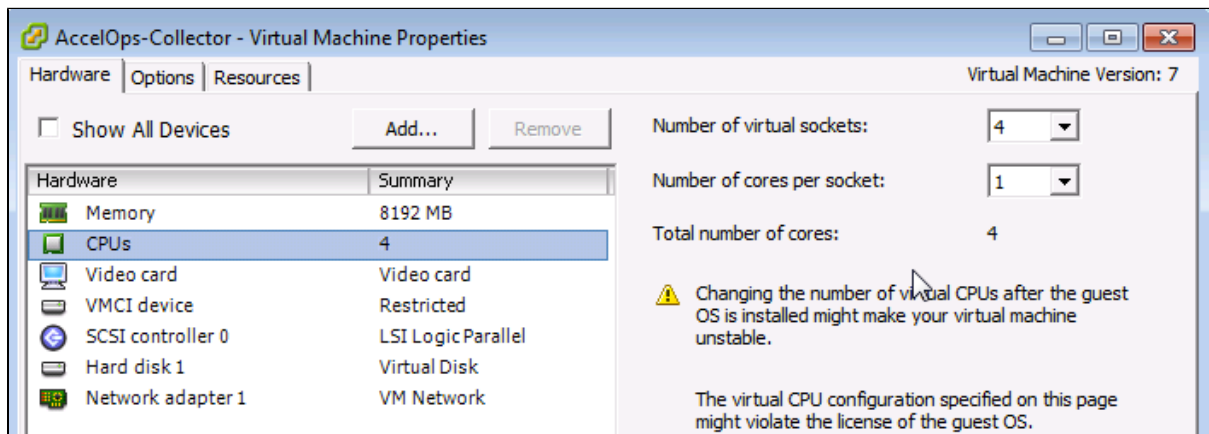
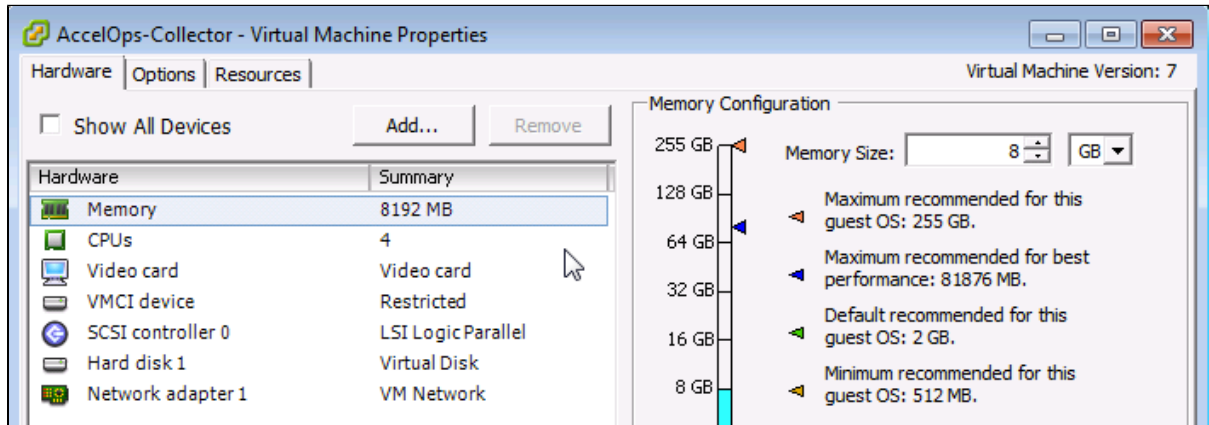
1. On Hardware tab, Click on Memory and verify the memory to at least 8 GB



**Allocate more memory for large deployments**

For large deployments you should allocate at least 24GB of memory. See the [section on Collector Hardware requirements](#) in the topic [Must Read Prior to Install](#) for more information.

2. Click on CPUs, verify the CPU cores to at least 4 cores



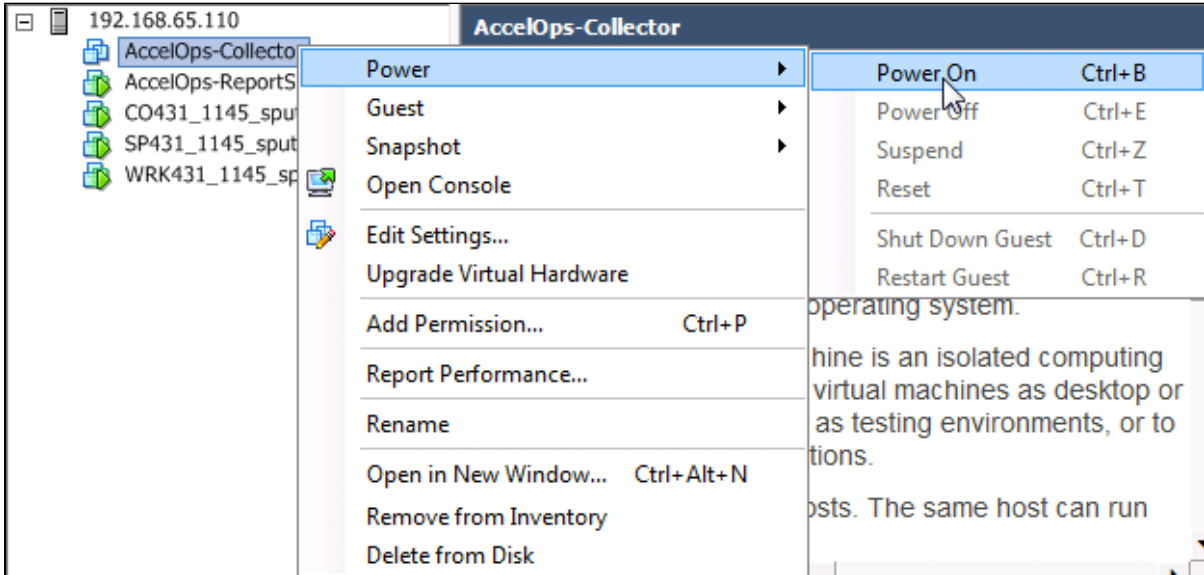
**Section 3: Configure the Collector from the VMware Console**

**Note**

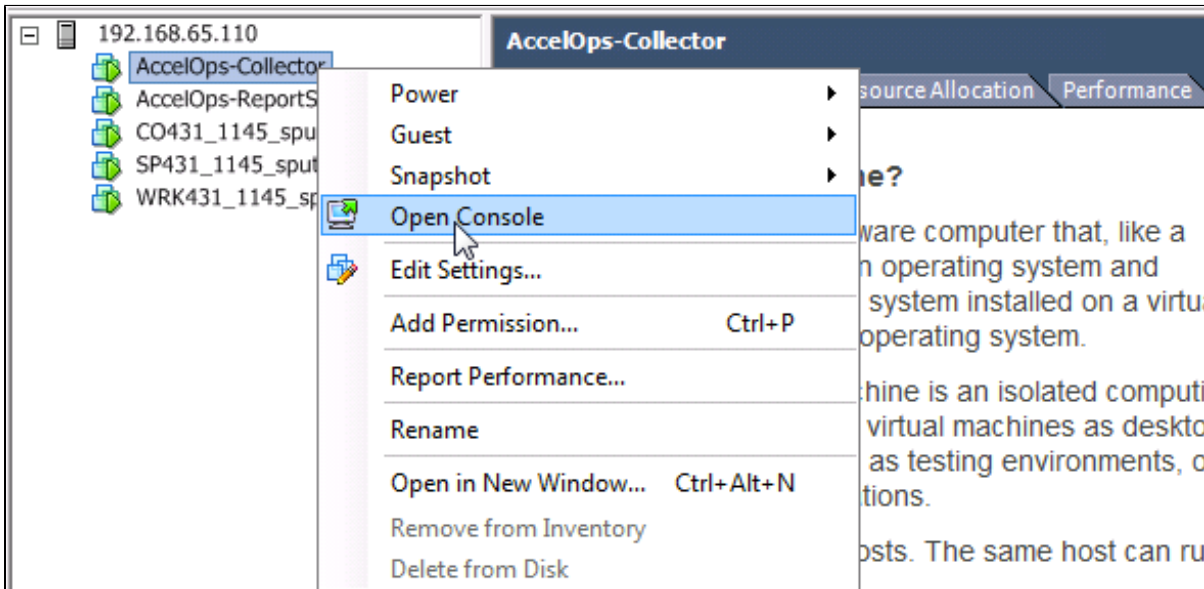
While in the VM console window do not press any control keys during the installation. Doing so might cause the installation to stop and not to precede any further. For example Ctrl – C or Ctrl – Z. If this occurs you must erase the VA and start over from the beginning.

**Step 1: Start Collector**

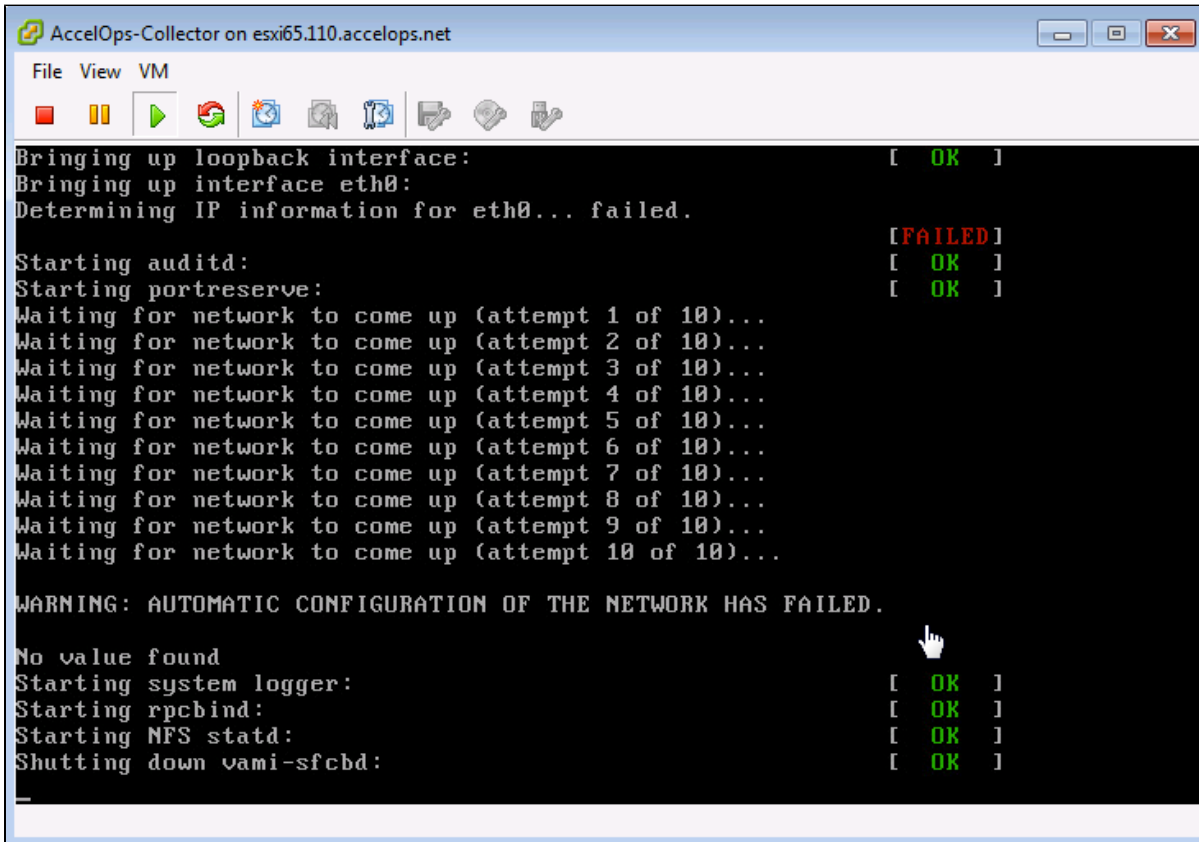
1. From the VMware vSphere client, Select recently imported Collector Virtual appliance and Power On (Right click > Power > Power On)



2. Open Collector VM Console (Right click > Open Console)



3. Network eth0 Failed message may appear on the screen during the loading, but this is to be expected the first time the AccelOps Collector is started.



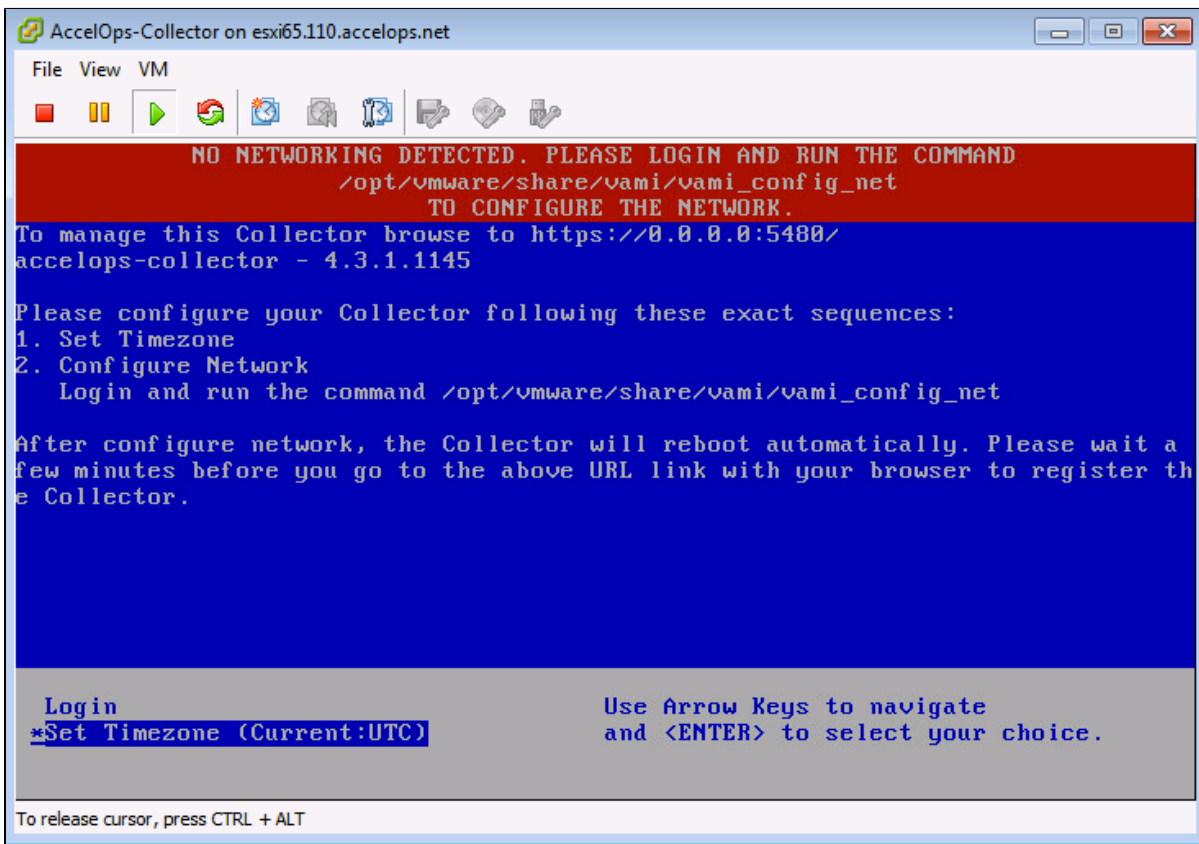
```
AccelOps-Collector on esxi65.110.accelops.net
File View VM
Bringing up loopback interface: [ OK ]
Bringing up interface eth0:
Determining IP information for eth0... failed. [ FAILED ]
Starting auditd: [ OK ]
Starting portreserve: [ OK ]
Waiting for network to come up (attempt 1 of 10)...
Waiting for network to come up (attempt 2 of 10)...
Waiting for network to come up (attempt 3 of 10)...
Waiting for network to come up (attempt 4 of 10)...
Waiting for network to come up (attempt 5 of 10)...
Waiting for network to come up (attempt 6 of 10)...
Waiting for network to come up (attempt 7 of 10)...
Waiting for network to come up (attempt 8 of 10)...
Waiting for network to come up (attempt 9 of 10)...
Waiting for network to come up (attempt 10 of 10)...
WARNING: AUTOMATIC CONFIGURATION OF THE NETWORK HAS FAILED.
No value found
Starting system logger: [ OK ]
Starting rpcbind: [ OK ]
Starting NFS statd: [ OK ]
Shutting down vami-sfcbd: [ OK ]
```

## Step 2: Configure the Timezone

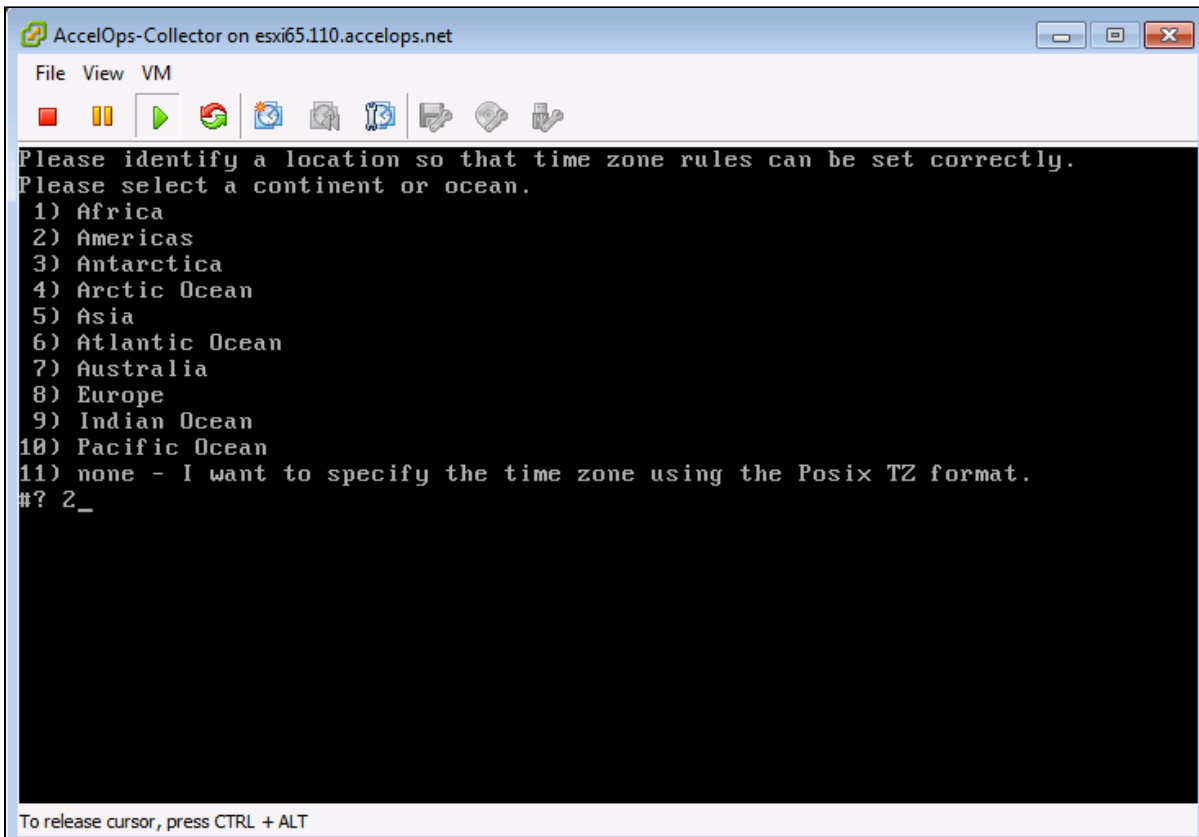
It is important to select the correct time zone that you are located in.

1. On VM console, Select 'Set Timezone' and press the 'Enter' key

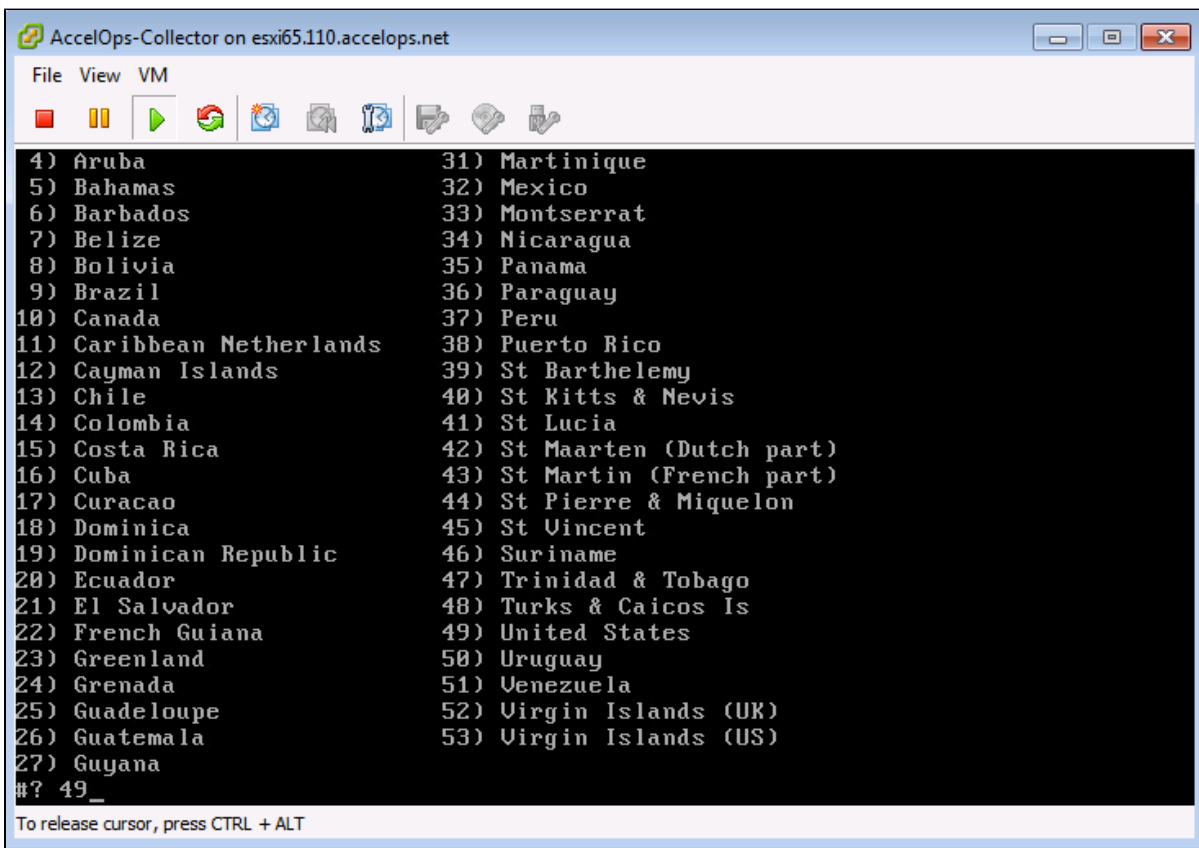




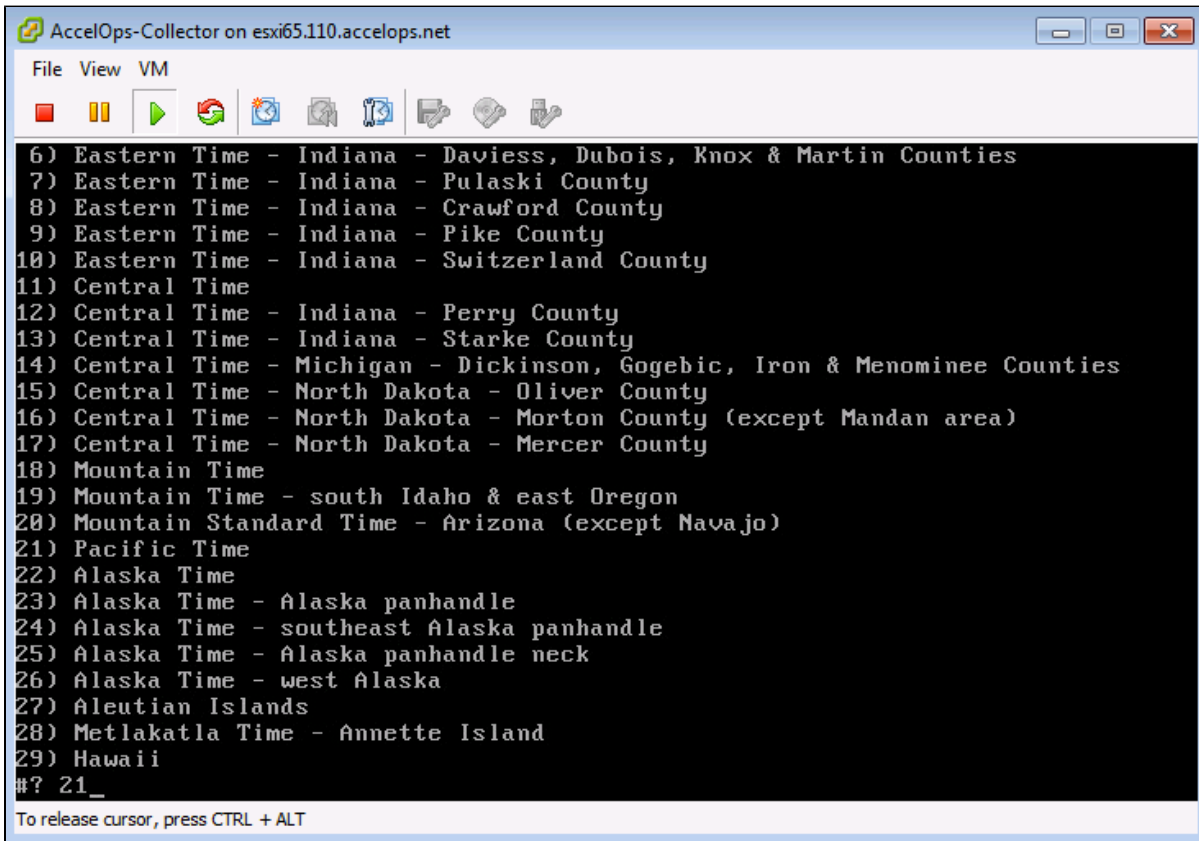
2. Select your Location and press the 'Enter' key



3. Select your Country and press the 'Enter' key



4. Select your Timezone and press the 'Enter' key

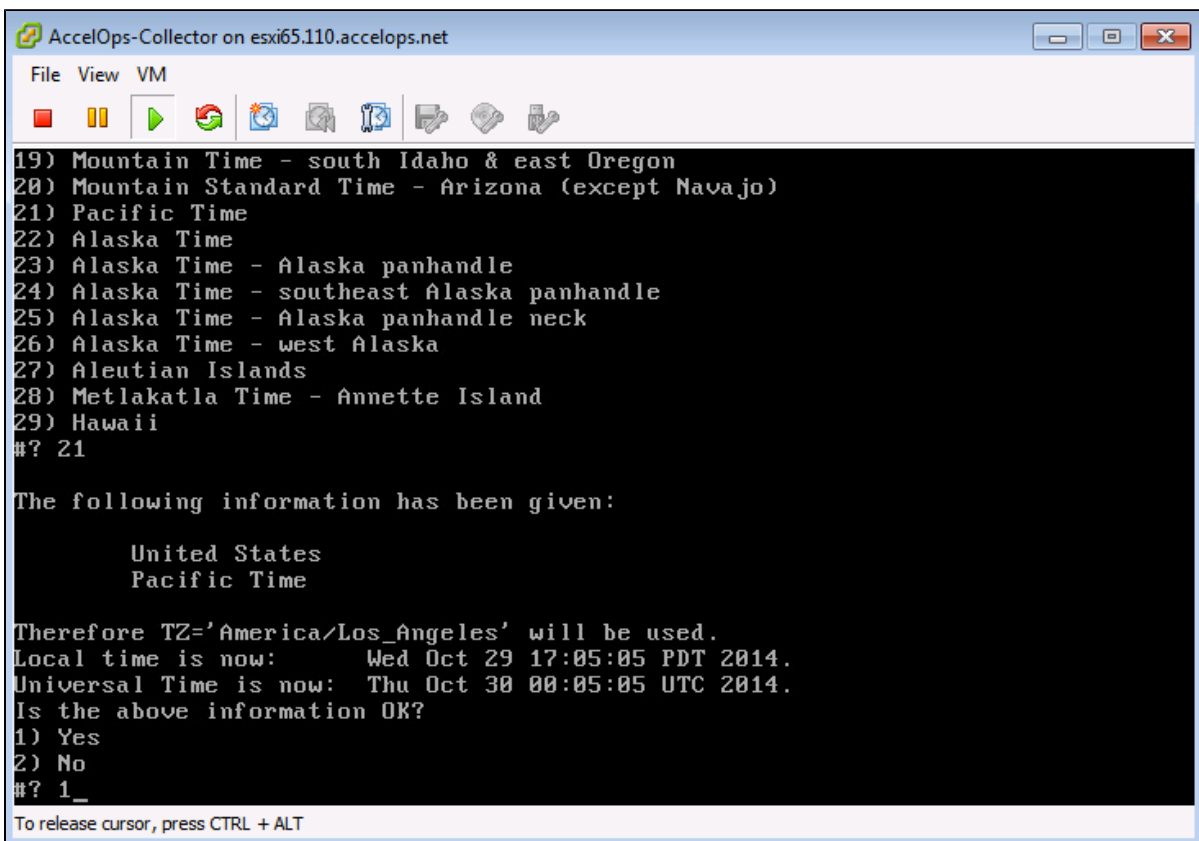


The screenshot shows a window titled "AccelOps-Collector on esxi65.110.accelops.net". The window has a menu bar with "File", "View", and "VM". Below the menu bar is a toolbar with various icons. The main content area is a black terminal window with white text listing 29 timezones. The list is as follows:

```
6) Eastern Time - Indiana - Daviess, Dubois, Knox & Martin Counties
7) Eastern Time - Indiana - Pulaski County
8) Eastern Time - Indiana - Crawford County
9) Eastern Time - Indiana - Pike County
10) Eastern Time - Indiana - Switzerland County
11) Central Time
12) Central Time - Indiana - Perry County
13) Central Time - Indiana - Starke County
14) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties
15) Central Time - North Dakota - Oliver County
16) Central Time - North Dakota - Morton County (except Mandan area)
17) Central Time - North Dakota - Mercer County
18) Mountain Time
19) Mountain Time - south Idaho & east Oregon
20) Mountain Standard Time - Arizona (except Navajo)
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - southeast Alaska panhandle
25) Alaska Time - Alaska panhandle neck
26) Alaska Time - west Alaska
27) Aleutian Islands
28) Metlakatla Time - Annette Island
29) Hawaii
#? 21_
```

At the bottom of the terminal window, it says "To release cursor, press CTRL + ALT".

5. Review and Confirm your Timezone selection



The screenshot shows a terminal window titled "AccelOps-Collector on esxi65.110.accelops.net". The terminal displays a list of time zones from 19 to 29, with option 21 selected. It then shows the selected time zone information and the local and universal times. The user has responded with '1' to confirm the configuration.

```
AccelOps-Collector on esxi65.110.accelops.net
File View VM
[Icons]
19) Mountain Time - south Idaho & east Oregon
20) Mountain Standard Time - Arizona (except Navajo)
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - southeast Alaska panhandle
25) Alaska Time - Alaska panhandle neck
26) Alaska Time - west Alaska
27) Aleutian Islands
28) Metlakatla Time - Annette Island
29) Hawaii
#? 21

The following information has been given:

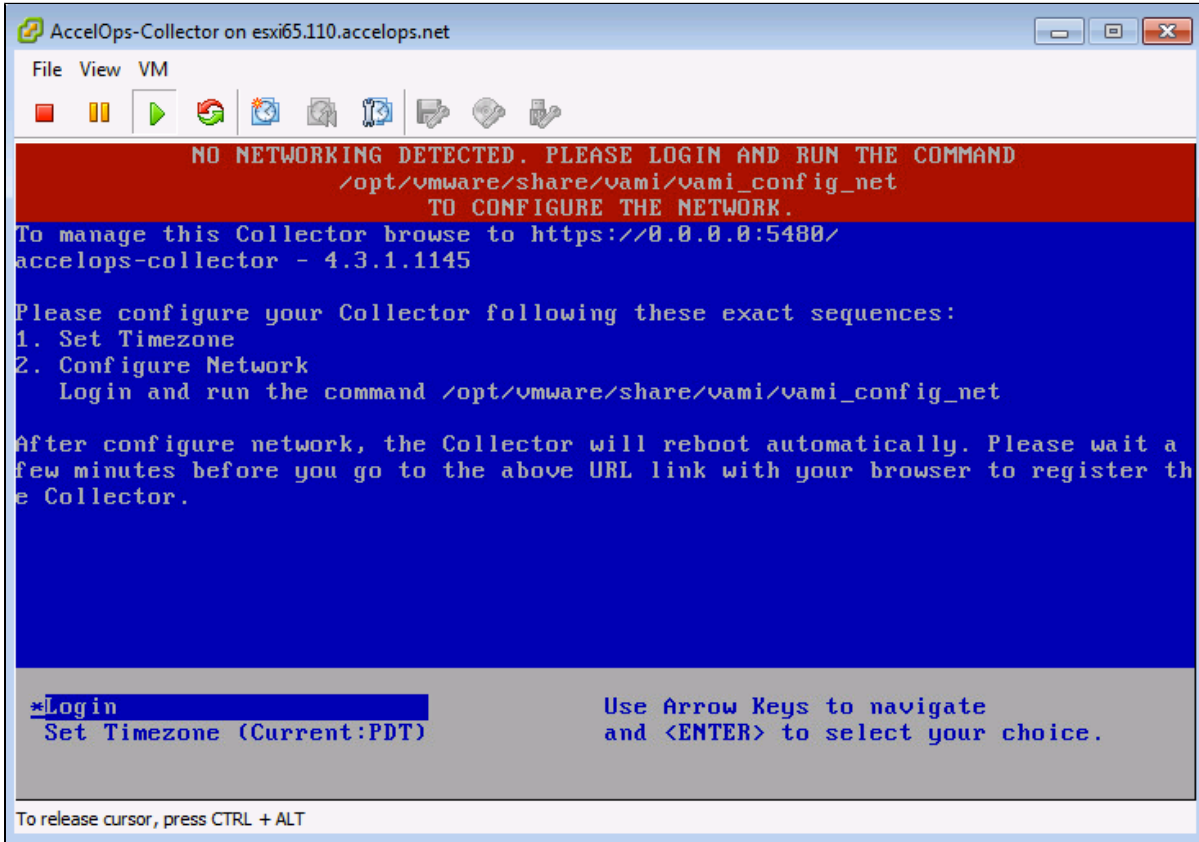
    United States
    Pacific Time

Therefore TZ='America/Los_Angeles' will be used.
Local time is now:    Wed Oct 29 17:05:05 PDT 2014.
Universal Time is now: Thu Oct 30 00:05:05 UTC 2014.
Is the above information OK?
1) Yes
2) No
#? 1_

To release cursor, press CTRL + ALT
```

### Step 3: Configure the Network

1. On VM console, Select 'Login' and press the 'Enter' key



- 2. Enter default log in information.  
 Login: root  
 Password: ProspectHills



- 3. Run vami\_config\_net script to configure network and installing AccelOps Collector

**Configure the Network**

```
/opt/vmware/share/vami/vami_config_net
```

```

AccelOps-Collector on esxi65.110.accelops.net
File View VM
localhost login: root
Password:
[root@localhost ~]# /opt/vmware/share/vami/vami_config_net

CHANGE NETWORK CONFIGURATION for eth0

Please enter the desired network parameters.
To exit, type q at any prompt.

IP Address [1]: _

```

4. Configure Network with Static IP address
  - a. The following information is required to configure Static IP address:
    - i. IP Address
    - ii. Netmask
    - iii. Gateway
    - iv. DNS Server(s)

#### Proxy Server

Do not configure the proxy server setting.

- b. Press the 'y' key to accept the changes
- c. Enter 'Host name' and press 'Enter' key

```

AccelOps-Collector on esxi65.110.accelops.net
File View VM
To exit, type q at any prompt.

IP Address [192.168.65.146]: 192.168.65.146
Netmask [255.255.252.0]: 255.255.252.0
Gateway [192.168.64.1]: 192.168.64.1
DNS Server 1 [192.168.0.10]: 192.168.0.10
DNS Server 2 [192.168.0.40]: 192.168.0.40
Is a proxy server necessary to reach the Internet? y/n [n]:

IP Address:      192.168.65.146
Netmask:        255.255.252.0
Gateway:        192.168.64.1
Proxy Server:
DNS Servers:    192.168.0.10, 192.168.0.40

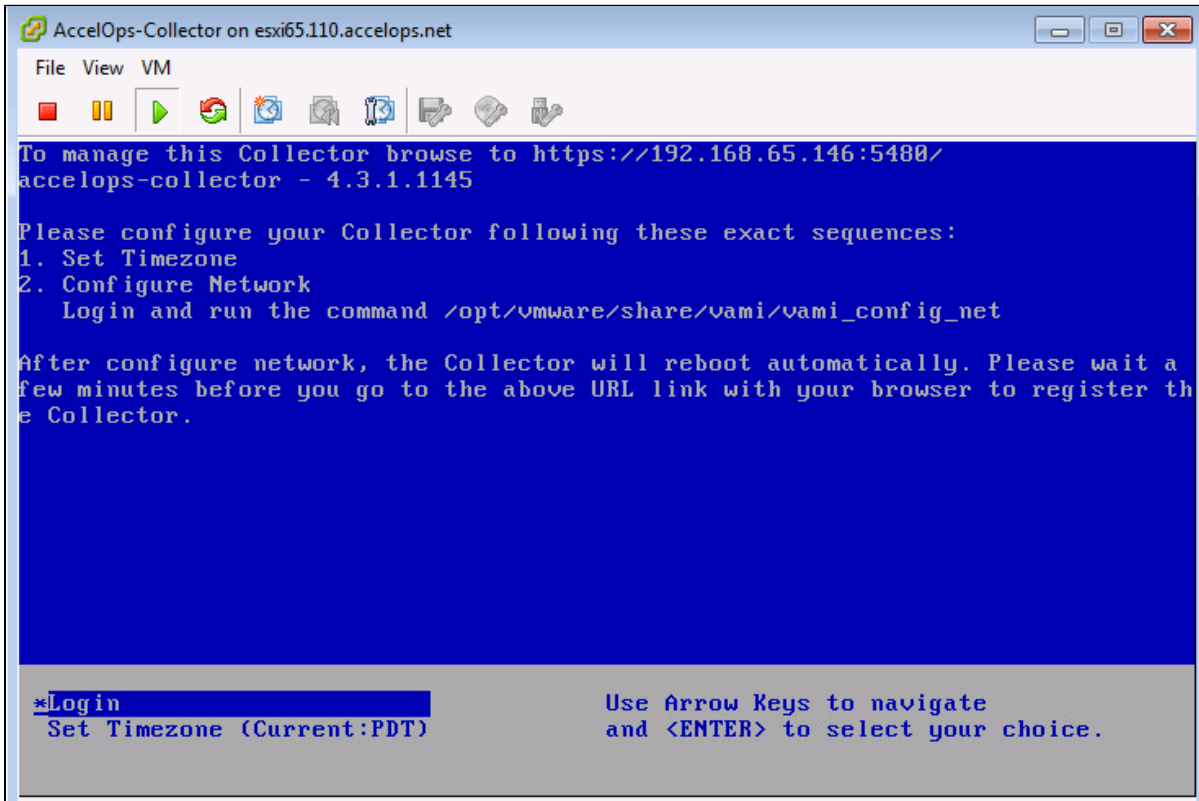
Is this correct? y/n [y]: y

Reconfiguring the network...
DNS server settings updated
Determining if ip address 192.168.65.146 is already in use for device eth0...
vami_login: no process killed
Network parameters successfully changed to requested values
Input hostname: collector.accelops.net_

```

5. After configuring network, Collector will reboot automatically

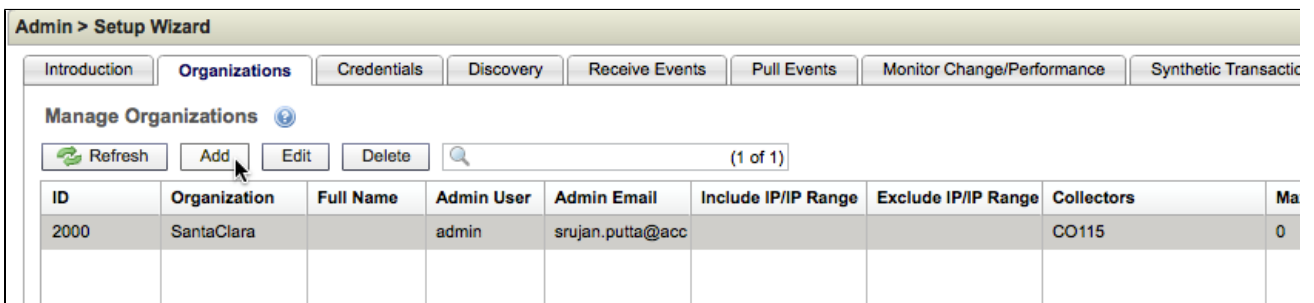
In 3 to 5mins, Collector will be configure successful.



## Section 4: Register the Collector to Supervisor

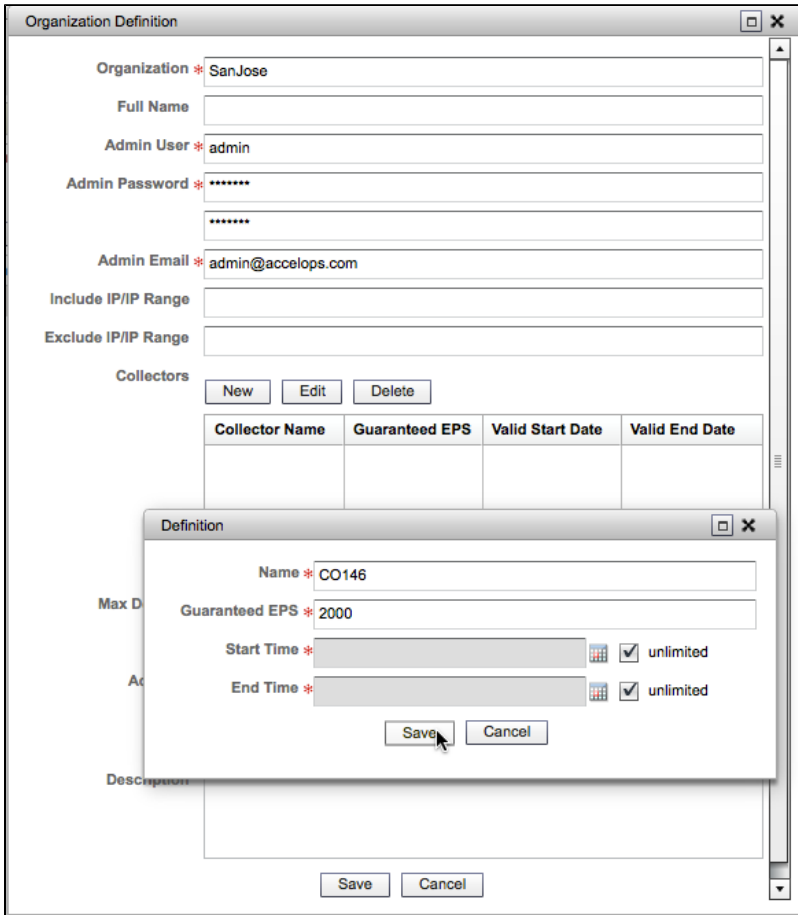
### Step 1: Create an Organization on Supervisor UI

1. Logon to Supervisor UI
2. Go to Admin > Setup Wizard > Organizations > click 'Add'

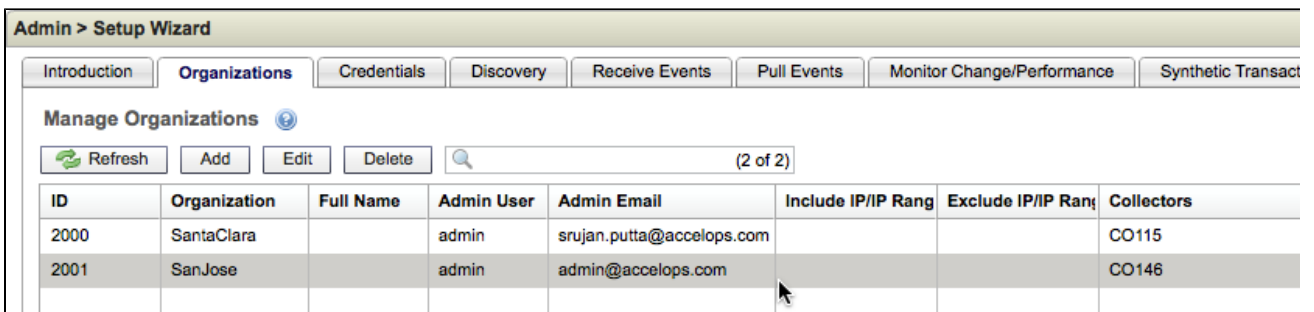


3. Enter Organization Name, Admin User, Admin Password
  - On Collectors section, click 'New' to add Collector enter Collector Name and license parameters
    - a. Enter Collector Name
    - b. Enter EPS limit
    - c. Select valid Start and End times
    - d. Click Save





4. Review recently added Organization and Collector names



**Step 2: Register the Collector to Supervisor**

1. In browser windows, enter 'https://<Collector-IP>:5480'

https://192.168.65.146:5480

Google

### Collector Setup

Name:

User ID:

Password:

Cust/Org ID:

Cloud URL:  Save

2. Enter Collector setup information
  - a. Name: Collector Name (defined in above Step 1.3)
  - b. User ID: Admin User (defined in above Step 1.3)
  - c. Password: Admin Password (defined in above Step 1.3)
  - d. Cust/Org ID: Organization Name (defined in above Step 1.3)
  - e. Cloud URL: Supervisor URL (eg: https://<Supervisor-IP>)
  - f. Click Save

https://192.168.65.146:5480

Google

### Collector Setup


Name:

User ID:

Password:

Cust/Org ID:

