



VOSS Automate Installation Guide

Release 24.2-PB1

August 21, 2025

Legal Information

- Copyright © 2025 VisionOSS Limited.
All rights reserved.
- This information is confidential. If received in error, it must be returned to VisionOSS ("VOSS"). Copyright in all documents originated by VOSS rests in VOSS. No portion may be reproduced by any process without prior written permission. VOSS does not guarantee that this document is technically correct or complete. VOSS accepts no liability for any loss (however caused) sustained as a result of any error or omission in the document.

DOCUMENT ID: 20250821043538

Contents

1	Introduction	1
1.1	Overview	1
1.2	Command Conventions	1
2	Deployment Topologies	2
2.1	VOSS Automate deployment topologies	2
2.2	Multi Data Center Deployments	17
2.3	Clustering Considerations	21
2.4	Network Communications between Nodes within the Cluster	21
2.5	Network Communications External to the Cluster	23
2.6	Cluster Commands	24
2.7	Geo-redundancy/Redundancy And DR Synopsis	25
3	Prepare to Install	26
3.1	Network Docker Container Range	26
3.2	Backup Size Considerations	27
3.3	Topology Hardware Specifications	27
4	Install VOSS Automate	35
4.1	Installation Process	35
4.2	Create a New VM Using the Platform-Install OVA	35
4.3	Using the screen command	40
4.4	Installation Logs	41
4.5	Commands to Determine the Primary Node Role in a Cluster	44
4.6	Unified Node Topology Installation	44
4.7	Modular Cluster Topology Installation	59
5	VOSS Automate Azure Installation	65
5.1	Azure Cloud Deployment	65
6	VOSS Automate AWS Installation	70
6.1	AWS Deployment	70
	Index	76

1. Introduction

1.1. Overview

Note: For upgrade steps, refer to the Upgrade Guide (major releases) or MOP (maintenance release patch bundles).

The VOSS Automate system can be deployed in either a test single node topology or a multi-node cluster with High Availability and Disaster Recovery capabilities. This guide is aimed at Technical and Operational personnel responsible for the deployment of a VOSS Automate system.

The guide describes the product in general and is not specific to a particular deployment/solution. For cluster installations, also refer to the Health Checks for Cluster Installations Guide.

Information may vary depending on the installation environment.

The guide provides an overview of the deployment of the VOSS Automate system on:

- VMware: [Create a New VM Using the Platform-Install OVA](#)
- MS Azure: [Azure Cloud Deployment](#)
- AWS: [AWS Deployment](#)

1.2. Command Conventions

For file transfers, either **sftp** or **scp** can be used. For example, either:

1. **sftp platform@<node_hostname>**
2. **cd media**
3. **put <upgrade_template_file>**

or:

scp <upgrade_template_file> platform@<node_hostname>:~/media

In this Guide, **scp** is used to show file transfer.

2. Deployment Topologies

2.1. VOSS Automate deployment topologies

2.1.1. Overview

Architecturally, VOSS Automate offers two main deployment topologies:

- *Unified node cluster topology*
- *Modular cluster deployment topology*

The VOSS Automate deployment topologies are comprised of a configuration of nodes.

In addition, the following deployment options are available:

- Cloud deployments
- VOSS Automate Cloudv1

2.1.2. Deployment topology nodes

Nodes are of distinct types: Web Proxy, Unified/Single node, Application and Database, that perform specific functions within a topology. Each node type comprises of one or more components (software subsystems).

Components:

- **Operating system:** Ubuntu, stripped down / hardened
- **Platform:** Docker, isolated components
- **Web server:** nginx, receives and forwards HTTP requests
 - Hosts static files: CSS, JS and images;
 - Load balance between unified nodes (UNs): round robin, configurable, e.g. 2 data centres
 - Detects inactive UN: removes from round robin
- **Database:** MongoDB (scalable, distributed), PostgreSQL (scalable)
- **Application:** JavaScript, Python, REST API, device drivers, workflow engine, transactions / queue engine, RBAC, search, bulk loader, and more ...

The table below shows the set of components in each node type.

Component	Node Types			
	Web Proxy	Uni-fied/Single node	Applica-tion	Database
Operating system	X	X	X	X
Platform	X	X	X	X
Web server	X	X		
Database		X		X
Application		X	X	

2.1.3. Unified node cluster topology

VOSS Automate's Unified Node Cluster topology comprises these options:

- Single-node cluster (cluster-of-one/standalone)
- Single-node cluster (cluster-of-one/standalone) with VMWare HA
- 2 Node with Web proxies
- 4 Node with Web proxies
- 6 Node with Web proxies

Important: Choose between a Unified Node deployment or a Modular Architecture deployment.

In a Unified Node deployment, VOSS Automate is deployed either as a single node cluster, 2 unified nodes, or a cluster of multiple nodes with High Availability (HA) and Disaster Recovery (DR) qualities.

Each node can be assigned one or more of the following functional roles:

- WebProxy - load balances incoming HTTP requests across unified nodes.
- Single node cluster - combines the Application and Database roles for use in a non-multi-clustered test environment.
- Unified - similar to the Single node cluster role Application and Database roles, but clustered with other nodes to provide HA and DR capabilities

The nginx web server is installed on the WebProxy, Single node cluster, and Unified node, but is configured differently for each role.

In a clustered environment containing multiple Unified nodes, a load balancing function is required to offer HA (High Availability providing failover between redundant roles).

VOSS Automate supports deployment of either the WebProxy node or a DNS load balancer. Here are some considerations in choosing a WebProxy node vs. DNS:

- The Proxy takes load off the Unified node to deliver static content (HTML/JAVA scripts). When using DNS or a third-party load balancer, the Unified node has to process this information.
- DNS does not know the state of the Unified node.
- The WebProxy detects if a Unified node is down or corrupt. In this case, the WebProxy will select the next Unified node in a round robin scheme.

We recommend that you run no more than two Unified nodes and one WebProxy node on a physical server (VMware server). Also recommended is that the disk subsystems be unique for each Unified node.

The following deployment topologies are defined:

- Test: a standalone, single node cluster with Application and Database roles combined. No High Availability/Disaster Recovery (HA/DR) is available.

Important: This deployment should be used for test purposes only.

- Production with Unified Nodes: in a clustered system, comprising:
 - 2, 3, 4 or 6 Unified nodes (each with combined Application and Database roles)
 - 0 to 4 (maximum 2 if 2 Unified nodes) WebProxy nodes offering load balancing. The WebProxy nodes can be omitted if an external load balancer is available.

Single-node cluster (cluster-of-one/standalone)

A single-node cluster (cluster-of-one/standalone) deployment should be used for test purposes only.



The table describes the advantages and disadvantages of a single-node cluster (cluster-of-one/standalone) deployment topology:

Advantages	Disadvantages
<ul style="list-style-type: none">• Smallest hardware footprint	<ul style="list-style-type: none">• No high availability or disaster recovery• Less throughput than clusters

Single-node cluster (cluster-of-one/standalone) with VMWare HA

The table describes the advantages and disadvantages of a single-node cluster (cluster-of-one/standalone) with VMWare HA deployment topology:

Advantages	Disadvantages
<ul style="list-style-type: none"> • Smallest hardware footprint • Disaster recovery available 	<ul style="list-style-type: none"> • Less throughput than clusters

Multinode cluster with unified nodes

To achieve Geo-Redundancy using the Unified nodes, consider the following:

- Either four or six Unified nodes - each node combining Application and Database roles - are clustered and split over two geographically disparate locations.
- Two Web Proxy nodes to provide High Availability that ensure an Application role failure is gracefully handled. More may be added if Web Proxy nodes are required in a DMZ.

It is strongly recommended *not* to allow customer end-users the same level of administrator access as the restricted groups of provider- and customer administrators. This is why Self-service web proxies as well as Administrator web proxies should be used.


Systems with Self-service only web proxies are *only* recommended where the system is customer facing, but where the customer does not administer the system themselves.

- Web Proxy and Unified nodes can be contained in separate firewalled networks.
- Database synchronization takes places between all Database roles, thereby offering Disaster Recovery and High Availability.
- For 6 unified nodes, all nodes in the cluster are active. For an 8 node cluster (with latency between data centers greater than 10ms) , the 2 nodes in the DR node are passive, in other words, the **voss workers 0** command has been run on the DR nodes.

Primary and fall-back Secondary Database servers can be configured manually. Refer to the Platform Guide for further details.