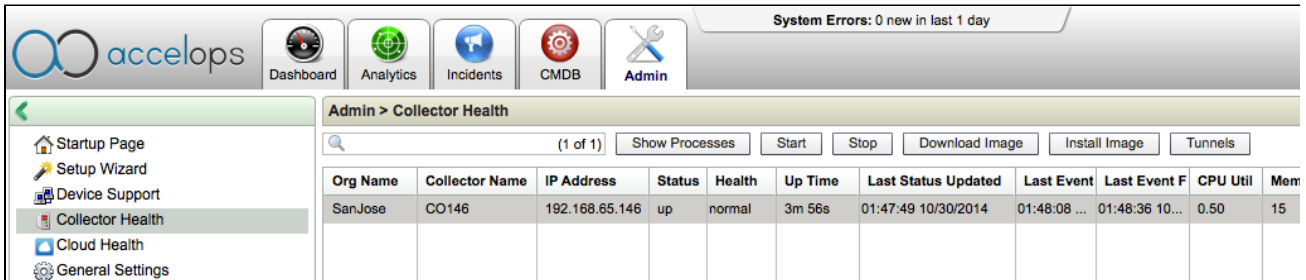
Cloud URL:  Save



3. Collector will reboot automatically after registration successful.

**Step 3: Check Collector health on Supervisor UI**

1. Go to Admin > Collector Health, make sure Collector status is Up and health is Normal.

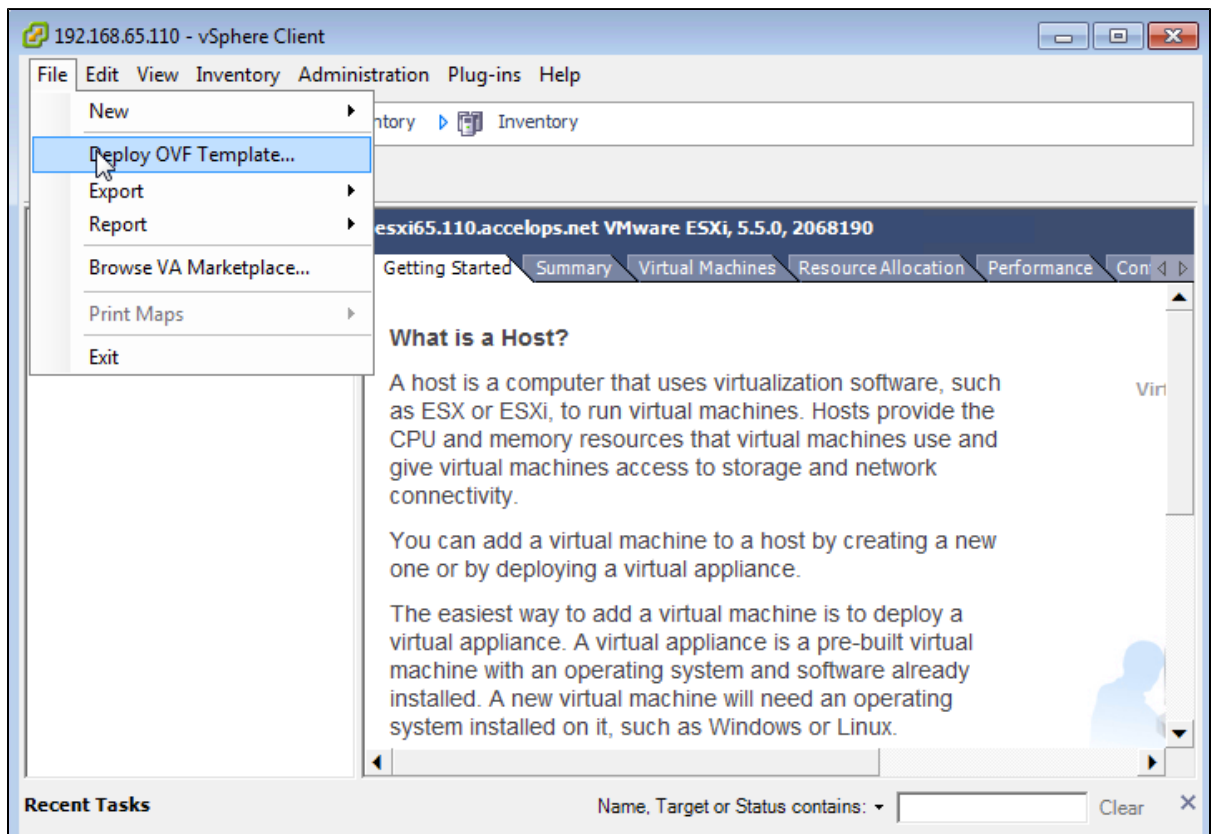


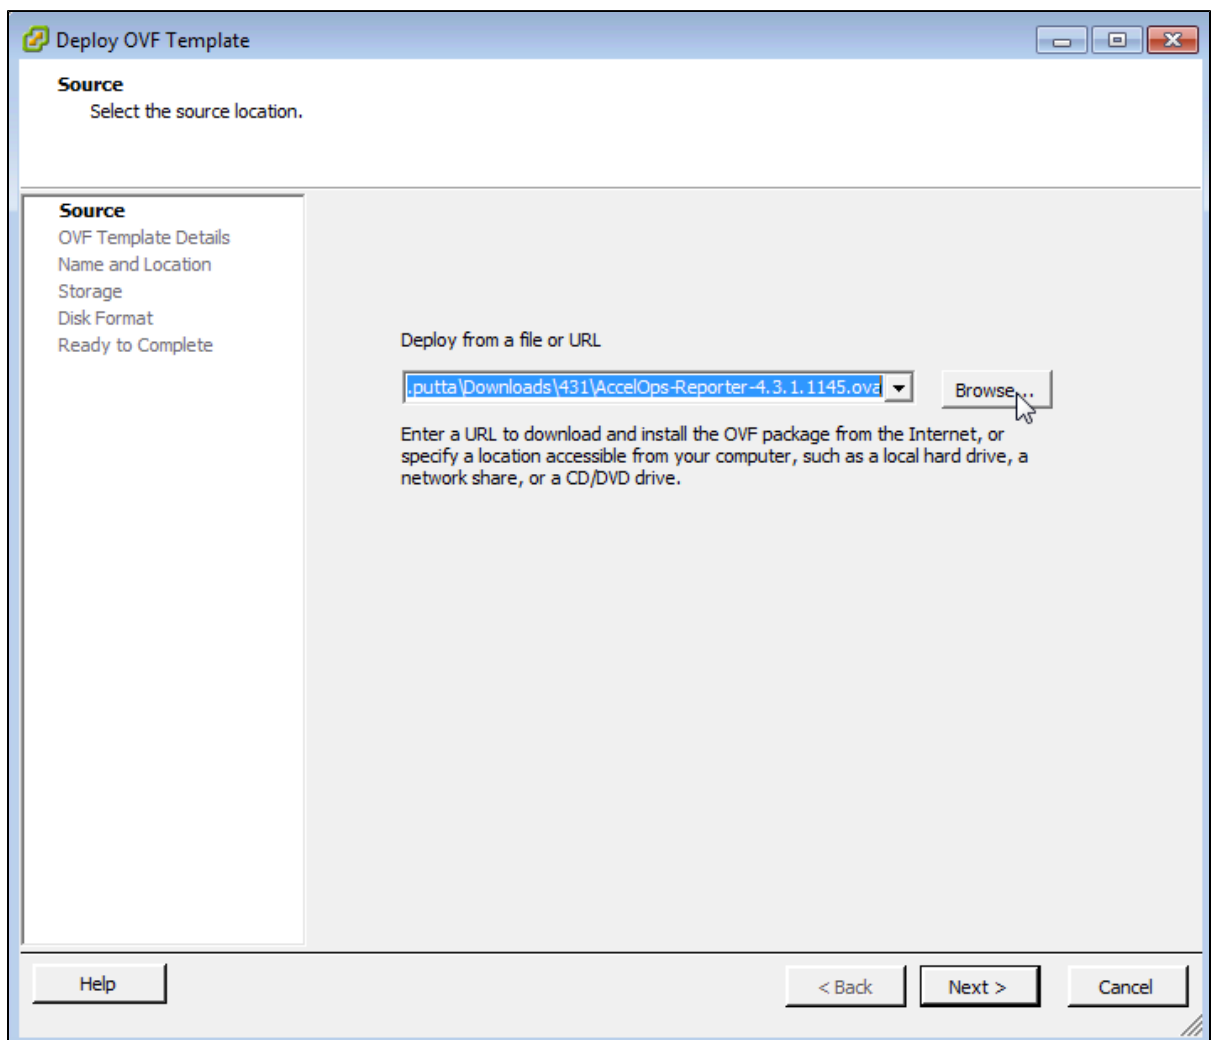
**Installing a Report Server node**

**Section 1: Import a Report Server image into VMware ESX server**

**Step 1: Begin importing Report Server**

1. Download Report server OVA package from AccelOps image server
2. Log into VMware vSphere Client
  - a. Click on File -> Select Deploy OVF Template
  - b. Click on Browse button to choose OVA file
  - c. Locate Report Server .ova file (Example: AccelOps-Reporter-4.3.1.1145.ova)



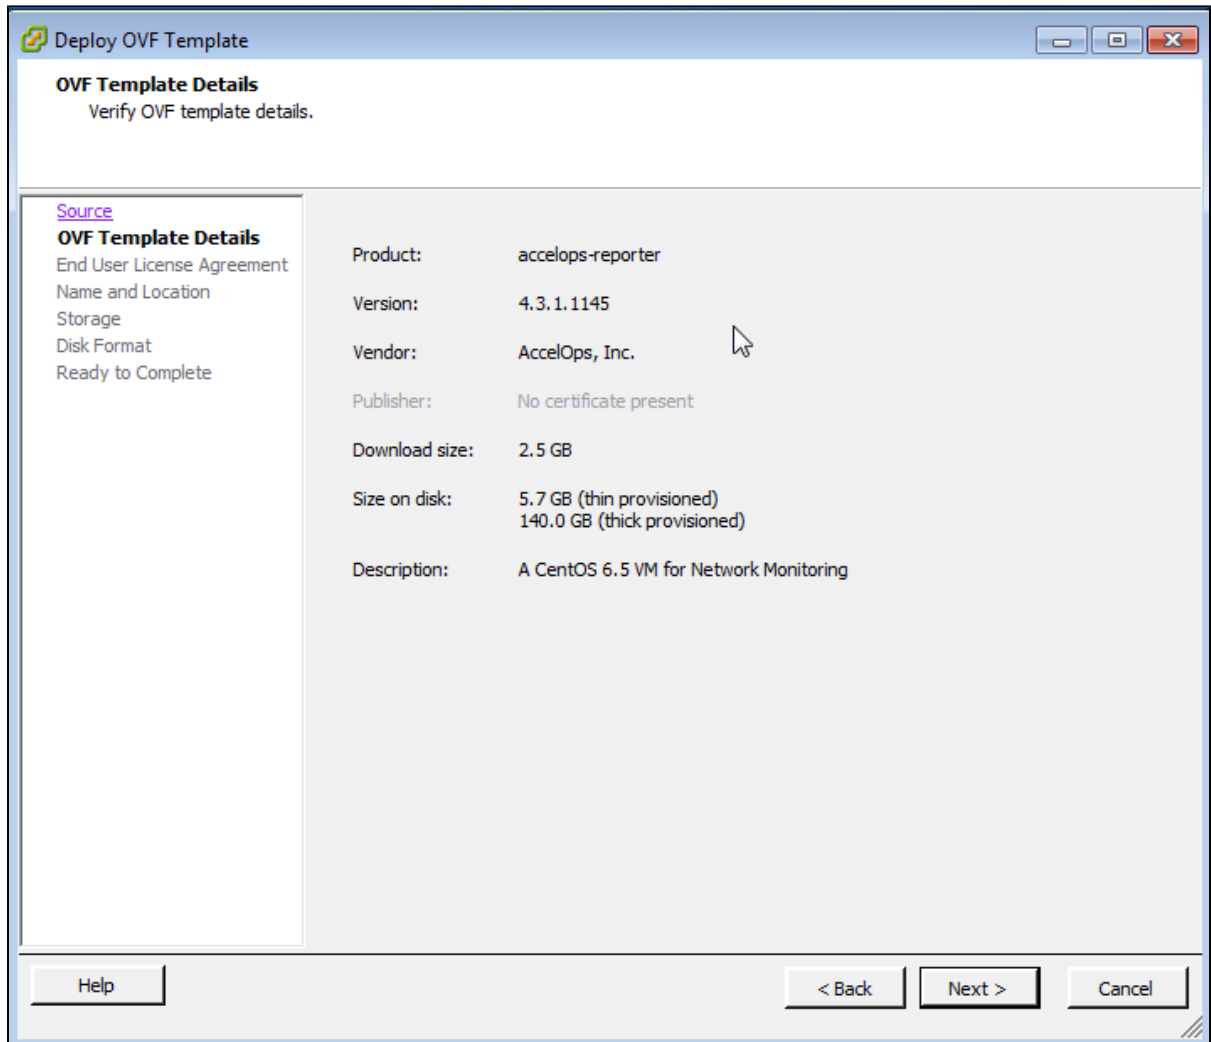


### **Step 2: Review the Virtual Appliance details**

OVF Template details will tell you how big the file is to import and how much disk space it will take up in your selected ESX server.

-

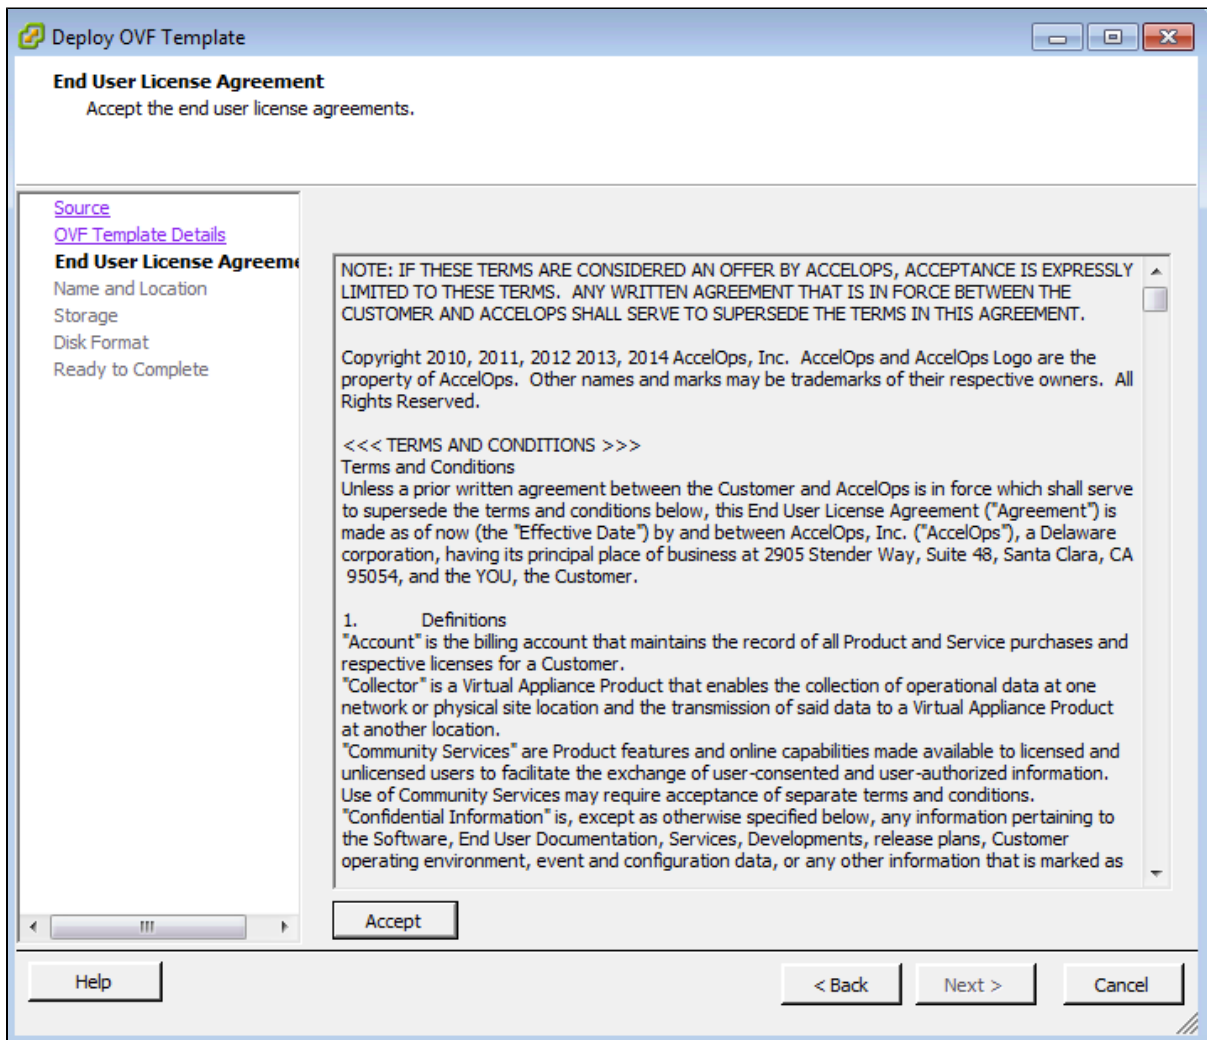
1. Click the 'Next' button



### **Step 3: End User License Agreement**

To continue you must accept the End User License Agreement.

1. Click on 'Accept' button on License Agreement
2. Click on 'Next' button



**Step 4: Name and Location**

Enter a name for the Virtual Appliance that is being imported. This name is independent of the host name and is only visible within the ESX environment.

1. Enter a Report Server name and Click on 'Next' button

**Deploy OVF Template**

**Name and Location**  
Specify a name and location for the deployed template

[Source](#)  
[OVF Template Details](#)  
[End User License Agreement](#)  
**Name and Location**  
Storage  
Disk Format  
Ready to Complete

Name:  
AccelOps-ReportServer

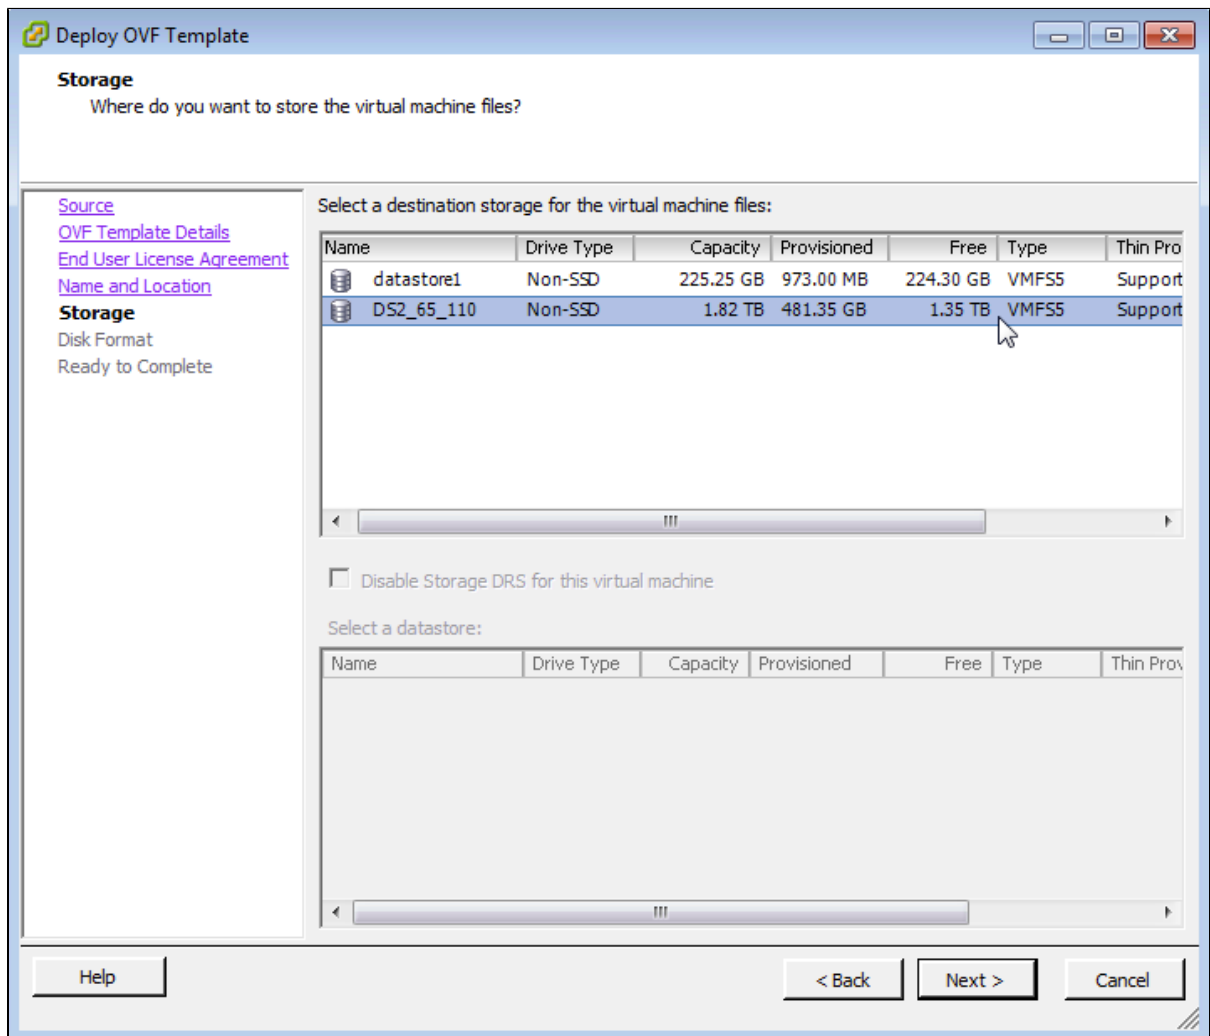
The name can contain up to 80 characters and it must be unique within the inventory folder.

Help < Back Next > Cancel

### Step 5: Storage

1. Select the desired storage location(Data store) and Click on 'Next' button





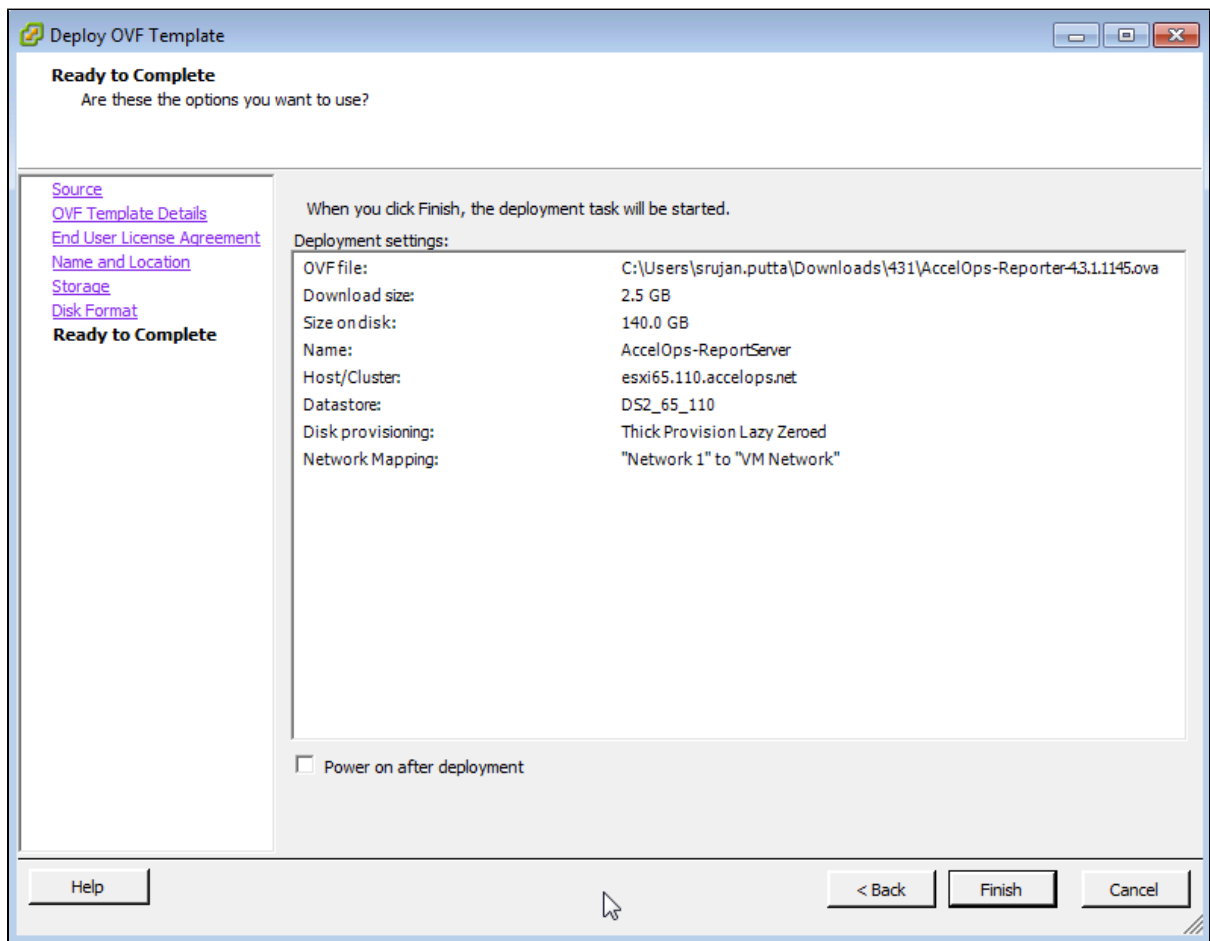
**Step 6: Disk Format**

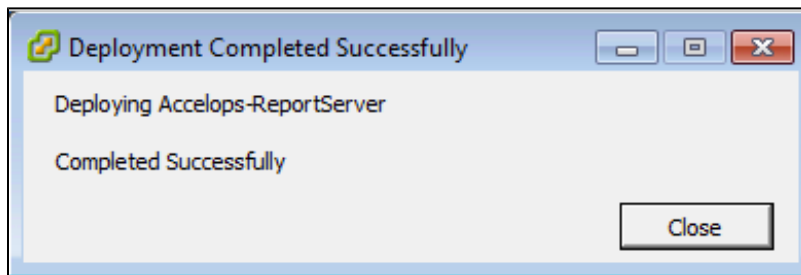
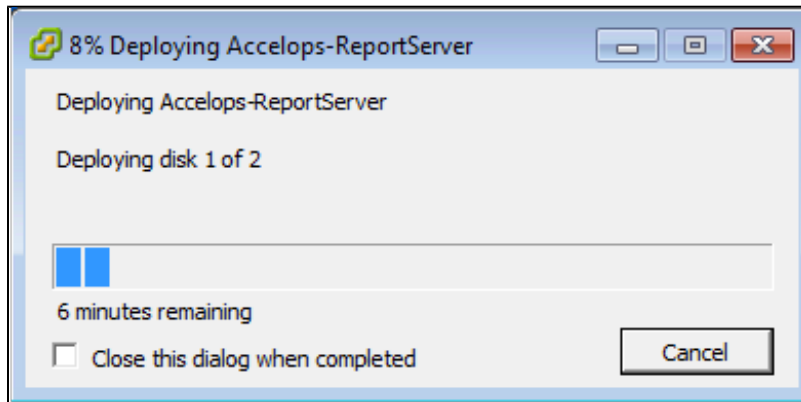
1. Select the Disk format as 'Thick Provision' (recommended) and Click on 'Next' button

The screenshot shows a window titled "Deploy OVF Template" with a sub-header "Disk Format". Below the sub-header is the question "In which format do you want to store the virtual disks?". On the left side, there is a navigation pane with links for "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Storage", and "Disk Format". The "Disk Format" section is currently selected and shows "Ready to Complete". The main area contains two input fields: "Datastore:" with the value "DS2\_65\_110" and "Available space (GB):" with the value "1381.4". Below these fields are three radio button options: "Thick Provision Lazy Zeroed" (which is selected), "Thick Provision Eager Zeroed", and "Thin Provision". At the bottom of the window, there are three buttons: "Help", "< Back", and "Next >" (with a mouse cursor over it), and a "Cancel" button.

### Step 7: Ready to Complete

1. Review the Report server resources and Click on 'Finish' button to begin importing Virtual Appliance
2. Virtual Appliance import will takes 7 to 10mins to complete. Do not turn off or reboot the system during this time.



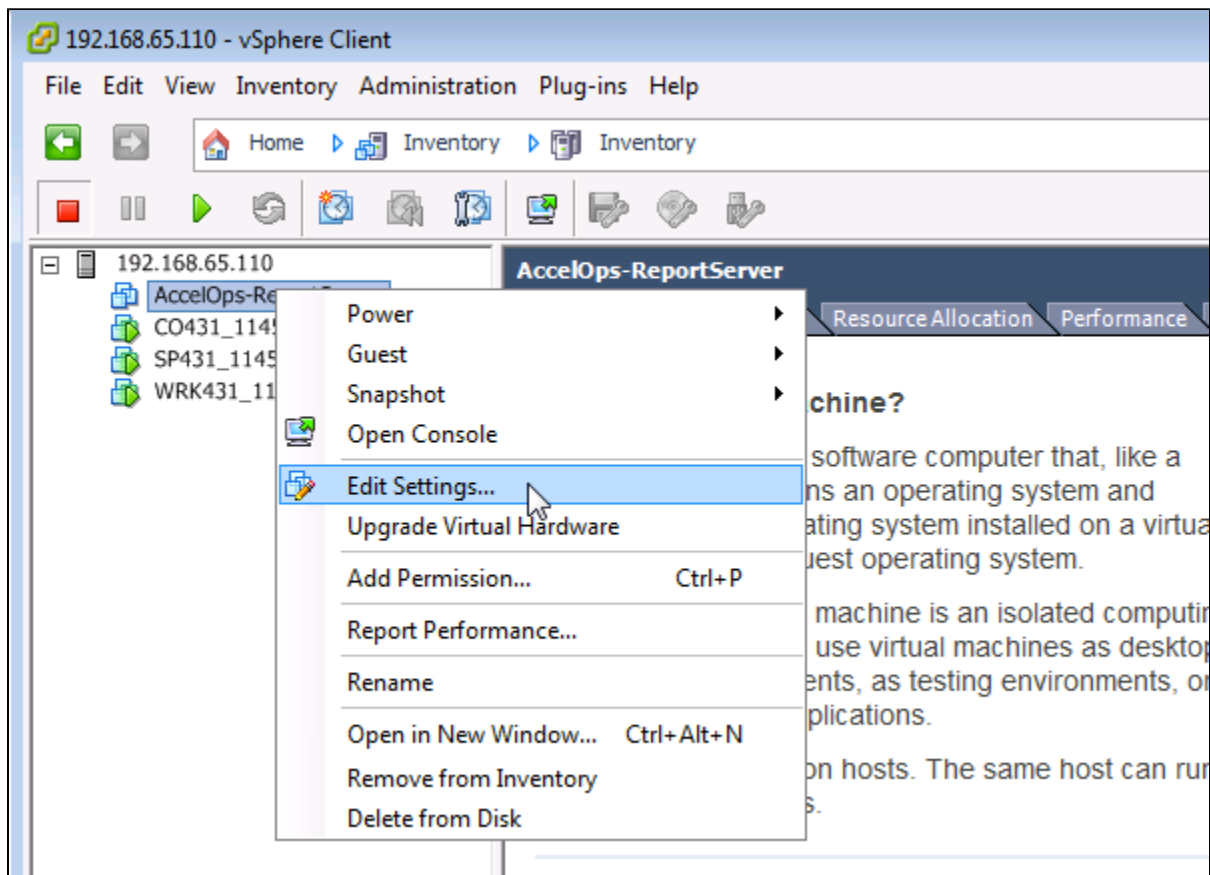


## Section 2: Edit the Report Server hardware settings

Prior to starting the Report Server for the first time it is required to make some modifications to the hardware setting for the Virtual Appliance.

### **Step 1: Edit Settings**

1. From the VMware vSphere client, Select the imported Report server Virtual appliance
2. Right mouse click and select Edit Settings



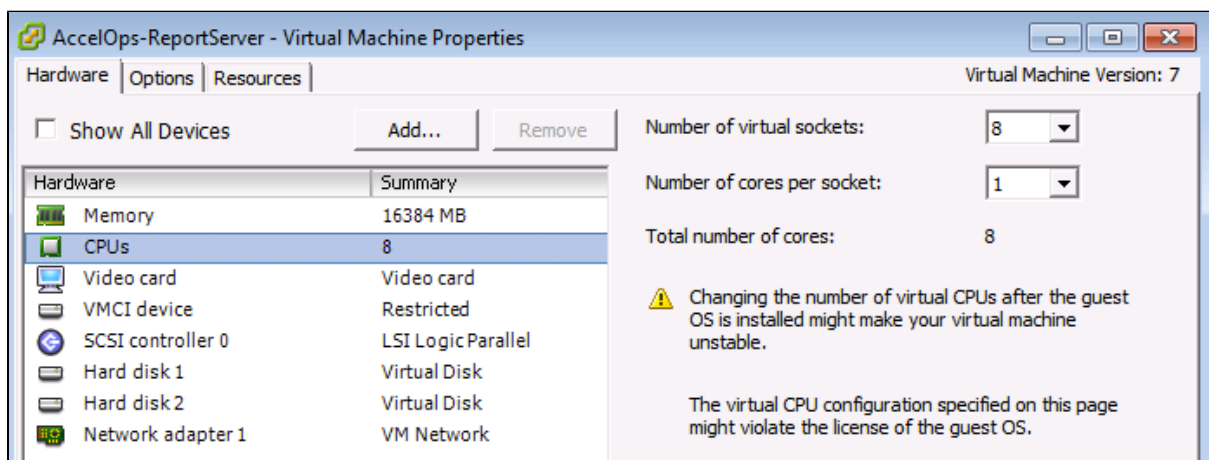
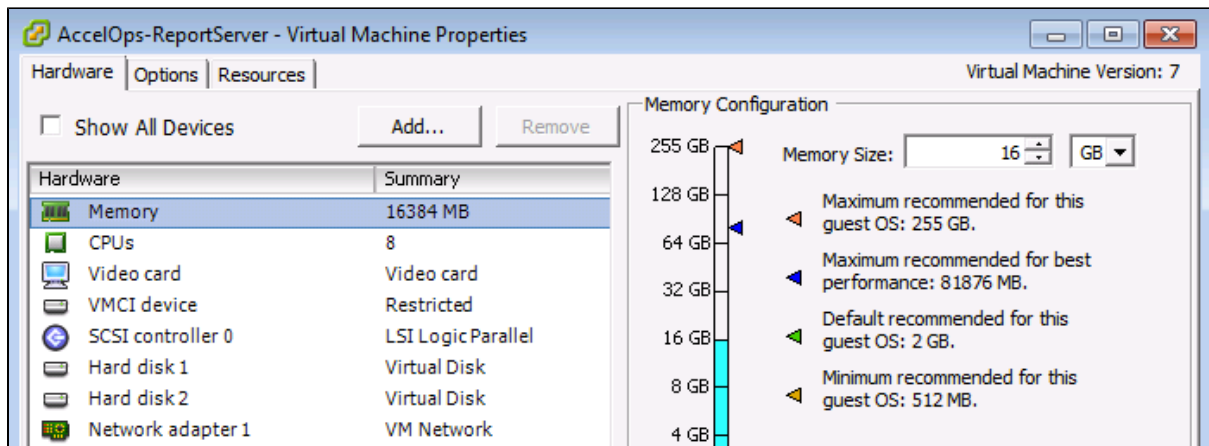
### Step 2: Verify Allocated Memory

1. On Hardware tab, Click on Memory and verify the memory to at least 16 GB

#### Allocate more memory for large deployments

For large deployments you should allocate at least 24GB of memory. See the [section on Report Server Hardware requirements](#) in the topic [Must Read Prior to Install](#) for more information

2. Click on CPUs, verify the CPU cores to at least 8 cores



### Section 3: Select mount point to Store Reports data

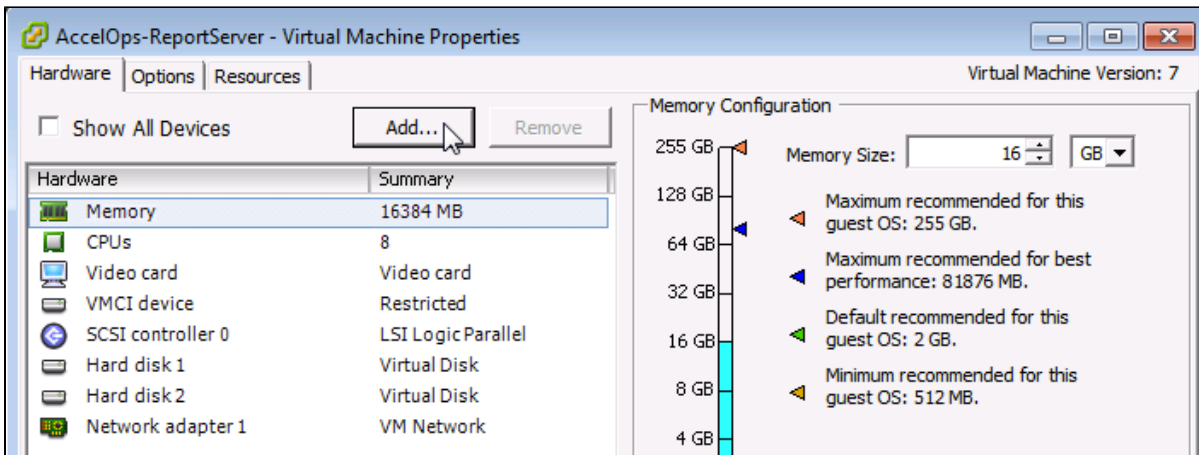
The AccelOps Report Server can be installed using either storage configured within the ESX server or NFS storage.

**Note**

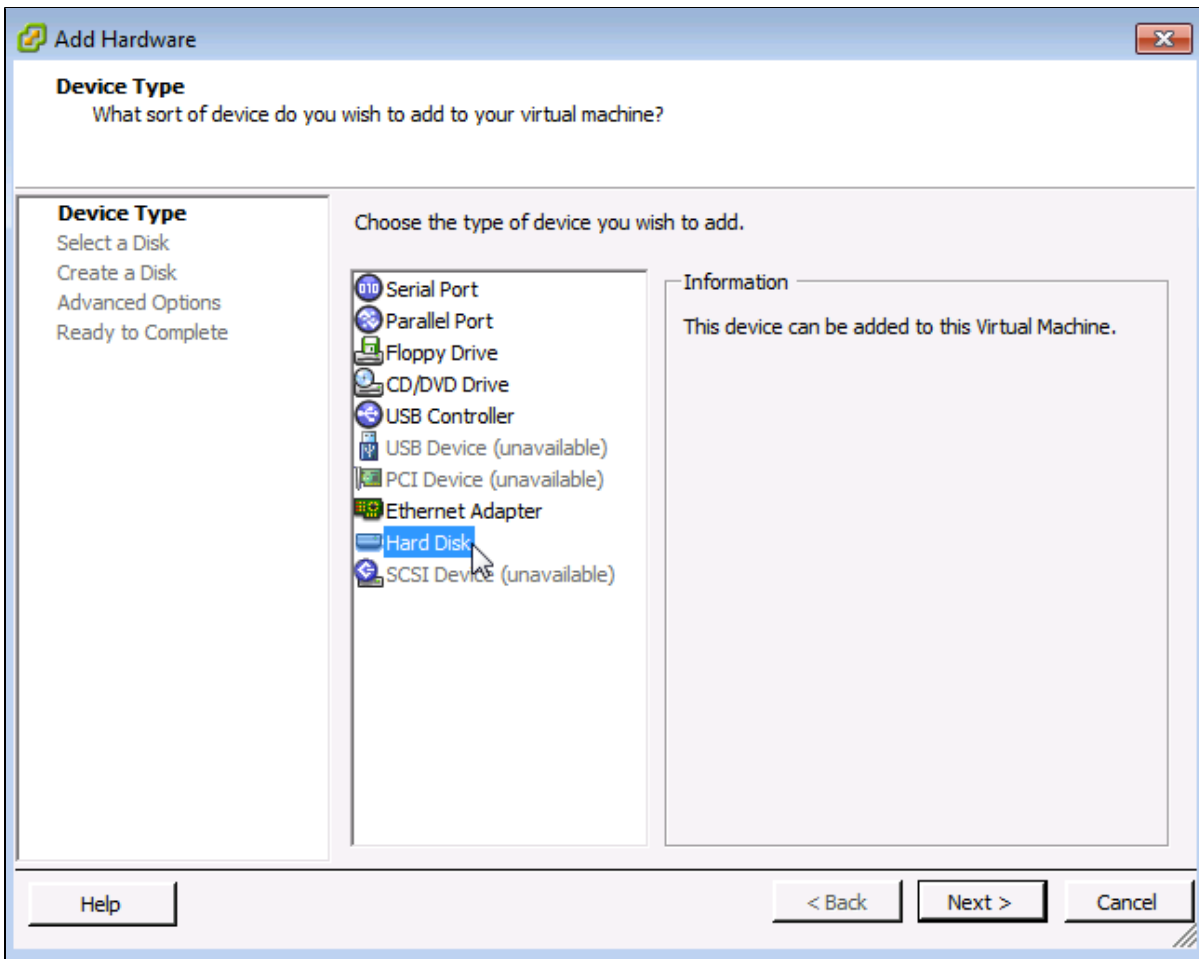
Skip Step 1 if AccelOps Super is used NFS storage to store the events data. Else, continue Step 1 to create local disk to store the reports data.

#### Step 1: Adding an Additional Hard Disk to store the Reports data

1. On Hardware tab, Click on the 'Add' button

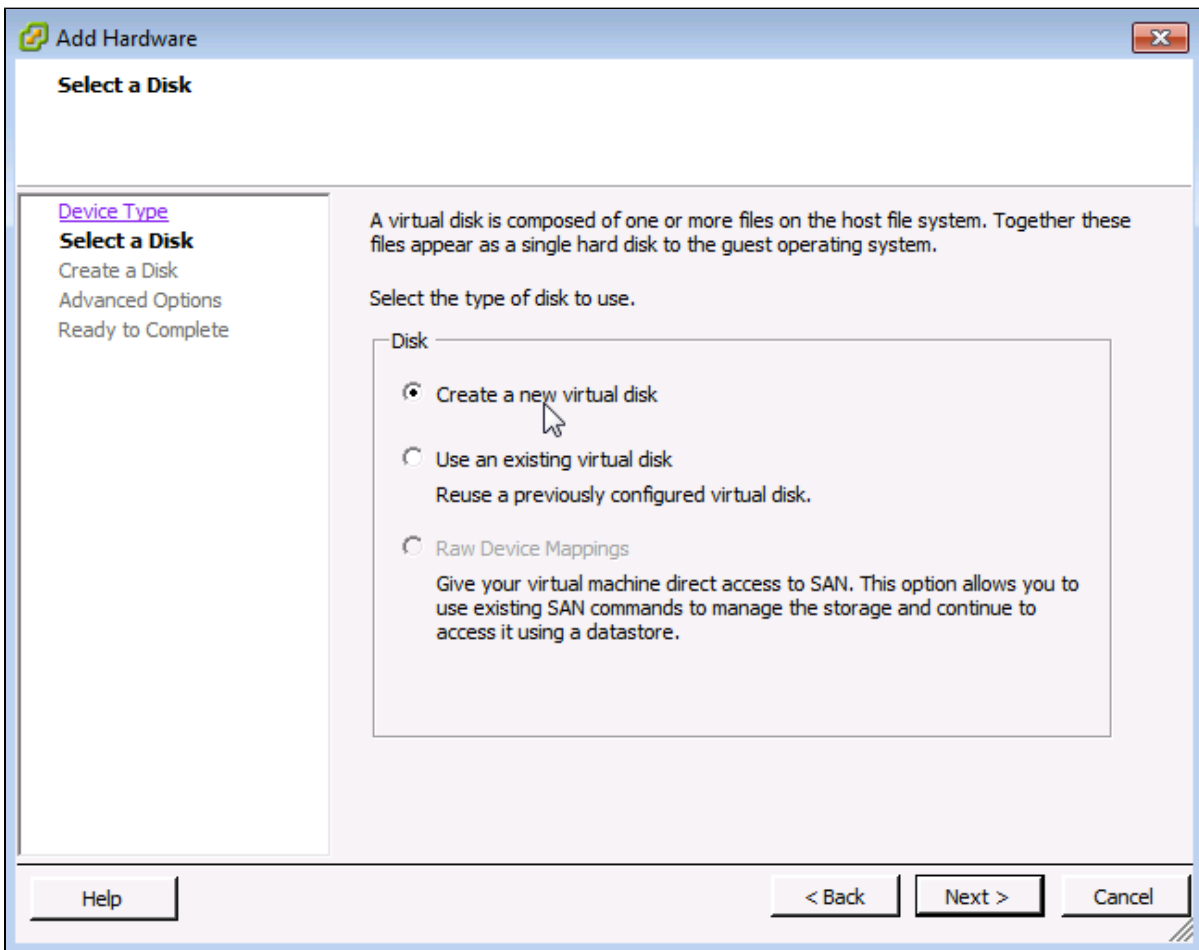


2. On Hardware pop-up, Select 'Hard Disk' and click on 'Next' button



3. Select 'Create a new virtual disk' and click on 'Next' button

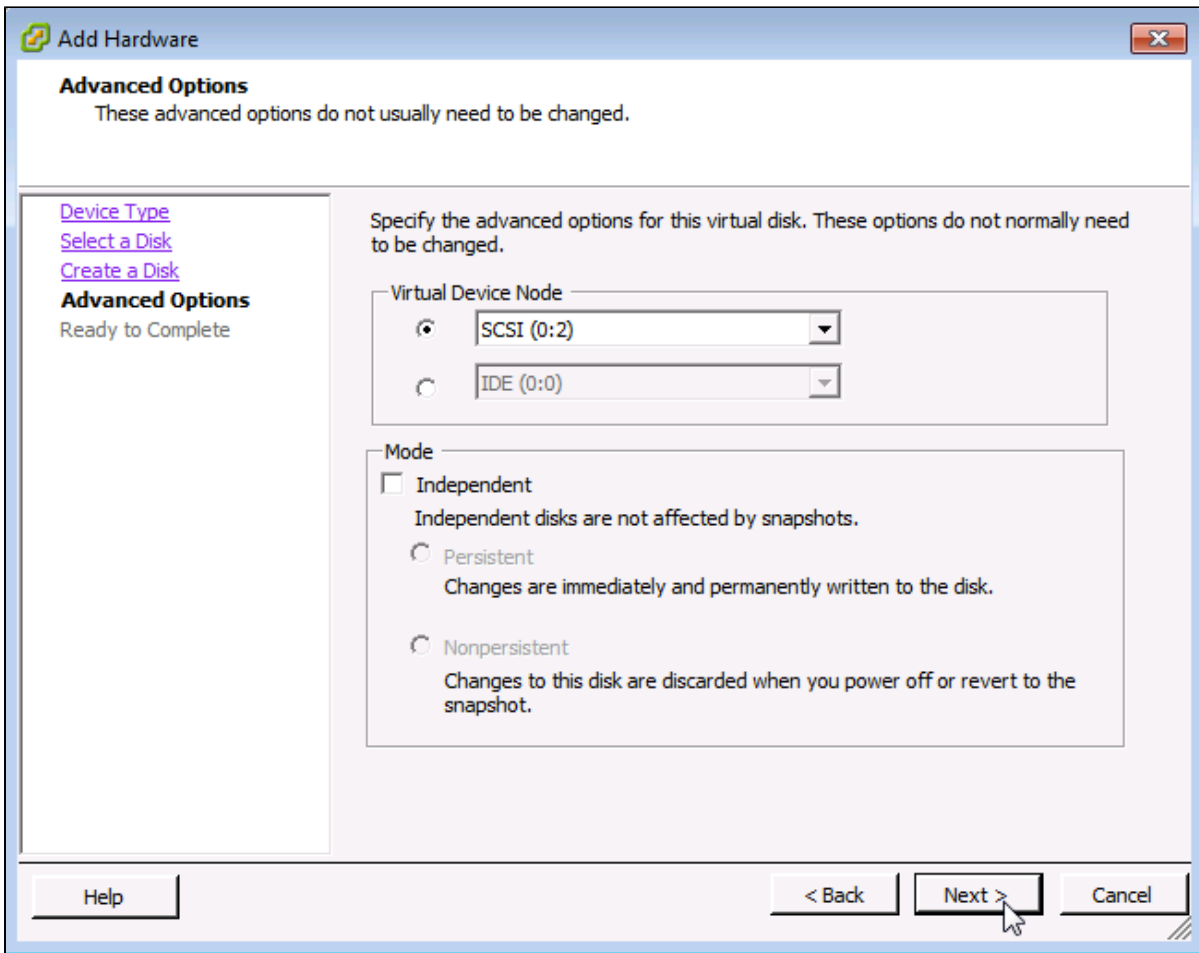




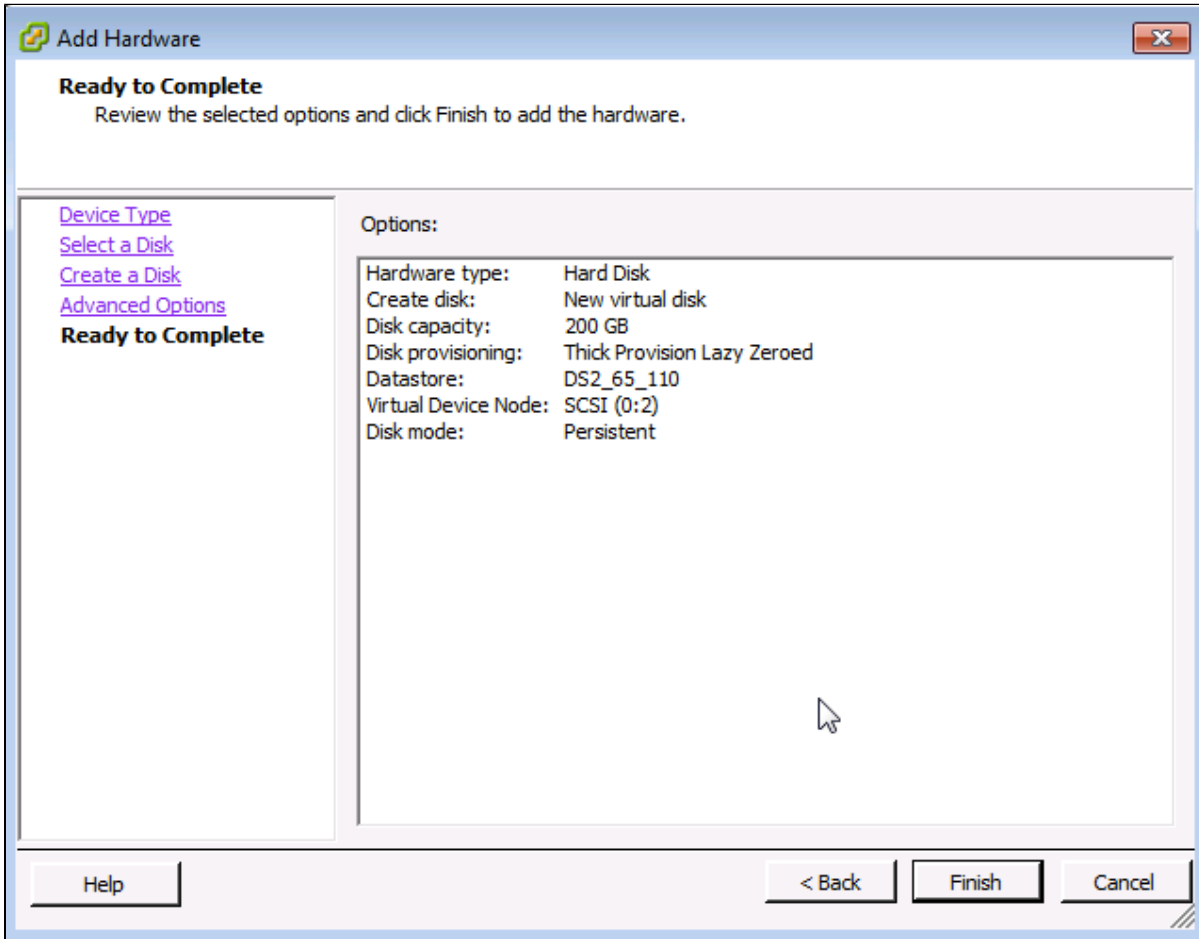
4. Enter the Disk size and click on 'Next' button  
Refer to Hardware Requirements for Report Server nodes

The screenshot shows a 'Create a Disk' dialog box within the 'Add Hardware' window. The dialog is titled 'Create a Disk' and has a subtitle 'Specify the virtual disk size and provisioning policy'. On the left side, there are links for 'Device Type' and 'Select a Disk', and a section for 'Create a Disk' with sub-sections 'Advanced Options' and 'Ready to Complete'. The main area is divided into three sections: 'Capacity' with a 'Disk Size' field set to '200' and a unit dropdown set to 'GB'; 'Disk Provisioning' with three radio button options: 'Thick Provision Lazy Zeroed' (selected), 'Thick Provision Eager Zeroed', and 'Thin Provision'; and 'Location' with two radio button options: 'Store with the virtual machine' and 'Specify a datastore or datastore cluster:' (selected). Below the second option is a text field containing 'DS2\_65\_110' and a 'Browse...' button. At the bottom of the dialog, there are three buttons: 'Help', '< Back', and 'Next >', with a mouse cursor pointing to the 'Next >' button. A 'Cancel' button is also present to the right of 'Next >'.