Example: 6-Node Cluster

The diagram illustrates an example of a 6-Node Cluster:




Example: 8 Node Cluster

The diagram illustrates an example of an 8-Node Cluster:




Example: 2 Web Proxy Nodes in a DMZ

The diagram illustrates an example of 2 Web proxy nodes in a DMZ:

**Example: 4 Web Proxy Nodes in a DMZ (2 admin, 2 Self-service)**

The diagram illustrates an example of 4 Web proxy nodes (2 admin, and 2 Self-service) in a DMZ:



2 Node cluster with unified nodes

To achieve Geo-Redundancy using the Unified nodes, consider the following:

- Two unified nodes - each node combining application and database roles - are clustered and optionally split over two geographically disparate locations.
- (Optional) Two web proxy nodes can be used. It may be omitted if an external load balancer is available.
- Web proxy and unified nodes can be contained in separate firewalled networks.
- Database synchronization takes place from primary to secondary unified nodes, thereby offering Disaster Recovery if the primary node fails.

- If the secondary unified node has *more than 10ms latency* with the primary unified node, it must be configured to be in the *same* geographical location.

Important:

With only two Unified nodes, with or without Web proxies, there is no High Availability. The database on the primary node is read/write, while the database on the secondary is read only.


Only redundancy is available.

- If the primary node fails, a manual delete of the primary node on the secondary and a cluster provision will be needed.
- If the secondary node fails, it needs to be replaced.

Refer to the topic on DR Failover and Recovery in a 2 Node Cluster in the Platform Guide.

Example: 2 Node Cluster

The diagram illustrates a 2 Node Cluster:



4 Node with web proxies

The table describes the advantages and disadvantages of a 4 Node with Web Proxies deployment topology:

Advantages	Disadvantages
<ul style="list-style-type: none"> • More disaster recovery scenarios supported • More throughput than 3 Node 	<ul style="list-style-type: none"> • More hardware than 3 Node

6 Node with web proxies

A 6 Node with Web proxies deployment topology:

- Is typically deployed for multi-data center deployments
- Supports Active/Standby

2.1.4. Modular cluster deployment topology

Overview

A Modular Cluster topology has separate Application and Database nodes:

- 3 Database nodes
- 1 - 8 Application Nodes
- Web Proxies

A Modular Cluster topology has the following advantages:

- Increased processing capacity
- Horizontal scaling by adding more Application nodes
- Improved database resilience with dedicated nodes and isolation from application
- Improved database performance by removing application load from the primary database

Important: Choose between a Unified Node Cluster deployment or a Modular Cluster deployment.

VOSS Automate is deployed as a modular cluster of multiple nodes with High Availability (HA) and Disaster Recovery (DR) qualities.

Each node can be assigned one or more of the following functional roles:

- WebProxy - load balances incoming HTTP requests across nodes.
- Application role node, clustered with other nodes to provide HA and DR capabilities
- Database role node, clustered with other nodes to provide HA and DR capabilities

The nginx web server is installed on the WebProxy and application role node, but is configured differently for each role.

Related Topics

Modular Architecture Multinode Installation in the Install Guide.

Migrate a Unified Node Cluster to a Modular Cluster in the Platform Guide.

A load balancing function is required to offer HA (High Availability providing failover between redundant roles).

VOSS Automate supports deployment of either the WebProxy node or a DNS load balancer. Consider the following when choosing a WebProxy node vs. DNS:

- The Proxy takes load off the application role node to deliver static content (HTML/JAVA scripts). When using DNS or a third-party load balancer, the application role node has to process this information.
- DNS does not know the state of the application role node.
- The WebProxy detects if an application role node is down or corrupt. In this case, the WebProxy will select the next application role node in a round robin scheme.

We recommend that you run no more than one application role node and one database role node and one WebProxy node on a physical server (VMware server). When selecting disk infrastructure, high volume data access by database role replica sets must be considered where different disk subsystems may be required depending on the performance of the disk infrastructure.

The following modular cluster topology is recommended (minimum):

Important: Single node cluster topologies are not available for modular cluster deployments.

- Production with nodes: in a clustered system of 2 data centers:
 - DC1 = primary data center containing primary database node (highest database weight)
 - DC2 = data recovery data center

The system comprises of the following nodes:

- 3 nodes with application roles (2 in DC1; 1 in DC2)
- 3 nodes with database roles (2 in DC1; 1 in DC2)
- Maximum 2 WebProxy nodes if 2 data centers; offering load balancing. The WebProxy nodes can be omitted if an external load balancer is available.