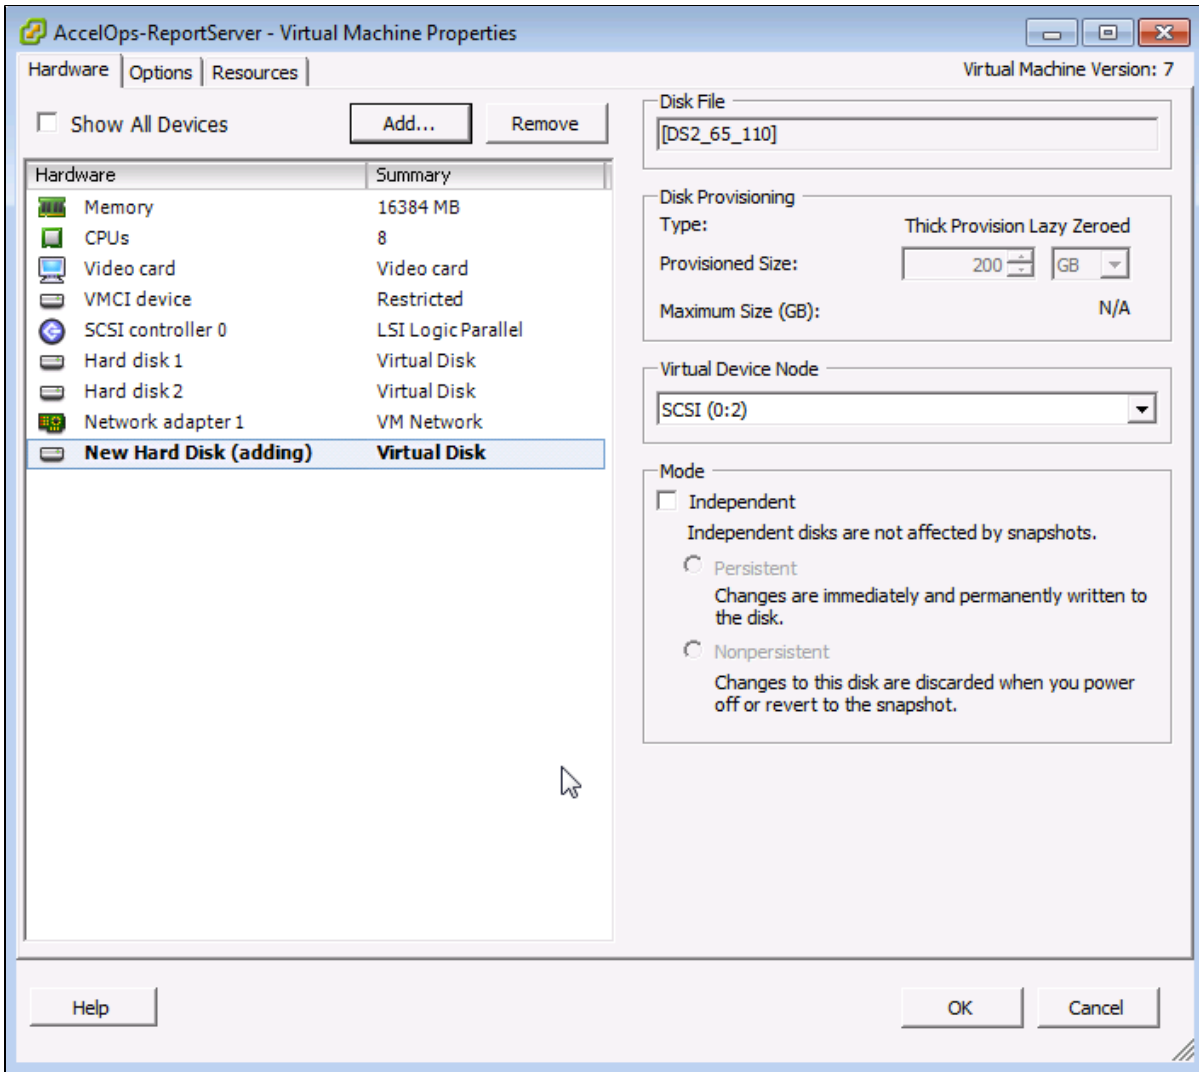
5. Ensure that 'Independent' is not checked and click on 'Next' button




6. Confirm the Options are correct and click the 'Finish' button



7. Click 'OK' on Virtual Machine properties pop-up



8. Wait for until Adding Disk task complete on vSphere client



**Note**

Skip Step 2 if you choose Step 1 to store the AccelOps event data in local Hard disk. Else, continue Step 2 to create NFS mount point to store the reports data.

**Step 2: Create new NFS mount point to store the Reports data**

1. Log into NFS server, Create new directory to store Reports data

**Note**

Do not use AccelOps Supervisor mount point location.

2. Refer to NFS documentation to check NFS directory permission

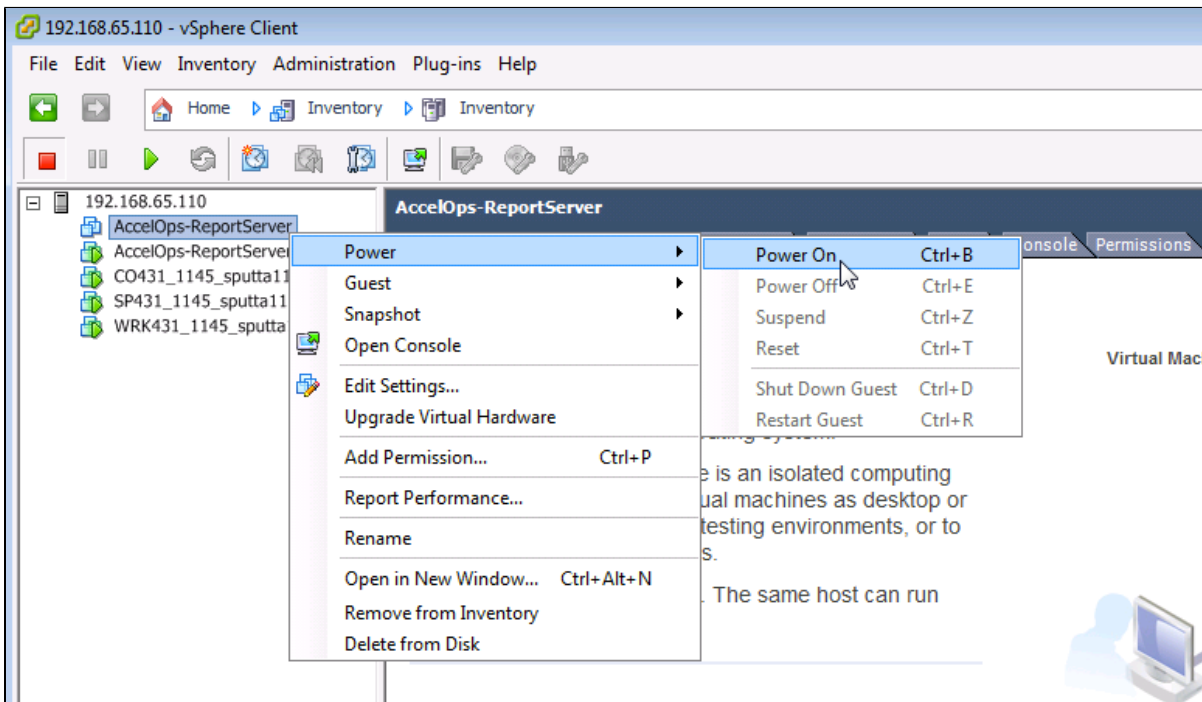
## Section 4: Configure the Report Server from the VMware Console

### Note

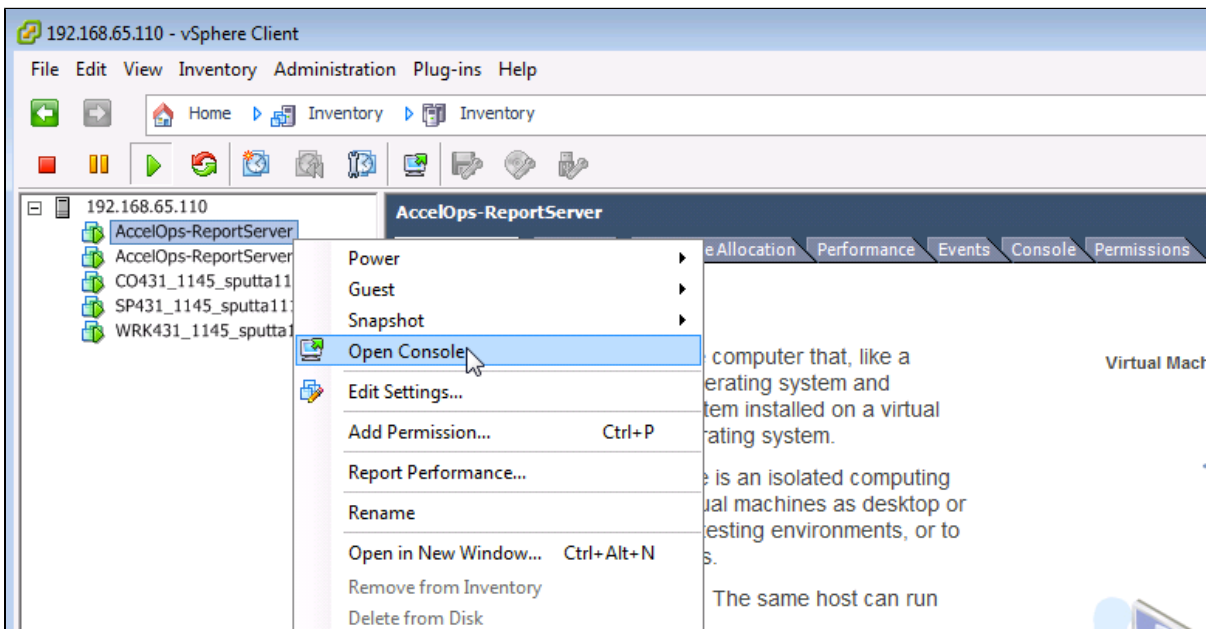
While in the VM console window do not press any control keys during the installation. Doing so might cause the installation to stop and not to precede any further. For example Ctrl – C or Ctrl – Z. If this occurs you must erase the VA and start over from the beginning.

### Step 1: Start Report Server

1. From the VMware vSphere client, Select recently imported Report server Virtual appliance and Power On (Right click > Power > Power On)



2. Open Report server VM Console (Right click > Open Console)



3. Network eth0 Failed message may appear on the screen during the loading, but this is to be expected the first time the AccelOps Report server is started.

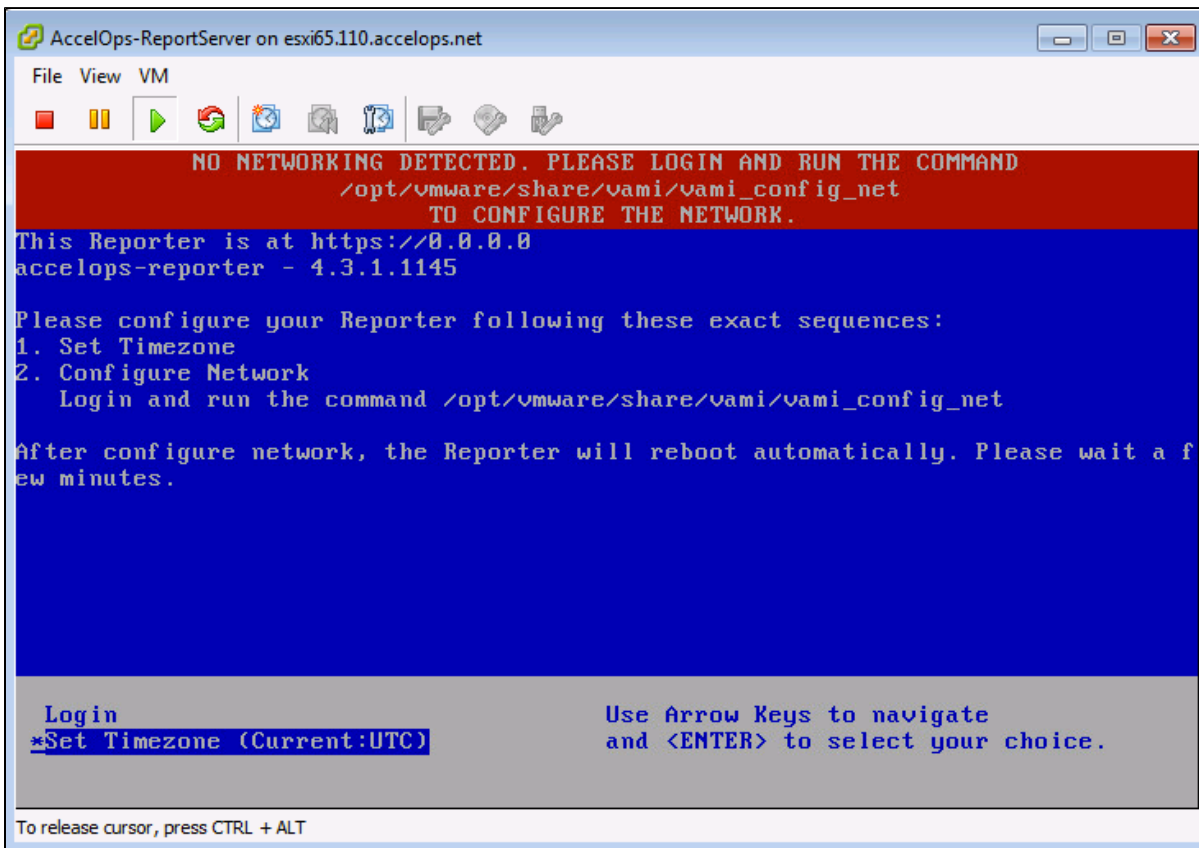
```
AccelOps-ReportServer on esxi65.110.accelops.net
File View VM
[Icons: Stop, Pause, Play, Refresh, Power, Network, Storage, Print, Help]
/dev/sda1: clean, 39/32768 files, 31115/131072 blocks [ OK ]
Remounting root filesystem in read-write mode: [ OK ]
Mounting local filesystems: [ OK ]
Enabling /etc/fstab swaps: [ OK ]
Entering non-interactive startup
Calling the system activity data collector (sadc)...
iptables: Applying firewall rules: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0:
Determining IP information for eth0... failed. [ FAILED ]
Starting auditd: [ OK ]
Starting portreserve: [ OK ]
Waiting for network to come up (attempt 1 of 10)...
Waiting for network to come up (attempt 2 of 10)...
Waiting for network to come up (attempt 3 of 10)...
Waiting for network to come up (attempt 4 of 10)...
Waiting for network to come up (attempt 5 of 10)...
Waiting for network to come up (attempt 6 of 10)...
Waiting for network to come up (attempt 7 of 10)...
Waiting for network to come up (attempt 8 of 10)...
Waiting for network to come up (attempt 9 of 10)...
Waiting for network to come up (attempt 10 of 10)...
```

### Step 2: Configure the Timezone

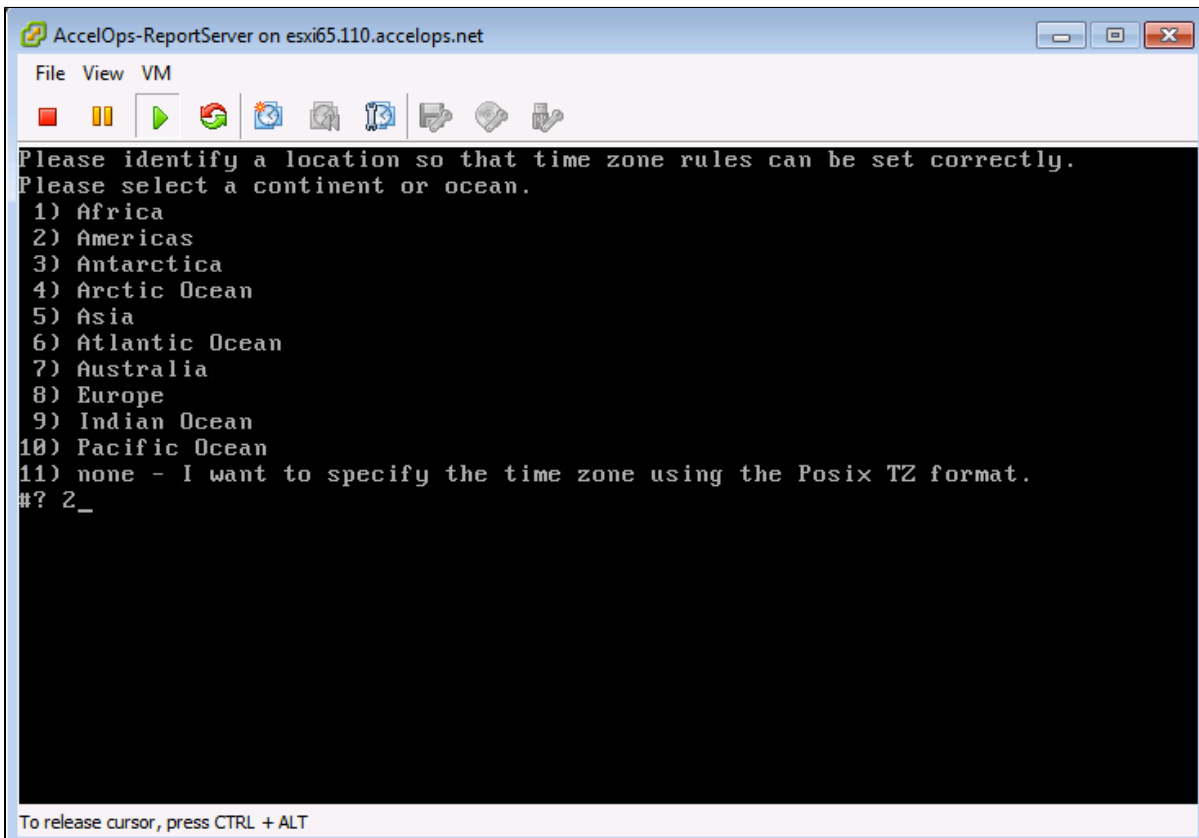
It is important to select the correct time zone that you are located in.

1. On VM console, Select 'Set Timezone' and press the 'Enter' key

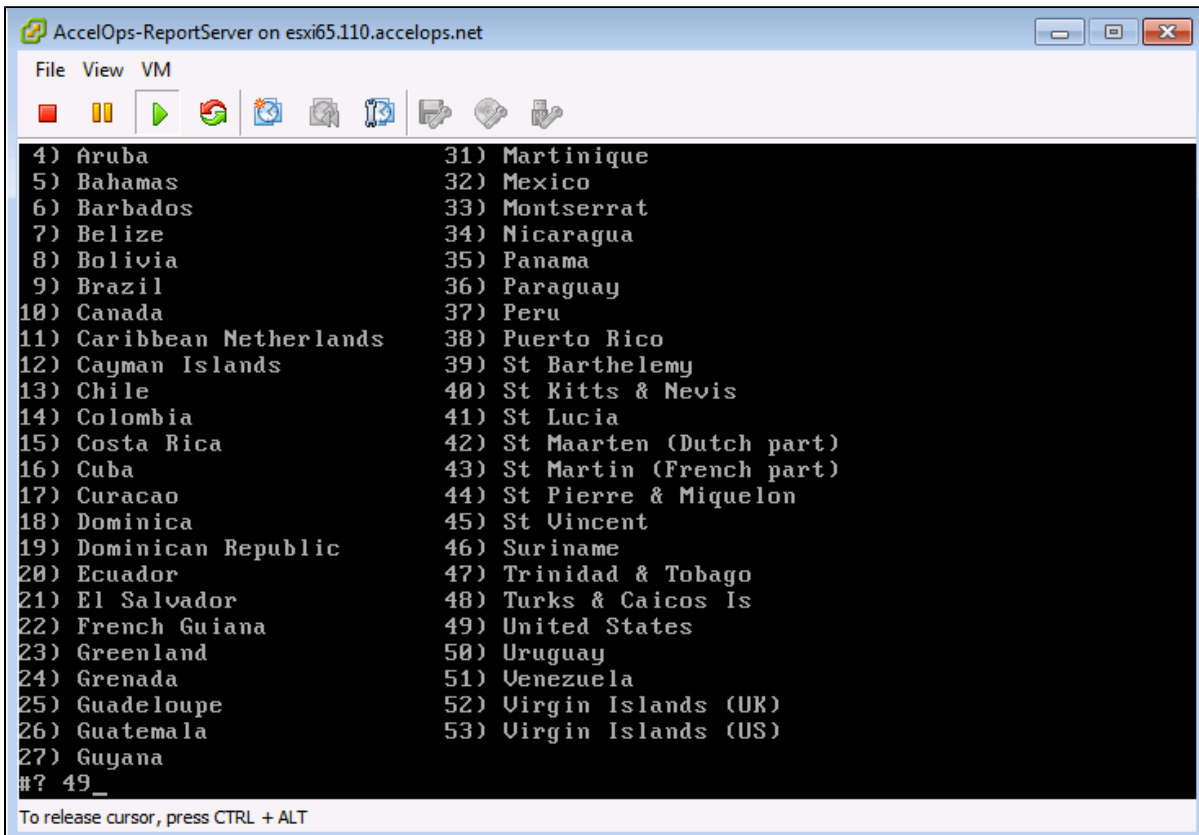




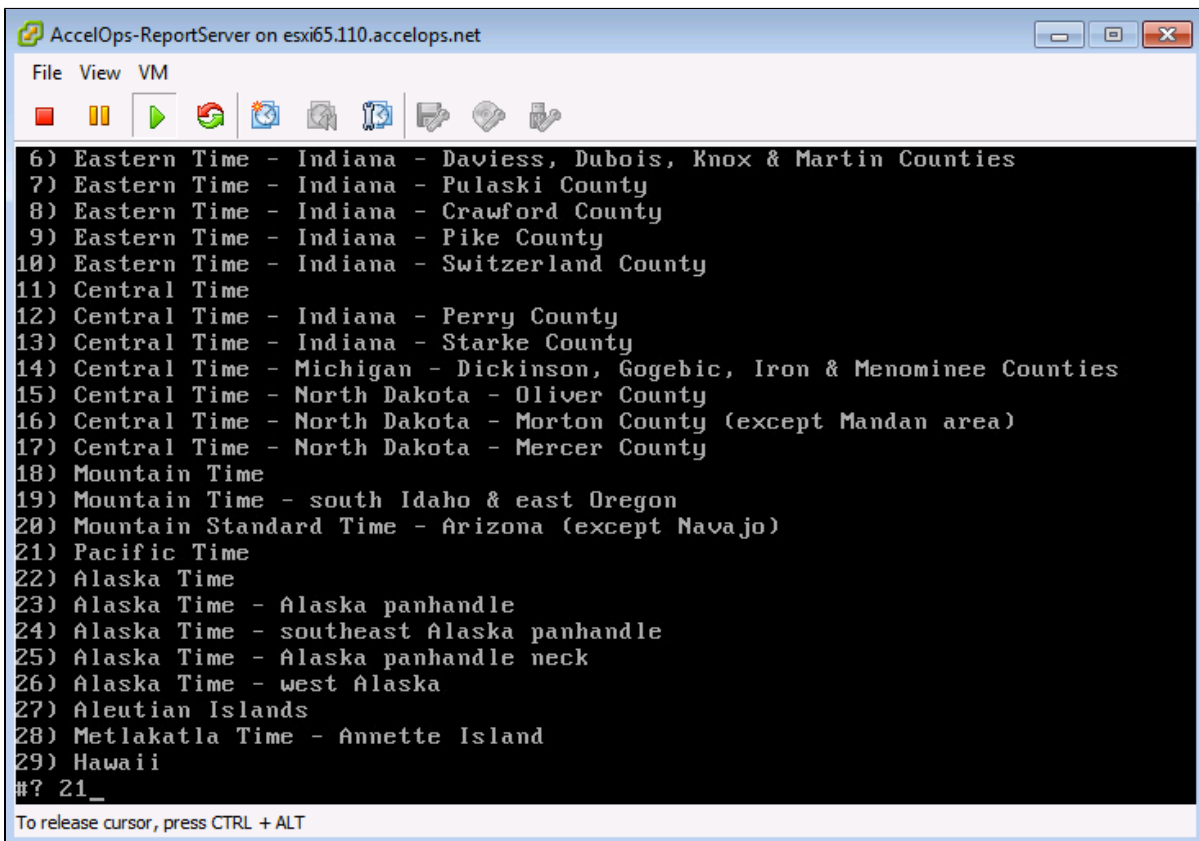
2. Select your Location and press the 'Enter' key



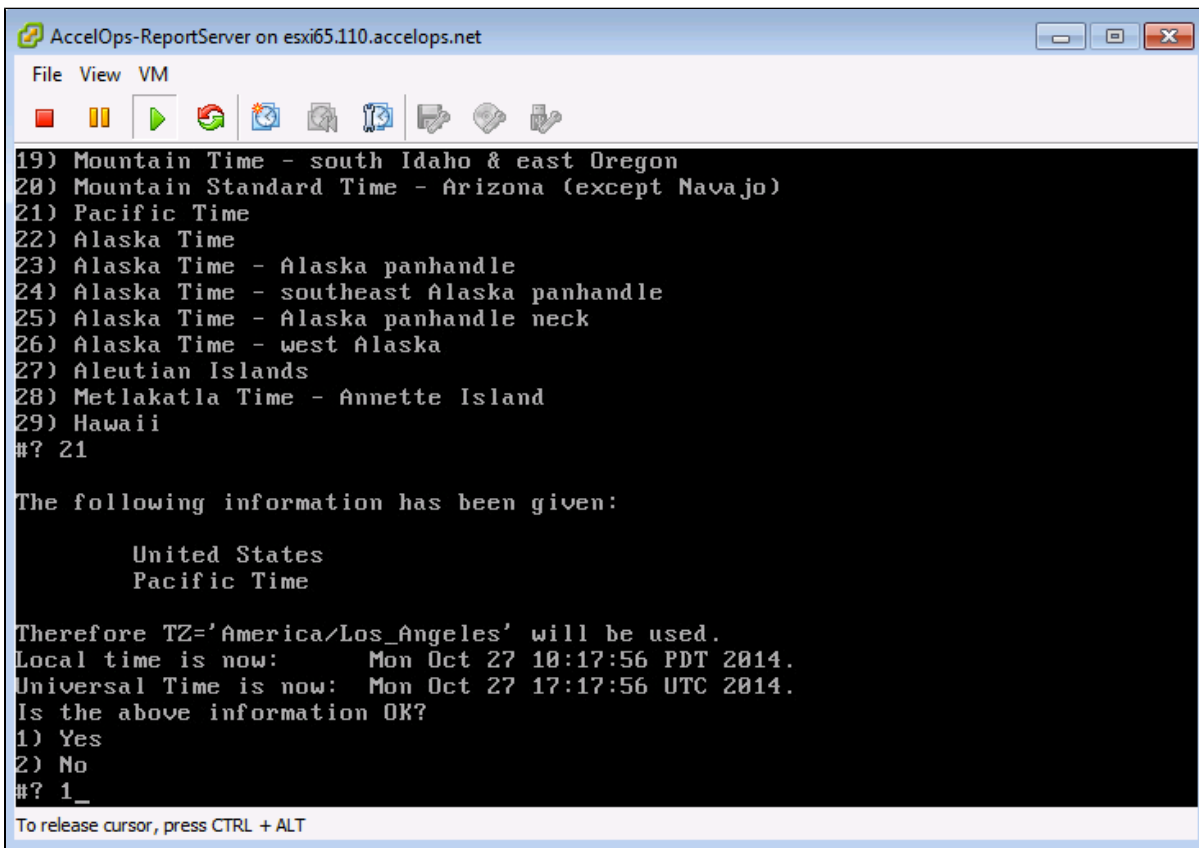
3. Select your Country and press the 'Enter' key



4. Select your Timezone and press the 'Enter' key



5. Review and Confirm your Timezone selection



The screenshot shows a terminal window titled "AccelOps-ReportServer on esxi65.110.accelops.net". The terminal displays a list of time zone options from 19 to 29. Option 21, "Pacific Time", is selected. The terminal then displays the following information:

```
19) Mountain Time - south Idaho & east Oregon
20) Mountain Standard Time - Arizona (except Navajo)
21) Pacific Time
22) Alaska Time
23) Alaska Time - Alaska panhandle
24) Alaska Time - southeast Alaska panhandle
25) Alaska Time - Alaska panhandle neck
26) Alaska Time - west Alaska
27) Aleutian Islands
28) Metlakatla Time - Annette Island
29) Hawaii
#? 21

The following information has been given:


    United States
    Pacific Time

Therefore TZ='America/Los_Angeles' will be used.
Local time is now:      Mon Oct 27 10:17:56 PDT 2014.
Universal Time is now:  Mon Oct 27 17:17:56 UTC 2014.
Is the above information OK?
1) Yes
2) No
#? 1_


To release cursor, press CTRL + ALT
```

### Step 3: Configure the Network

1. On VM console, Select 'Login' and press the 'Enter' key



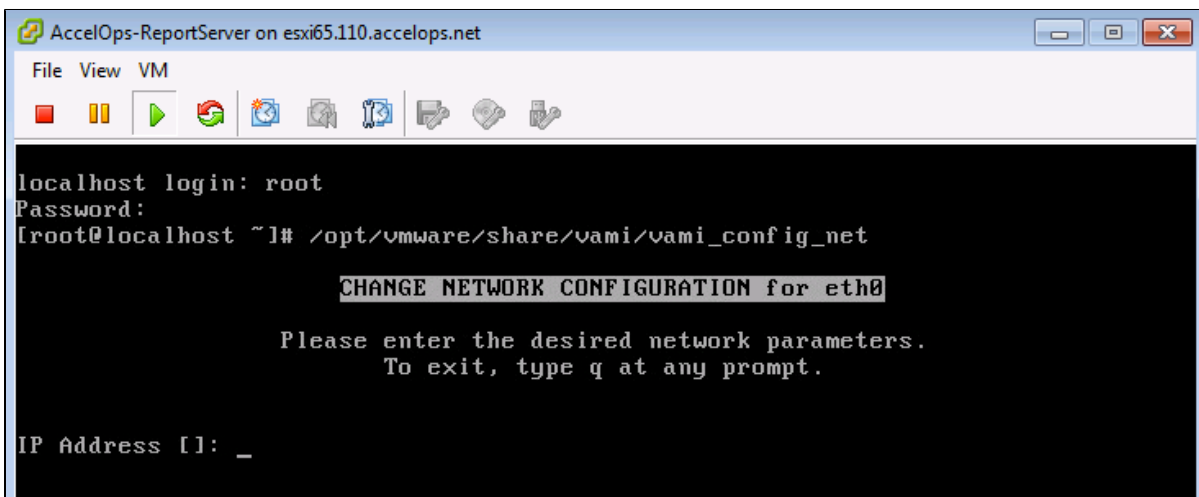
2. Enter default log in information.  
Login: root  
Password: ProspectHills



3. Run vami\_config\_net script to configure network and installing AccelOps Report server

**Configure the Network**

`/opt/vmware/share/vami/vami_config_net`



4. Configure Network with Static IP address
  - a. The following information is required to configure Static IP address:
    - i. IP Address
    - ii. Netmask
    - iii. Gateway
    - iv. DNS Server(s)

**Proxy Server**

Do not configure the proxy server setting.

- b. Press the 'y' key to accept the changes

- c. Enter 'Host name' and press 'Enter' key

```

AccelOps-ReportServer on esxi65.110.accelops.net
File View VM
Gateway []: 192.168.64.1
DNS Server 1 []: 192.168.0.10
DNS Server 2 [192.168.0.10]: 192.168.0.40
Is a proxy server necessary to reach the Internet? y/n [n]:

IP Address:      192.168.65.145
Netmask:        255.255.252.0
Gateway:        192.168.64.1
Proxy Server:
DNS Servers:    192.168.0.10, 192.168.0.40

Is this correct? y/n [y]: y

Reconfiguring the network...
DNS server settings updated
Determining if ip address 192.168.65.145 is already in use for device eth0...
vami_login: no process killed
Network parameters successfully changed to requested values
Input hostname: reportserver.accelops.net

```

#### Step 4: Select Mount Point to store your data

The AccelOps Report Server can be installed using either storage configured within the ESX server or NFS storage.

Choose only following Step 4.1 or 4.2 based on [Section 3](#)

1. Enter your Local hard disk mount point

##### Note

Skip Step 4.1 if AccelOps Super is used NFS storage to store the events data. Else, continue below steps.

##### Local Storage Mount point

/dev/sdc



The screenshot shows a terminal window titled "AccelOps-ReportServer on esxi65.110.accelops.net". The terminal output is as follows:

```

Mount Point For Your Data: local storage(/dev/sdc) or NFS (ip:/dir) with right p
ermission to proceed or use Ctrl+C to go back to console.
Mount Point: /dev/sdc
mke2fs 1.41.12 (17-May-2010)
/dev/sdc is entire device, not just one partition!
Proceed anyway? (y,n) y
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
3604480 inodes, 14417920 blocks
720896 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
440 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424

Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: _

```

At the bottom of the terminal window, it says "To release cursor, press CTRL + ALT".

2. Enter your NFS mount point

**Note**

Skip Step 4.2 if you choose Step 4.1 to store Reports data in local Hard disk. Else, continue below steps.

**NFS Mount point**

<NFS\_Server\_IP\_Address>:/<Directory\_Path>

```

Mount Point For Your Data: local storage(/dev/sdc) or NFS (ip:/dir) with right p
ermission to proceed or use Ctrl+C to go back to console.
Mount Point: 192.168.67.168:/data/RS/431
192.168.67.168:/data/RS/431 mounted
Running initialization of Reporter host

```

To release cursor, press CTRL + ALT

3. After mount point, Report server will reboot automatically  
In 5 to 10mins, Report server will be configure successful.

```

AccelOps-ReportServer on esxi65.110.accelops.net
File View VM
This Reporter is at https://192.168.65.145
accelops-reporter - 4.3.1.1145

Please configure your Reporter following these exact sequences:
1. Set Timezone
2. Configure Network
   Login and run the command /opt/vmware/share/vami/vami_config_net

After configure network, the Reporter will reboot automatically. Please wait a f
ew minutes.

Login
*Set Timezone (Current:PDT)
Use Arrow Keys to navigate
and <ENTER> to select your choice.

```

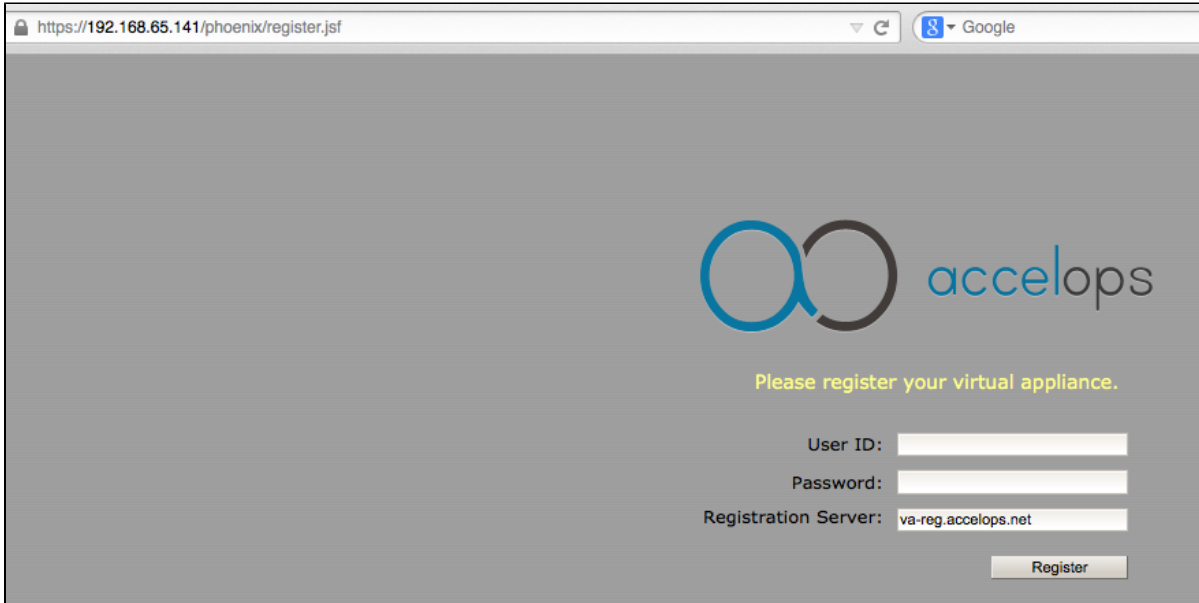
## Section 5: Register the Report Server to Supervisor

### Note

A new license is required to enable Report Server. Contact AccelOps support for Report Server License. You can skip this step if a Report Server license is already registered with the Supervisor.

### Step 1: Re-register AccelOps Supervisor with new license

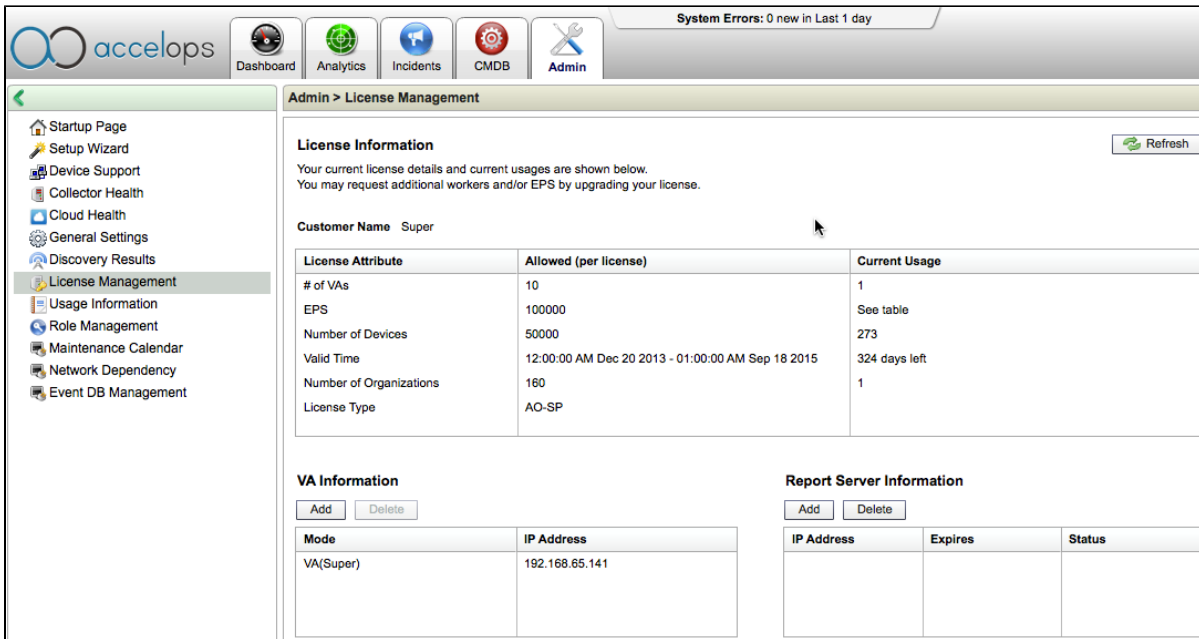
1. In browser windows, enter 'https://<Supervisor IP>/phoenix/register.jsf'



2. Enter license user name and password

**Step 2: Add Report Server on Supervisor UI**

1. Go to Admin > License Management. You will see a section called 'Report Server Information'



2. Click 'Add' to register the Report Server Node
  - a. Enter Report Server IP Address
  - b. Enter Read only credentials that will be used by the AccelOps Visualization Analytics Server to read the data in the PostgreSQL tables in the Report Server

**Add Report Server**

Report Server IP Address:

Database Username:

Database Password:

Confirm Password:

**Report Server Information**

IP Address	Expires	Status

- Report Server registration will takes 3mins to complete when Supervisor's CMDB size is <1 GB. Else, it takes more time to complete registration based Supervisor's CMDB size. You can 'Run in Background' when Supervisor's CMDB size is > 1GB.

**Adding Report Server: 192.168.65.145**

**Report Server Information**


IP Address	Expires	Status

- Click 'OK' on Confirmation pop-up. The Supervisor node will configure the Report Server

**Report Server Information**

IP Address	Expires	Status
192.168.65.145	09/05/2015	Active

- Make sure Report Server appears in Admin > Cloud health

Admin > Cloud Health										
 Refresh    Last Updated: 17:38:55										
Name	IP Address	Health	Module Role	Version	Build Date	Cores	Load Averag	CPU	Swap Size	Swap I
VA141	192.168.65.141	Normal	supervisor	4.3.1.1145	10:04:55 AM Oct 15 20...	8	0.26,0.27,0.3	0.88	6 GB	57.33 I
reportserver.accelops.net	192.168.65.145	Normal	report server	4.3.1.1145	10:04:55 AM Oct 15 20...	3	0.02,0.03,0	2.33	6 GB	0 KB
reportserver.accelops.net ( 192.168.65.145 ) - report server: Process Details										
Process Name	Status	Up Time	CPU	Event Rate	Physical Mem	Virtual Mem	SharedStore ID			
phMonitorReportServe	Up	2h 32m 48s	0%	0/s	24 MB	909 MB	0			
rsyslogd	Up	2h 55m	0%	0/s	1 MB	243 MB	0			
Report DB	Up	2h 34m 1s	0%	0/s	17 MB	446 MB	0			
pgsql DB	Up	2h 38m 19s	0%	0/s	17 MB	446 MB	0			

## Installing in KVM

The following steps are needed to set up AccelOps Virtual Appliance under Redhat Kernel Virtualization Manager (KVM). The procedures are identical for both Service Provider and non-Service Provider versions; the difference is in the licenses applied which will make multiple organizations to appear.

This section describes the basic installation steps on KVM. Please follow the detailed steps for configuring AO Collector, Super, Worker as described under ESX.

- Prerequisites
- [Prerequisites](#)
- [Preliminary step: setup bridge network for KVM](#)
- [Deploying AccelOps on KVM](#)

### Prerequisites

To install AccelOps under KVM, you will need this information:

- Static IP address
  - Determine the static IP address and subnet mask for your virtual appliance.
- Enter the AccelOps host name within your local DNS server
- Proxy server information (IP address and port number)
  - The authenticated proxy server is not supported in this version of AccelOps. Turn off authentication on proxy server for this host or completely disable the proxy for this host.
- Determine the IP address of NFS mount point and NFS share name

### Preliminary step: setup bridge network for KVM

If AccelOps is the first guest on KVM, then a bridge network may be required to enable network connectivity. For details see [here](#).

For the instructions below, **em1** connected as management network and **em4** connected to your local area network.

1. In KVM host go to `/etc/sysconfig/network-scripts/`
2. Create bridge network config file `ifcfg-br0` (Sample config for br0)

```
DEVICE=br0
BOOTPROTO=none
NM_CONTROLLED=yes
ONBOOT=yes
TYPE=Bridge
NAME="System br0"
```

3. Edit network config file `ifcfg-em4`

```
DEVICE=em4
BOOTPROTO=shared
NM_CONTROLLED=no
ONBOOT=yes
TYPE=Ethernet
UUID="24078f8d-67f1-41d5-8eea-xxxxxxxxxxxxx"
IPV6INIT=no
USERCTL=no
DEFROUTE=yes
IPV4_FAILURE_FATAL=yes
NAME="System em4"
HWADDR=F0:4D:00:00:00:00
BRIDGE=br0
```

4. Restart network service
- 5.

### Deploying AccelOps on KVM

1. Download AO KVM package from image server. Verify checksum of the package. Uncompress the file.
2. Start Virtual Machine Manager

? Unknown Attachment

3. Right click host and select New

? Unknown Attachment

4. Create new VM

? Unknown Attachment

5. Choose the decompressed KVM disk image

? Unknown Attachment

6. Allocate resource based on the system requirement

? Unknown Attachment

7. Confirm and create new VM

? Unknown Attachment

8. Open new VM

? Unknown Attachment

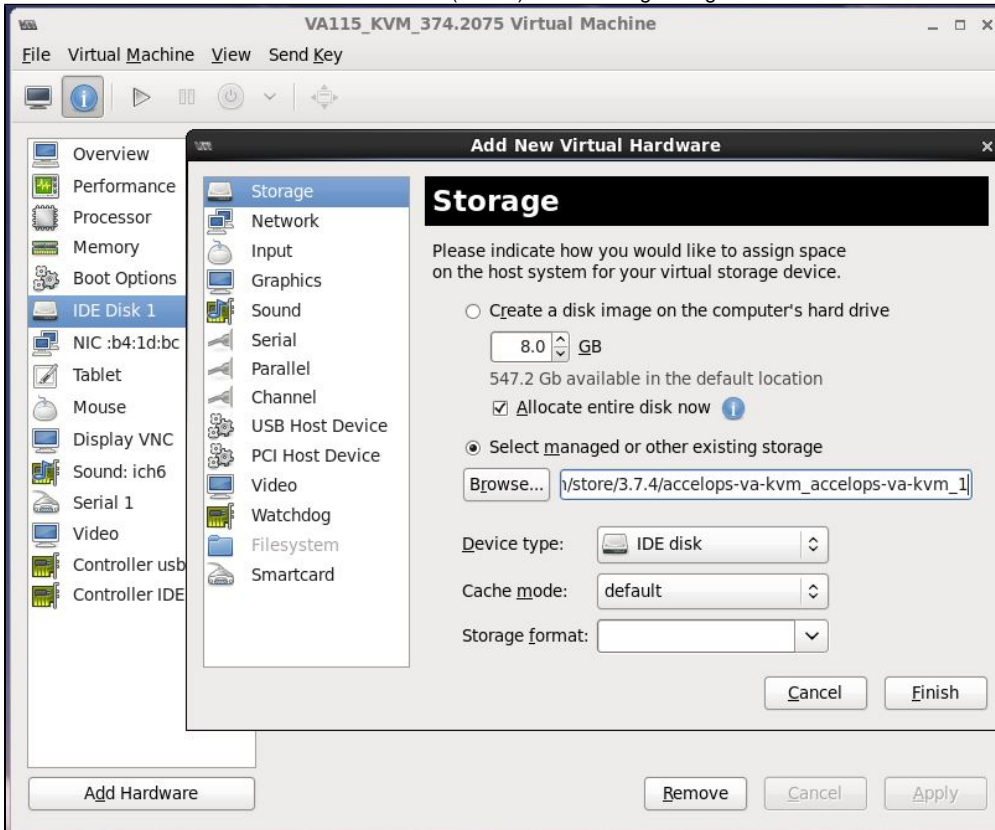
9. Change network setting bridge via "br0"



- 10. Change Storage Disk bus to IDE



- 11. Add new Hardware disk and select CMDB disk (60 GB) from existing storage



- 12. **Optional:** add local eventDB storage (VA mode only)





13. Power on AccelOps VA
14. Configure AccelOps system from console by following the steps described for ESX hypervisor.

## Installing in HyperV

The following steps are needed to set up AccelOps Virtual Appliance under Microsoft Hyper-V virtualization platform. The procedures are identical for both Service Provider and non-Service Provider versions; the difference is in the licenses applied which will make multiple organizations to appear.