Multinode modular cluster with application and database nodes

To achieve Geo-Redundancy using Application and Database nodes, consider the following:

- Six Application and Database nodes - 3 nodes with an application role and 3 nodes with a database role - are clustered and split over two geographically disparate locations.
- Two Web Proxy nodes to provide High Availability that ensure an Application role failure is gracefully handled. More may be added if Web Proxy nodes are required in a DMZ.

It is strongly recommended *not* to allow customer end-users the same level of administrator access as the restricted groups of provider- and customer administrators. This is why Self-service web proxies as well as Administrator web proxies should be used.

Systems with Self-service only web proxies are *only* recommended where the system is customer facing, but where the customer does not administer the system themselves.


- Web Proxy, application and database nodes can be contained in separate firewalled networks.

- Database synchronization takes place between all database role nodes, thereby offering Disaster Recovery and High Availability.
- All nodes in the cluster are active.

Primary and fall-back Secondary Database servers can be configured manually. Refer to the Platform Guide for further details.


Example: 6 Node Cluster

The diagram illustrates an example of a 6 Node Cluster:




Example: 2 Web Proxy Nodes in a DMZ

The diagram illustrates an example of 2 Web Proxy Nodes in a DMZ:

**Example: 4 Web Proxy Nodes in a DMZ**

The diagram illustrates an example of 4 Web Proxy Nodes in a DMZ (2 admin, 2 Self-service):



2.1.5. Cloud deployments

VOSS Automate supports the following Cloud deployments:

- Microsoft Azure
- Amazon Web Services (AWS)

Although Google Cloud Platform (GCP) is not officially supported, contact us to discuss your requirements.


The advantages of a Cloud deployment topology:

- Leverage cloud tooling, such as proxies (which can be used instead of VOSS Web Proxy)

2.1.6. VOSS Automate Cloudv1 (SaaS)

VOSS Automate Cloudv1 is a Software-as-a-Service (SaaS) offering hosted on a shared VOSS Automate instance within Microsoft Azure.

VOSS manages this instance, which seamlessly integrates with a customer's Unified Communications (UC) platform, Microsoft Exchange, Microsoft Active Directory, and third-party applications like ServiceNow and Identity Providers (IdPs) for Single Sign-On (SSO) authentication.



2.2. Multi Data Center Deployments

Multinode clusters can be deployed in both Active/Active or Active/Standby configurations. Active/Active configurations have all the Unified nodes enabled for transaction processing. Active/Standby configurations have only the Unified nodes in the primary data center (the data center containing the Unified node with the primary database) enabled for transaction processing.

In order to run an Active/Active configuration, the latency (RTT) between data centers must not exceed 20ms. For higher latencies an Active/Standby configuration must be used.

To switch to an Active/Standby configuration, do the following on all Unified nodes in the secondary data center (the data center *not* containing the Unified node with the primary database):

1. log into platform:

```
ssh platform@<ip_address>
```

2. set the worker count on the nodes in the secondary data center to zero:

```
voss workers 0
```

Setting the number of workers is persistent; in other words, this setting will still apply after upgrades and system restart.

For an Active/Standby configuration, the proxy server web weights should be set to the Unified nodes on the primary data center. This is done with the **web weight add** command.

The web weight specifies the routing and relative counts of the initial HTTP request from the Web Proxy to a Unified Node. The initial request could be a request such as a transaction, or for example a GET request. Consider web weights configuration as shown below:

```
172.0.0.158:443: 1
172.0.0.159:443: 1
172.0.0.161:443: 1
172.0.0.162:443: 1
173.0.0.163:443: 0
173.0.0.164:443: 0
```

This configuration means that the servers 173.0.0.163 and 173.0.0.164 serve as backup servers and requests are only routed to these if the other servers are not available. While the other servers are available, an equal number of requests are routed to them in a round-robin manner. Refer to the Best Practices Guide for more details on deployment models and web weight settings.

Consider an example where:

Note: The examples below show system command output where the Phone Based Registration application is running.