This section describes the basic installation steps on Hyper-V. Please follow the detailed steps for configuring AO Collector, Super, Worker as described under ESX

## Supported Versions

AccelOps have been tested to run on Hyper-V on Microsoft Windows 2012.

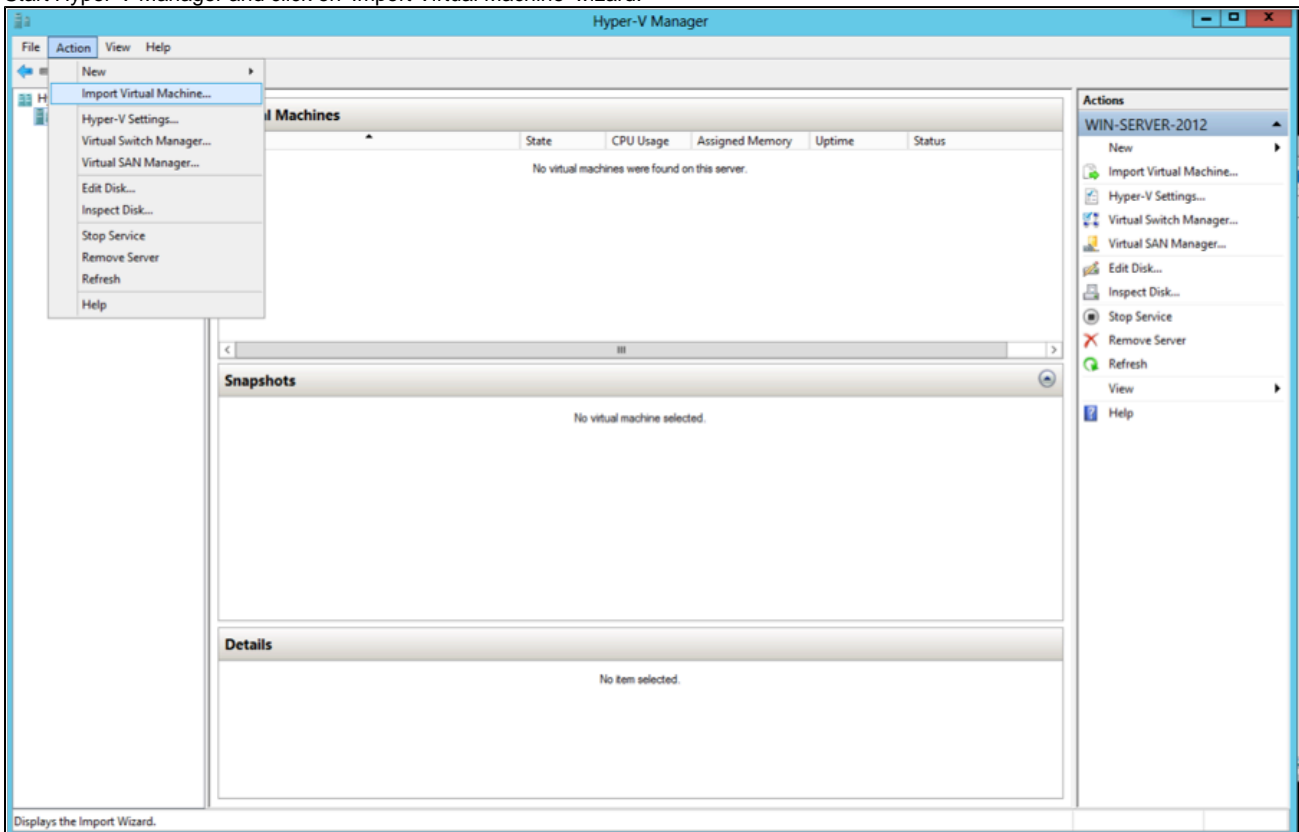
## Prerequisites

To install AccelOps under HyperV, you will need this information:

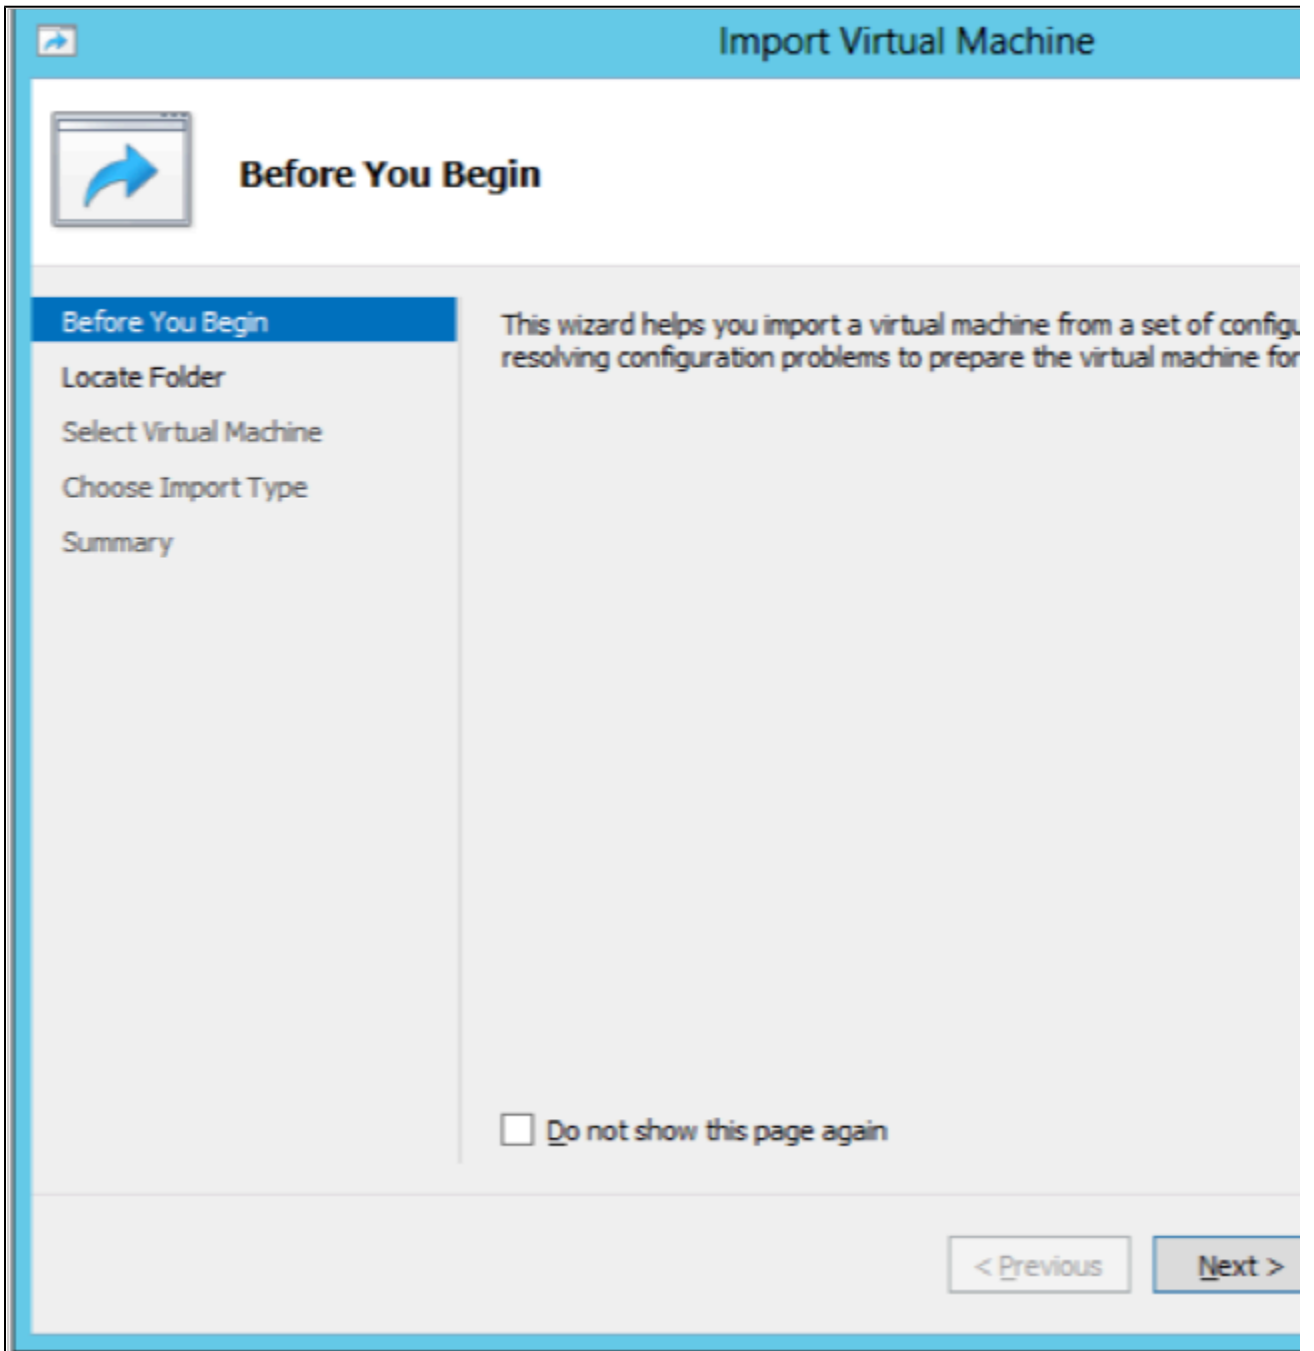
- Static IP address
  - Determine the static IP address and subnet mask for your virtual appliance.
- Enter the AccelOps host name within your local DNS server
- Proxy server information (IP address and port number)
  - The authenticated proxy server is not supported in this version of AccelOps. Turn off authentication on proxy server for this host or completely disable the proxy for this host.
- Determine the IP address of NFS mount point and NFS share name

## Deploying a virtual appliance on Hyper-V

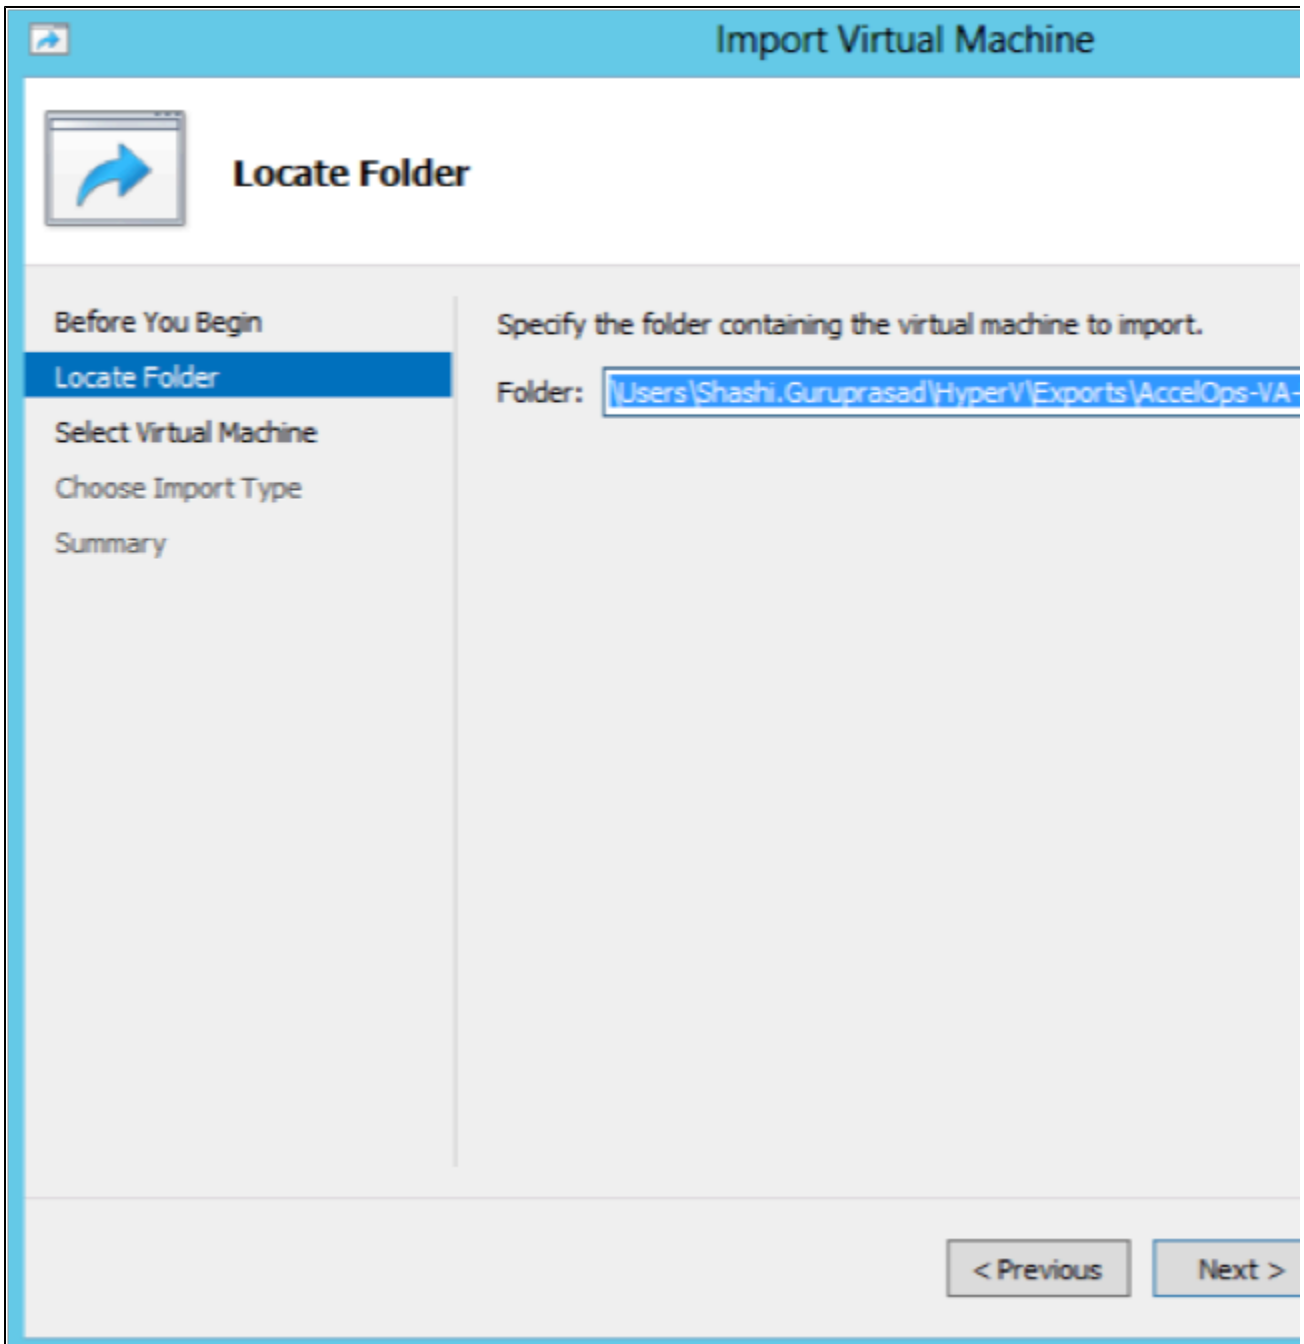
1. Download AO HyperV package from image server. Verify checksum of the package. Un-compress the file.
2. Start Hyper-V Manager and click on 'Import Virtual Machine' wizard.



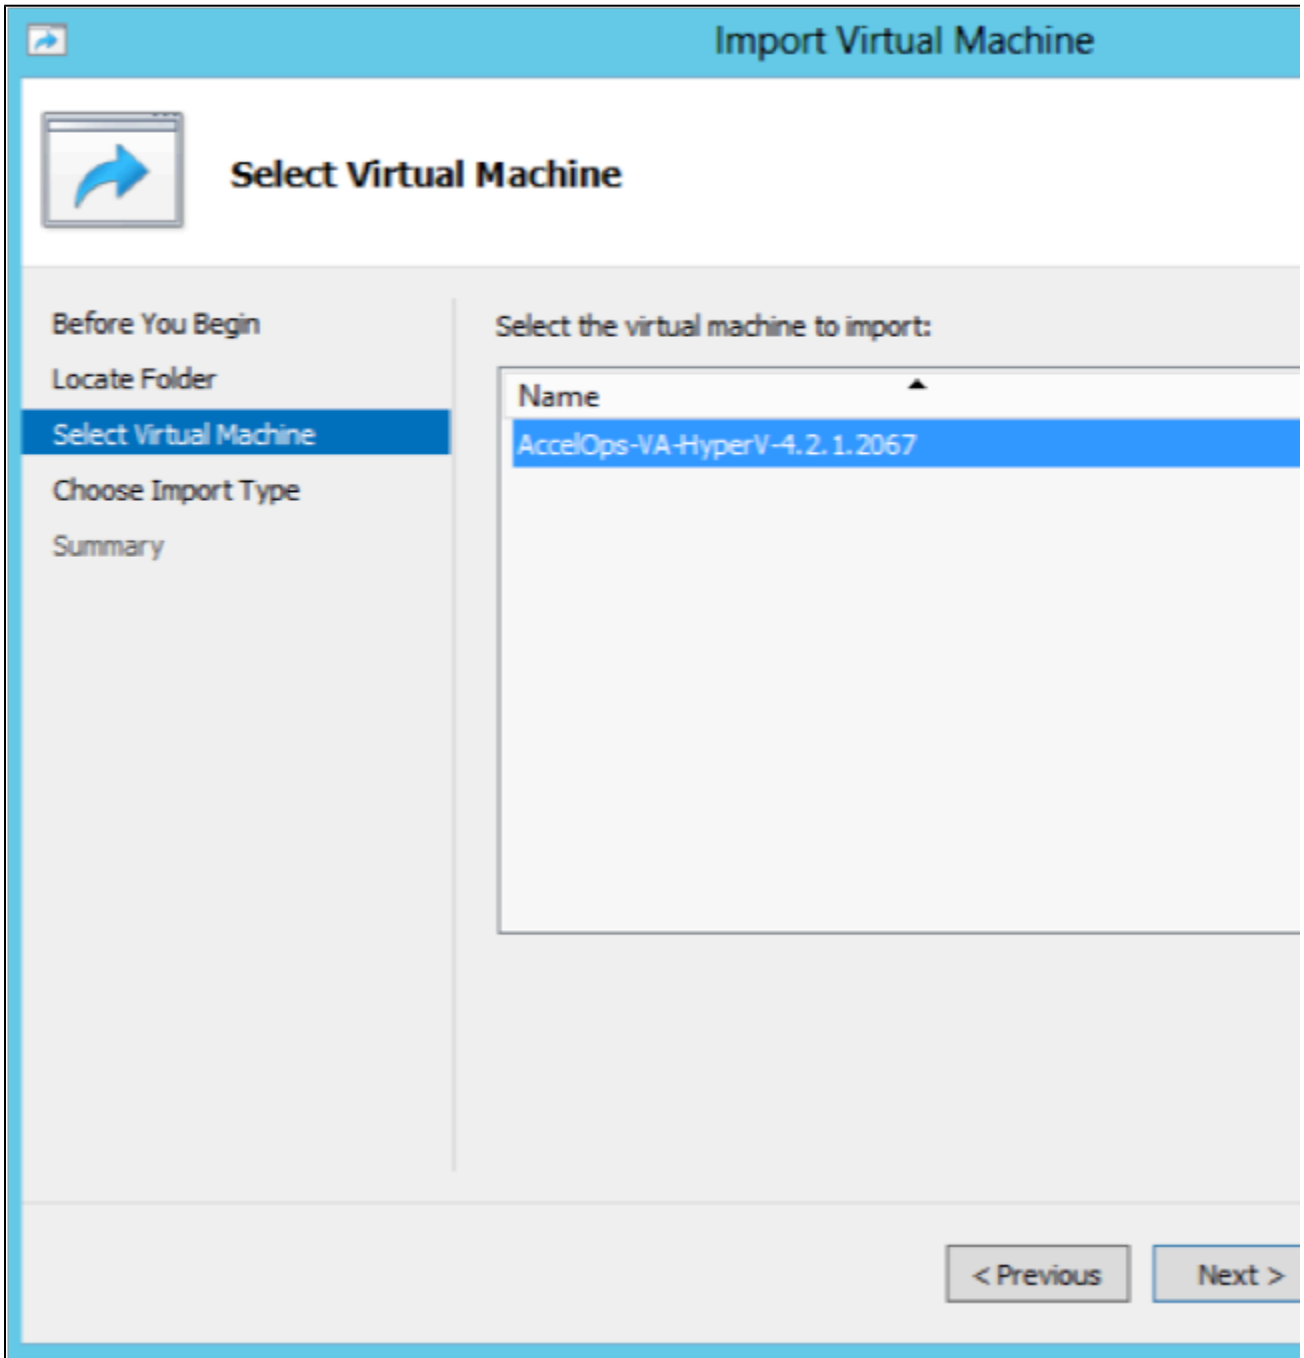
3.



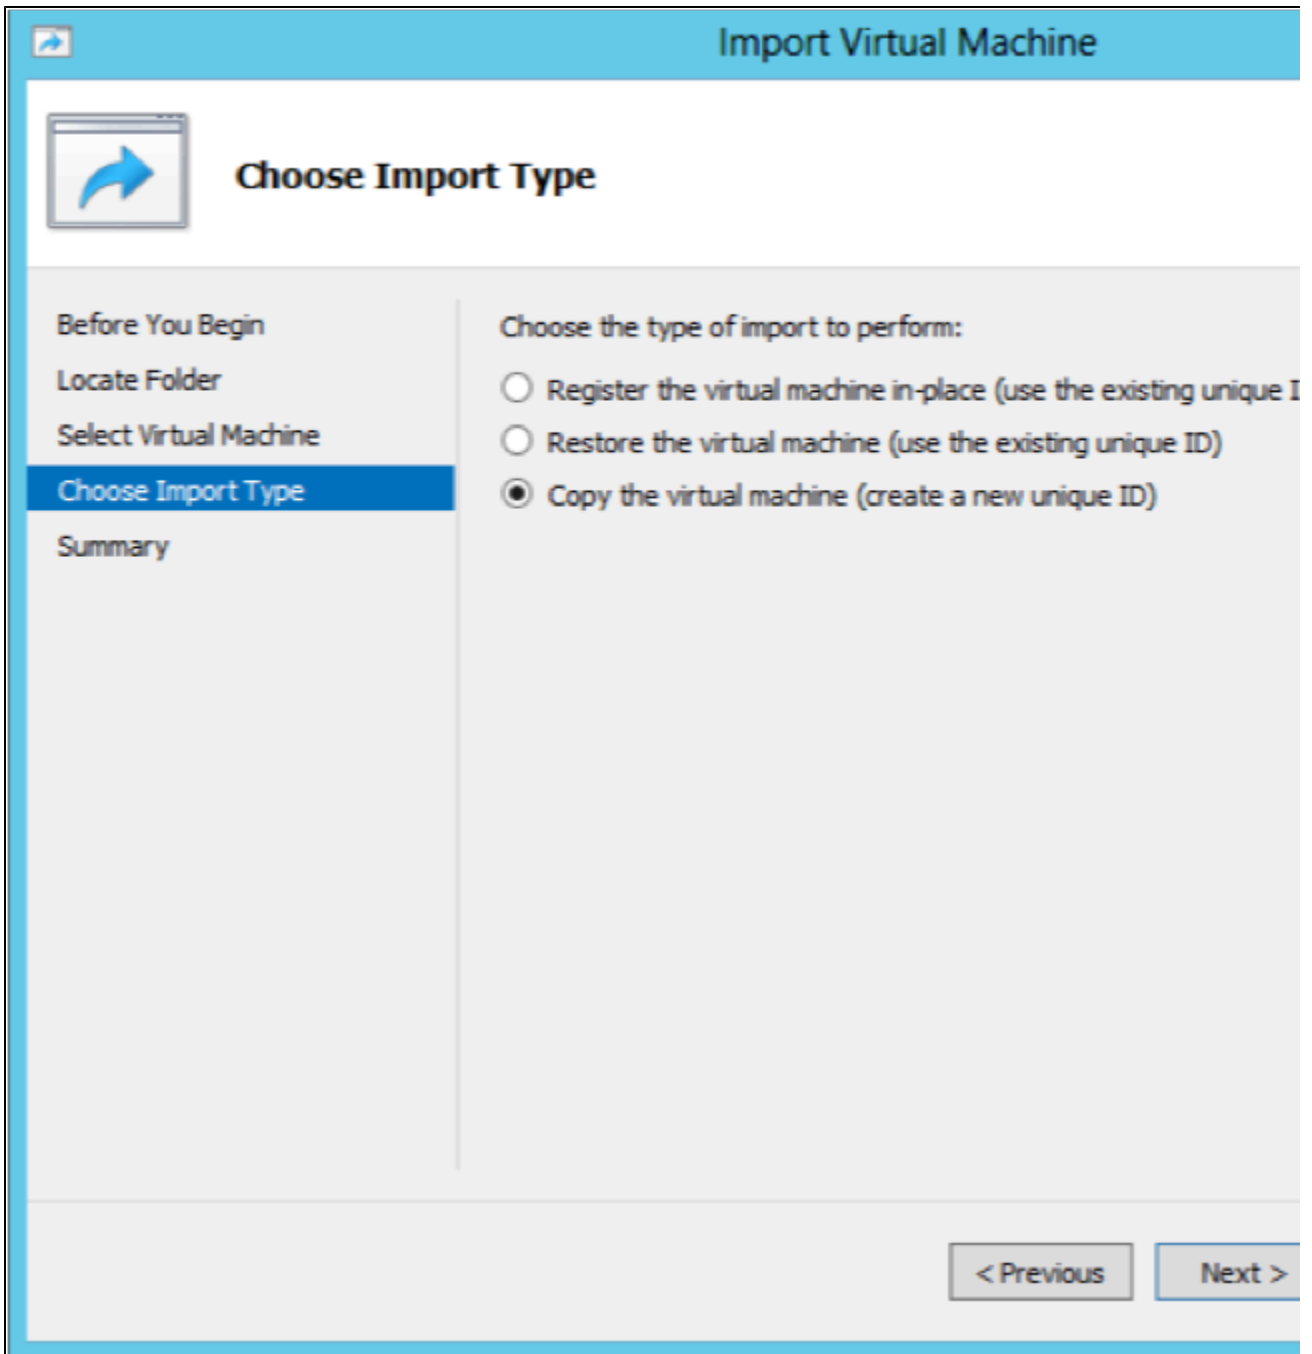
4. Specify the folder containing AccelOps virtual machine



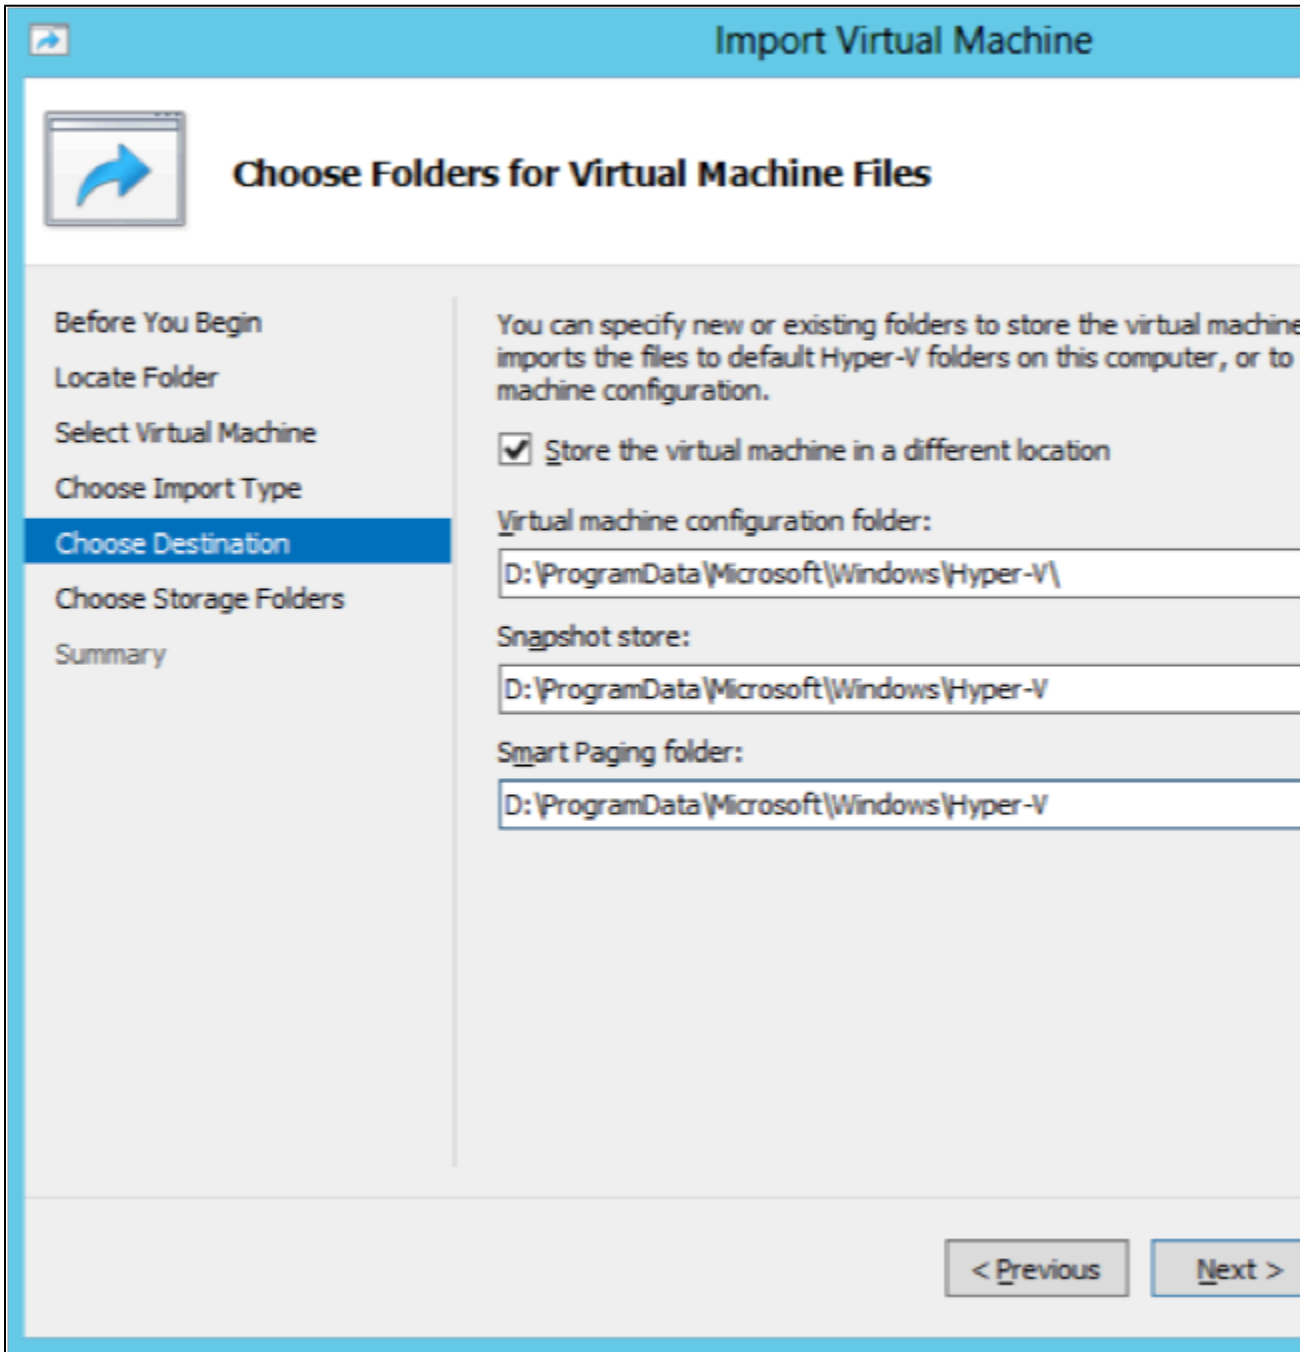
5. Select the proper AccelOps 4.2.1 image



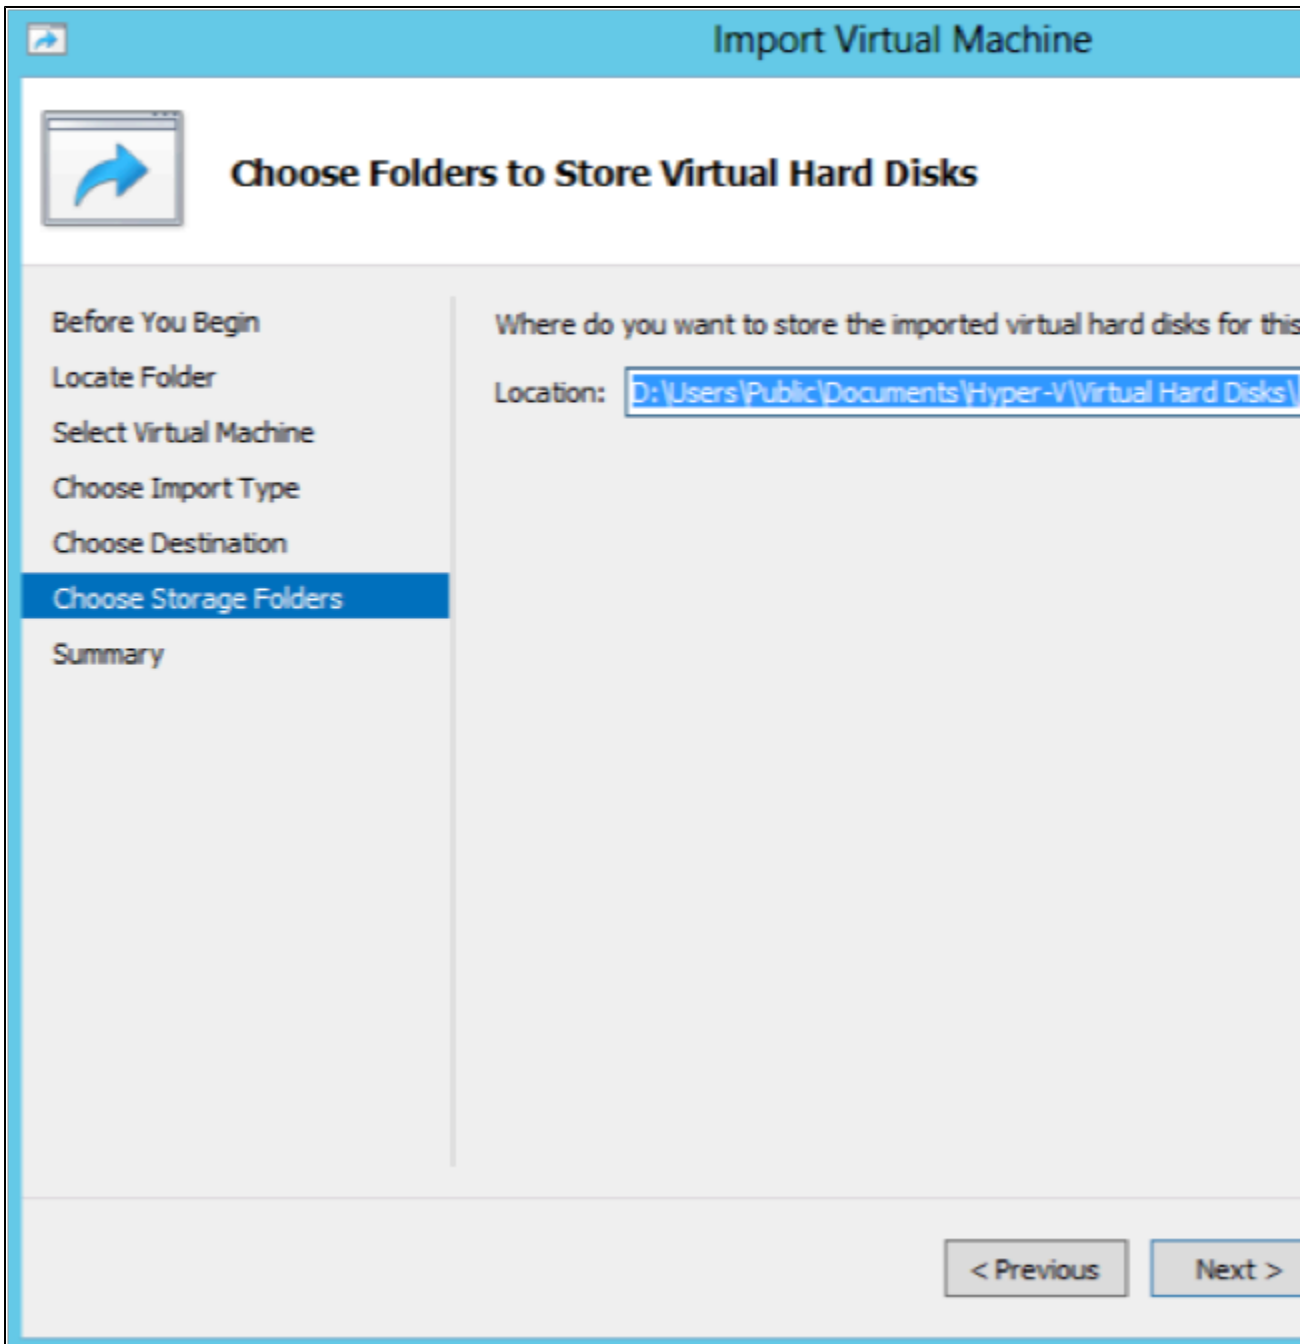
6. Start copying the virtual machine



7. Choose folders for Virtual Machine files

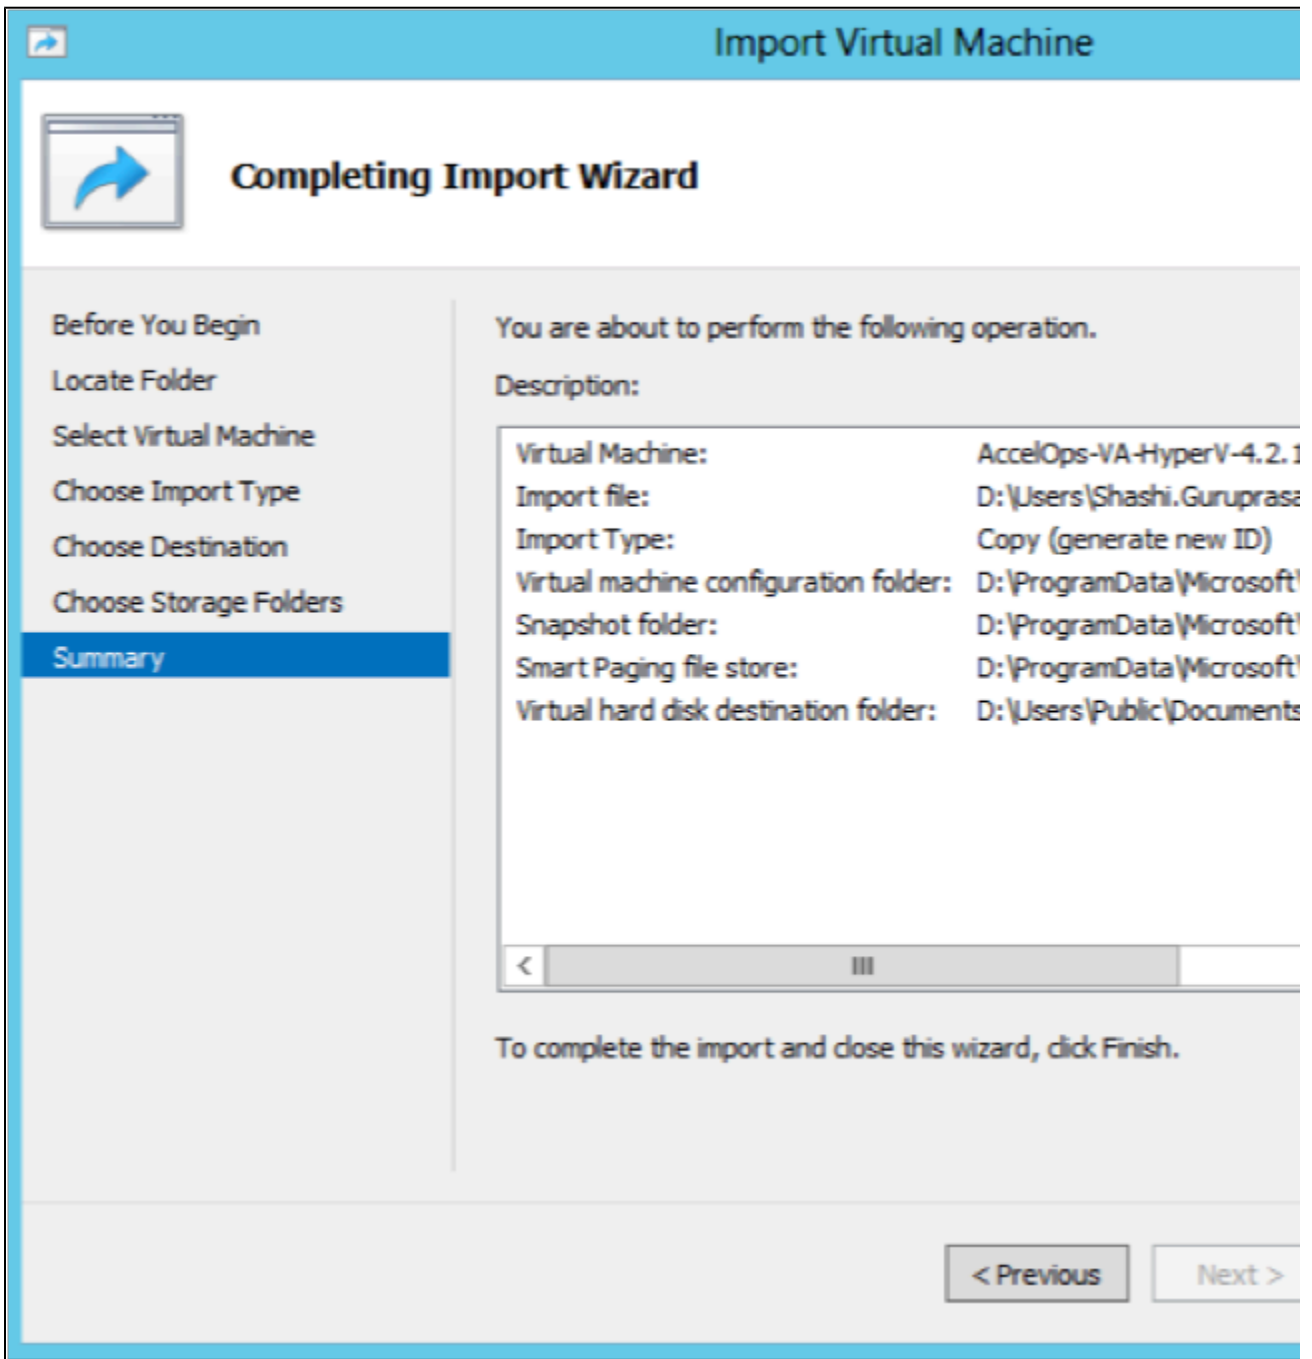


8. Choose folders to store virtual hard disks




9. Verify the information and complete the import by clicking Finish





10. Before you power on the VM, determine whether you plan on using NFS to store events (i.e. EventDB) or use a local disk. If you plan on using a local disk, a dynamically expanding disk of 100GB is already included in IDE controller 1. If this size is sufficient, proceed to complete the remaining installation steps. Use /dev/sdc1 as the device when you are asked for the path for EventDB. If you need a larger sized local disk, you can convert this disk to a larger one using 'Edit Disk' and selecting 'Convert' to create a new hard disk. After this step is complete, replace the third disk on IDE controller 1 with the new one that you create. If you are going to use NFS for EventDB, follow the same steps as described for ESX hypervisor
11. Power on AccelOps VA and Connect to the console (right click the VM and click 'Connect')



12. Configure AccelOps system from console by following the steps described for ESX hypervisor.

**Note:** AccelOps HyperV images uses dynamically expanding VHD disks for the root and CMDB partitions, and a dynamically expanding VHDX disk for EventDB. Dynamically expanding disks are used to keep the exported HyperV image within reasonable limits. For production deployments, Microsoft recommends using a 'Fixed type' disk for VHD format (root and CMDB partition) and 'Dynamic type' for VHDX format (EventDB). Please refer to 'Performance Tuning Guidelines for Windows Server 2012 (or R2)'

## Installing in AWS

The following steps are needed to set up AccelOps Virtual Appliance in Amazon Web Services (AWS). The procedures are identical for both Service Provider and non-Service Provider versions; the difference is in the licenses applied which will make multiple organizations to appear.

The following are the Amazon Machine Instance ids for the various relevant releases

Region Code	Region Name	4.2.1 Super/Worker AMI-id	4.2.1 Collector AMI-id	4.2.3 Super/Worker AMI-id	4.2.3 Collector AMI-id
us-east-1	US East (Northern Virginia)	ami-cc7e9ba4	ami-49190320	ami-2872a340	ami-5a954532
us-west-1	US West (Northern California)	ami-9c4443d9	ami-a24443e7	ami-17242852	ami-2b24286e
us-west-2	US West (Oregon) Region	ami-2fc6b21f	ami-aa1f6b9a	ami-d90145e9	ami-2f02461f
eu-west-1	EU (Ireland) Region	ami-bf08c7c8	ami-a708c7d0	ami-aa3ce3dd	ami-903ce3e7

- Note: SVN password reset issue after system reboot for AccelOps 3.7.6 customers in AWS Virtual Private Cloud (VPC)
- Setup a VPC (Virtual Private Cloud) in AWS for use by AccelOps VMs
- Determine Local or NFS Storage for EventDB
- Setup NFS Server in AWS for use by AccelOps VMs
- Performing consistent snapshots of EBS volumes that host EventDB and CMDB
- Setup AO Supervisor and Worker nodes in AWS
- (Optional - If using Local storage) Configure Local Storage for EventDB
- (Optional - if Worker is needed) Setup AO Worker node in AWS
- (Optional - if Collector is needed) Setup AO Collector node in AWS
- Deploy AccelOps Collector AMI
- Validate AccelOps Collector connectivity and health
- Register AccelOps Collector
- Setup AccelOps Report Server node

### Note: SVN password reset issue after system reboot for AccelOps 3.7.6 customers in AWS Virtual Private Cloud (VPC)

AccelOps uses SVN to store monitored device configurations. In AWS VPC setup, we have noticed that AccelOps SVN password gets changed if the system reboots - this prevents AccelOps from storing new configuration changes and viewing old configurations. The following procedure can be used to reset the SVN password to AccelOps factory default so that AccelOps can continue working correctly.

**This script needs to be run only once.**


1. Logon to Super
2. Copy the attached "ao\_svnpwd\_reset.sh" script to Super on EC2+VPC deployment
3. Stop all backend processes before running script by issuing the following command: **phtools --stop all**
4. Run following command to change script permissions: "**chmod +x ao\_svnpwd\_reset.sh**"
5. Execute "ao\_svnpwd\_reset.sh" as root user: **./ao\_svnpwd\_reset.sh**
6. The system will reboot
7. Check SVN access to make sure that old configurations can be viewed.

### Setup a VPC (Virtual Private Cloud) in AWS for use by AccelOps VMs

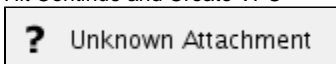
The user **must** setup a VPC in AWS and deploy AccelOps cluster (or VA) within a VPC rather than classic-EC2. In classic-EC2, starting or stopping an instance changes the private IP address of the instance. This not only invalidates the license but also causes AccelOps software to stop working.

1. Visit the [VPC dashboard](#) page and start VPC Wizard. In the example below, we create a VPC with a single public subnet (i.e. each VM has an elastic IP attached to it). Customers may use other configurations depending on their requirements of security and public access requirements (e.g: customer collectors need access to AccelOps cluster over public Internet as opposed to reaching AccelOps over a


VPN)



- Hit Continue and Create VPC







- You should get a 'success' screen as follows



- Once this is done, you will see the following defined: VPC, subnet, route table, Internet gateway, DHCP option set, default Network ACLs, and a 'default' security group. The only change needed to this default VPC is to modify the 'default' security group to add inbound rules for ssh(22) and https(443). The example screenshots below show this configuration (user does not need to perform any action except for the security group). For better security, Users can customize the inbound rules and limit the IP addresses from where this VPC can be accessed from.

Alternatively, a VPN connection can be setup to customer's private data center in which the default inbound rules ensure that the VPC is firewalled off from public Internet. This user guide does not explain such a setup. Please refer to Amazon's documentation to set up VPN connections.



## Determine Local or NFS Storage for EventDB

If the aggregate EPS for your AccelOps installation requires a cluster (VA or SP + Worker nodes), then you **must** setup an NFS server. The next section describes setting up an NFS Server. If your storage requirements for EventDB is more than 1TB, it is recommended that you use an NFS server where you can configure LVM+RAID0 which is described in the next section. Although it is possible to setup a similar LVM+RAID0 on AccelOps node itself, this has not been tested.

To calculate storage requirements, here is an example. At 5000 EPS, you can calculate daily storage requirements to about 22-30GB (300k events take roughly 15-20MB on average in compressed format stored in eventDB). So, in order to have 6 months of data available for querying, you need to have 4-6TB of storage.

If you only need one AccelOps node and your storage requirements are lower than 1TB and is not expected to ever grow beyond this limit, you can avoid setting up an NFS server and use a local EBS volume for EventDB. For this option, skip the next section, and install AccelOps as described in section 'Setup AO Supervisor and Worker nodes in AWS'.

## Setup NFS Server in AWS for use by AccelOps VMs


To setup an NFS server in AWS that is also highly available across availability zone failures, several architectures and partner options are presented by an AWS Solutions Architect in this [talk \(40 min youtube video\)](#) and [link to slides](#).

In this section, we describe setting up a simple NFS server using EBS volumes. **Note: If you are running these machines in production, it is significantly cheaper to use EC2 Reserved Instances (1 or 3 year) as opposed to on-demand instances.**


Note that EBS volumes have 10 times higher durability guarantees than traditional disk drives because data is replicated within an availability

zone for [component failures \(RAID equivalent\)](#), so adding another layer of RAID on them does not provide you with higher durability guarantees. EBS has an annual failure rate (AFR) of 0.1 to 0.5%. In order to have higher durability guarantees, it is necessary to take periodic snapshots of the volumes. Snapshots are stored in AWS S3 which has [99.999999999% durability](#) (via synchronous replication of data across multiple data centers) and 99.99% availability.


1. To configure the NFS Server for AccelOps eventDB, launch an instance and choose 64-bit HVM Amazon Linux AMI for the NFS server and follow the wizard. The reason to choose HVM image over the default Paravirtualized (PV) image is that HVM image automatically includes drivers to support [enhanced networking](#) which uses SR-IOV for networking that results in higher performance (packets per second), lower latency, and lower jitter.



2. Choose instance type. For this example, I chose c3.2xlarge (current generation). Choose any instance type that has the capabilities 'EBS optimized' and 'High' or '10 gigabit' network performance. Click 'Next: Configure Instance Details'




3. Choose the VPC and Subnet you just created in the previous section. Make sure you choose a placement group that you created previously or create a new one. Also select EBS optimized instance




4. Click 'Next: Add Storage'. Add EBS volumes up to the capacity you need for AccelOps eventDB. For example, at 5000 EPS, you can calculate daily storage requirements to about 22-30GB (300k events take roughly 15-20MB on average in compressed format stored in eventDB). So, in order to have 6 months of data available for querying, you need to have 4-6TB of storage. On AWS, the maximum EBS volume is sized 1TB. In order to have larger disks, you need to create software RAID-0 volumes. You can attach at most 8 such volumes to an instance which gets you 8TB with RAID-0. Note that there is no advantage in using a different RAID configuration other than RAID-0 because it does not increase durability guarantees. In order to ensure much better durability guarantees, plan on performing regular snapshots which store the data in S3 (See next section 'Performing consistent snapshots of EBS volumes that host EventDB and CMDB' for details on how to do this). Since RAID-0 stripes data across these volumes, the aggregate IOPS you get will be the sum of the IOPS on individual volumes.


Click 'Next: Tag Instance'



5. Choose an appropriate name for the NFS instance, and click 'Next: Configure Security Group'




6. Select the default VPC security group that you configured previously. Click 'Next: Review and Launch'



7. Review the information on the review page and click 'Launch'

8. You will need to select a key-pair for communicating to the NFS server and select 'Launch Instances'



9. Make sure your NFS server comes up correctly.

10. Select the NFS server instance and click 'connect' and you will get information to ssh into the instance

11. Configure NFS mount point access give AO internal IP full access.

```

# Update the OS and libraries with the latest patches
$ sudo yum update -y

$ sudo yum install -y nfs-utils nfs-utils-lib lvm2

$ sudo su -
# echo Y | mdadm --verbose --create /dev/md0 --level=0 --chunk=256
--raid-devices=4 /dev/sdf /dev/sdg /dev/sdh /dev/sdi
# mdadm --detail --scan > /etc/mdadm.conf
# cat /etc/mdadm.conf
# dd if=/dev/zero of=/dev/md0 bs=512 count=1
# pvcreate /dev/md0
# vgcreate VolGroupData /dev/md0
# lvcreate -l 100%vg -n LogVolDataMd0 VolGroupData
# mkfs.ext4 -j /dev/VolGroupData/LogVolDataMd0
# echo "/dev/VolGroupData/LogVolDataMd0 /data ext4 defaults 1 1"
>> /etc/fstab
# mkdir /data
# mount /data
# df -kh

# vi /etc/exports
/data 10.0.0.0/24(rw,no_root_squash)
# exportfs -ar
# chkconfig --levels 2345 nfs on
# chkconfig --levels 2345 rpcbind on
# service rpcbind start
Starting rpcbind: [ OK ]
# service nfs start
Starting NFS services: [ OK ]
Starting NFS mountd: [ OK ]
Stopping RPC idmapd: [ OK ]
Starting RPC idmapd: [ OK ]
Starting NFS daemon: [ OK ]

```

## Performing consistent snapshots of EBS volumes that host EventDB and CMDB

In order to get high durability guarantees of AccelOps data, it is necessary to periodically create EBS snapshots (hourly, daily, weekly), and these are stored in S3. The EventDB is typically hosted as a RAID-0 volume of several EBS volumes (as given in example in previous section). In order to reliably snapshot these EBS volumes together, a script called `ec2-consistent-snapshot` can be used to briefly freeze the volumes and create a snapshot. Another script called `ec2-expire-snapshots` can be used to schedule cron jobs to delete old snapshots that are no longer needed. CMDB is hosted on a much smaller EBS volume and the same scripts can be used to snapshot this volume.

The details of how download these scripts and setup periodic snapshots and expiration is described in the following blog post:

<http://twigmon.blogspot.com/2013/09/installing-ec2-consistent-snapshot.html>

The actual scripts can be downloaded from these github projects

<https://github.com/alestic/ec2-consistent-snapshot>

<https://github.com/alestic/ec2-expire-snapshots>

## Setup AO Supervisor and Worker nodes in AWS

1. Logon to AWS EC2 console using your AWS account and go to EC2 dashboard
2. If you have not created a VPC, please do that as described in the beginning of this document
3. Click Launch Instance

- Select "Community AMIs" and search for **the following AMI ids** and click Select


The following AMIs are for release 4.2.1 and 4.2.3

Region Code	Region Name	4.2.1 Super/Worker AMI-id	4.2.3 Super/Worker AMI-id
us-east-1	US East (Northern Virginia)	ami-cc7e9ba4	ami-2872a340
us-west-1	US West (Northern California)	ami-9c4443d9	ami-17242852
us-west-2	US West (Oregon) Region	ami-2fc6b21f	ami-d90145e9
eu-west-1	EU (Ireland) Region	ami-bf08c7c8	ami-aa3ce3dd




- Select 'Compute Optimized' instances, and preferably select one of c3 instances with 'High' network performance or 10 gigabit performance. C3 instances (current generation) run on the latest Intel Xeons that AWS provides as of AccelOps 3.7.6. For this guide, I chose c3.xlarge and click Continue.


**Note: If you are running these machines in production, it is significantly cheaper to use EC2 Reserved Instances (1 or 3 year) as opposed to on-demand instances.**




- Click Next: Configure Instance Details




- Click 'Next: Add Storage' - In this example, provisioned IOPS was selected for all 3 nodes in the cluster even though only Super node's CMDB data needs higher IOPS. For workers, Standard IOPS is sufficient. You can always launch with Standard IOPS first, and then add a separate EBS volume for CMDB separately with higher provisioned IOPS. As described in section 'Determine Local or NFS Storage for EventDB', if you are using local storage for EventDB, allocate a new EBS volume sized 50GB to 1TB (depending on storage requirement) with provisioned IOPS at 2000 and set it to device /dev/xvdi




- Click 'Next: Tag Instance'. The name you give here will be assigned to all the instances you launch together in this wizard run. You will then have to rename the instances to its role (Super, worker1, worker2, etc.)



- Click 'Next: Configure Security Group'. Choose the VPC's 'default' security group. AccelOps needs access to ports 443 (HTTPS) for GUI and API access, SSH(22) for remote management which are already included in the default security group. This group allows all traffic between instances within the VPC.




- Click 'Review and Launch' to be taken to a review page. Make sure all the information provided is correct. Click 'Launch'




- You will be asked to select an existing key pair or create a new key pair to connect to these instances via SSH. If you use an existing key pair, make sure you have access to it. If you are creating a new key pair, download the private key and store them in a secure location accessible from the machine from where you usually connect to these AWS instances. Then click 'Launch Instances'




- Click 'View Instances'



- Make sure your instances are running correctly. They will all be named 'AccelOps Cluster' to begin with. Rename them appropriately to distinguish between Super, Worker1, Worker2 as the screenshot below shows.



- Create VPC based Elastic IPs and attach them to your nodes so that the public IPs do not change when you stop and start instances.



? Unknown Attachment

15. You will now need to ssh to AccelOps SP. From the EC2 dashboard, select the instance, click Connect, select "Connect with a standalone SSH client" and locate the command example from the sample screenshot below. To login to AccelOps nodes, replace ec2-user with root as user name.

? Unknown Attachment

16. SSH to the Super node. **Note: If you are installing a single AccelOps node with Local EventDB storage, please refer to section '(Optional - If using local storage) Configuring Local EventDB Storage' to perform additional steps**
17. Run "**cd /opt/phoenix/deployment/jumpbox/aws**"
18. Run the script **pre-deployment.sh** to configure host name and NFS mount point. Host name can be obtained from the EC2 dashboard, selecting the instance, right clicking on the instance, selecting "Connect with a standalone SSH client" and locating Public DNS.
- Agree with License Agreements
  - Enter host name
  - Enter the mount point -
    - <NFS Server IP>:/data (For <NFS Server IP>, use the 10.0.0.X IP address of the NFS Server running within the VPC) **q**
    - If using Local EventDB Storage as configured in section '(Optional - If using local storage) Configuring Local EventDB Storage', then use /dev/xvdi
19. The system will reboot
20. Logon to AccelOps SP; run "**cd /opt/phoenix/deployment/jumpbox/aws**" and run the script **deployment.sh**
21. The system will reboot
22. Logon to AccelOps SP; run "**cd /opt/phoenix/deployment/jumpbox/aws**" and run the script **post-deployment.sh**
23. The system will reboot is now ready
24. Register the virtual appliance. Launch browser to AccelOps VA using the DNS name in step 16. Register using the username and password provided by AccelOps.

? Unknown Attachment

25. Once registration succeeds, log on using the default username and password.  
User ID: admin  
Password: admin\*1

? Unknown Attachment

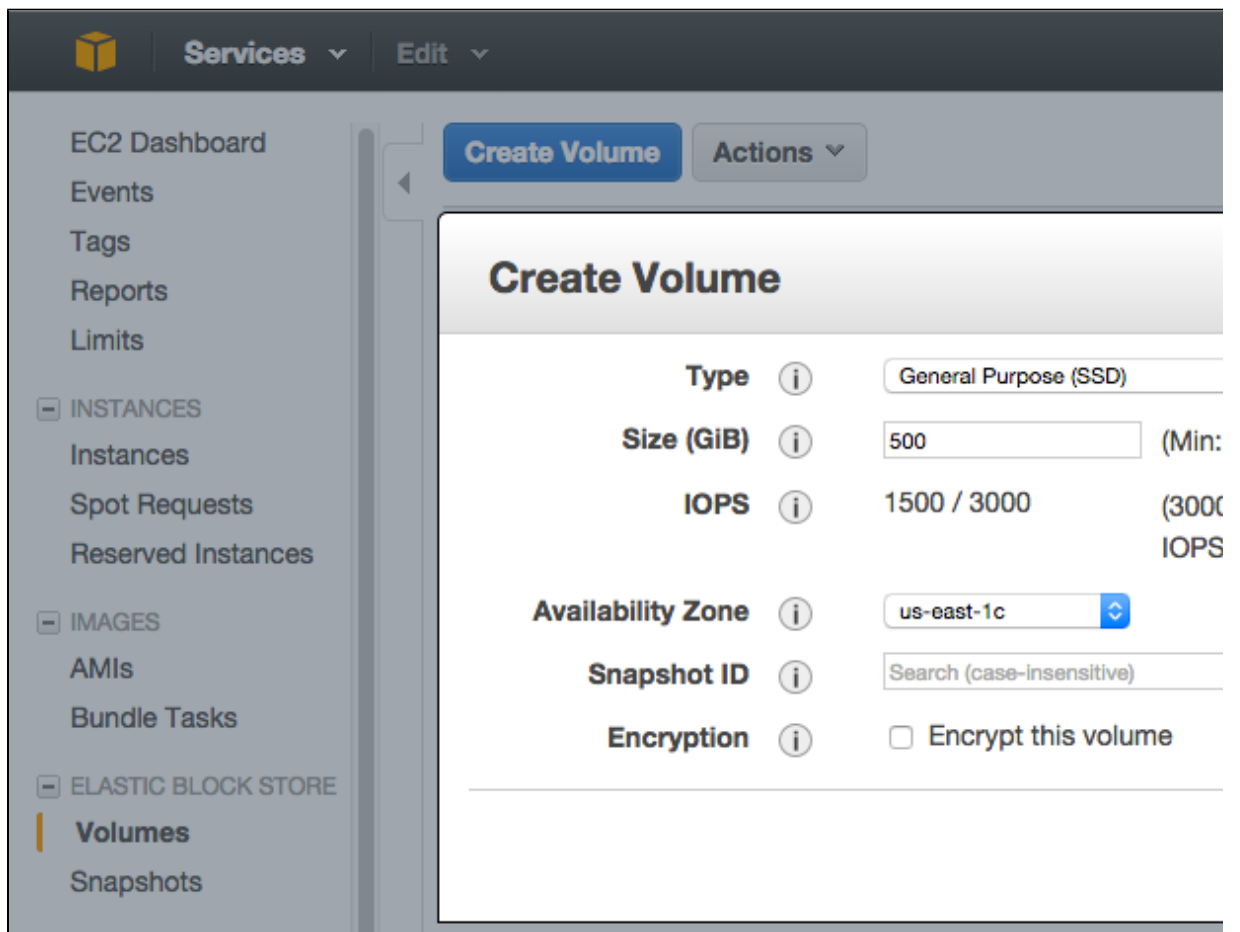
26. The installation is now complete. The health of the system can be checked by ssh-ing to the device (as in step 14) and running phstatus command

? Unknown Attachment

### (Optional - If using Local storage) Configure Local Storage for EventDB

- Login to AWS, Go to Ec2 > Volumes > Create Volume
  - Allocate EBS volume sized 100 GB to 1 TB (depending on storage requirement)
  - Select 'Availability Zone' same as AccelOps Supervisor instance launched region
  - Click on Create Volume





The screenshot displays the AWS Management Console interface for creating an EBS volume. The left sidebar shows the navigation menu with 'Volumes' selected under 'ELASTIC BLOCK STORE'. The main panel is titled 'Create Volume' and includes a 'Create Volume' button and an 'Actions' dropdown. The configuration fields are as follows:

Field	Value
Type	General Purpose (SSD)
Size (GiB)	500 (Min: 1)
IOPS	1500 / 3000 (3000 IOPS)
Availability Zone	us-east-1c
Snapshot ID	Search (case-insensitive)
Encryption	<input type="checkbox"/> Encrypt this volume

2. Select recently created Volume and Attach to Supervisor instance
  - a. Select Supervisor instance name
  - b. Change device name to **'/dev/xvdi'** and click 'Attach' button

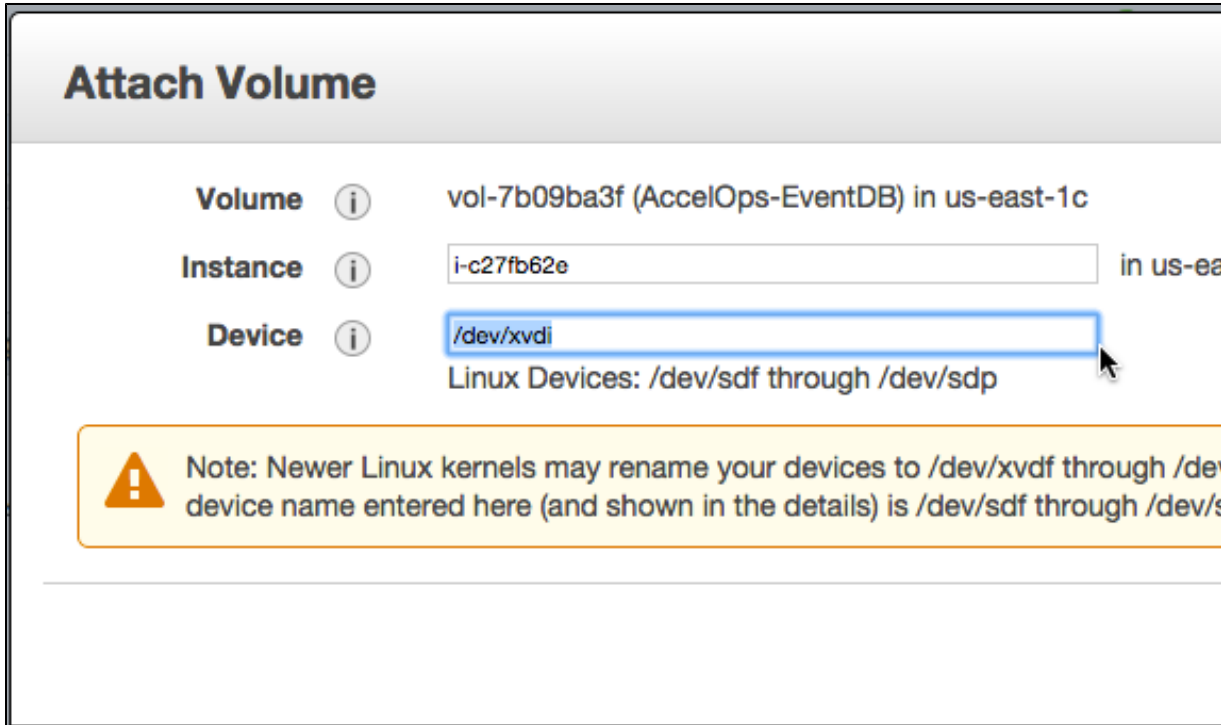
**Create Volume** **Actions** ^

Filter by tags

- Delete Volume
- Attach Volume**
- Detach Volume
- Force Detach Volume
- Create Snapshot
- Change Auto-Enable IO Setting
- Add/Edit Tags

<input type="checkbox"/>	Name		Volume Type	Snapshot	Created	Avail	
<input type="checkbox"/>			gp2	snap-93d7692a	November 20, 2014...	us-ee	
<input type="checkbox"/>			gp2		November 20, 2014...	us-ee	
<input type="checkbox"/>			io1 (1800)	snap-89d76930	November 20, 2014...	us-ee	
<input type="checkbox"/>		vol-931baad7	60 GiB	io1 (1800)	snap-94d7692d	November 20, 2014...	us-ee
<input checked="" type="checkbox"/>	AccelOps-EventDB	vol-7b09ba3f	500 GiB	gp2	November 21, 2014...	us-ee	

**Volume:** | vol-7b09ba3f (AccelOps-EventDB)



**(Optional - if Worker is needed) Setup AO Worker node in AWS**

1. If you have not already launched worker instances, then launch an AO Worker instance following steps 1-16 as described in previous section.
2. Follow steps 17-18 as in AO Supervisor case for basic setup. The system will now reboot and is ready for use.
3. Logon to AO Supervisor node and register the Worker to the Super
4. The installation is now complete. The health of the system can be checked by ssh-ing to the device (as in step 15) and running phstatus command

**(Optional - if Collector is needed) Setup AO Collector node in AWS**

1. First launch an AO Collector instance following steps 1-14 as in AO Supervisor case, **except use the following AMI id** for Collector and choose a large instance (2 core, 4GB).

Region Code	Region Name	4.2.1 Collector AMI-id	4.2.3 Collector AMI-id
us-east-1	US East (Northern Virginia)	ami-49190320	ami-5a954532
us-west-1	US West (Northern California)	ami-a24443e7	ami-2b24286e
us-west-2	US West (Oregon) Region	ami-aa1f6b9a	ami-2f02461f
eu-west-1	EU (Ireland) Region	ami-a708c7d0	ami-903ce3e7

2. Register Collector to Supervisor
  - a. Logon to AccelOps Supervisor

- b. Create an organization to which the collector AMI will belong to. Create a collector for that organization. Assign the collector IP to the collector's public IP.
- c. Register the collector
  - i. Locate a Windows machine on AWS
  - ii. Open a Remote desktop session from your PC to that Windows machine on AWS.
    - 1. Within the remote desktop session, launch a browser session to the private address of AccelOps Collector WMI
    - 2. Fill in the registration information and click Save (Register)




- 3. Once registration completes, the collector will reboot
- 4. Connect to AccelOps Collector AMI via SSH. Make sure all processes are up.

## Deploy AccelOps Collector AMI


This section discusses how to deploy AccelOps Collector as an Amazon Machine Instance for monitoring Amazon Web Services (AWS) EC2 instances.

- 1. Logon to Amazon Web Services.
- 2. Find public AccelOps Collector AMI by searching the ami ids as per the following table

Region Code	Region Name	4.2.3 Collector AMI-id
us-east-1	US East (Northern Virginia)	ami-5a954532
us-west-1	US West (Northern California)	ami-2b24286e
us-west-2	US West (Oregon) Region	ami-2f02461f
eu-west-1	EU (Ireland) Region	ami-903ce3e7



- 3. Launch AccelOps Collector AMI
  - a. Select the AccelOps Collector AMI and click Launch



- b. Fill in the following information
  - i. Select Large instance
    - Placeholder box: ? Unknown Attachment
  - ii. Use default for both Kernel ID and RAM Disk ID
    - Placeholder box: ? Unknown Attachment
  - iii. Create and download key pair
    - Placeholder box: ? Unknown Attachment
  - iv. Create Security Group to allow specific protocols
    - Placeholder box: ? Unknown Attachment
    - Placeholder box: ? Unknown Attachment
  - v. Review and Launch
    - Placeholder box: ? Unknown Attachment
    - Placeholder box: ? Unknown Attachment

- c. Click "Instances" and see the new AMI (instance Id i-f5ff7490) initializing..
  - Placeholder box: ? Unknown Attachment

- d. Once initialized, the AccelOps Collector AMI instance show like this
  - Placeholder box: ? Unknown Attachment

- e. Note the private and public IP addresses for the collector

### Validate AccelOps Collector connectivity and health

1. Connect to AccelOps Collector AMI via SSH
  - a. Logon to Amazon Web Services
  - b. Go to EC2 > My instances
  - c. Click on the AccelOps Collector AMI
  - d. Choose Connect and the help window appears. Follows the instructions on that popup to connect to AccelOps AMI via SSH



2. Run "phstatus" - it will show the following processes but they will all be down since the collector is not yet registered.




### Register AccelOps Collector

1. Logon to AccelOps Supervisor
2. Create an organization to which the collector AMI will belong to. Create a collector for that organization. Assign the collector IP to the collector's public IP.
3. Register the collector
  - a. Locate a Windows machine on AWS
  - b. Open a Remote desktop session from your PC to that Windows machine on AWS.
  - c. Within the remote desktop session, launch a browser session to the private address of AccelOps Collector WMI
  - d. Fill in the registration information and click Save (Register)



4. Once registration completes, the collector will reboot
5. Connect to AccelOps Collector AMI via SSH. Make sure all processes are up




### Setup AccelOps Report Server node

1. Logon to AWS EC2 console using your AWS account and go to EC2 dashboard
2. If you have not created a VPC, please do that as described in the beginning of this document
3. Go to EC2 Dashboard and Click on 'Launch Instance' button
4. Select "Community AMIs" and search for **the following AMI ids** and click Select

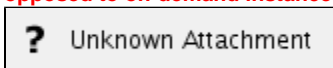
The following AMIs are for release 4.2.1 and 4.2.3

Region Code	Region Name	4.3.1 Report Server AMI-id
us-east-1	US East (Northern Virginia)	<b>ami-70851f18</b>
us-west-1	US West (Northern California)	
us-west-2	US West (Oregon) Region	
eu-west-1	EU (Ireland) Region	

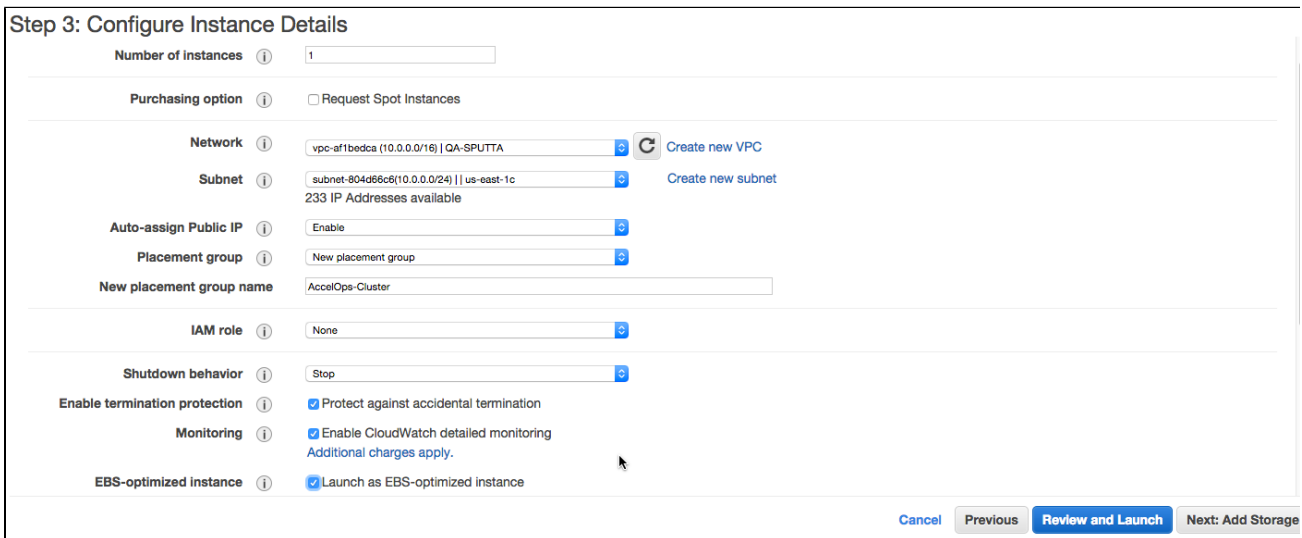


5. Select 'Compute Optimized' instances, and preferably select one of c3 instances with 'High' network performance or 10 gigabit performance. C3 instances (current generation) run on the latest Intel Xeons that AWS provides as of AccelOps 3.7.6. For this guide, I chose c3.4xlarge and click Continue.

**Note: If you are running these machines in production, it is significantly cheaper to use EC2 Reserved Instances (1 or 3 year) as opposed to on-demand instances.**



6. Click Next: Configure Instance Details



7. Click 'Next: Add Storage' - Select Provisioned IOPS. As described in section 'Determine Local or NFS Storage for ReportDB', if you are using local storage for ReportDB, allocate a new EBS volume sized 50GB to 1TB (depending on storage requirement) with provisioned IOPS at 2000 and set it to device /dev/xvdi

## Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance using the storage options in Amazon EC2.

Type <small>(i)</small>	Device <small>(i)</small>	Snapshot <small>(i)</small>	Size (GiB) <small>(i)</small>	Volume Type <small>(i)</small>
Root	/dev/sda1	snap-a2c35910	80	Provisioned IOPS (SSD)
EBS	/dev/xvdj	snap-a1c35913	60	Provisioned IOPS (SSD)

**Add New Volume**

- Click 'Next: Tag Instance'. The name you give here will be assigned to the 'Report Server instance'.

## Step 5: Tag Instance

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Report Server.

Key <small>(127 characters maximum)</small>	Value
Name	Report Server

**Create Tag** (Up to 10 tags maximum)

- Click 'Next: Configure Security Group'. Choose the VPC's 'default' security group. AccelOps needs access to ports 443 (HTTPS) for GUI and API access, SSH(22) for remote management which are already included in the default security group. This group allows all traffic between instances within the VPC.

**? Unknown Attachment**

- Click 'Review and Launch' to be taken to a review page. Make sure all the information provided is correct. Click 'Launch'

## Step 7: Review Instance Launch



### AccelOps-Reporter-4.3.1.1157-EBS-Update - ami-70851f18

AccelOps-Reporter-4.3.1.1157-EBS-Update

Root Device Type: ebs    Virtualization type: paravirtual

#### ▼ Instance Type

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage
c3.2xlarge	28	8	15	2 x 80

#### ▼ Security Groups

Security Group ID	Name
sg-7e04d51b	default

#### All selected security groups inbound rules

Security Group ID	Type ⓘ	Protocol ⓘ
sg-7e04d51b	Custom TCP Rule	TCP
sg-7e04d51b	SSH	TCP

- You will be asked to select an existing key pair or create a new key pair to connect to these instances via SSH. If you use an existing key pair, make sure you have access to it. If you are creating a new key pair, download the private key and store them in a secure location accessible from the machine from where you usually connect to these AWS instances. Then click 'Launch Instances'



## Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Choose an existing key pair

**Select a key pair**

accelops

I acknowledge that I have access to the selected private key file (accelops.pem), and that without this file, I won't be able to log into my instance.

[Cancel](#) [Launch Instances](#)

12. Click 'View Instances'

## Launch Status

### ✓ Your instance is now launching

The following instance launch has been initiated: [i-8a870166](#) [View launch log](#)



### Get notified of estimated charges

[Create billing alerts](#) to get an email notification when estimated charges on your AWS bill ex

## How to connect to your instance

Your instance is launching, and it may take a few minutes until it is in the **running** state, when it will be ready to connect. You can stop or terminate your instance.

Click **View Instances** to monitor your instance's status. Once your instance is in the **running** state, you can connect to it.

### ▼ Here are some helpful resources to get you started

- [How to connect to your Linux instance](#)
- [Amazon EC2: User Guide](#)
- [Learn about AWS Free Usage Tier](#)
- [Amazon EC2: Discussion Forum](#)

While your instances are launching you can also

[Create status check alarms](#) to be notified when these instances fail status checks. (Additional charges may apply)

[Create and attach additional EBS volumes](#) (Additional charges may apply)

[Manage security groups](#)

13. Make sure your instances are running correctly. They will all be named 'AccelOps Cluster' to begin with. Rename them appropriately to distinguish between Super, Worker1, Worker2 as the screenshot below shows.

? Unknown Attachment

14. Create VPC based Elastic IPs and attach them to your nodes so that the public IPs do not change when you stop and start instances.

? Unknown Attachment

? Unknown Attachment

15. You will now need to ssh to AccelOps SP. From the EC2 dashboard, select the instance, click Connect, select "Connect with a standalone SSH client" and locate the command example from the sample screenshot below. To login to AccelOps nodes, replace ec2-user with root as user name.

? Unknown Attachment

16. SSH to the Super node. **Note: If you are installing a single AccelOps node with Local EventDB storage, please refer to section '(Optional - If using local storage) Configuring Local EventDB Storage' to perform additional steps**
17. Run "**cd /opt/phoenix/deployment/jumpbox/aws**"
18. Run the script **pre-deployment.sh** to configure host name and NFS mount point. Host name can be obtained from the EC2 dashboard, selecting the instance, right clicking on the instance, selecting "Connect with a standalone SSH client" and locating Public DNS.
- a. Agree with License Agreements
  - b. Enter host name
  - c. Enter the mount point -
    - i. <NFS Server IP>:/data (For <NFS Server IP>, use the 10.0.0.X IP address of the NFS Server running within the VPC) **q**
    - l
    - ii. If using Local EventDB Storage as configured in section '(Optional - If using local storage) Configuring Local EventDB Storage', then use /dev/xvdi
19. The system will reboot
20. Logon to AccelOps SP; run "**cd /opt/phoenix/deployment/jumpbox/aws**" and run the script **deployment.sh**
21. The system will reboot
22. Logon to AccelOps SP; run "**cd /opt/phoenix/deployment/jumpbox/aws**" and run the script **post-deployment.sh**
23. The system will reboot is now ready
24. Register the virtual appliance. Launch browser to AccelOps VA using the DNS name in step 16. Register using the username and password provided by AccelOps.

? Unknown Attachment

25. Once registration succeeds, log on using the default username and password.  
User ID: admin  
Password: admin\*1

? Unknown Attachment

26. The installation is now complete. The health of the system can be checked by ssh-ing to the device (as in step 14) and running phstatus command

? Unknown Attachment

# Visual Analytics

Visual Analytics is an add-on for AccelOps that lets you create individual visualizations of AccelOps report data, as well as dashboards containing multiple visualization charts. AccelOps Visual Analytics has three components:

1. The AccelOps Report Server, which syncs with and replicates AccelOps reports in near-real time.
2. The Visual Analytics Server from Tableau Software, which enables the publication and distribution of your visualizations.
3. The Visual Analytics Desktop, which is your primary tool for creating visualizations

Topics in this section will walk you through the process of installing and configuring all three components, as well as a brief introduction to creating single and multiple sheet workbooks of visualizations. For more detailed information about Tableau Software's Visual Analytics Server and Visual Analytics Desktop, including examples of creating sheets and workbooks, you should consult the [Product Support section of the Tableau Software website](#).

- [AccelOps Visual Analytics: Overview](#)
- [Installation and Configuration of AccelOps Visual Analytics](#)
- [Creating and Managing Workbooks](#)
- [Working with AccelOps Report Server](#)

# AccelOps Visual Analytics: Overview

- [Architecture](#)
- [Set Up](#)
- [Using Visual Analytics](#)

## Architecture

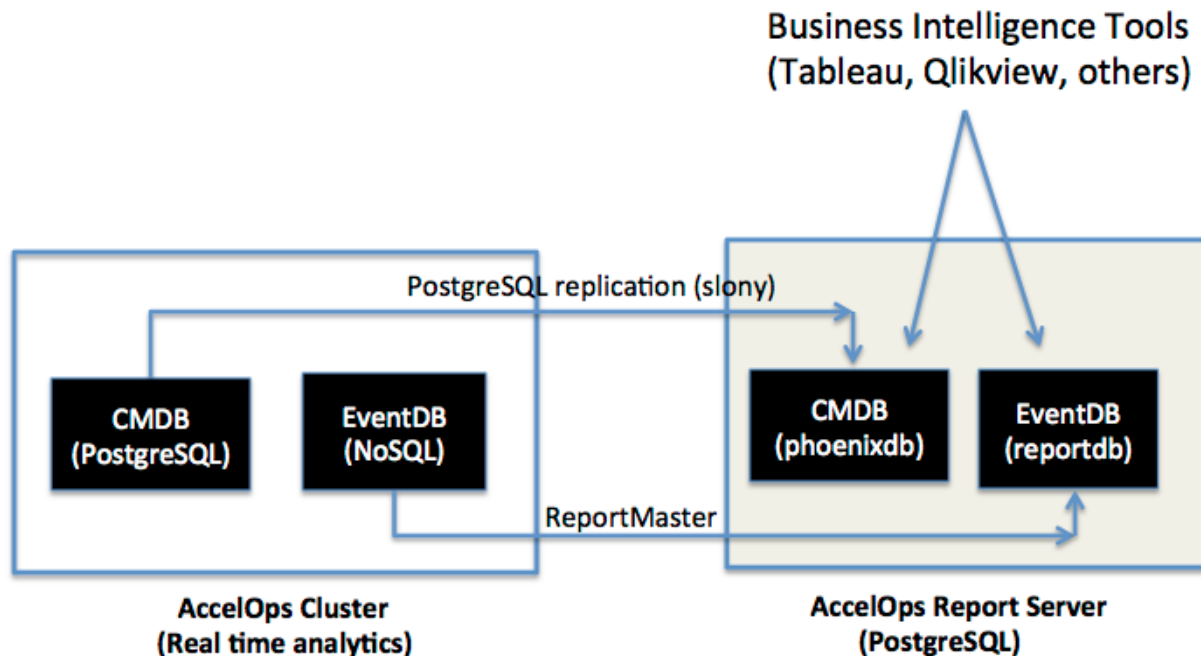
With AccelOps Visual Analytics, you can now create visual representations of the data that is stored in AccelOps. This includes:

- **Structured data** stored in the AccelOps CMDB relational PostgreSQL database, such as:
  - Discovered information about devices, systems, applications and users
  - Identity and location
  - Incidents and notifications
- **Unstructured data** such as logs, events, performance metrics etc. that is stored in the AccelOps EventDB NoSQL database that is accessible by the cluster members (Supervisor, Workers) over NFS.

For fast visual analytics without compromising the real time performance of AccelOps cluster, the data is exported to a separate virtual machine, the **AccelOps Report Server**, running PostgreSQL. The Report Server contains two databases that are queried by AccelOps Visual Analytics:

- **phoenixdb**  
This database contains the entire AccelOps CMDB and is populated via asynchronous PostgreSQL replication (slony) in near-real time.
- **reportdb**  
This database contains the results of event queries.

You can find more information about AccelOps Report Server in the topic [Report Server Architecture: phoenixdb and reportdb](#) and its related topics.



## Set Up

There are three components to Accelops Visual Analytics:

1. Accelops Report Server
2. Visual Analytics Server (Powered by Tableau)
3. Visual Analytics Desktop (Powered by Tableau)

You must first install Report Server as described in [Installing and Configuring AccelOps Report Server](#). After installing Visual Analytics Server on

a Windows server, and installing Visual Analytics Desktop on a Windows or Mac OS X device, you then connect the two systems as described in [Installing and Configuring Visual Analytics Desktop](#). When this connection is established, it automatically triggers the remote registration and configuration of the AccelOps Report Server, including replication of the CMDB and EventDB data from the AccelOps Cluster to the AccelOps Report Server, as well as the user account required for access to the original databases. Registration of the Report Server and replication of the AccelOps cluster data may take some time depending on the size of the original CMDB. Registration is complete when the replication process catches up with the latest data in the system. From that point on replication from CMDB to AccelOps Report Server takes place in near real time, letting you run Visual Analytics queries against CMDB data that has been replicated to the Report Server's phoenixdb.

You can find full information about setting up all components of AccelOps Visual Analytics in the section [Installation and Configuration of AccelOps Visual Analytics](#)

## Using Visual Analytics

Using AccelOps Visual Analytics involves first syncing reports contained in the primary AccelOps application to the AccelOps Report Server.

1. Log in to the AccelOps application.
2. Select a report that can be synced to the Report Server.  
Currently only reports that contain a **Group By** condition can be synced. Both system and user-created reports can be synced as long as there is a Group By condition.
3. Click **Sync**.

When the sync process initiates, the Supervisor node dynamically creates a table within the Report Server reportdb database. When the sync is established, it will run every five minutes, and the last five minutes of data in the synced report will be pushed to the corresponding table. This lets you run Visual Analytics on event data stored in the Report Server reportdb database.

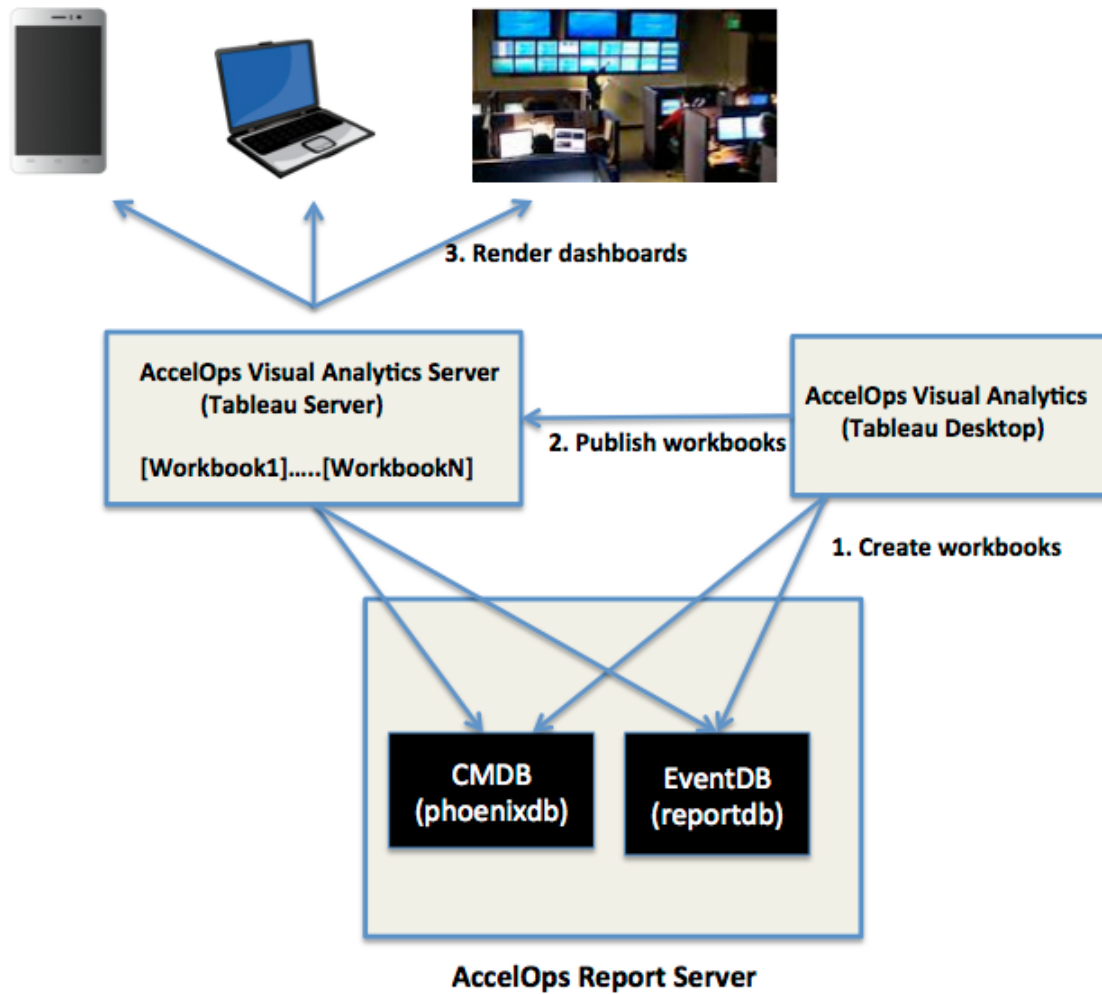
Viewing charts and dashboards in Visual Analytics involves using Visual Analytics Desktop to create workbooks that are connected to a datasource within Report Server (reportdb), and then providing users with access to those workbooks.

1. Install Visual Analytics Server on a Windows Server.
2. Load workbooks created in Visual Analytics Desktop onto Visual Analytics Server.
3. Create users and assign them access to specific workbooks on Visual Analytics Server.
4. Log on to Visual Analytics Server from a Windows or Mac OS X device, such as a phone, tablet, or PC, to view the contents of workbooks.

A few basic workbooks are provided by Accelops. To create your own workbook:

1. Install Visual Analytics Desktop onto your desktop.
2. Connect to the appropriate data source (phoenixdb or reportdb) to create a workbook.
3. Publish the workbook to the Visual Analytics Server.

You can find more information about workbooks in the section [Creating and Managing Workbooks](#).



**(PostgreSQL)**



# Installation and Configuration of AccelOps Visual Analytics

The process for installing and configuring AccelOps Visual Analytics involves three steps:

1. [Installing and configuring the AccelOps Report Server.](#)
2. [Installing and configuring the Visual Analytics Server.](#)
3. [Installing and configuring the Visual Analytics Desktop.](#)

You can find requirements for all three components in [Requirements for AccelOps Visual Analytics.](#)

## Requirements for AccelOps Visual Analytics

- AccelOps Visual Analytics Report Server
  - Hardware Requirements for Report Server nodes
- Virtual Analytics Server and Virtual Analytics Desktop (Tableau Software)
  - Visual Analytics Server Requirements
    - Supported Windows Servers
    - Memory, Cores, and Disk Space Recommendations
    - Accounts and Permissions
    - Configuration Information
    - Ports
    - Drivers
  - Visual Analytics Desktop Requirements
    - Supported operating systems
    - Memory and disk space recommendations
    - Drivers

### AccelOps Visual Analytics Report Server

You install Visual Analytics Report Server as a node running under the AccelOps All-in-One Virtual Appliance, and these requirements assume that you have already set up and installed AccelOps. If you are working with a fresh install of AccelOps that includes Report Server, see the topics under [Fresh Installation](#) for complete requirements and installation instructions for the AccelOps Virtual Appliance.

You must install Visual Analytics Report Server on a dedicated machine.

### Hardware Requirements for Report Server nodes

Model	Quantity	Host SW	Processor	Memory	OS/App Storage	Reports Data Storage (1 year)
Report Server	1	ESX	8 Core 3 GHz, 64 bit	16 GB	200GB (80GB OS/App; 60GB CMDB )	?

### Virtual Analytics Server and Virtual Analytics Desktop (Tableau Software)

These are the basic requirements for installing Visual Analytics Server and Visual Analytics Desktop, both of which are powered by [Tableau Software](#). For more detailed information about Tableau Server, see the [Tableau Server Administrator's Guide](#). For more information about Tableau Desktop, see the [Tableau Desktop User's Guide](#).

### Visual Analytics Server Requirements

#### Supported Windows Servers

Tableau Server is available in 32-bit and 64-bit versions. We recommend running the 64-bit version of Tableau Server on a 64-bit operating system.

- Windows Server 2003 R2 SP2 or higher
- Windows Server 2008
- Windows Server 2008 R2
- Windows Server 2012 and 2012R2

#### Memory, Cores, and Disk Space Recommendations

The minimum requirements for running the 64-bit version of Tableau Server are 4 cores and 8 GB of RAM.

These are the **minimum** recommendations based on the number of users on the server:

Deployment Type	Server Users	CPU	RAM	Free Disk Space
Small	<25	4-core	8 GB	20 GB
Medium	<100	8-core	32 GB	100 GB

Enterprise	>100	16-core	32 GB or more	100 GB or more
------------	------	---------	---------------	----------------

## Accounts and Permissions

- **Administrative account**

The account under which you install Tableau Server must have permission to install software and services.

- **Run As account (optional)**

A Run As User account for the Tableau Server service to run under is useful if you're using NT Authentication with data sources, or if you're planning on using SQL Server impersonation. For more information, see [Run As User](#) and [SQL Server Impersonation](#) in the Tableau Administrator's Guide.

## Configuration Information

When you install and configure Tableau Server you may be asked for the following information:

Option	Description	Your Information
<b>Server Account</b>	The server must have a user account that the service can use. The default is the built-in Windows Network Service account. If you use a specific user account you'll need the domain name, user name, and password.	Username: Password: Domain:
<b>Active Directory</b>	Instead of using Tableau's built-in user management system, you can authenticate through Active Directory. If so, you'll need the <a href="#">fully-qualified domain name</a> .	Active Directory Domain:
<b>Open port in Windows firewall</b>	When selected, Tableau Server will open the port used for http requests in the Windows Firewall software to allow other machines on your network to access the server.	<input type="checkbox"/> - Yes <input type="checkbox"/> - No

## Ports

- **TCP/IP, IIS and port 80**

Tableau Server's gateway listens on port 80, which is also used by Internet Information Services (IIS) by default. If you are installing Tableau Server on a machine that's also running IIS, you should modify Tableau's gateway port number to avoid conflict with IIS. Tableau Server also requires several TCP/IP ports to be available to the server. See [TCP/IP Ports](#) and [Edit the Default Ports](#) in the Tableau Administrator's Guide for details.

## Drivers

You may need to install additional database drivers. Download drivers from [www.tableausoftware.com/support/drivers](http://www.tableausoftware.com/support/drivers).

## Visual Analytics Desktop Requirements

### Supported operating systems

Tableau Desktop is available in 32-bit and 64-bit versions and can be used with these operating systems:

- Windows 7
- Windows 8
- Windows 8.1
- Mac OSX

### Memory and disk space recommendations

The minimum requirements for Tableau Desktop are 4 GB of RAM and 10 GB disk free space.