1. Primary data center has unified nodes with IP addresses and ports:

- 172.0.0.158:443
- 172.0.0.159:443
- 172.0.0.161:443
- 172.0.0.162:443

A unified node typically shows the following output for the **web weight list** command:

```
$ web weight list
Default service weights

voss-deviceapi:
  phoneservices:
    172.0.0.158:8412: 1
  selfservice:
    172.0.0.158:5000: 1
  voss-deviceapi:
    172.0.0.158:9902: 1
  voss-portal:
    172.0.0.158:6001: 1
```

2. Secondary data center has unified nodes with IP addresses and ports:

- 173.0.0.163:443
- 173.0.0.164:443

The defaults of the **web weight list** command run on the *proxy servers* is as below:

1. Primary data center proxy server:

```
$ web weight list
Default service weights

  phonebasedreg:
    phoneservices:
      172.0.0.158:443: 1
      172.0.0.159:443: 1
      172.0.0.161:443: 1
      172.0.0.162:443: 1
      173.0.0.163:443: 0
      173.0.0.164:443: 0
    voss-deviceapi:
      selfservice:
        172.0.0.158:443: 1
        172.0.0.159:443: 1
        172.0.0.161:443: 1
        172.0.0.162:443: 1
        173.0.0.163:443: 0
        173.0.0.164:443: 0
      voss-deviceapi:
        172.0.0.158:443: 1
        172.0.0.159:443: 1
        172.0.0.161:443: 1
        172.0.0.162:443: 1
        173.0.0.163:443: 0
        173.0.0.164:443: 0
```

2. Secondary data center proxy server:

```
$ web weight list
Default service weights

  phonebasedreg:
    phoneservices:
      172.0.0.158:443: 0
      172.0.0.159:443: 0
      172.0.0.161:443: 0
      172.0.0.162:443: 0
      173.0.0.163:443: 1
      173.0.0.164:443: 1
    voss-deviceapi:
      selfservice:
        172.0.0.158:443: 0
        172.0.0.159:443: 0
        172.0.0.161:443: 0
        172.0.0.162:443: 0
        173.0.0.163:443: 1
        173.0.0.164:443: 1
      voss-deviceapi:
        172.0.0.158:443: 0
        172.0.0.159:443: 0
        172.0.0.161:443: 0
```

(continues on next page)

(continued from previous page)

```
172.0.0.162:443: 0
173.0.0.163:443: 1
173.0.0.164:443: 1
```

In order to ensure that the secondary data center is configured for a Standby mode, change the web weights to show userweights as seen in the output on the *secondary* data center proxy server below:

```
$ web weight list
Default service weights

  phonebasedreg:
    phoneservices:
      172.0.0.158:443: 0
      172.0.0.159:443: 0
      172.0.0.161:443: 0
      172.0.0.162:443: 0
      173.0.0.163:443: 1
      173.0.0.164:443: 1
    voss-deviceapi:
      selfservice:
        172.0.0.158:443: 0
        172.0.0.159:443: 0
        172.0.0.161:443: 0
        172.0.0.162:443: 0
        173.0.0.163:443: 1
        173.0.0.164:443: 1
      voss-deviceapi:
        172.0.0.158:443: 0
        172.0.0.159:443: 0
        172.0.0.161:443: 0
        172.0.0.162:443: 0
        173.0.0.163:443: 1
        173.0.0.164:443: 1

Customized service weights
  userweights:
    172.0.0.158:443: 1
    172.0.0.159:443: 1
    172.0.0.161:443: 1
    172.0.0.162:443: 1
    173.0.0.163:443: 0
    173.0.0.164:443: 0
```

As can be seen, the load balancing web weights have been changed to the unified nodes on the *primary* data center.

2.3. Clustering Considerations

The cluster contains multiple nodes which can be contained in separate firewalled networks.

Network ports need to be opened on firewalls to allow inter-node communication – these are described in more detail in the Platform Guide.

All communication between nodes is encrypted.

Node type	Ports
WebProxy	22 (ssh & sftp), 80 (http), 161 & 162 (snmp), 443 & 8443 (https)
Unified	22 (ssh & sftp), 80 (http), 161 & 162 (snmp), 443 & 8443 (https), 27020 & 27030 (database)

- 22/ssh is used for remote administration
- 80 and 443 is used for the web server
- 161 and 162 are used for sending and receiving snmp
- 8443 is used for inter-cluster communication
- 27020 and 27030 is used for database queries and replication

2.4. Network Communications between Nodes within the Cluster

The cluster contains multiple nodes which can be contained in separate firewalled networks. Network ports need to be opened on firewalls to allow inter-node communication.

All communication between nodes is encrypted.

The following details are all based on the default settings. These can vary depending on the application setup and network design (such as NAT) of the solution, so may need adjustment accordingly. Where a dependant is noted, this is fully dependant on the configuration with no default.