## Drivers

You may need to install additional database drivers. Download drivers from [www.tableausoftware.com/support/drivers](http://www.tableausoftware.com/support/drivers).

## Installing and Configuring AccelOps Report Server

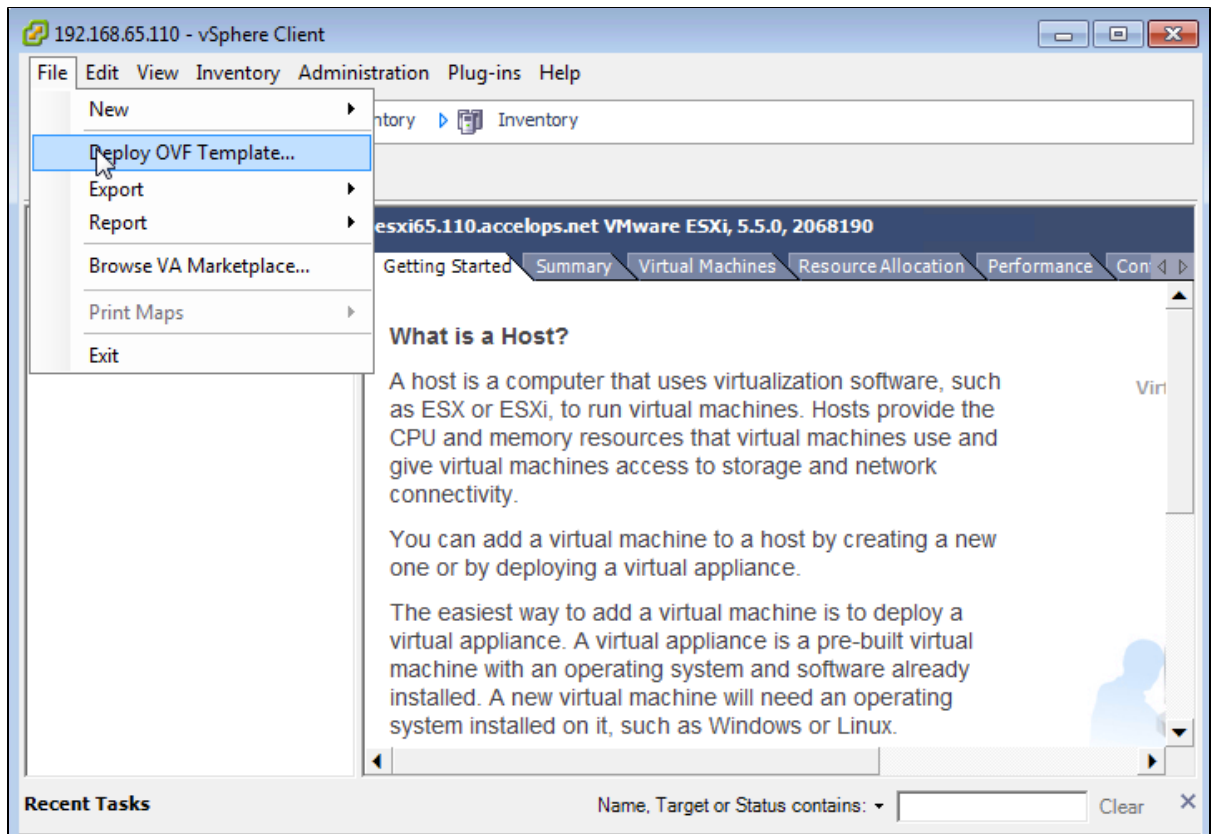
These instructions are for installing Report Server on VMWare ESX, and assume that you have already installed and configured an AccelOps environment. For instructions for a complete AccelOps install, see the topics under [Fresh Installation](#). Information on installing and configuring Report Server with other hypervisors will be included in upcoming releases.

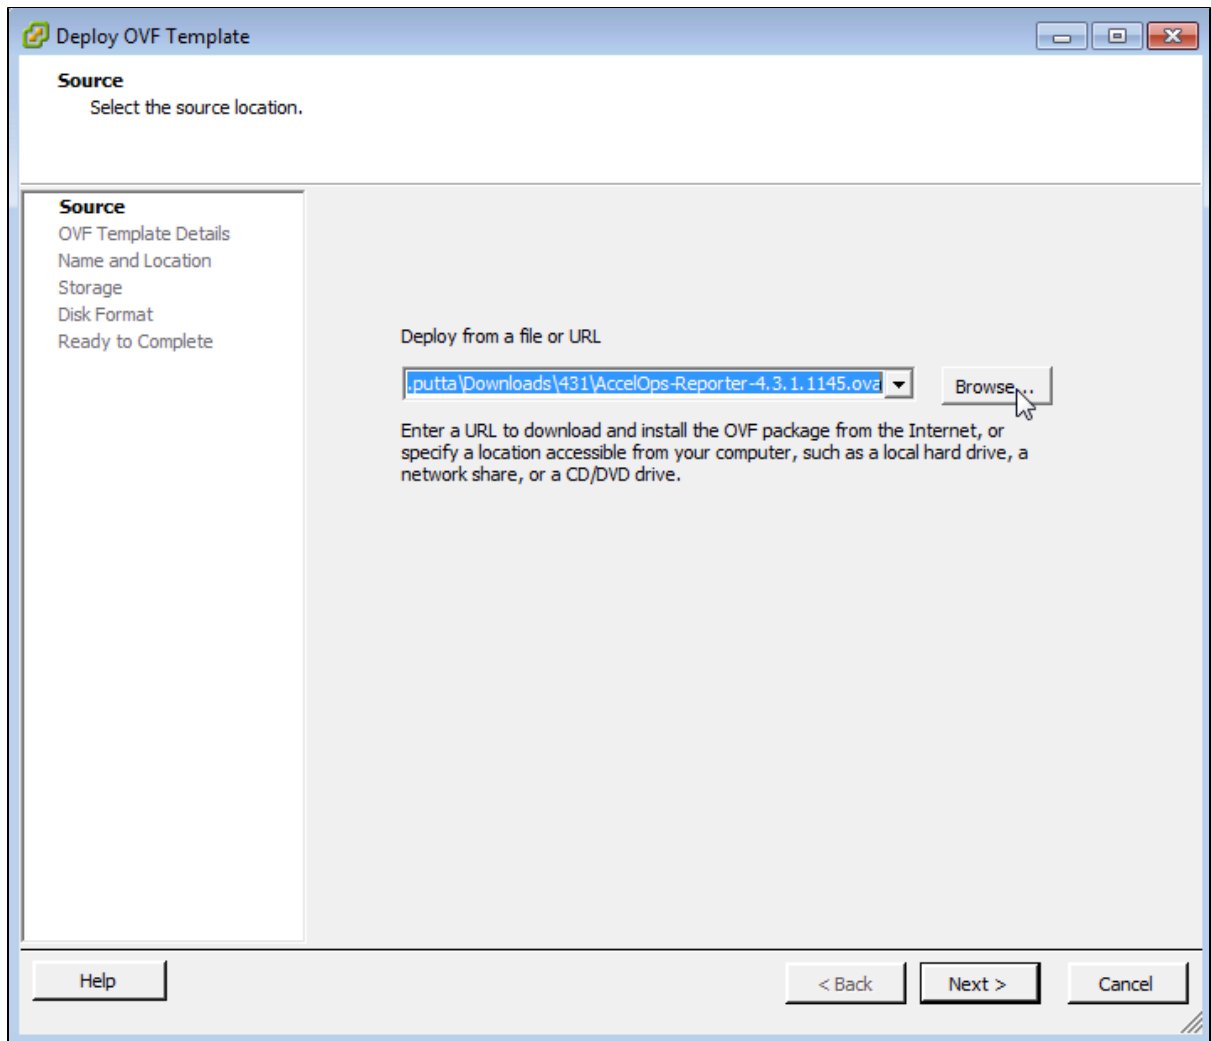
- Section 1: Import a Report Server image into VMware ESX server
  - Step 1: Begin importing Report Server
  - Step 2: Review the Virtual Appliance details
  - Step 3: End User License Agreement
  - Step 4: Name and Location
  - Step 5: Storage
  - Step 6: Disk Format
  - Step 7: Ready to Complete
- Section 2: Edit the Report Server hardware settings
  - Step 1: Edit Settings
  - Step 2: Verify Allocated Memory
- Section 3: Select mount point to Store Reports data
  - Step 1: Adding an Additional Hard Disk to store the Reports data
  - Step 2: Create new NFS mount point to store the Reports data
- Section 4: Configure the Report Server from the VMware Console
  - Step 1: Start Report Server
  - Step 2: Configure the Timezone
  - Step 3: Configure the Network
  - Step 4: Select Mount Point to store your data
- Section 5: Register the Report Server to Supervisor
  - Step 1: Re-register AccelOps Supervisor with new license
  - Step 2: Add Report Server on Supervisor UI

### Section 1: Import a Report Server image into VMware ESX server

#### Step 1: Begin importing Report Server

1. Download Report server OVA package from AccelOps image server
2. Log into VMware vSphere Client
  - a. Click on File -> Select Deploy OVF Template
  - b. Click on Browse button to choose OVA file
  - c. Locate Report Server .ova file (Example: AccelOps-Reporter-4.3.1.1145.ova)

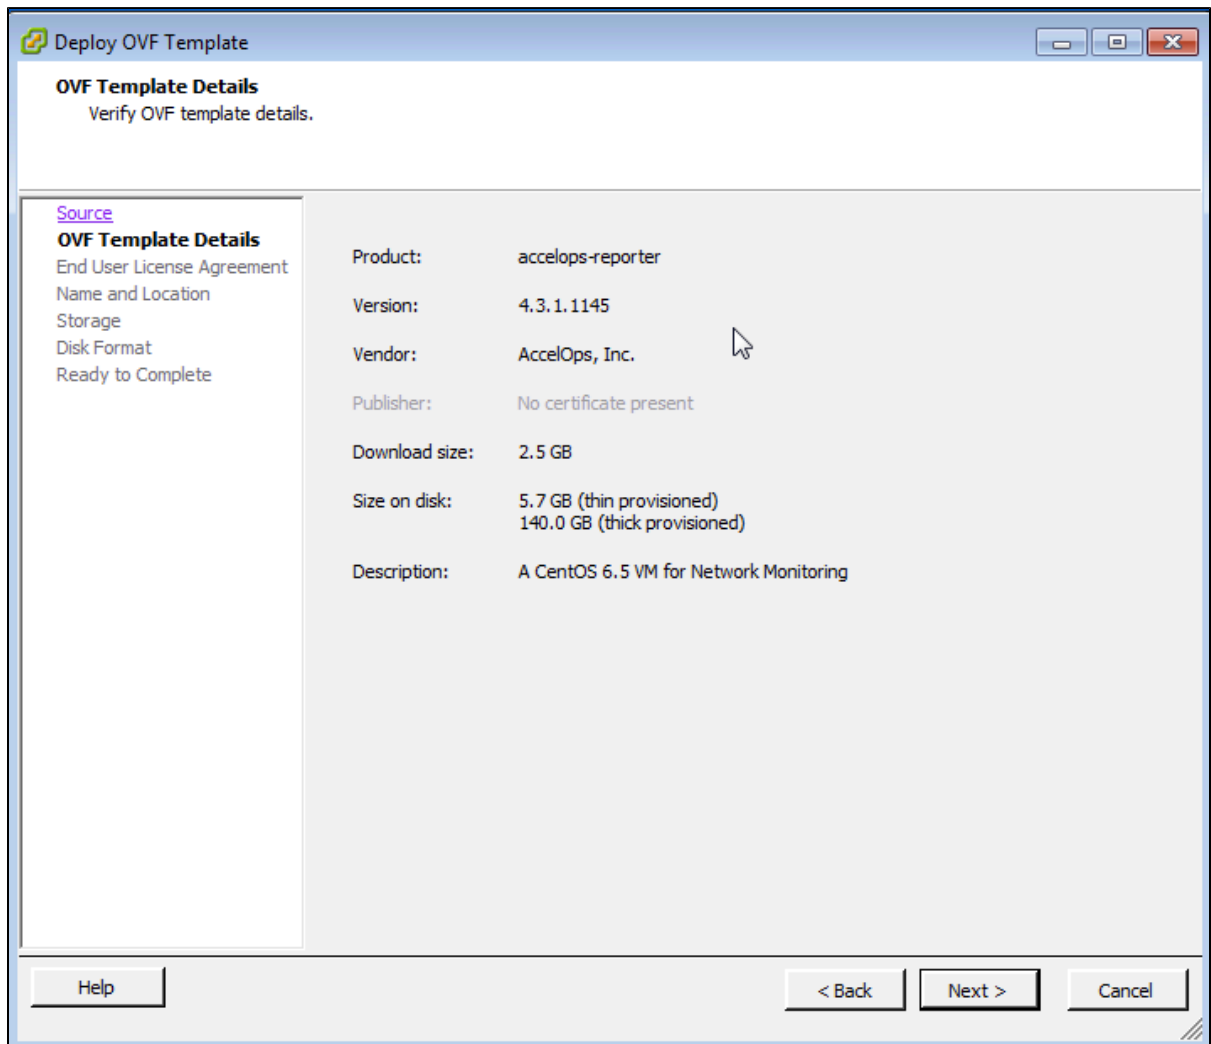




### Step 2: Review the Virtual Appliance details

OVF Template details will tell you how big the file is to import and how much disk space it will take up in your selected ESX server.

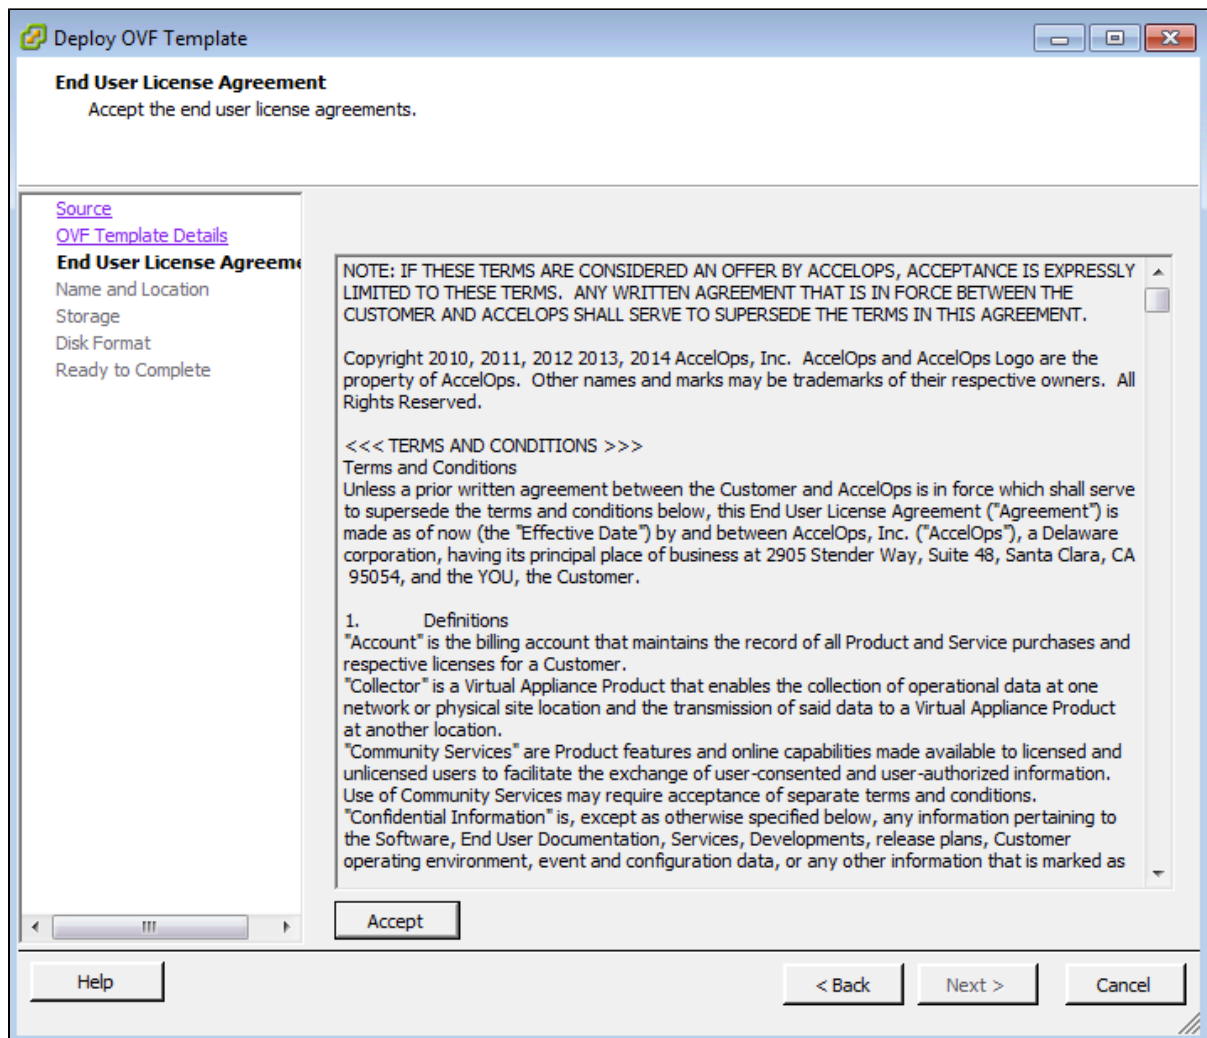
- 1. Click the 'Next' button



**Step 3: End User License Agreement**

To continue you must accept the End User License Agreement.

1. Click on 'Accept' button on License Agreement
2. Click on 'Next' button

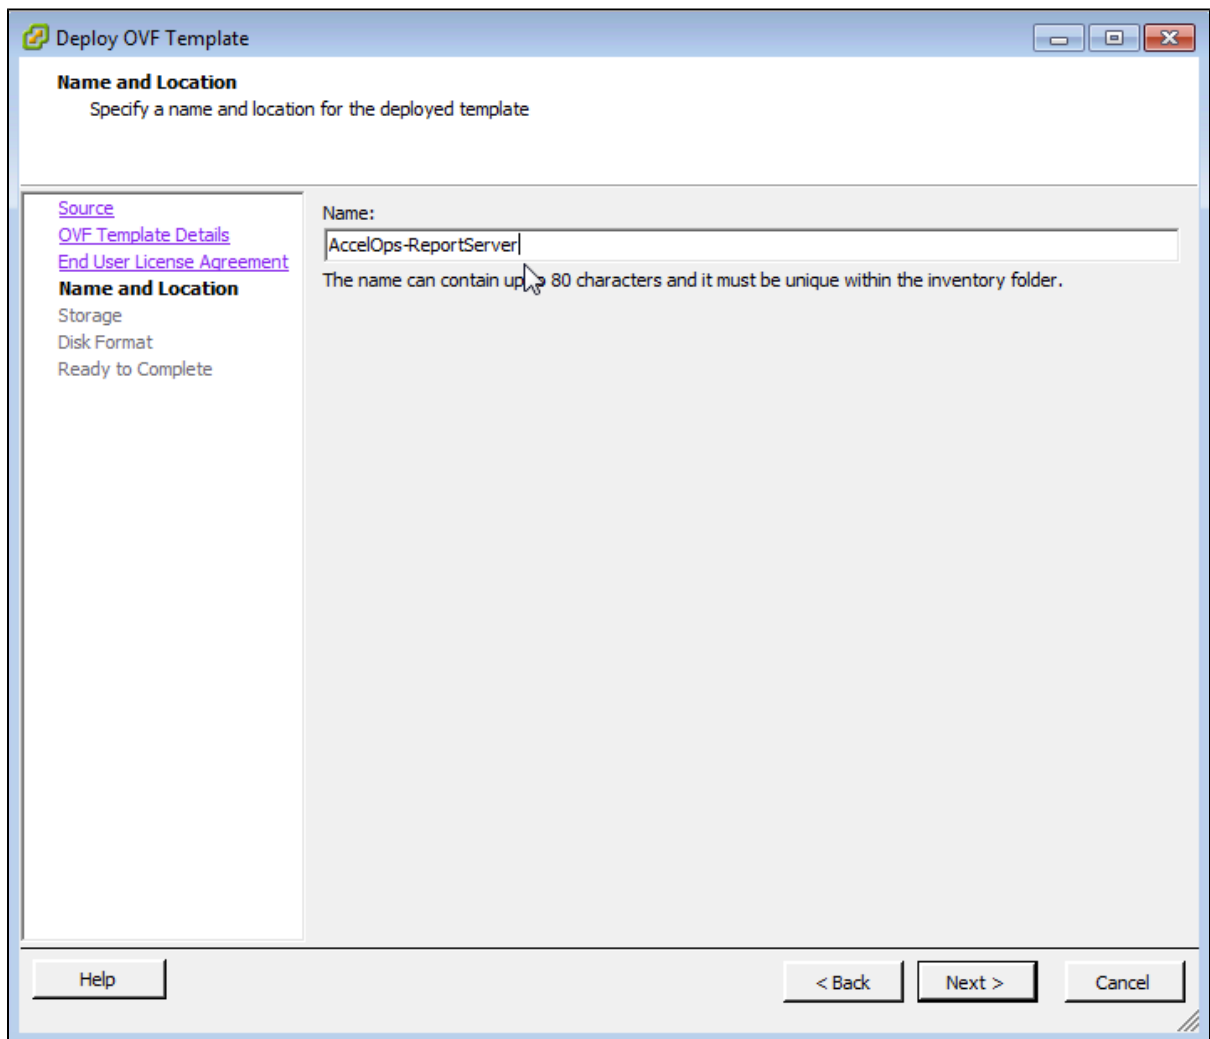


#### Step 4: Name and Location

Enter a name for the Virtual Appliance that is being imported. This name is independent of the host name and is only visible within the ESX environment.

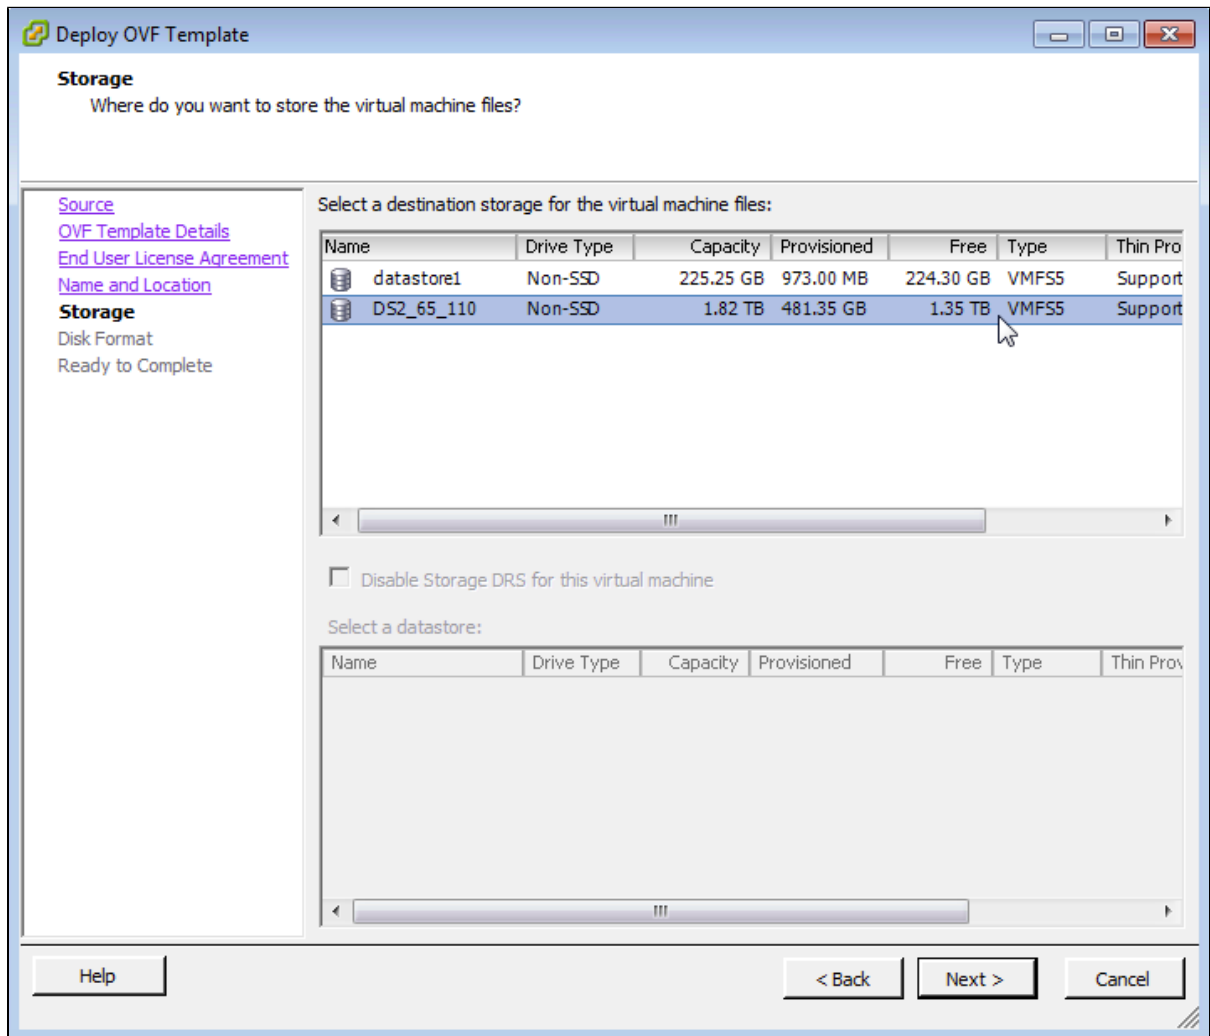
1. Enter a Report Server name and Click on 'Next' button





**Step 5: Storage**

1. Select the desired storage location(Data store) and Click on 'Next' button



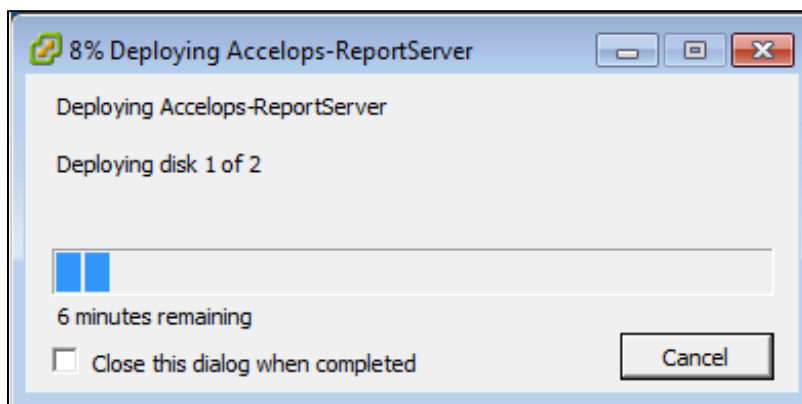
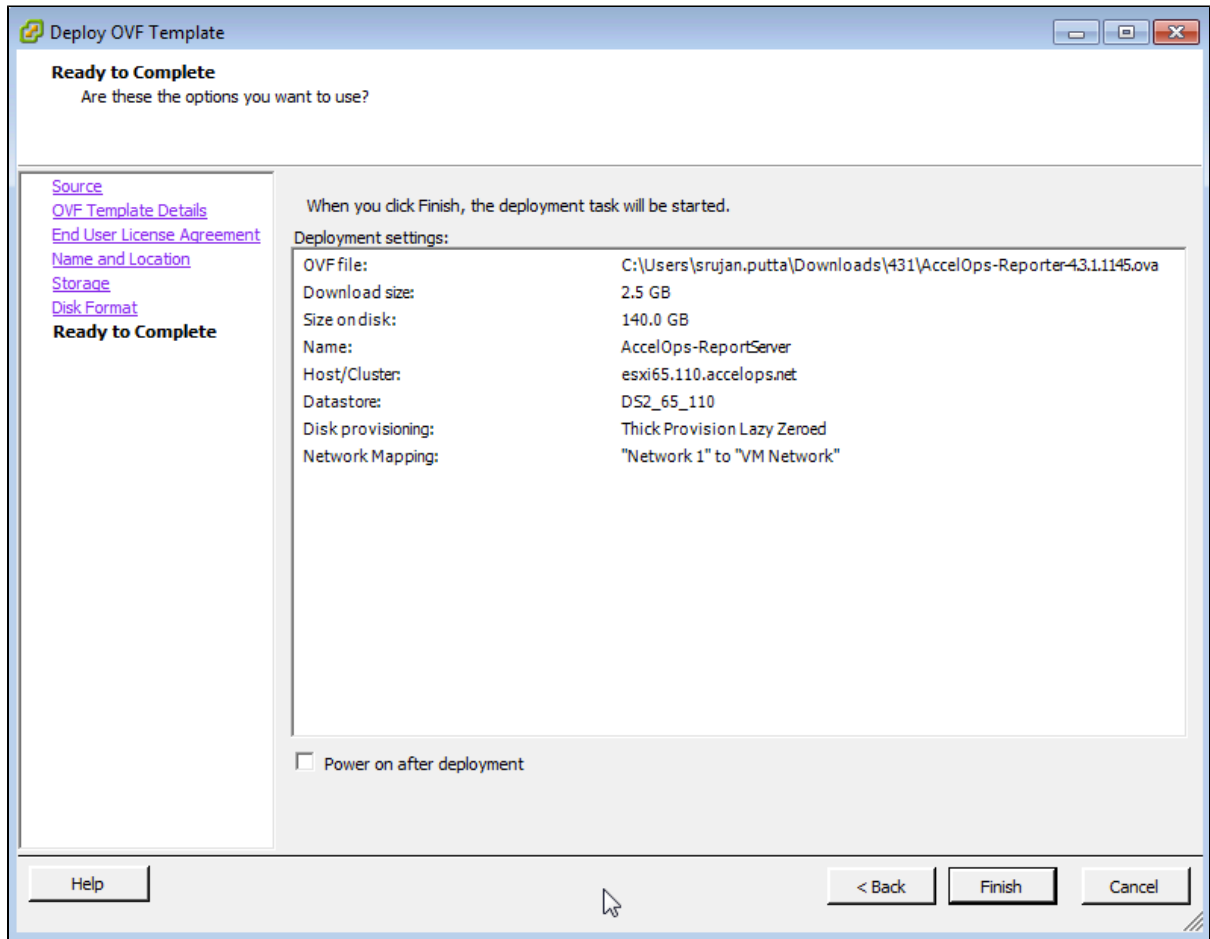
**Step 6: Disk Format**

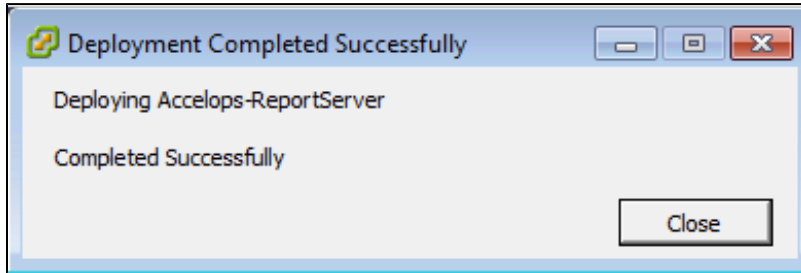
1. Select the Disk format as 'Thick Provision' (recommended) and Click on 'Next' button

The screenshot shows a window titled "Deploy OVF Template" with a "Disk Format" sub-header. Below the header is the question "In which format do you want to store the virtual disks?". On the left, a sidebar contains links for "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Storage", and "Disk Format" (which is highlighted and labeled "Ready to Complete"). The main area contains two input fields: "Datastore:" with the value "DS2\_65\_110" and "Available space (GB):" with the value "1381.4". Below these are three radio button options: "Thick Provision Lazy Zeroed" (selected), "Thick Provision Eager Zeroed", and "Thin Provision". At the bottom, there are three buttons: "Help", "< Back", and "Next >" (with a mouse cursor over it), and a "Cancel" button.

**Step 7: Ready to Complete**

1. Review the Report server resources and Click on 'Finish' button to begin importing Virtual Appliance
2. Virtual Appliance import will takes 7 to 10mins to complete. Do not turn off or reboot the system during this time.



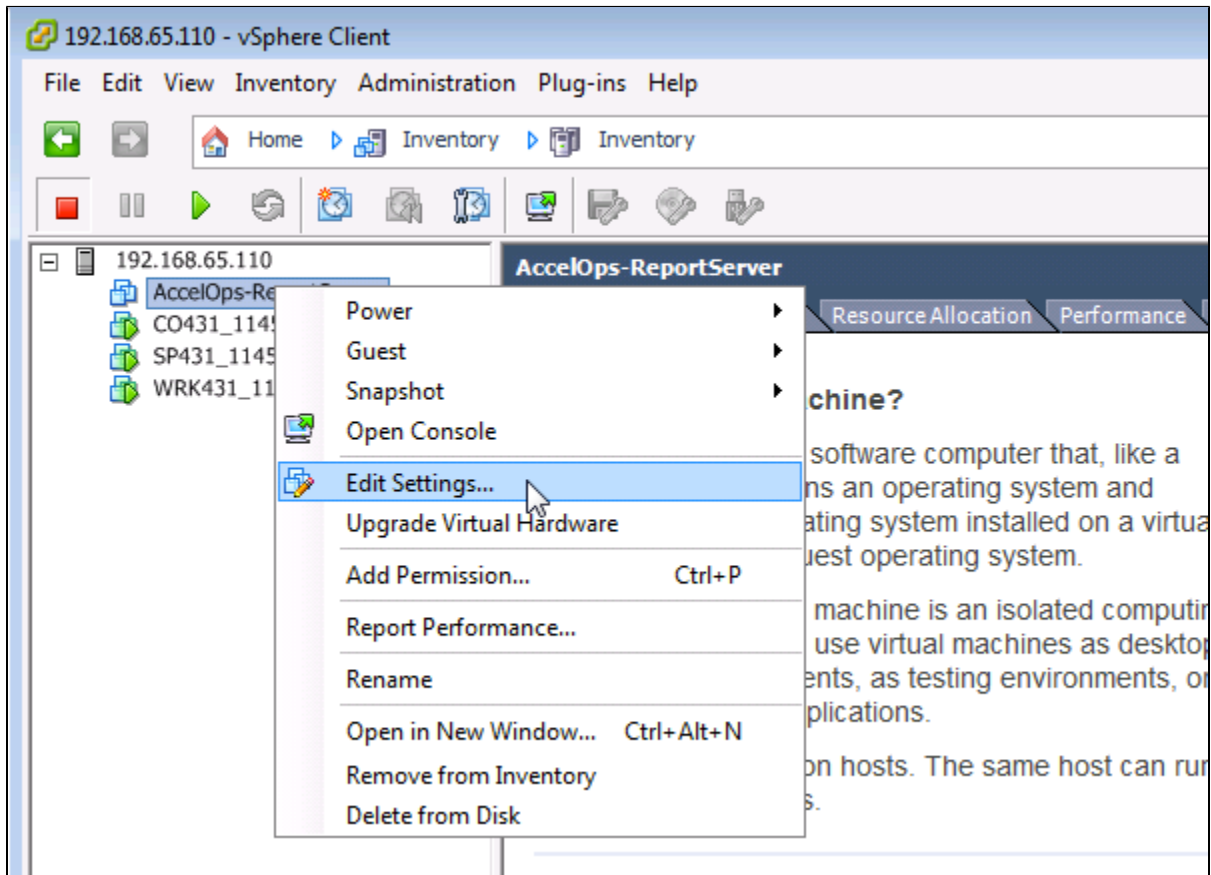


## Section 2: Edit the Report Server hardware settings

Prior to starting the Report Server for the first time it is required to make some modifications to the hardware setting for the Virtual Appliance.

### Step 1: Edit Settings

1. From the VMware vSphere client, Select the imported Report server Virtual appliance
2. Right mouse click and select Edit Settings



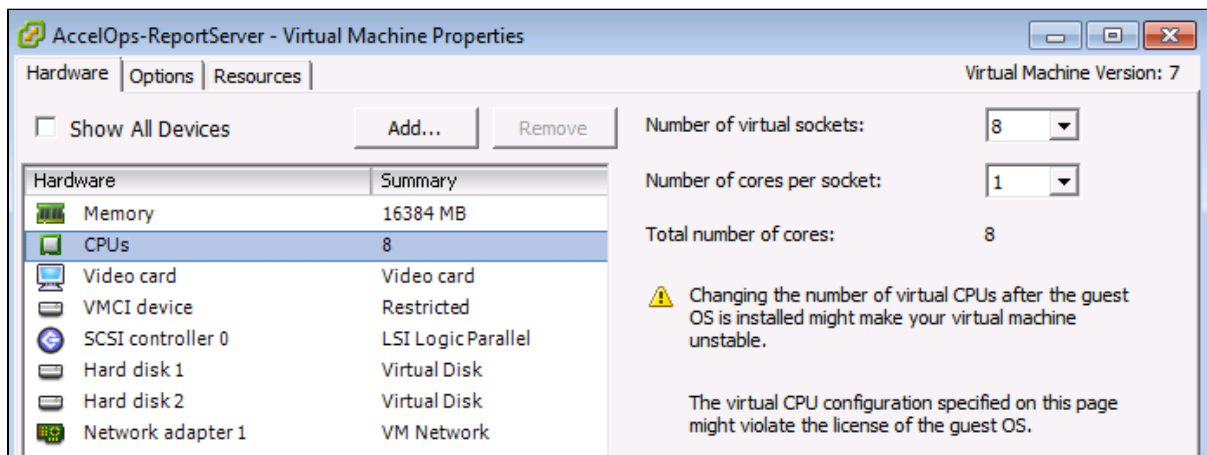
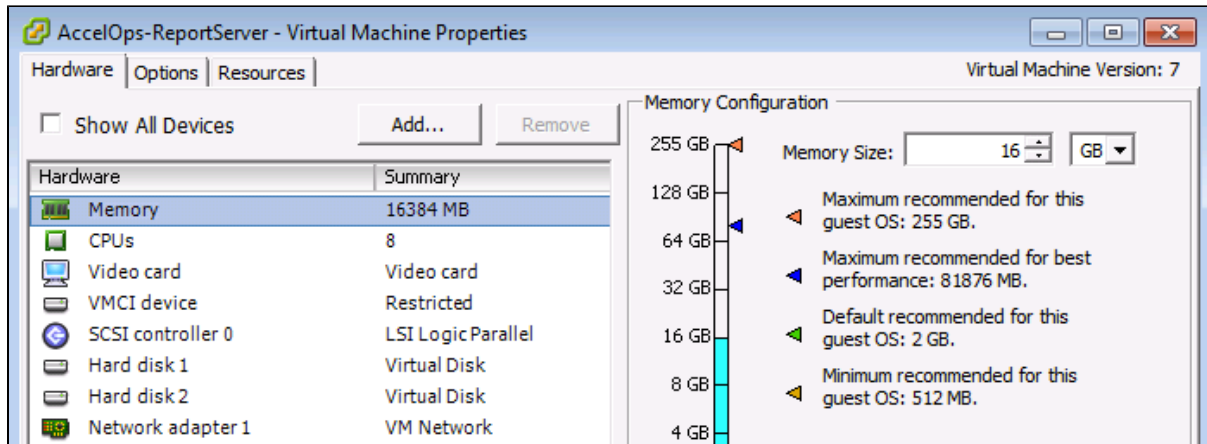
### Step 2: Verify Allocated Memory

1. On Hardware tab, Click on Memory and verify the memory to at least 16 GB

**Allocate more memory for large deployments**

For large deployments you should allocate at least 24GB of memory. See the section on Report Server Hardware requirements in the topic [Must Read Prior to Install](#) for more information

2. Click on CPUs, verify the CPU cores to at least 8 cores



**Section 3: Select mount point to Store Reports data**

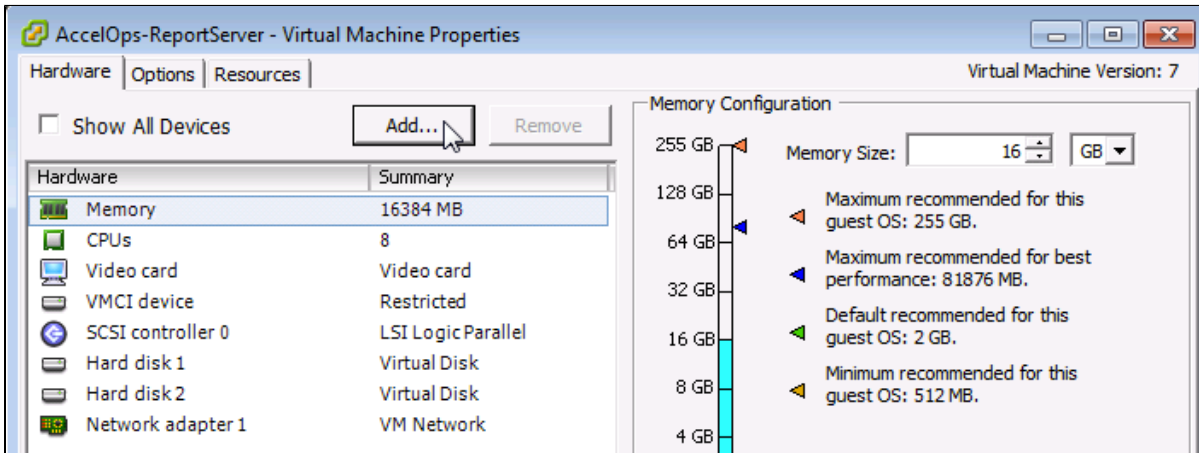
The AccelOps Report Server can be installed using either storage configured within the ESX server or NFS storage.

**Note**

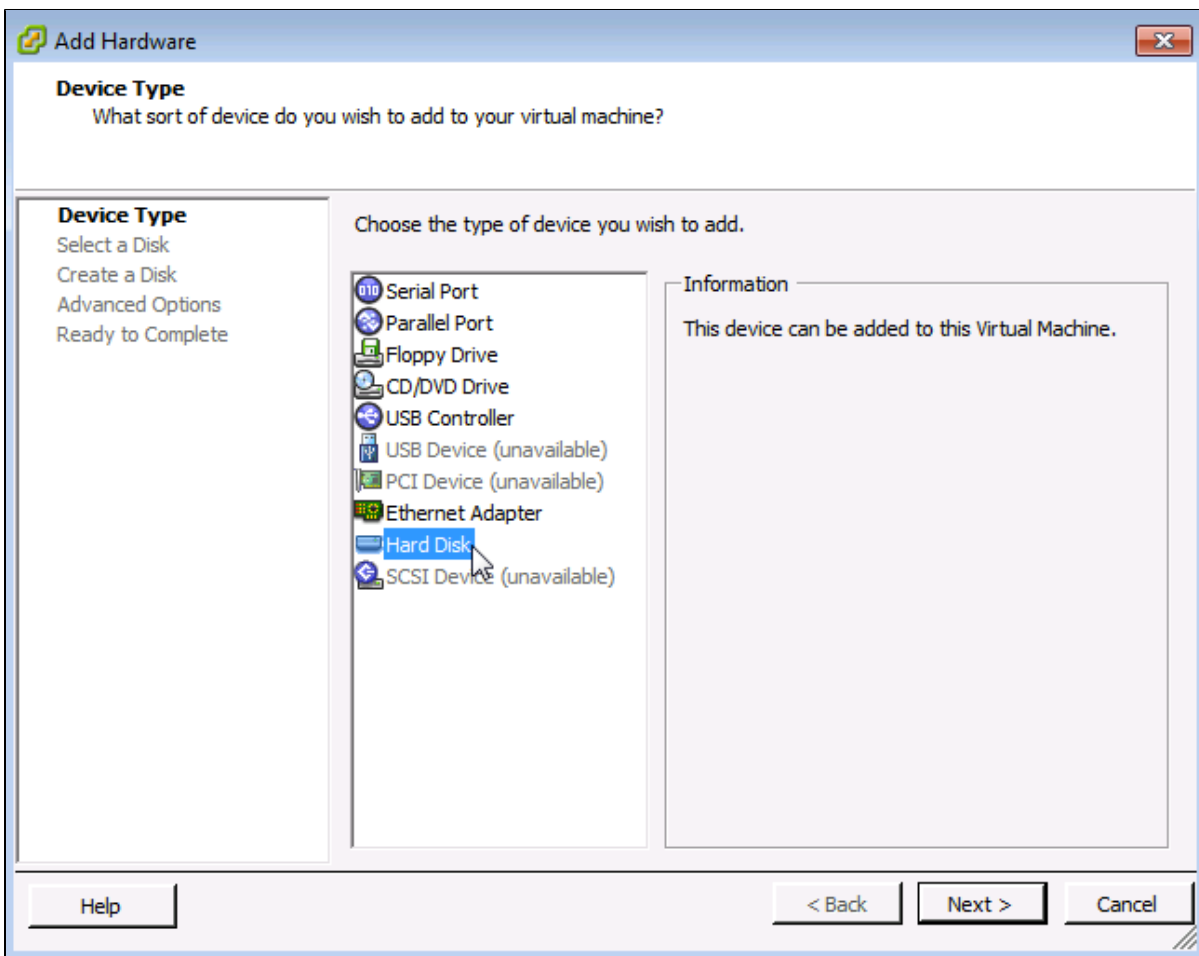
Skip Step 1 if AccelOps Super is used NFS storage to store the events data. Else, continue Step 1 to create local disk to store the reports data.