These communications are all related to communications between application nodes within the cluster. There are a few different deployment models so the details below cover the different models and relevant ports. So review and implement according to the deployment model in use.

Note that Standalone is only a single node so this section is not relevant for that deployment model.

- Proxy to Proxy Node

This is relevant if the proxy node is present in the system.

Communication	Protocol	Port
Cluster Communications	HTTPS	TCP 8443 bi-directional

- Proxy to Unified/Application Node

This is relevant if the proxy node is present in the system.

Communication	Protocol	Port
User access	HTTPS	TCP 443
Cluster Communications	HTTPS	TCP 8443 bi-directional

- Unified Node to Unified node

This is relevant to the communications between the unified nodes (application and database combined). If the application and database nodes are split, then see the relevant application and database node details below. Database arbiters run on port 27030.

Communication	Protocol	Port
Database access	database	TCP 27020 and 27030 bi-directional
Cluster Communications	HTTPS	TCP 8443

- Application node to Application node

This is relevant to the communications between application nodes in the system. This is only relevant where the database node is separate from the application node (in other words, not Unified node).

Communication	Protocol	Port
Cluster communications	HTTPS	TCP 8443 bi-directional

- Application Node to Database node

This is relevant to the communications between the application node and the database node. This is relevant if the database node is separate from the application node. Database arbiters run on port 27030.

Communication	Protocol	Port
Database access	database	TCP 27020, 27030, 5432 bi-directional
Cluster Communications	HTTPS	TCP 8443

- Database Node to Database node

This is relevant to the communications between the application node and the database node. This is relevant if the database node is separate from the application node. Database arbiters run on port 27030.

Communication	Protocol	Port
Database access	database	TCP 27020 and 27030 bi-directional
Cluster Communications	HTTPS	TCP 8443

2.5. Network Communications External to the Cluster

The following details are all based on the default settings. These can vary depending on the application setup and network design (such as NAT) of the solution, so may need adjustment accordingly. Where a dependant is noted, this is fully dependant on the configuration with no default.

These communications are all related to communications with devices external to the cluster.

2.5.1. Outbound Communications to Devices from the Application/Unified Nodes

Communication	Protocol	Port
Cisco Unified Communications Manager (UCM)	HTTPS	TCP 8443
Cisco Unity Connection (CUXN)	HTTPS	TCP 443
Webex	HTTPS	TCP 443
LDAP directory	LDAP	TCP/UDP 389 and/or 636(TLS/SSL)
MS PowerShell Proxy Node	HTTPS	TCP 5986
Microsoft 365 (Graph API)	HTTPS	TCP 443
Zoom	HTTPS	TCP 443

2.5.2. Outbound to External Systems from the Proxy Node

Communication	Protocol	Network Protocol and Port
API Sync and Async responses	HTTPS	TCP 443
Northbound Notification messages	HTTPS	dependant
Microsoft Teams / Microsoft Exchange	HTTPS	443
VOSS Cloud Licensing Service	HTTP HTTPS	80 443

2.5.3. Outbound to External Systems from All Nodes

Communication	Protocol	Port
SNMP	SNMP	TCP/UDP 162
SFTP as required for backup destinations	SFTP	TCP 22
NTP	NTP	UDP 123

2.5.4. Inbound Communications From External Systems to the Proxy Node

Communication	Protocol	Port
Web Access	HTTPS	TCP 443
API Request	HTTPS	TCP 443

2.5.5. Inbound Communications to All Nodes

Communication	Protocol	Port
SSH and SFTP for management and files transfers	SFTP/SSH	TCP/UDP 22

2.5.6. On-line Help Links to External Documentation

In order to have access to the online help website URL, you may need to request that the website be made accessible by your network administrator.

2.6. Cluster Commands

Note: From release 21.1 onwards, a standalone topology is considered a single node cluster (cluster-of-one). This means that commands such as **cluster provision**, **cluster list**, **cluster status** and **cluster maintenance-mode** should be run on a standalone topology.