**Step 1: Adding an Additional Hard Disk to store the Reports data**

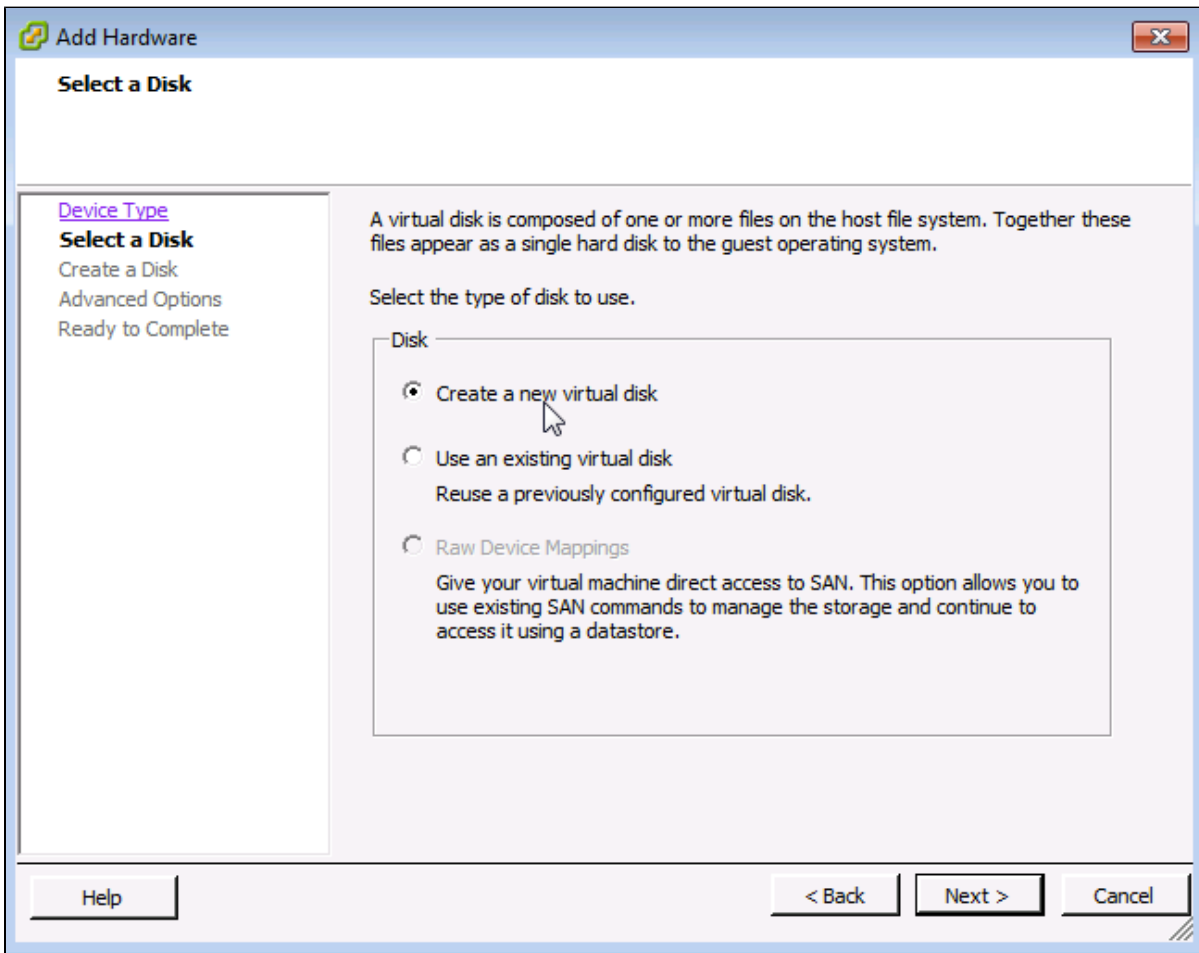
1. On Hardware tab, Click on the 'Add' button



2. On Hardware pop-up, Select 'Hard Disk' and click on 'Next' button



3. Select 'Create a new virtual disk' and click on 'Next' button



4. Enter the Disk size and click on 'Next' button  
Refer to Hardware Requirements for Report Server nodes



**Add Hardware**

**Create a Disk**  
Specify the virtual disk size and provisioning policy

[Device Type](#)  
[Select a Disk](#)  
**Create a Disk**  
Advanced Options  
Ready to Complete

**Capacity**  
Disk Size: 200 GB

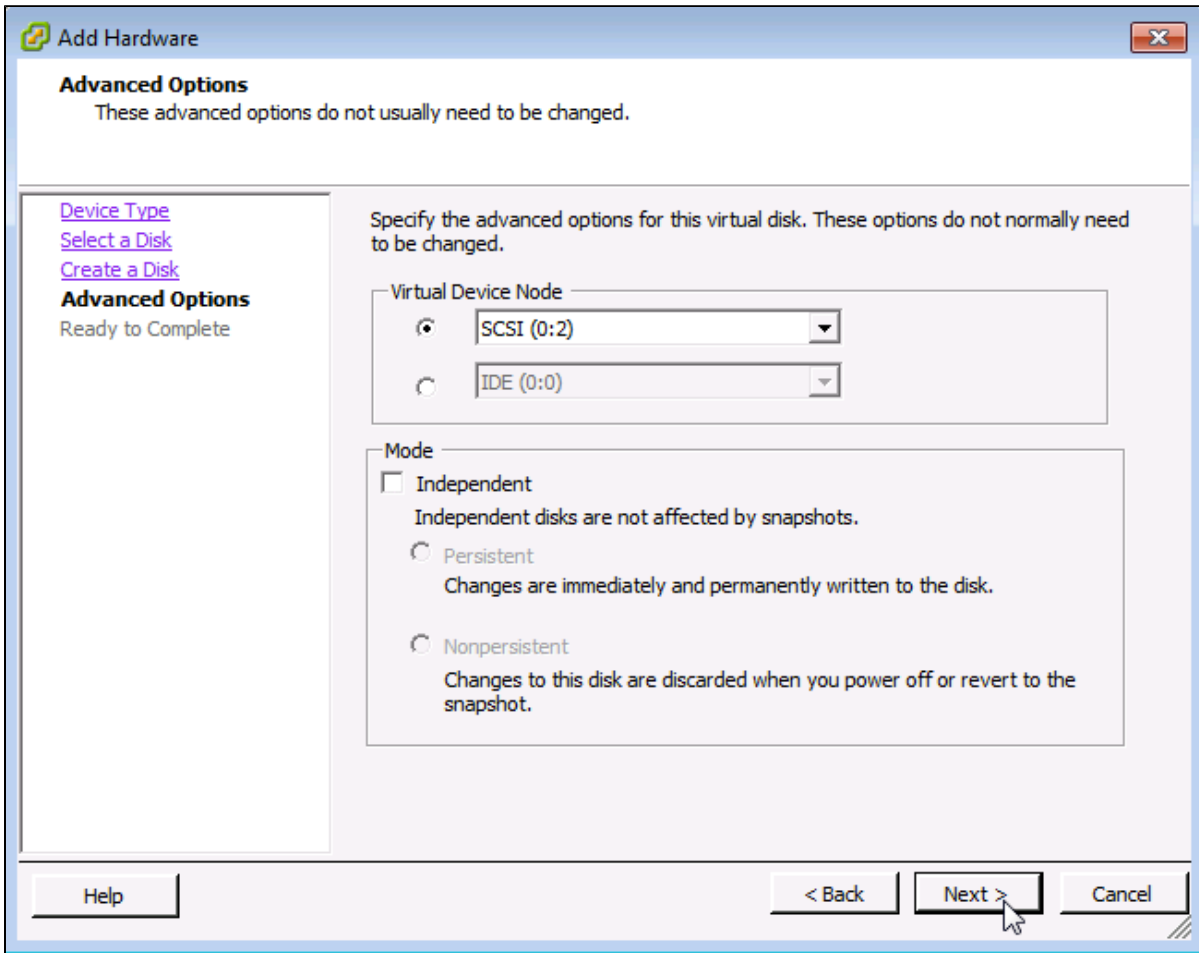
**Disk Provisioning**

- Thick Provision Lazy Zeroed
- Thick Provision Eager Zeroed
- Thin Provision

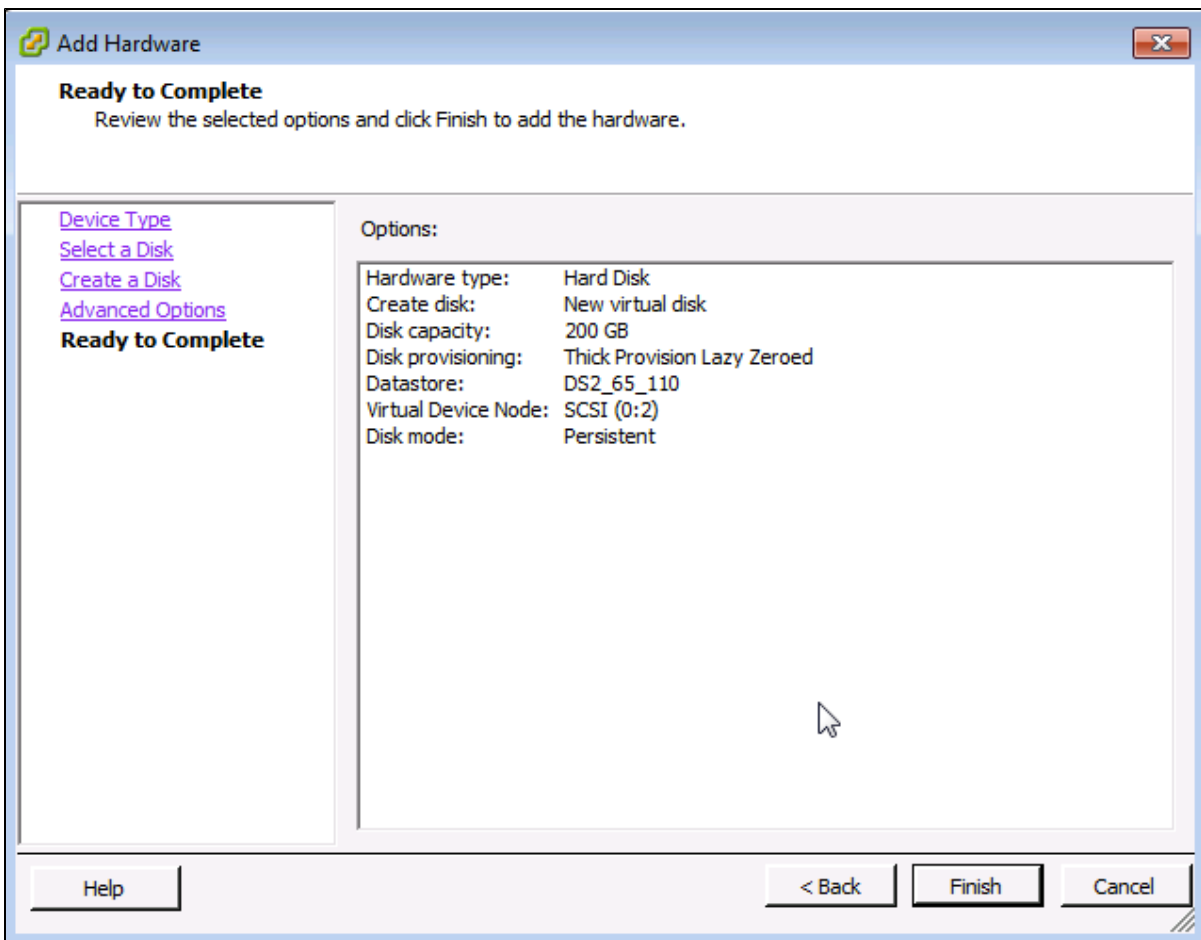
**Location**

- Store with the virtual machine
- Specify a datastore or datastore cluster:  
DS2\_65\_110

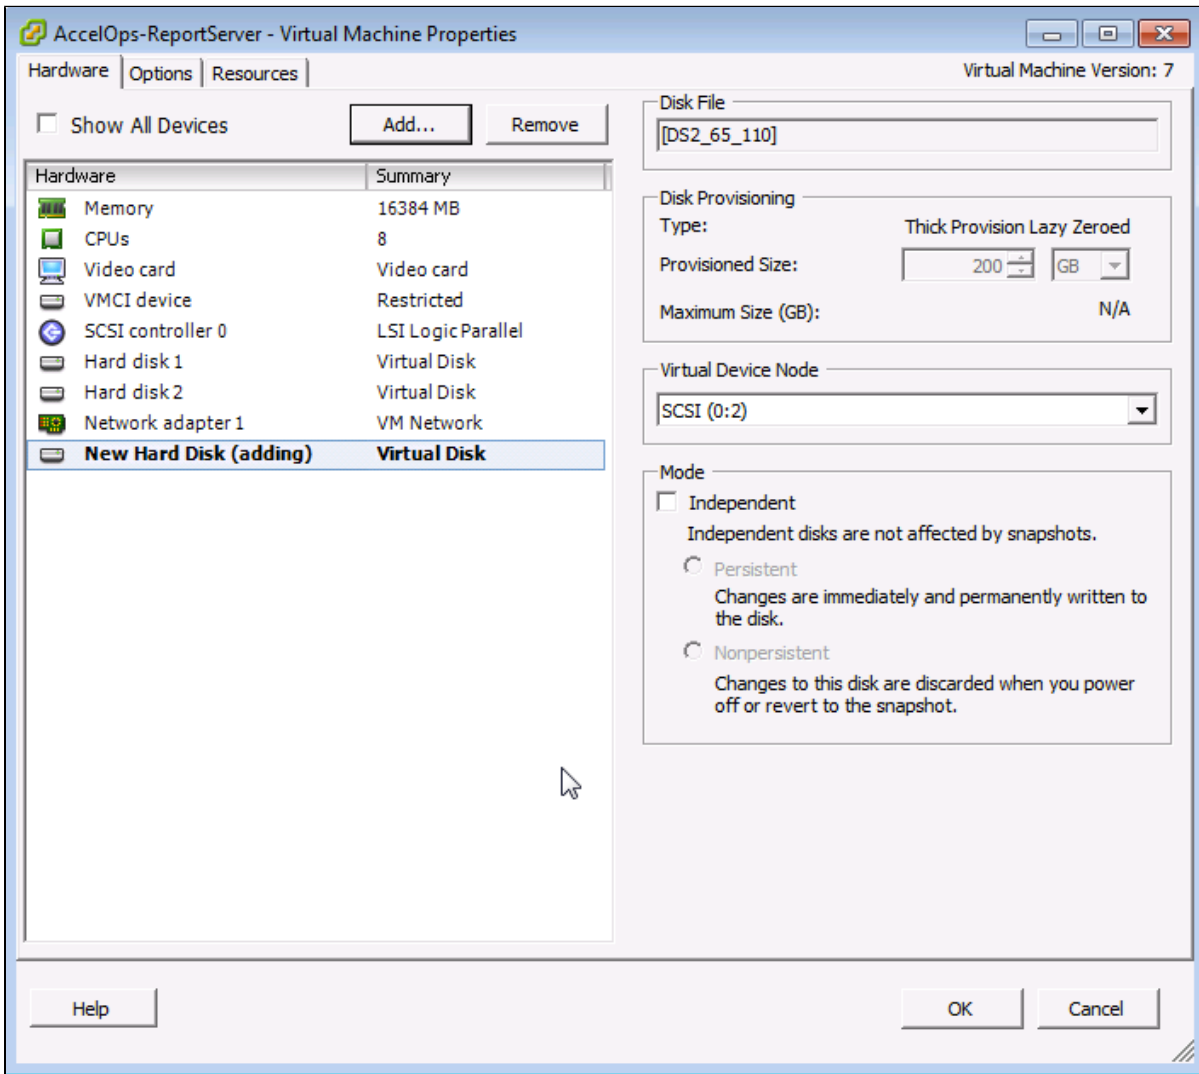
5. Ensure that 'Independent' is not checked and click on 'Next' button




6. Confirm the Options are correct and click the 'Finish' button



7. Click 'OK' on Virtual Machine properties pop-up



8. Wait for until Adding Disk task complete on vSphere client



**Note**

Skip Step 2 if you choose Step 1 to store the AccelOps event data in local Hard disk. Else, continue Step 2 to create NFS mount point to store the reports data.

**Step 2: Create new NFS mount point to store the Reports data**

1. Log into NFS server, Create new directory to store Reports data

**Note**  
Do not use AccelOps Supervisor mount point location.

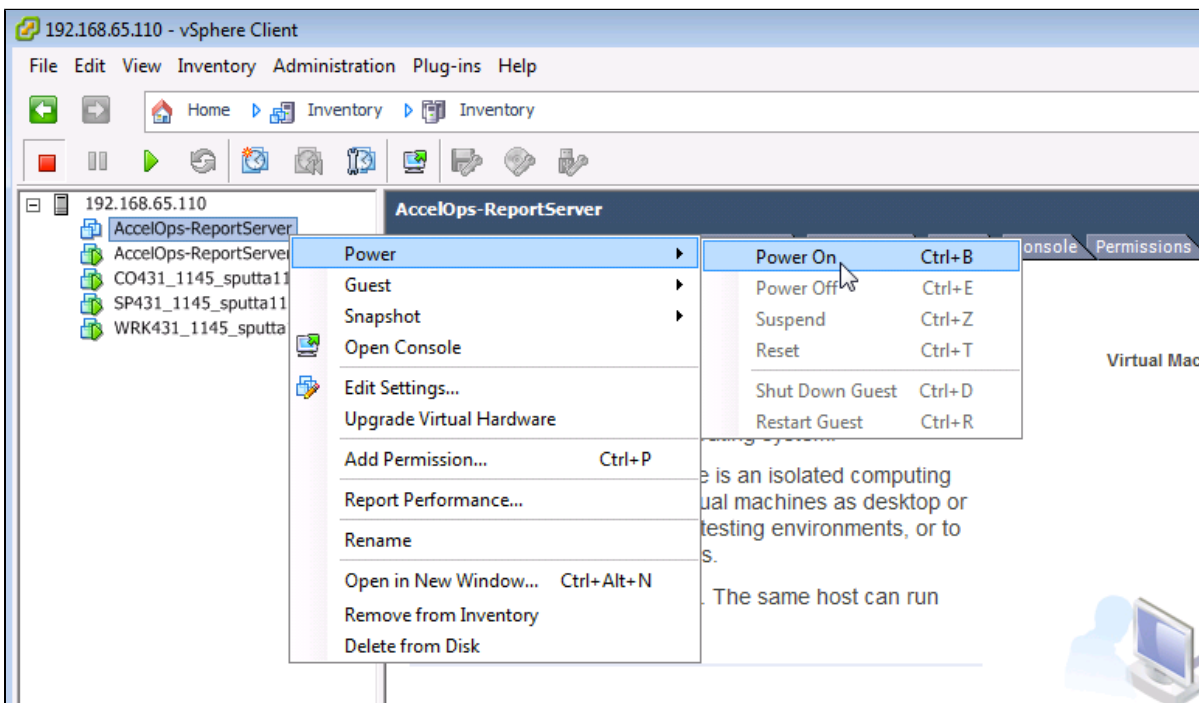
2. Refer to NFS documentation to check NFS directory permission

**Section 4: Configure the Report Server from the VMware Console**

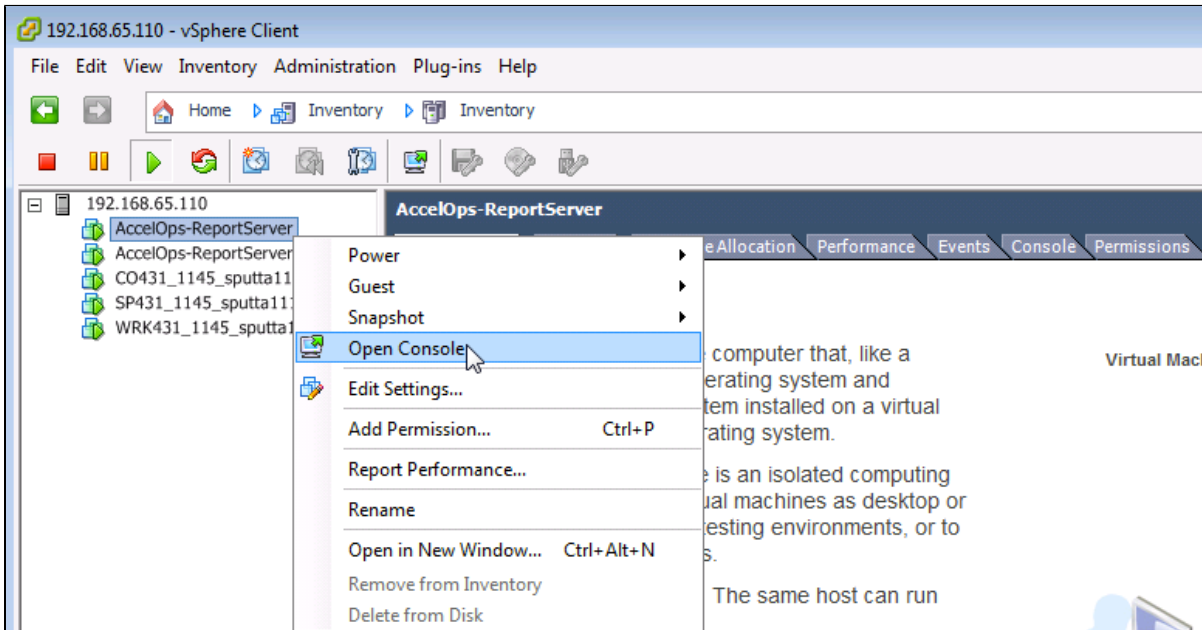
**Note**  
While in the VM console window do not press any control keys during the installation. Doing so might cause the installation to stop and not to precede any further. For example Ctrl – C or Ctrl – Z. If this occurs you must erase the VA and start over from the beginning.

**Step 1: Start Report Server**

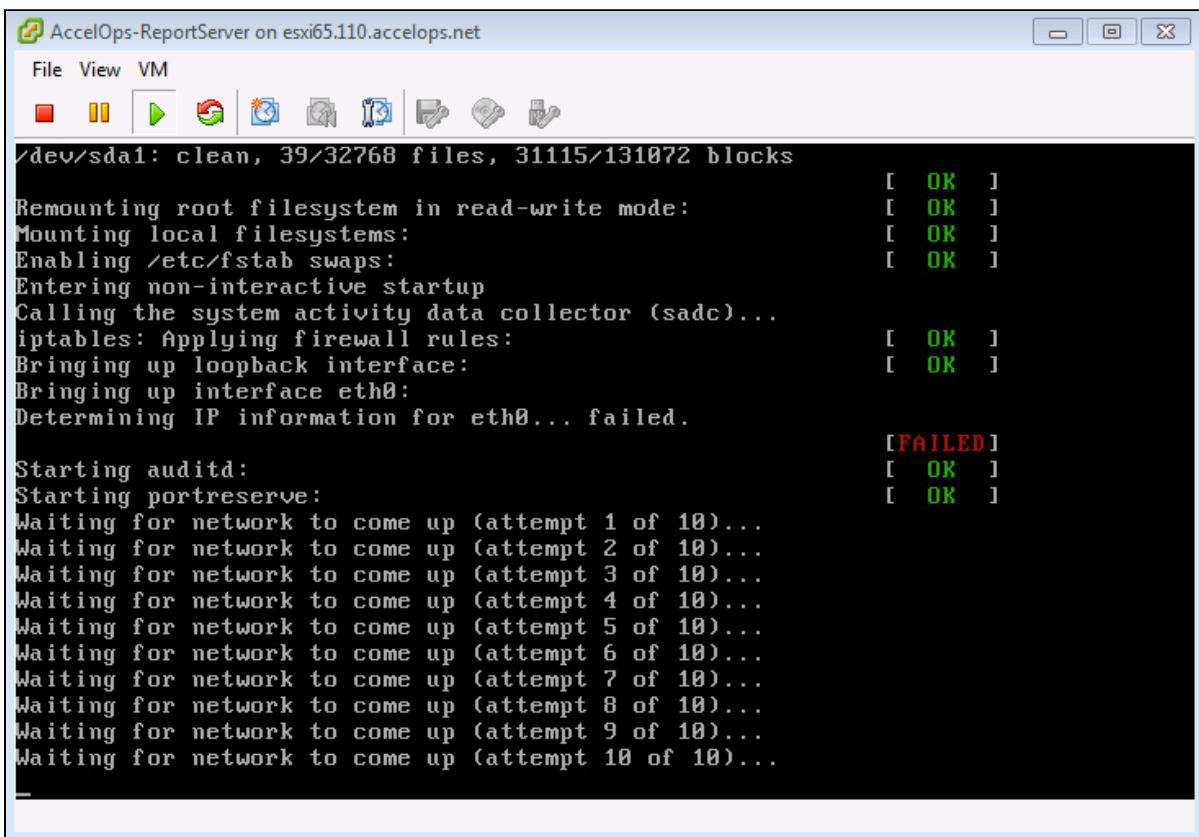
1. From the VMware vSphere client, Select recently imported Report server Virtual appliance and Power On (Right click > Power > Power On)



2. Open Report server VM Console (Right click > Open Console)



3. Network eth0 Failed message may appear on the screen during the loading, but this is to be expected the first time the AccelOps Report server is started.




```
AccelOps-ReportServer on esxi65.110.accelops.net
File View VM
/dev/sda1: clean, 39/32768 files, 31115/131072 blocks
Remounting root filesystem in read-write mode: [ OK ]
Mounting local filesystems: [ OK ]
Enabling /etc/fstab swaps: [ OK ]
Entering non-interactive startup
Calling the system activity data collector (sadc)...
iptables: Applying firewall rules: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0:
Determining IP information for eth0... failed. [ FAILED ]
Starting auditd: [ OK ]
Starting portreserve: [ OK ]
Waiting for network to come up (attempt 1 of 10)...
Waiting for network to come up (attempt 2 of 10)...
Waiting for network to come up (attempt 3 of 10)...
Waiting for network to come up (attempt 4 of 10)...
Waiting for network to come up (attempt 5 of 10)...
Waiting for network to come up (attempt 6 of 10)...
Waiting for network to come up (attempt 7 of 10)...
Waiting for network to come up (attempt 8 of 10)...
Waiting for network to come up (attempt 9 of 10)...
Waiting for network to come up (attempt 10 of 10)...
```

## Step 2: Configure the Timezone

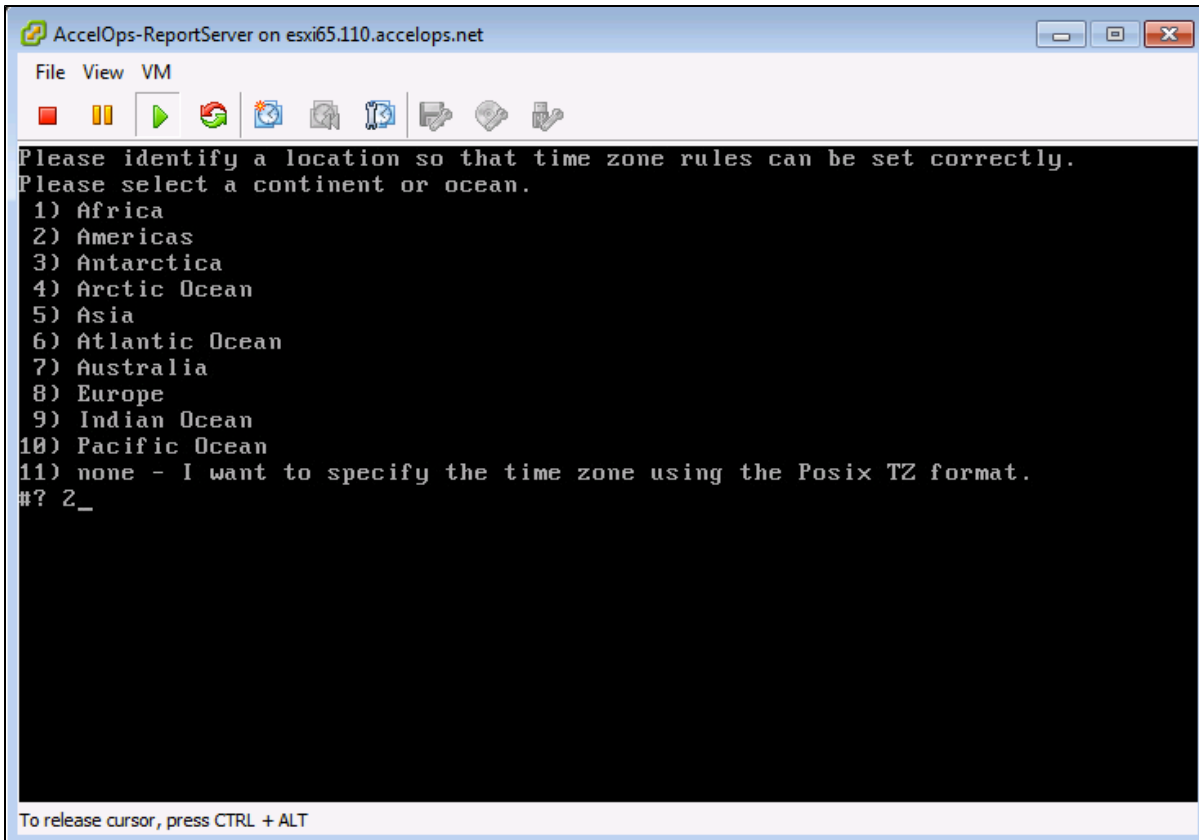
It is important to select the correct time zone that you are located in.

1. On VM console, Select 'Set Timezone' and press the 'Enter' key




2. Select your Location and press the 'Enter' key

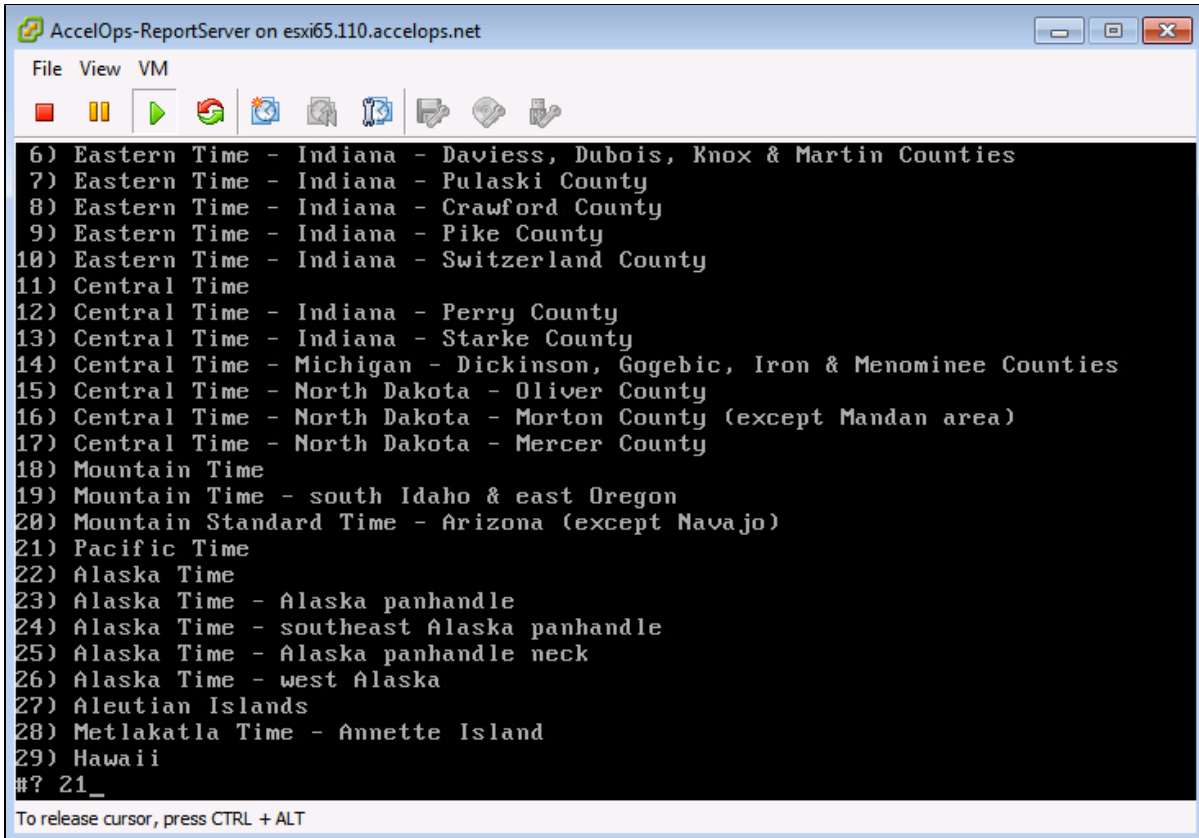




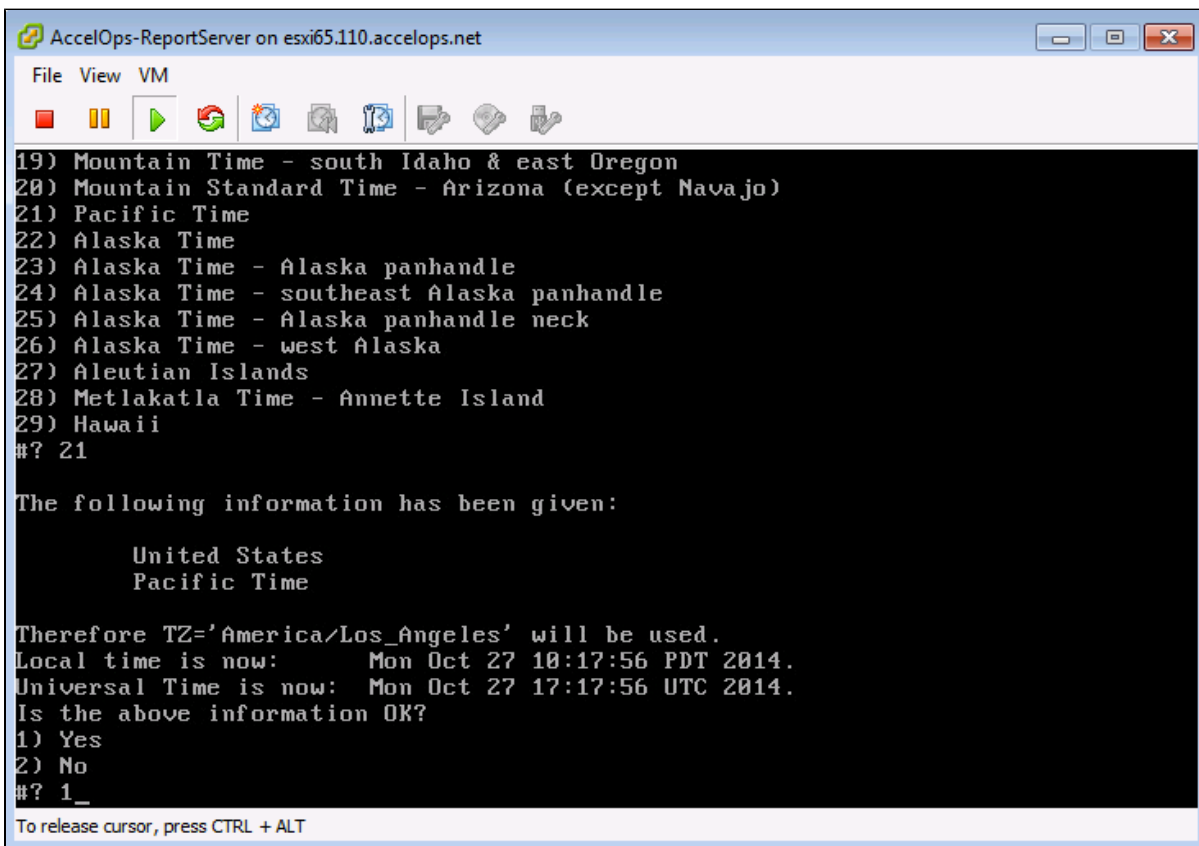
3. Select your Country and press the 'Enter' key



4. Select your Timezone and press the 'Enter' key




5. Review and Confirm your Timezone selection




**Step 3: Configure the Network**

1. On VM console, Select 'Login' and press the 'Enter' key



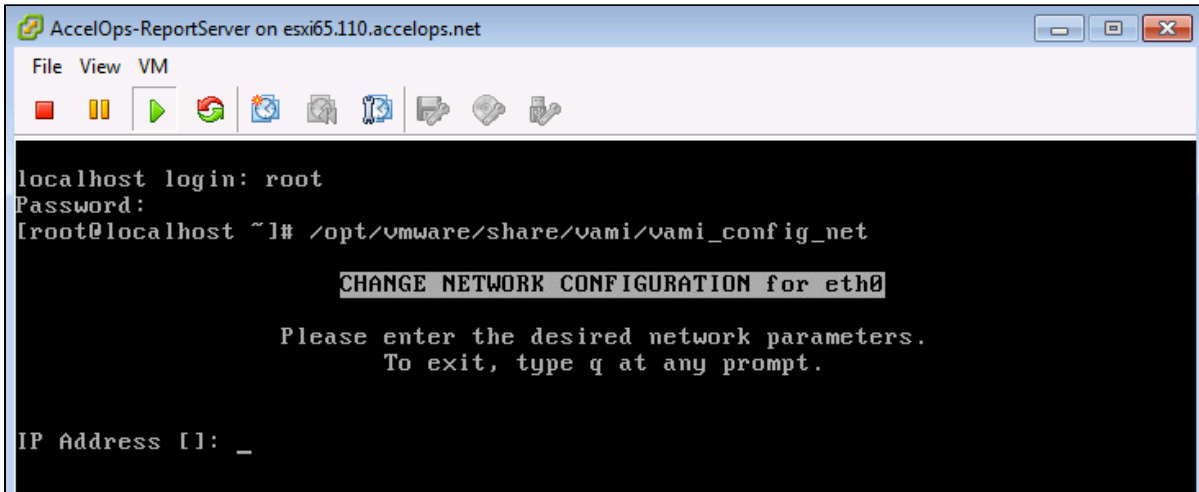
2. Enter default log in information.  
Login: root  
Password: ProspectHills



3. Run vami\_config\_net script to configure network and installing AccelOps Report server

```

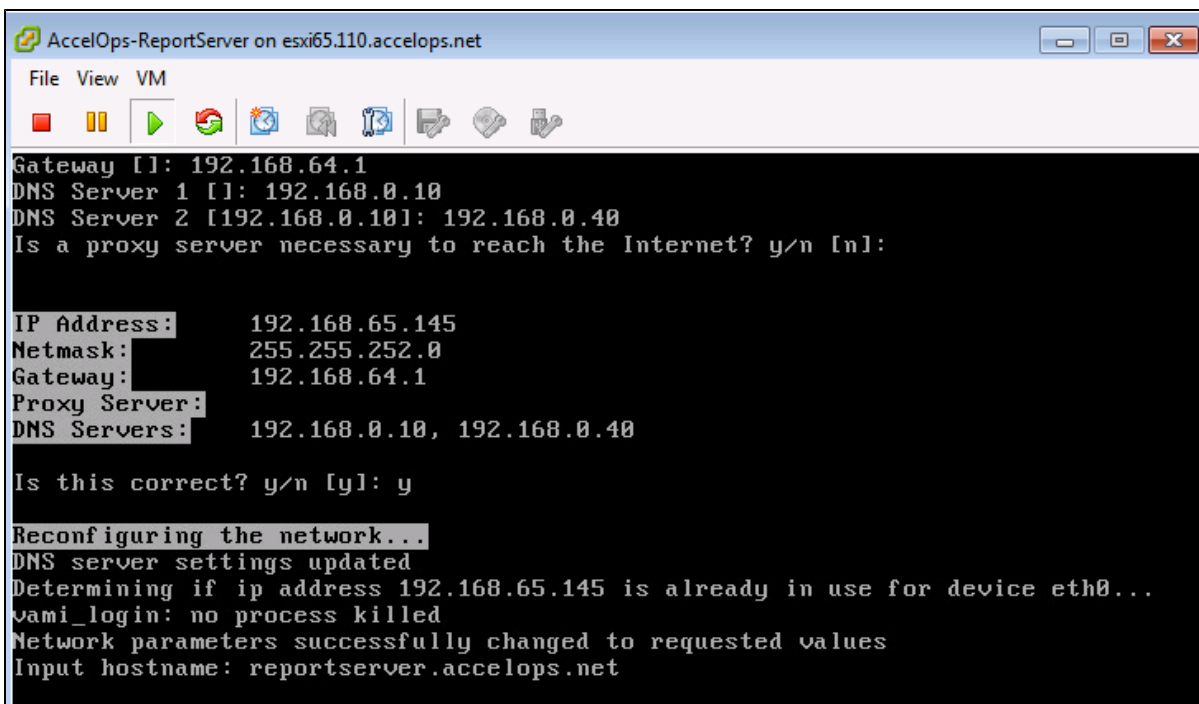
Configure the Network
/opt/vmware/share/vami/vami_config_net
    
```



4. Configure Network with Static IP address
  - a. The following information is required to configure Static IP address:
    - i. IP Address
    - ii. Netmask
    - iii. Gateway
    - iv. DNS Server(s)

**Proxy Server**  
Do not configure the proxy server setting.

- b. Press the 'y' key to accept the changes
- c. Enter 'Host name' and press 'Enter' key



Step 4: Select Mount Point to store your data

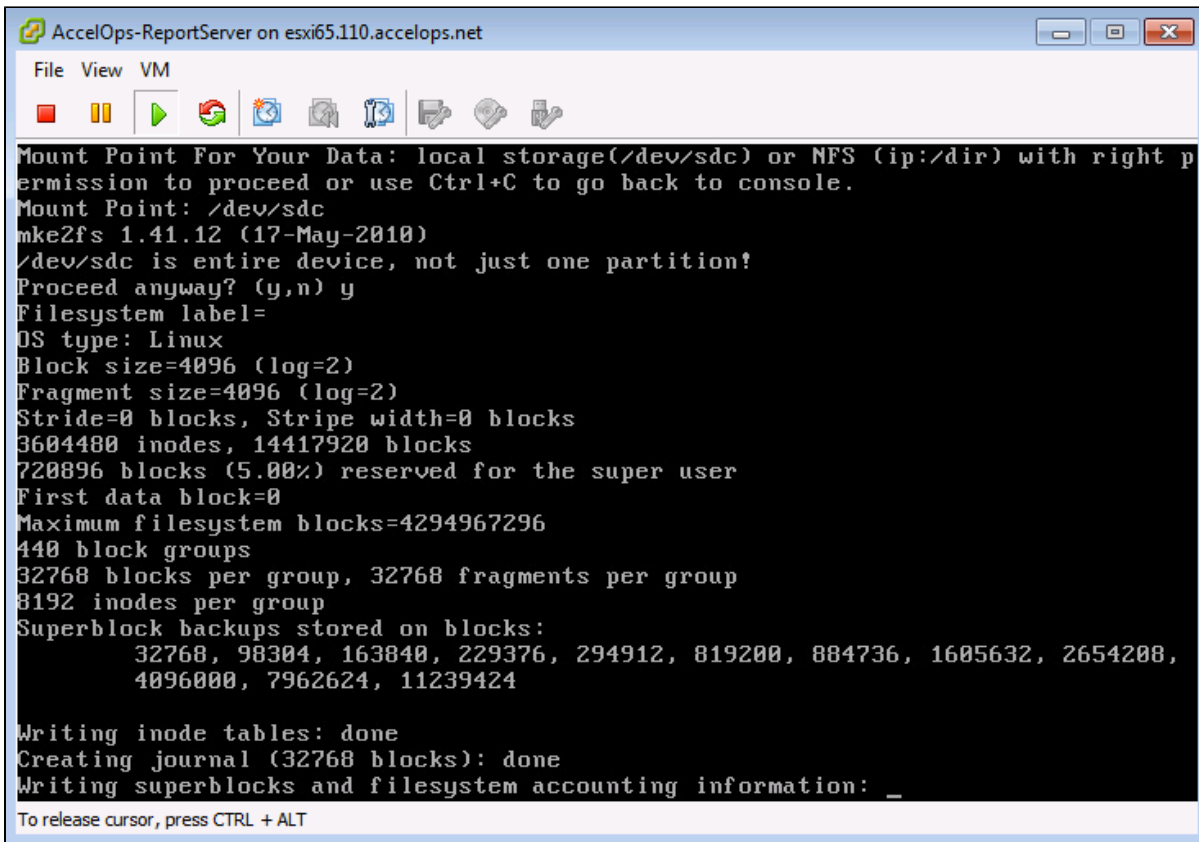
The AccelOps Report Server can be installed using either storage configured within the ESX server or NFS storage.

Choose only following Step 4.1 or 4.2 based on [Section 3](#)

1. Enter your Local hard disk mount point

**Note**  
Skip Step 4.1 if AccelOps Super is used NFS storage to store the events data. Else, continue below steps.

**Local Storage Mount point**  
/dev/sdc



2. Enter your NFS mount point

**Note**  
Skip Step 4.2 if you choose Step 4.1 to store Reports data in local Hard disk. Else, continue below steps.

**NFS Mount point**  
<NFS\_Server\_IP\_Address> : /<Directory\_Path>

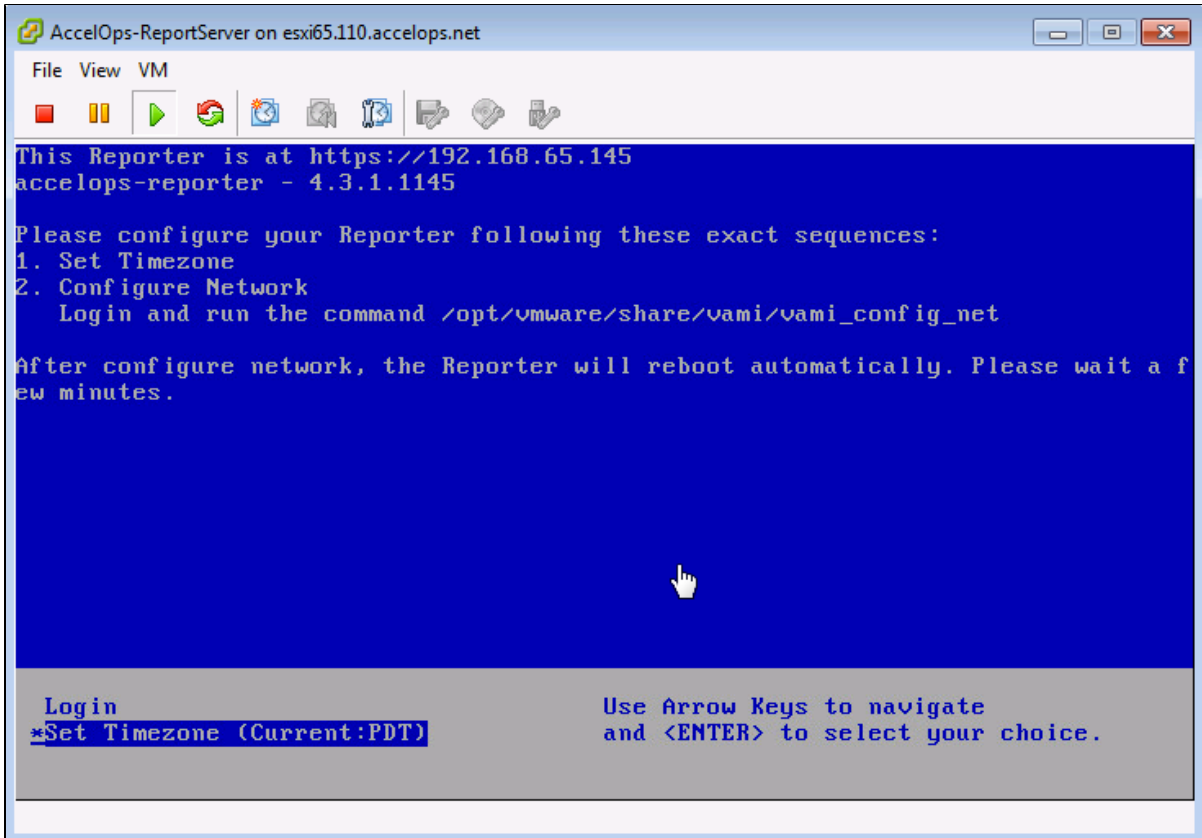
```

Mount Point For Your Data: local storage(/dev/sdc) or NFS (ip:/dir) with right p
ermission to proceed or use Ctrl+C to go back to console.
Mount Point: 192.168.67.168:/data/RS/431
192.168.67.168:/data/RS/431 mounted
Running initialization of Reporter host

```

To release cursor, press CTRL + ALT

3. After mount point, Report server will reboot automatically  
In 5 to 10mins, Report server will be configure successful.



```

AccelOps-ReportServer on esxi65.110.accelops.net
File View VM
This Reporter is at https://192.168.65.145
accelops-reporter - 4.3.1.1145

Please configure your Reporter following these exact sequences:
1. Set Timezone
2. Configure Network
   Login and run the command /opt/vmware/share/vami/vami_config_net

After configure network, the Reporter will reboot automatically. Please wait a f
ew minutes.

Login
*Set Timezone (Current:PDT)
Use Arrow Keys to navigate
and <ENTER> to select your choice.

```

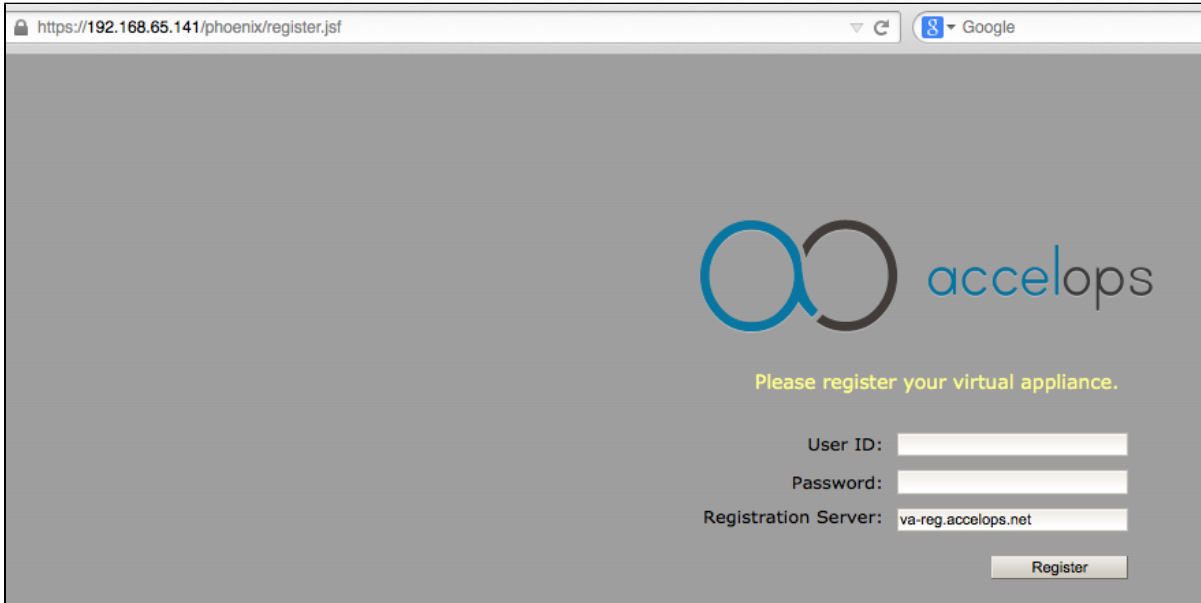
## Section 5: Register the Report Server to Supervisor

### Note

A new license is required to enable Report Server. Contact AccelOps support for Report Server License. You can skip this step if a Report Server license is already registered with the Supervisor.