The following Command Line Interface console display shows the available commands for clustering.

```
cluster add <ip>                - add a new node to join the existing cluster
cluster check <verbose>         - Display pre-upgrade readiness.
                                For each node, test if ports
                                are available, the time taken to connect, drive space
                                percentage and lastly checks if NTP is running
cluster del <ip>                - remove a node from the existing cluster
cluster job kill <pid>          - Kill a detached job <pid>
cluster job list                 - List detached jobs in the cluster
cluster job reconnect <pid>     - Reconnect to a detached job <pid>
cluster list                     - display the list of nodes associated with the cluster
cluster maintenance-mode       - Display the status, start or stop maintenance
                                <start|stop|status> mode across the cluster
cluster prepnode                - Prepares the system so that it can be joined to a
↳ cluster                       or delta install scripts
cluster primary                 - Check if the system is considered the primary by
↳ patch
```

(continues on next page)

(continued from previous page)

```

cluster primary role          - Check if the system is considered the
    <application|database>    'primary' in the given role
cluster provision             - perform cluster-wide provisioning
    [datacentre <location>]
    [role <role>]
cluster run <where> <command> - run the command on a particular host.
                                <where> can either be a name prefix, ip, role, or
                                ↪ 'all'
cluster status                - display the status of the cluster
cluster upgrade <iso/url>     - upgrade all applications from iso image <iso-name>.
    [datacentre <location>]    <iso-name> can be a URL for upgrading
    [backup <location>]        from a remote server.
cluster where <application>   - determine on which nodes the application is installed

```

2.7. Geo-redundancy/Redundancy And DR Synopsis

High Availability (HA) is an approach to IT system design and configuration that ensures VOSS Automate is operational and accessible during a specified time frame. This is achieved using redundant hardware and resources. If there is a failure, an automatic failover will occur to a second node.

Web server proxy nodes perform load-balancing between application roles, so that load is distributed. During provisioning, the web server proxy is provided with all the IP addresses of the application nodes. The web server software then does load balancing among these nodes, according to its configuration. If a node fails to respond in a set time, the proxy will send the transaction to another node. This means that in the event that an Application role is lost, the WebProxy will transparently bypass the faulty Application role.

The proxy web server that is configured to be located in the primary site normally load balances to the two unified nodes in the primary site. The proxy web server falls back to the two nodes in the Disaster Recovery site if the nodes in the primary site are down. The web proxy nodes in the secondary site defaults load balancing to the two unified nodes configured for the secondary site.

Data is replicated between Database roles, and role failure is recoverable. This is done using the database replication facility. Automatic failover between Database roles occurs while there is greater than 50% Database role availability. Once there is insufficient role availability, the system needs to be manually re-provisioned.

HA can be increased by adding nodes to the cluster. Application performance and availability can be increased by adding additional application role servers.

Backups can be scheduled to run automatically across the cluster. Backups include application data, configuration and software. Backups can take place to both local disk and remote network location. Every node upgrade is preceded by a snapshot backup which allows any upgrade to be rolled back. Refer to the Platform Guide for details.

3. Prepare to Install

Note: From release 21.4 onwards, VOSS Automate allows for the registration and update of product licenses within the application. A licensing service is installed during installation or upgrade and a license token is associated with the platform on which it is installed.

3.1. Network Docker Container Range

Important: When troubleshooting network issues, verify that the address range is not in use.

If it is in use, modify the Private Address Space using the **network container range add <private IP>** command as described below.

RFC-1918 states that the following three blocks of the IP address space are reserved for private internets:

10.0.0.0	-	10.255.255.255	(10/8 prefix)
172.16.0.0	-	172.31.255.255	(172.16/12 prefix)
192.168.0.0	-	192.168.255.255	(192.168/16 prefix)

This subnet block address range can be modified to another Private Address Space if needed.

Use the command **network container range list** to see the current Private Address Space.

For example:

```
$ network container range list
range: 10.1.2.1/24
```

Use the command **network container range add <private IP>** to modify the Private Address Space.

Important: A valid Private Address IP is required as input.

The range /24 is appended to the IP. For example, if 192.168.0.6 is used, the Private Address range 192.168.0.0/24 is used.

In a clustered environment, you could use **cluster run all network container range add <private IP>**, but if required, the Private Address Space can be also set to be different on each node by running the **add** command on each individual node.

For example:

```
$ network container range add 192.168.2.3
You are about to restart all services. Do you wish to continue?y
Application processes stopped. (note this line changes dynamically)

Reconfiguring applications....
Application processes started. (note this line changes dynamically)
```

3.2. Backup Size Considerations

The default backup partition size is 50GB for the default 250GB database partition size. These are the default partition installation sizes.

Important: Backups should be created and restored in a **screen** session - see: [Using the screen command](#).

To determine the required space for a specific backup partition, carry out and consider the following:

1. Run **backup create <location-name>** from the CLI. The command output indicates the required space needed to do the backup.