### Step 1: Re-register AccelOps Supervisor with new license

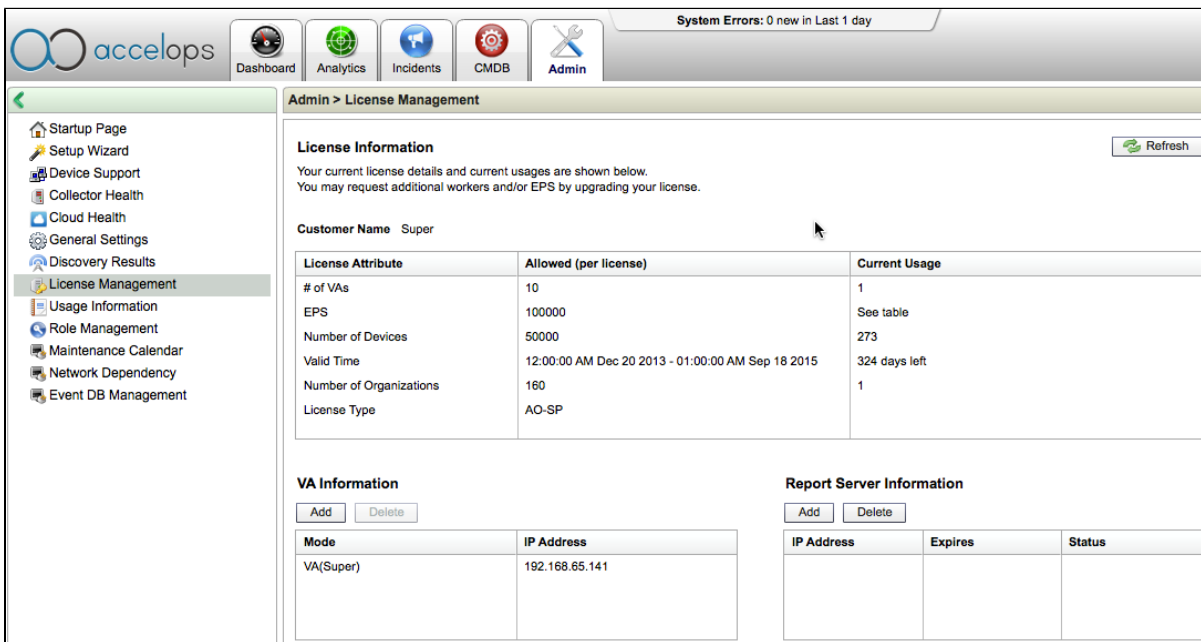
1. In browser windows, enter 'https://<Supervisor IP>/phoenix/register.jsf'



2. Enter license user name and password

**Step 2: Add Report Server on Supervisor UI**

1. Go to Admin > License Management. You will see a section called 'Report Server Information'



2. Click 'Add' to register the Report Server Node
  - a. Enter Report Server IP Address
  - b. Enter Read only credentials that will be used by the AccelOps Visualization Analytics Server to read the data in the PostgreSQL tables in the Report Server

**Add Report Server**

Report Server IP Address:

Database Username:

Database Password:

Confirm Password:

**Report Server Information**

IP Address	Expires	Status

- Report Server registration will takes 3mins to complete when Supervisor's CMDB size is <1 GB. Else, it takes more time to complete registration based Supervisor's CMDB size. You can 'Run in Background' when Supervisor's CMDB size is > 1GB.

Adding Report Server: 192.168.65.145

Run in Background

**Report Server Information**


IP Address	Expires	Status

- Click 'OK' on Confirmation pop-up. The Supervisor node will configure the Report Server

<b>Report Server Information</b>		
<input type="button" value="Add"/> <input type="button" value="Delete"/>		
IP Address	Expires	Status
192.168.65.145	09/05/2015	Active



5. Make sure Report Server appears in Admin > Cloud health

Admin > Cloud Health										
 Refresh		Last Updated: 17:38:55								
Name	IP Address	Health	Module Role	Version	Build Date	Cores	Load Averag	CPU	Swap Size	Swap I
VA141	192.168.65.141	Normal	supervisor	4.3.1.1145	10:04:55 AM Oct 15 20...	8	0.26,0.27,0.3	0.88	6 GB	57.33 f
reportserver.accelops.net	192.168.65.145	Normal	report server	4.3.1.1145	10:04:55 AM Oct 15 20...	3	0.02,0.03,0	2.33	6 GB	0 KB
reportserver.accelops.net ( 192.168.65.145 ) - report server: Process Details										
Process Name	Status	Up Time	CPU	Event Rate	Physical Mem	Virtual Mem	SharedStore ID			
phMonitorReportServe	Up	2h 32m 48s	0%	0/s	24 MB	909 MB	0			
rsyslogd	Up	2h 55m	0%	0/s	1 MB	243 MB	0			
Report DB	Up	2h 34m 1s	0%	0/s	17 MB	446 MB	0			
pgsql DB	Up	2h 38m 19s	0%	0/s	17 MB	446 MB	0			

## Installing and Configuring Visual Analytics Server

- [Prerequisites](#)
- [Installation](#)
- [Activation](#)
- [Configuration](#)

### Prerequisites

Before you begin installing Visual Analytics Server, make sure you have read the section on **Visual Analytics Server** in [Requirements for AccelOps Visual Analytics](#). This contains information on the Administrator Account and Ports that you will need during the configuration process. The Visual Analytics Server is powered by Tableau Server, and you may want to also consult the [Tableau Server Administration Guide](#) before you begin the installation process.

### Installation

1. Download the Visual Analytics Server installation file [from Tableau Software](#).
2. Double-click the installation file to launch the Visual Analytics Server Setup Wizard.
3. When the Setup Wizard launches, click **Next** to begin the installation process.
4. Enter a **Destination Location** where you want to install the server files, and then click **Next**.
5. When the system verification process completes, click **Next**.
6. Enter a location for the Start Menu folder, or use the default location, and then click **Next**.
7. Click **Install** to complete the installation process.
8. Click **Next** to begin the server activation process.

### Activation

1. If you are evaluating Visual Analytics Server, click **Start trial now**. Otherwise, click **Activate the product** to enter a license key.
2. If you enter a license key, click **Activate**.
3. Click **Continue** to launch the Visual Analytics Server configuration process.

### Configuration

1. In the Configuration dialog, enter a **User Name** and **Password** for the domain admin account that you will use to administer the Visual Analytics Server.
2. If necessary, enter a **Gateway** port through which you will connect to the server over HTTP.
3. Click **OK**.  
The initialization process will launch and complete within several minutes.
4. Click **Finish** to complete the configuration process.
5. Launch the Visual Analytics Server user interface by entering the URI for the server in a browser window.  
The URI will be in the format of `http://<Windows_Server_IP_Address>:<Port_Number_Used_In_Step_2>`
6. Sign in to the server by entering the credentials for the domain admin account that you created in Step 1, and then click **Sign In**.
7. Click the **Admin** tab and select **Maintenance**.
8. Under **Status**, check to make sure that all systems are up and running.

You are now ready to install Visual Analytics Desktop. After you have completed the Visual Analytics Desktop installation process and connect to Report Server for the first time to create a sheet, as described in [Creating a Single Sheet Workbook](#), you will also establish the connection between AccelOps Report Server and Visual Analytics Server.

## Installing and Configuring Visual Analytics Desktop

- [Prerequisites](#)
- [Installation and Activation](#)
- [Connecting Visual Analytics Desktop to Visual Analytics Server](#)

### Prerequisites

1. Make sure you have read the section on Visual Analytics Desktop in [Requirements for AccelOps Visual Analytics](#).
2. You should set up the Visual Analytics Server before you install the Visual Analytics Desktop, as you will connect Visual Analytics Desktop and Visual Analytics Server using the Domain Admin account and IP address that you used to configure Visual Analytics Server.
3. Download and install the [Postgres SQL data source driver](#) from Tableau Software and install it on your desktop.

### Installation and Activation

1. Download the Visual Analytics Desktop installation file [from Tableau Software](#).  
The downloader will automatically detect the operating system of the device you are using to download, so you should be sure to initiate the download from the machine where you want to install Analytics Desktop.
2. Double-click the installation file to launch the Setup Wizard.
3. In the Setup Wizard, click **Install**.
4. In the **User Account Control** dialog, click **Yes**.  
The installation will complete within a few minutes, and the Activation dialog will launch.
5. Click **Start Free Trial** to begin evaluation of Analytics Desktop, or click **Activate the Product** if you have a license key.
6. Click **Continue** to launch Visual Analytics Desktop.

### Connecting Visual Analytics Desktop to Visual Analytics Server

1. Launch Visual Analytics Desktop.
2. In the **Server** menu, select **Sign In**.
3. Enter the IP address and HTTP port for the Visual Analytics Server.
4. Click **Connect**.
5. Enter the user name and password for the Admin user you created during the [Visual Analytics Server set up process](#), and then click **Sign In**.  
Your Visual Analytics Desktop application will now be connected to the Visual Analytics Server.

## Creating and Managing Workbooks

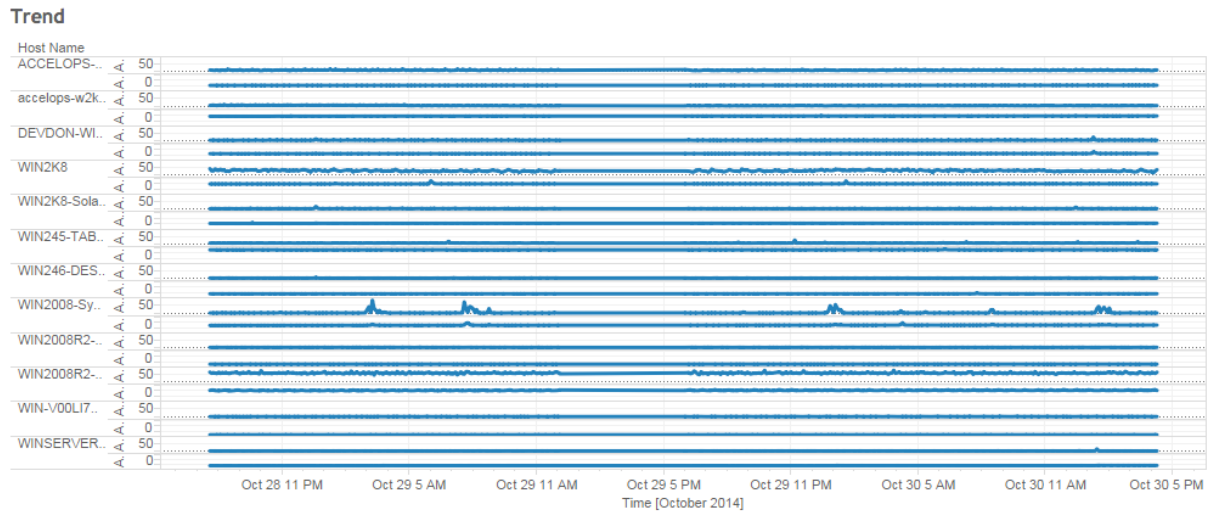

This section contains information on using Visual Analytics Desktop to create sheets and workbooks that are based on AccelOps reports, and then publishing them for others to use.

- [Viewing Workbooks](#)
- [Creating and Publishing Workbooks](#)
  - [Creating a Single Sheet Workbook](#)
  - [Creating a Multiple Sheet Workbook](#)
  - [Using AccelOps Workbooks with Tableau Visual Analytics Desktop and Server](#)
- [Adding Users to Workbooks](#)

## Viewing Workbooks

1. Log in to Visual Analytics Server.
2. Click the **Content** tab and select **Workbooks**.
3. Click on a workbook.
 

The workbook along with the various worksheets are displayed.
4. Select a workbook or worksheet.
5. You will be prompted for credentials that will allow the workbook or worksheet to access database information. Enter the Admin credential that you used to set up Accelops Report Server and click **OK**.
6. When your credential is accepted, the chart associated with the selected workbook or worksheet will be displayed.



## Creating and Publishing Workbooks

Workbooks are collections of AccelOps reports that have been synced to AccelOps Report Server, and which are then the basis for charts and dashboards that can be published to Visual Analytics Server for access by other users. Information in this section describes how to create single and multiple sheets of report information, and then make them accessible to other users.

- [Creating a Single Sheet Workbook](#)
- [Creating a Multiple Sheet Workbook](#)
- [Using AccelOps Workbooks with Tableau Visual Analytics Desktop and Server](#)

## Creating a Single Sheet Workbook

These instructions demonstrate how to create a single-sheet workbook that will chart the CPU and memory utilization trend for various servers. This example uses the **Servers by CPU, Memory** report and its associated table, but any report with a table in the reportdb database can also be used. The [Tableau Desktop online Help](#) also contains extensive information about building sheets and workbooks with the Tableau Desktop editor, which powers the AccelOps Visual Analytics Desktop.

- [Prerequisites](#)
- [Procedure](#)
  - [Create the Sheet](#)
  - [Create the Workbook](#)
  - [Publish the Workbook](#)

### Prerequisites

- Follow the instructions in [Syncing an AccelOps Report with Report Server](#) to sync the report you want to use for your worksheet.
- You will need to know the name of the parent table for your synced report. Follow the instructions in [Viewing reportdb Organization](#) to find the table that corresponds to your report.

### Procedure

#### Create the Sheet

1. Launch AccelOps Visual Analytics Desktop.
2. Connect to AccelOps Report Server with the **Username** and **Password** that you used during Report Server installation.  
For **Database**, enter **reportdb**.  
For **Port**, enter **30000**.

#### Connecting to Port 30000

It's important to make sure you enter the correct port to connect to the reportdb database. If you leave this option blank you will connect to the default PostgreSQL database port, which will connect you with phoenixdb instead of reportdb. For more information about the databases contained in Report Server, see [Report Server Architecture: phoenixdb and reportdb](#).

3. Under **Tables**, select the parent table for your report.  
For the steps following, we will use the **Servers by CPU, Memory** table and its associated columns.
4. Drag the table to the **View** pane and click **Update Now**.  
The data in the table will load into the pane below. Note that the table columns match closely to the Report Display Columns in AccelOps.
5. For **Connection**, select **Live**.
6. Click **Go to Worksheet**.  
In the worksheet view you will see that a set of **Dimensions** and **Measures** are populated for the table.
7. Under **Measures**, select **Report Time** and drag it to the **Dimensions** section to create Report Time as a calculated measurement.
8. Under **Dimensions**, right-click on **Report Time** to edit the calculation formula and convert it to a human-readable format from UNIX time. The formula should look like `DATEADD('second',INT([Report Time]),#1969-12-31 16:00:00#)`

You may also want to rename **Report Time** to **Time** to make it easier to read on the resulting chart.

9. Drag **Report Time** from **Dimensions** to **Columns**.
10. Under **Columns**, right-click on **Report Time** and select **Exact Date**.  
You should now see dates and time increments in your chart as the X-axis.
11. Under **Measures**, select and drag **AVG(cpuUtil)** and **AVG(memUtil)** to **Rows**.
12. Set the aggregation of both **AVG(cpuUtil)** and **AVG(memUtil)** to **AVG**.  
For example, **AVG(AVG(cpuUtil))** and **AVG(AVG(memUtil))**.  
You should now see both measures on the Y-axis of your chart.
13. Under **Dimensions**, drag **Host Name** to the **Color** section under **Marks**.  
Each host will be assigned a color and added to the chart.
14. Change the chart display name for **AVG(cpuUtil)** and **AVG(memUtil)** by clicking on each in the Y-axis to launch the Edit Y-Axis dialog.  
You can now edit the **Title** and **Range**, as well as other attributes, for each measure.
15. Under **Data**, click on the data source to open the Options menu, then click **Refresh**.
16. Rename the sheet by clicking on the data source to open the Options menu, then select **Rename** and enter a new name.

Your sheet is now complete. Hover your mouse over a trend line to view information about a specific host.

#### Create the Workbook



1. Click the **Dashboard** tab on the bottom of the Sheet editor to open the **Dashboard** editor.
2. Under **Dashboard**, select an appropriate **Size** and screen resolution.
3. Under **Dashboard**, select the sheet and drag it into the display pane.
4. Open the **Dashboard** options menu and select **Rename**.  
Change the name of the dashboard from **Server CPU/Memory Trend** to **Server Performance**.
5. In the **File** menu, select **Save**.

### ***Publish the Workbook***

1. In the **Server** menu, select **Sign In...**
2. Enter the IP address and port number for the Visual Analytics Server.
3. Enter the **Username** and **Password** for the Visual Analytics Server admin user, and then click **Sign In**.
4. In the **Server** menu, select **Publish Workbook**.
5. Enter attributes for the workbook, such the associated **Project**, **Name**, **View Permissions**, and **Views to Share**.  
See [Adding Users to Workbooks](#) for more information about user permissions for workbooks.
6. Click **Publish**.



## Creating a Multiple Sheet Workbook

These instructions demonstrate how to create a multiple-sheet workbook that will contain a set of charts related to Network Health. This example uses the **Network Devices by Ping RTT**, **Network Interfaces By Utilization**, and **Network Devices By CPU, Memory** reports, but any report with an associated table and views in the reportdb database could be used. The [Tableau Desktop online Help](#) also contains extensive information about building sheets and workbooks with the Tableau Desktop editor, which powers the AccelOps Visual Analytics Desktop.

- [Prerequisites](#)
- [Procedure](#)
  - [Create a View](#)
  - [Create a Workbook that Uses the View](#)
  - [Create the Workbook](#)
  - [Publish the Workbook](#)

### Prerequisites

- Follow the instructions in [Syncing an AccelOps Report with Report Server](#) to sync the reports you want to use for your worksheet.
- You will need to know the name of the parent table for your synced reports. Follow the instructions in [Viewing reportdb Organization](#) to find the table that corresponds to your report.

### Procedure

#### **Create a View**

Each report you want to include in your workbook corresponds to a table in the AccelOps reportdb. These tables need to be joined to cross-link the information that will appear in your workbook. In the case of a Network Health workbook that includes the sheets Network Devices by Ping RTT, Network Interfaces By Utilization, and Network Devices By CPU, Memory, the joining keys are **host name** and **time**.

1. Follow the instructions in [Viewing reportdb Organization](#) to find the parent tables for the reports you want to join.  
For each report there is one parent table and multiple child tables containing data for a particular month.
2. Create a SQL statement in pgAdmin to join the tables.  
In this example data is captured for one day. This enables quick generation of the data visualization.

```

SELECT cpu.report_time, cpu."hostName", cpu."hostIpAddr", cpu."AVG(cpuUtil)",
cpu."AVG(memUtil)",
    uptime."SUM(sysDownTime)", uptime."AVG(avgDurationMSec)",
uptime."LAST(sysUpTime)",
    uptime."SUM(pollIntv)", util."intfName", util."intfAlias",
    util."AVG(inIntfUtil)" AS "totalAvgInIntfUtil", util."AVG(outIntfUtil)" AS
"totalAvgOutIntfUtil",
    util."AVG(recvBitsPerSec)" AS "totalAvgRecvBitsPerSec",
    util."AVG(sentBitsPerSec)" AS "totalAvgSentBitsPerSec",
    util."AVG(outQLen)", util."AVG(intfSpeed64)"

FROM "Network Devices By CPU, Memory_1278492569_1" cpu,
    "Network Devices by Ping RTT_2021056235_1" uptime,
    "Network Interfaces By Utilization_382117475_1" util

WHERE ((cpu.report_time * 1000)::double precision * '00:00:00.001'::interval +
'1969-12-31 16:00:00-08'::timestamp with time zone) >= (now() - 1::double
precision * '1 day'::interval)
    AND ((uptime.report_time * 1000)::double precision *
'00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time zone) >=
(now() - 1::double precision * '1 day'::interval)
    AND ((util.report_time * 1000)::double precision * '00:00:00.001'::interval
+ '1969-12-31 16:00:00-08'::timestamp with time zone) >= (now() - 1::double
precision * '1 day'::interval)
    AND cpu.report_time = uptime.report_time AND cpu."hostName" =
uptime."hostName" AND uptime.report_time = util.report_time AND uptime."hostName"
= util."hostName";

```

3. Click the **Play** icon in pgAdmin to execute the query.  
Make sure the output pane contains data that is the result of the query execution.
4. Modify the SQL statement to create a view.  
Add this command at the top of the SQL statement:

```
CREATE OR REPLACE VIEW ph_network_health_view AS
```

Add this command at the bottom of the SQL statement:

```
grant select on ph_network_health_view TO public;
```

Your complete SQL statement should look like this:

```

CREATE OR REPLACE VIEW ph_network_health_view AS

SELECT cpu.report_time, cpu."hostName", cpu."hostIpAddr", cpu."AVG(cpuUtil)",
cpu."AVG(memUtil)",
    uptime."SUM(sysDownTime)", uptime."AVG(avgDurationMSec)",
uptime."LAST(sysUpTime)",
    uptime."SUM(pollIntv)", util."intfName", util."intfAlias",
    util."AVG(inIntfUtil)" AS "totalAvgInIntfUtil", util."AVG(outIntfUtil)" AS
"totalAvgOutIntfUtil",
    util."AVG(recvBitsPerSec)" AS "totalAvgRecvBitsPerSec",
    util."AVG(sentBitsPerSec)" AS "totalAvgSentBitsPerSec",
    util."AVG(outQLen)", util."AVG(intfSpeed64)"

FROM "Network Devices By CPU, Memory_1278492569_1" cpu,
    "Network Devices by Ping RTT_2021056235_1" uptime,
    "Network Interfaces By Utilization_382117475_1" util

WHERE ((cpu.report_time * 1000)::double precision * '00:00:00.001'::interval +
'1969-12-31 16:00:00-08'::timestamp with time zone) >= (now() - 1::double
precision * '1 day'::interval)
    AND ((uptime.report_time * 1000)::double precision *
'00:00:00.001'::interval + '1969-12-31 16:00:00-08'::timestamp with time zone) >=
(now() - 1::double precision * '1 day'::interval)
    AND ((util.report_time * 1000)::double precision * '00:00:00.001'::interval
+ '1969-12-31 16:00:00-08'::timestamp with time zone) >= (now() - 1::double
precision * '1 day'::interval)
    AND cpu.report_time = uptime.report_time AND cpu."hostName" =
uptime."hostName" AND uptime.report_time = util.report_time AND uptime."hostName"
= util."hostName";

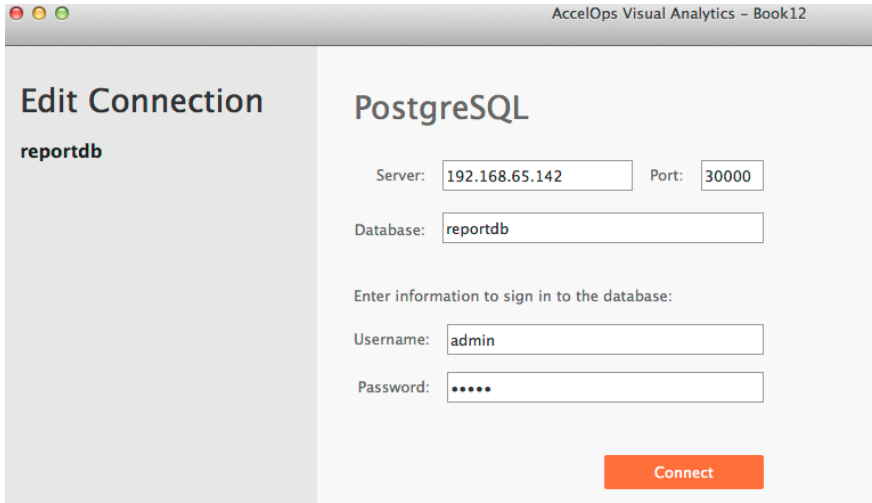
grant select on ph_network_health_view TO public;

```

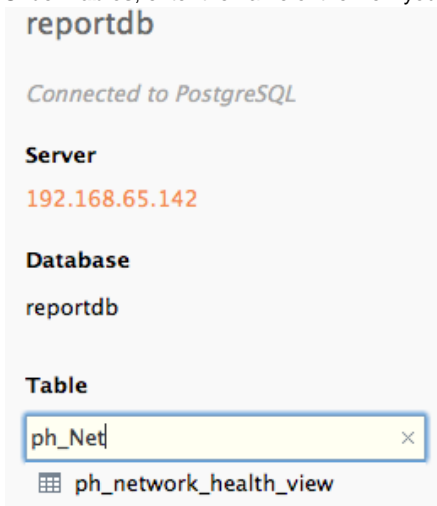
5. In pgAdmin, click the **Play** icon to execute the statement.
6. Using pgAdmin, navigate to the **Views** and make sure the `ph_network_health_view` has been created.
7. Right-click on `ph_network_health_view` to open the **Options** menu, then select **View Data > View Last 100 Rows** to make sure the view contains data.

### Create a Workbook that Uses the View

1. Launch AccelOps Visual Analytics Desktop.
2. Connect to AccelOps Report Server with the **Username** and **Password** that you used during Report Server installation.  
For **Database**, enter `reportdb`.  
For **Port**, enter `30000`.



- Under **Tables**, enter the name of the view you created in the search box to locate the view.



- Drag the view into the **Join** pane and click **Update Now**. The data in the view will load into the pane below.

ph\_network\_health\_view (reportdb) Live Extract 0 Add...

Connected to PostgreSQL

**Server**  
192.168.65.142

**Database**  
reportdb

**Table**  
ph\_Net  
ph\_network\_health\_view

ph\_network\_health\_view

Copy Go to Worksheet Show hidden fields Rows 2,358

Report Time #	Host Name Abc	Host Ip Addr Abc	AVG(cpuUtil) #	AVG(memUtil) #	SUM(sysDownTime) #	AVG(avgDurationMSec) #	LAST(sysUpTime) #	SUM(pollr #
1,416,870,299	SJ-Main-Cat65...	192.168.20.1	21.0000	31.24410	0	0	31,652,714	
1,416,870,299	SJ-Main-Cat65...	192.168.20.1	21.0000	31.24410	0	0	31,652,714	
1,416,870,299	SJ-Main-Cat65...	192.168.20.1	21.0000	31.24410	0	0	31,652,714	
1,416,870,299	SJ-Main-Cat65...	192.168.20.1	21.0000	31.24410	0	0	31,652,714	
1,416,870,299	SJ-Main-Cat65...	192.168.20.1	21.0000	31.24410	0	0	31,652,714	
1,416,870,299	ph-network-3...	192.168.64.1	25.0000	34.98160	0	1	6,560,735	

New Custom SQL


- For **Connection**, select **Live**.
- Click **Go to Worksheet**. In the worksheet view you will see that a set of **Dimensions** and **Measures** are populated for the view.

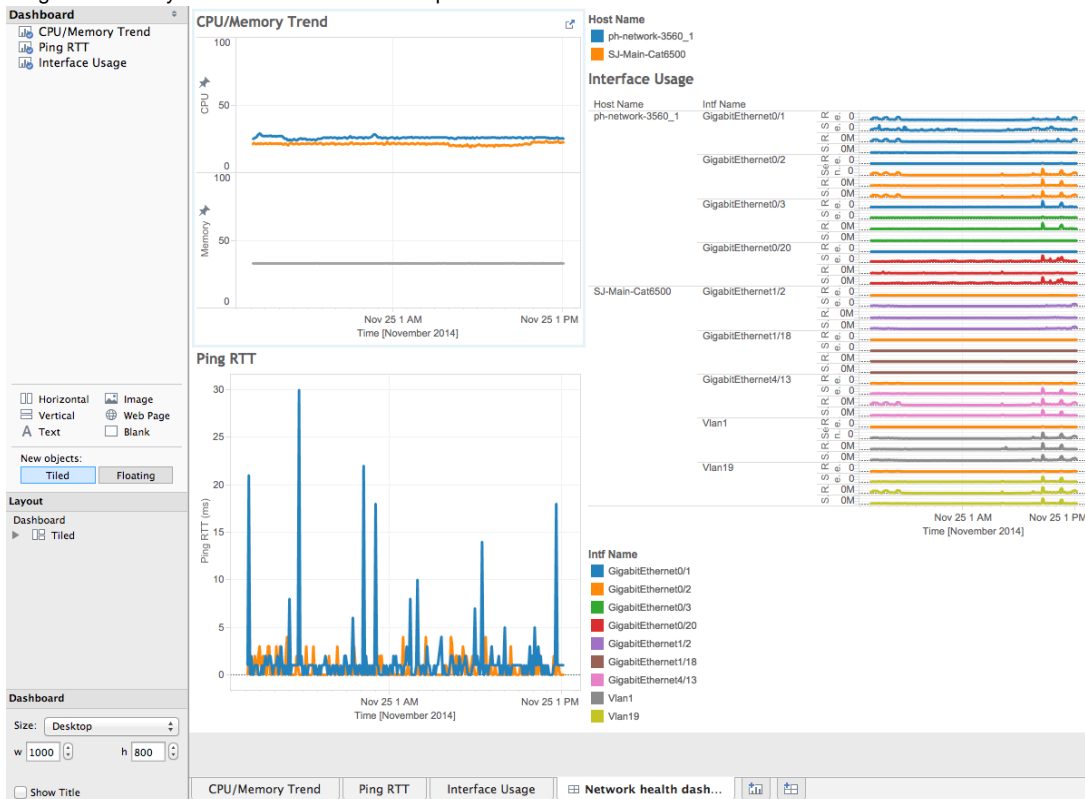


An example worksheet showing CPU and Memory Utilization with several Dimensions and Measures populated from the original table.

- For each report in your workbook you can now create an individual sheet, as described in [Creating a Single Sheet Workbook](#).

**Create the Workbook**

- Click the **Dashboard** tab  on the bottom of the Sheet editor to open the **Dashboard** editor.
- Drag each sheet you've created into the **Join** pane.



An example of three worksheets loaded into the Dashboard Join pane.

- Under **Dashboard**, select an appropriate **Size** and screen resolution.

4. Open the Dashboard **Options** menu and select **Rename**.
5. In the **File** menu, select **Save**.

### **Publish the Workbook**

1. In the **Server** menu, select **Sign In...**
2. Enter the IP address and port number for the Visual Analytics Server.
3. Enter the **Username** and **Password** for the Visual Analytics Server admin user, and then click **Sign In**.
4. In the **Server** menu, select **Publish Workbook**.
5. Enter attributes for the workbook, such as the associated **Project**, **Name**, **View Permissions**, and **Views to Share**. See [Adding Users to Workbooks](#) for more information about user permissions for workbooks.

Project:

Name:

Add Tags:

User/Group	Role
▶ All Users	(Custom)
Owner	Editor

Views to Share

Sheet

CPU/Memory Trend

Ping RTT

Interface Usage

Network health dashboard

Options

Show Sheets as Tabs

Show Selections



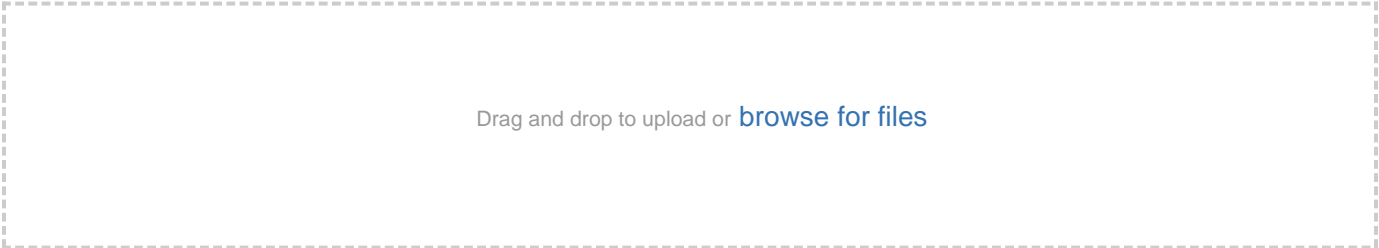
Buttons: Add..., Edit..., Remove, All, None, Authentication..., Cancel, Publish

6. Click **Publish**.

## Using AccelOps Workbooks with Tableau Visual Analytics Desktop and Server

You can use any of the workbooks provided by AccelOps, which are attached to this page, to create visualizations of AccelOps data.

1. Download a workbook attached to this page to your local device where Tableau Visual Analytics Desktop is installed.
2. In Visual Analytics Desktop, go to **File > Open...**
3. Browse to the file you downloaded and open it.
4. You can make any changes you want to the workbook, but you can upload it to the server and start using it as is. Follow the instructions in the **Publish the Workbook** section of [Creating a Single Sheet Workbook](#) to publish to the Tableau Visual Analytics Sever, and add user permissions as described in [Adding Users to Workbooks](#).

File	Modified 
<a href="#">&gt;  Accelops Incidents.twbx</a>	Dec 08, 2014 by Phil Gochenour
 <p>Drag and drop to upload or <a href="#">browse for files</a></p>	

## Adding Users to Workbooks

Only the workbook publisher can give access to specific users during report creation time. As the AccelOps Visual Analytics Server Administrator, you can add users to the system and view which workbooks users can access.

- [Adding Users to Visual Analytics Server](#)
- [Viewing User Access to Workbooks](#)

## Adding Users to Visual Analytics Server

1. Log in to AccelOps Visual Analytics Server.
2. In the **Admin** tab click **Users**.
3. Click **Add**.
4. Enter the user name as it appears in Active Directory.
5. Select the **License Level** for the user and assign **User Rights** as necessary.
6. Click **OK**.

## Viewing User Access to Workbooks

1. Log in to Visual Analytics Server.
2. In the **Admin** tab click **Users**.
3. Select a user name to see the workbooks that the user can access.



## Working with AccelOps Report Server

This section contains information on AccelOps Report Server architecture, viewing and querying CMDB and Event data in contained in the Report Server databases, and database maintenance.

- [Report Server Architecture: phoenixdb and reportdb](#)
- [Working with CMDB Data in AccelOps Report Server](#)
  - [Viewing phoenixdb Organization](#)
  - [Querying Incident Data in AccelOps Report Server](#)
    - [Reference: Attribute Columns in the ph\\_incident\\_view Table](#)
    - [Sample Incident Queries](#)
  - [Querying Other CMDB Tables in AccelOps Report Server](#)
    - [Querying Device Vendor and Model Distribution for Discovered Devices](#)
    - [Querying Discovered Devices](#)
- [Working with Event Data in AccelOps Report Server](#)
  - [Viewing reportdb Organization](#)
  - [Syncing an AccelOps Report with Report Server](#)
  - [Deleting a Report from AccelOps Report Server](#)
  - [Modifying an Existing Report in AccelOps Report Server](#)
- [Database Maintenance](#)
  - [Checking the Size of Tables](#)
  - [Creating Indices for Faster Queries](#)
  - [Deleting Old Tables](#)

## Report Server Architecture: phoenixdb and reportdb

AccelOps Report Server contains two databases:

- **phoenixdb**  
This database contains the entire AccelOps CMDB and is populated via asynchronous PostgreSQL replication (slony) in near-real time.
- **reportdb**  
This database contains the results of event queries.

Topics in this section describe how to view the tables in these databases, and how those tables are organized. For viewing the tables, we recommend using the pgAdmin PostgreSQL database utility, which you can download from [the pgAdmin website](#).

## Working with CMDB Data in AccelOps Report Server

Data from the AccelOps CMDB database is populated to the AccelOps Report Server and stored in the Report Server phoenixdb. This section contains information on how to view the organization of phoenixdb, and write queries against the data it contains.

- [Viewing phoenixdb Organization](#)
- [Querying Incident Data in AccelOps Report Server](#)
  - [Reference: Attribute Columns in the ph\\_incident\\_view Table](#)
  - [Sample Incident Queries](#)
- [Querying Other CMDB Tables in AccelOps Report Server](#)
  - [Querying Device Vendor and Model Distribution for Discovered Devices](#)
  - [Querying Discovered Devices](#)

## Viewing phoenixdb Organization

This database contains the contents of the entire AccelOps CMDB database, including incidents.

1. In the [pgAdmin utility](#), go to **File > Add Server**.
2. In the **New Server Registration** dialog, enter connection details for AccelOps Report Server.  
For **Maintenance DB**, select **phoenixdb**.  
For **Username** and **Password**, use the read-only user name and password that you created when you provisioned the Report Server.
3. Click **OK**.  
When the connection to the AccelOps Report Server is established, phoenixdb will load in the Object browser. There are approximately 197 tables in phoenixdb, which are replicated from the AccelOps cluster.
4. Select a table to view, then right-click to open the **Options** menu.
5. In the **Options** menu, select **View Data**, and then select an option for which rows you want to view.  
For example, to view the contents of the `ph_device` table, which contains CMDB information about discovered devices, you would select and then right click on `ph_device`, then select **View Data > View All Rows**.

You can also use this method to examine Views and other objects in the phoenixdb database.

## Querying Incident Data in AccelOps Report Server

There are two ways to look at the incident data inside AccelOps Report Server:

- **Incident Tables** (`ph_incident` and `ph_incident_detail`)  
Contains the incidents
- **Incident View** (`ph_incident_view`)  
This is a database view that adds other context to the incident tables by joining with other tables in the database. Added information includes location and business service. Some information is parsed out for easier query, such as host names and IP address fields from `incident_source`, and `incident_target` fields in `ph_incident` are parsed out as separate fields in `ph_incident_view`.

This topic describes how to view the data contained in **Incident View**.

1. Follow the instructions in [Viewing phoenixdb Organization](#) to access the phoenixdb database in AccelOps Report Server.
2. Go to **Views > ph\_incident\_view > Columns** to view the table columns.
3. Go to **Views > ph\_incident\_view > View Data > View Last 100 Rows** to view the incidents.

## Reference: Attribute Columns in the ph\_incident\_view Table

Column Name	Format	Description
incident_id	integer	Unique id for an incident
cust_org_id	integer	Customer Id (for AO-SP)
first_seen_time	integer	The time when the incident was first seen. The format is UNIX time but with milliseconds granularity. It is defined as the number of milliseconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970
last_seen_time	integer	The time when the incident was last seen. The format is UNIX time but with milliseconds granularity. It is defined as the number of milliseconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970
incident_et	string	Incident event type id e.g. PH_RULE_SERVER_HW_CRITICAL
incident_status	integer	0: Active 1: Auto Cleared 2: Manually Cleared 3: System Cleared
incident_count	integer	The number of times this exact incident (with the same parameters: source, destination etc has happened)
biz_name	string	Associated business service name
severity	integer	Numerical severity of the incident - range 0-10
severity_cat	string	Incident severity category: 0-4: LOW, 5-8: MEDIUM and 9-10: HIGH
orig_device_ip	string	IP address of the device that reported the incident
ph_incident_category	string	Category of infrastructure affected by this incident: possible values: Network, Server, Storage, Virtualization, Application, Internal
incident_src	string	Incident Source string formatted as a list of <Attribute>:Value; e.g. srcIpAddr:10.1.1.1,srcName:JoeLaptop
src_ip_addr	string	Source IP parsed out from incident_src field
src_name	string	Source Name parsed out from incident_src field
src_device_location	string	(Geo) Location display name string for the object specified in incident_src
src_country	string	(Geo) Country name string for the object specified in incident_src
src_state	string	(Geo) State name for the object specified in incident_src
src_building	string	(Geo) Building name for the object specified in incident_src
src_floor	string	(Geo) Floor for the object specified in incident_src
src_latitude	double	(Geo) Latitude for the object specified in incident_src
src_longitude	double	(Geo) Longitude for the object specified in incident_src
incident_target	string	Incident Destination string formatted as a list of <Attribute>:Value; e.g. "destIpAddr:10.1.1.1,destName:JoeLaptop" or "hostIpAddr:10.1.1.1,hostName:JoeLaptop"
dest_ip_addr	string	Destination IP parsed out from incident_target field
dest_name	string	Destination Name parsed out from incident_target field
dest_device_location	string	(Geo) Location display name string for the object specified in incident_target
dest_country	string	(Geo) Country name string for the object specified in incident_target
dest_state	string	(Geo) State name for the object specified in incident_target
dest_building	string	(Geo) Building name for the object specified in incident_target
dest_floor	string	(Geo) Floor for the object specified in incident_target
dest_latitude	double	(Geo) Latitude for the object specified in incident_target

dest_longitude	double	(Geo) Longitude for the object specified in incident_target
host_ip_addr	string	Host IP address parsed out from incident_target field
host_name	string	Host Name parsed out from incident_target field
host_device_location	string	(Geo) Location display name string for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_country	string	(Geo) Country name string for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_state	string	(Geo) State name for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_building	string	(Geo) Building name for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_floor	string	(Geo) Floor for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_latitude	double	(Geo) Latitude for the object specified in incident_target - populated if incident_target contains hostIpAddr
host_longitude	double	(Geo) Longitude for the object specified in incident_target - populated if incident_target contains hostIpAddr
vm_name	string	VM Name if incident involves a Virtual machine - populated if incident_target contains vmName
user_attr	string	User name if incident involves user, i.e. incident_target contains user
target_user_attr	string	Target user name if incident involves user, i.e. incident_target contains targetUser
ldap_domain	string	Domain if incident involves user, i.e. incident_target contains domain
computer	string	Computer name incident_target contains computer
target_computer	string	Target Computer name incident_target contains targetComputer
incident_details	string	Incident Details containing evidence on why the incident triggered e.g. Triggered Event Count = 90 or AVG(CPUUtil) = 90 etc

## Sample Incident Queries

- [Show Incident Categories with Severity and Frequency Occurrence](#)
- [Show Incident Location](#)

### Show Incident Categories with Severity and Frequency Occurrence

This query will show which parts of the infrastructure are triggering events.

1. Follow the instructions in [Viewing phoenixdb Organization](#) to access the phoenixdb in AccelOps Report Server.
2. Under **Views**, select `ph_incident_view`.
3. In `pgAdmin`, click on the **SQL** icon in the menu bar to open the SQL query window.
4. Enter this SQL query:

```
SELECT ph_incident category, incident_et, severity_cat, src_ip_addr, host_name, COUNT(*)
FROM ph_incident_view
GROUP BY ph_incident category, incident_et, severity_cat, src_ip_addr, host_name
ORDER BY COUNT(*) DESC;
```

5. When the query executes, you will see a list of matching incidents in the Output Pane.

### Show Incident Location

1. Follow the instructions in [Viewing phoenixdb Organization](#) to access the phoenixdb in Accelops Report Server.
2. Under **Views**, select `ph_incident_view`.
3. In `pgAdmin`, click on the **SQL** icon in the menu bar to open the SQL query window.
4. Enter this SQL query:

```
SELECT host_device_location, severity_cat, ph_incident_category, COUNT(*)
FROM ph_incident_view
GROUP BY host_device_location, ph_incident_category, severity_cat
ORDER BY host_device_location ASC, severity_cat ASC COUNT(*) DESC;
```

5. When the query executes, you will see a list of incidents and their locations in the Output Pane.



## Querying Other CMDB Tables in AccelOps Report Server

- [Querying Device Vendor and Model Distribution for Discovered Devices](#)
- [Querying Discovered Devices](#)

## Querying Device Vendor and Model Distribution for Discovered Devices

## Querying Discovered Devices

## Working with Event Data in AccelOps Report Server

Data from the AccelOps EventDB database is populated to the AccelOps Report Server and stored in the Report Server reportdb. This section contains information on how to view the organization of reportdb, and write queries against the data it contains.

- [Viewing reportdb Organization](#)
- [Syncing an AccelOps Report with Report Server](#)
- [Deleting a Report from AccelOps Report Server](#)
- [Modifying an Existing Report in AccelOps Report Server](#)

## Viewing reportdb Organization

This database contains the reports that are synced from the AccelOps cluster.

1. In the [pgAdmin utility](#), go to **File > Add Server**.
2. In the **New Server Registration** dialog, enter connection details for AccelOps Report Server.  
For **Maintenance DB**, select **reportdb**.  
For **Username** and **Password**, use the read-only user name and password that you created when you provisioned the Report Server.
3. Click **OK**.  
When the connection to the Report Server is established, reports will load in the Object browser.
4. Select a table to view, then right-click to open the **Options** menu.
5. In the **Options** menu, select **View Data**, and then select an option for which rows you want to view.

## Syncing an AccelOps Report with Report Server

1. Log in to AccelOps.
2. Go to **Analytics > Reports**.
3. Select a report.  
Any reports with a **Sync** checkbox can be synced. Run the report to make sure it contains some data.
4. For each report you want to sync, select the **Sync** checkbox.

**AO-SP**  
In the **Sync Details** dialog, select the organizations whose data needs to be synced.

5. Click **OK**.
6. After several minutes, follow the instructions in [Viewing reportdb Organization](#) to view the reportdb database.
7. Under **Tables**, you should now see the synced reports.

### Table Structure for Synced Reports

When you sync an AccelOps report to AccelOps Report Server, two pairs of tables are created in reportdb, one pair for each organization in the case of AO-SP. For each organization, multiple tables are created:

1. A **parent table** containing data for all months: the table name is of the form `<Report Name>_<ID>_<custId>`
2. A **child table** for the current month: `<Report Name>_<ID>_<custId>_<yYYYYmMM>` where YYYY is the year and MM is the month.

Queries should be written using the parent table. To see data in the parent table, follow the instructions in [Viewing reportdb Organization](#). The reportdb database fields are generated from the display fields in AccelOps report definitions. Only the field `report_time` is added to the Report Server table definitions to capture the time when the particular report is generated. For example, if you synced the report **Network Devices by CPU, Memory**, you would see these fields:

Field	Description
<code>report_time</code>	UNIX time at which the report is generated. <b>Unix time</b> (or <b>POSIX time</b> or <b>Epoch time</b> ) is a system for describing instants in time, defined as the number of seconds that have elapsed since 00:00:00 Coordinated Universal Time (UTC), Thursday, 1 January 1970 not counting leap seconds.
<code>hostName</code>	Host Name of the device for which CPU and memory are being measured
<code>hostIpAddr</code>	Access IP of the device for which CPU and memory are being measured
<code>AVG(cpuUtil)</code>	Average of all the CPU utilization metrics within the last 5 minutes ending with <code>report_time</code>
<code>AVG(memUtil)</code>	Average of all the CPU utilization metrics within the last 5 minutes ending with <code>report_time</code>

## Deleting a Report from AccelOps Report Server

1. Log in to AccelOps.
2. In **Analytics > Reports > Synced Reports**, select the report you want to delete.
3. In the **Sync Details** dialog, clear the **Sync** option for the report, and then click **OK**.  
The report will no longer be synced with Report Server. You can verify this by making sure the **Sync** option is not selected for the report on the **Analytics > Reports > Synced Reports** page. You can now delete the report from AccelOps Report Server.
4. Log in to AccelOps Report Server via SSH and navigate to the directory `/opt/phoenix/deployment/jumpbox`.
5. Run the `phreportdbmanager.py` command, along with the table name and date as arguments, to delete the report.

```
phreportdbmanager.py --remove tablenames="Network Devices By CPU,
Memory_1278492569_1" reporttimes=2014-10
```

### Viewing the Names of Reports in Report Server

Use the `pgAdmin` utility to view the names of all tables and reports in Report Server, as described in [Viewing reportdb Organization](#).

When the deletion process completes, you will see a command line output like this:

```
[admin@RS142 jumpbox]$ pwd
/opt/phoenix/deployment/jumpbox
[admin@RS142 jumpbox]$ phreportdbmanager.py
Usage: --remove reporttimes=yyyy-mm[,yyyymm,...]
Usage: --remove tablenames='tablename[|tablename[| ...]]' reporttimes=yyyy-mm[,yyyymm,...]
Usage: --archive reporttimes=yyyy[-mm][,yyyymm[-mm],...] archivepath=path-to-archiving-directory
[admin@RS142 jumpbox]$ phreportdbmanager.py --remove tablenames="Network Devices By CPU, Memory_1278492569_1" reporttimes=2014-10
Connected to reportdb
dropping child tables for report table "Network Devices By CPU, Memory_1278492569_1"
dropping child table "Network Devices By CPU, Memory_1278492569_1_y2014m10"
dropped child table "Network Devices By CPU, Memory_1278492569_1_y2014m10"
[admin@RS142 jumpbox]$
```

6. After you have deleted the table containing the report information, you will need to delete the parent table, which will now be empty of content, using the same `phreportdbmanager.py` command.

## Modifying an Existing Report in AccelOps Report Server

Suppose a system rule is synced and exported to AccelOps Report Server. When you modify that report in AccelOps, you must rename it, at which point it becomes a user report. When you then sync that report for AccelOps Report Server, a new table is created on the AccelOps Report Server.

Suppose now that you have a user-defined report that is already synced to the AccelOps Report Server, but you modify it inline in AccelOps, which means that you have changed the report conditions without changing the report name. This will cause a change in the table, but a new table will not be created. Here are some examples of inline modifications, and how they affect the structure of the table as well as the data collected in the table:

Modification	Effect
GROUP BY field added	The corresponding table has the new GROUP BY field, but only newer data populates the field
GROUP BY field removed	There is no change in the corresponding table, and newer data does not populate the field
GROUP BY field changed	For example, the field <code>srcIpAddr</code> is changed to <code>destIpAddr</code> . Both fields are retained, but newer data populates <code>destIpAddr</code> .
Aggregated fields added	The corresponding table has the new field, but only newer data populate that field
Aggregated field removed	There is no change in the corresponding table, and newer data does not populate the field
Aggregated Field Changed	For example, <code>AVG(cpuUtil)</code> is changed to <code>MAX(cpuUtil)</code> . Both fields are retained, but newer data populates <code>MAX(cpuUtil)</code> .



## Database Maintenance

- Checking the Size of Tables
- Creating Indices for Faster Queries
- Deleting Old Tables

## Checking the Size of Tables

## Creating Indices for Faster Queries

## Deleting Old Tables