If the current backup partition size is too small, the command will fail and suggest the size of the partition required. The size can be increased in size with the **drives add** command.

See the Drive Control topic in the Platform Guide.

If there is sufficient space but only a size check is required, the backup command can be canceled (Ctrl-C), if needed.

2. Run **voss db_collection_stats all** to show the size of the current database. This command validates the size of the database. This total will be smaller than the suggested backup size.

- A local backup requires free space of at least twice the size of the database. Preferably add another 30% of this. At any given time it is therefore recommended that available backup space should be approximately twice the database size plus 30% of the database size.

For remote backups, the backup requires free space of at least the size of the database plus 30% of this.

- Database growth over time needs to be considered and allowed for in the backup partition size.
- Space for multiple local backups also needs to be considered and added to the calculated backup partition size.

3.3. Topology Hardware Specifications

3.3.1. VOSS Automate Hardware Specifications

Overview

Note: For details around the open source software components used in Automate, see the *Open Source License Usage Guide*.

Virtualized hardware and resource oversubscription

It is recommended that no more than two Unified nodes and one Web Proxy node be run on a physical server (VMware server) and that the disk subsystems are unique for each Unified node.

VOSS Automate virtual machines should maintain a 1:1 ratio between virtual RAM and Disk hardware and physical hardware, in other words:

- 1 GB of virtual RAM (vRAM) must map to 1 GB of physical RAM
- 1 GB of virtual Disk (vDisk) storage must map to 1 GB of physical storage

For virtual CPU (vCPU), hyper-threading is supported.

Unified node hardware specifications

Single-node cluster (cluster-of-one) hardware specification

This section provides the virtual machine specification for a single node cluster deployment topology in VOSS Automate.

Node type	Quantity	VM	Memory	CPU	Disk	Network
Single node cluster	1	At least VMware 11, compatible with ESXi 6.0 and up	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	Total: 440 GB as allocated below. 370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40 GB for our apps • 50 GB for compressed backups • 250 GB for database 70 GB database disk to be added after upgrade or installation. Refer to the Upgrade Guide with ISO and Template or Installation Guide.	1 Gbit/s minimum

For Memory and CPU, the Resource Allocation Reservation on VMware is indicated in the table. Hyper-threading is supported.

Note:

- From release 24.1, allowance should be made for an additional 70GB database disk to be added after upgrade or installation. This disk will be used for Insights sync.

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Add disks in the AWS or MS Azure cloud hosted platform](#).

- If memory allocations are customized, ensure that the memory reservation remains equal to the allocated memory in order to prevent possible negative side-effects due to memory reclamation.

For VMware details, refer to the VMware Support topic in the Compatibility Matrix.

The maximum number of users for a single node cluster is 50,000.

Multinode cluster hardware specification

Virtual machine requirements are specified in the table below.

Node type	Quantity	VM	Memory	CPU	Disk	Network
Unified	4 or 6	At least VMware 11, compatible with ESXi 6.0 and up	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	Total: 440 GB as allocated below. 370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40 GB for our apps • 50 GB for compressed backups • 250 GB for database 70 GB database disk to be added after upgrade or installation. Refer to the Upgrade Guide with ISO and Template or Installation Guide.	1 Gbit/s minimum
WebProxy	2	At least VMware 11, compatible with ESXi 6.0 and up	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

For Memory and CPU, the Resource Allocation Reservation on VMware is indicated in the table. Hyper-threading is supported.

Note:

- From release 24.1, allowance should be made for an additional 70GB database disk to be added after upgrade or installation. This disk will be used for Insights sync.

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Add disks in the AWS or MS Azure cloud hosted platform](#).

- If memory allocations are customized, ensure that the memory reservation remains equal to the allocated memory in order to prevent possible negative side-effects due to memory reclamation.

The OS disk requirement is fixed and logs are rotated to ensure that 10 GB is sufficient. 40 GB for applications is a generous allocation and does not scale with the number of users.

The Database storage partition is sized to support 250 K users. Database backups are compressed and the partition is sized to ensure that sufficient space available to support backup of 250 GB database.

Note: To change the TRANSACTION_LOG cap size to greater than 10GB at larger providers for operational reasons, for example for diagnostics and a longer transaction replay window, the **voss db_collection_cap TRANSACTION_LOG <10-50GB>** command can be used from the command line.

Refer to Database Commands in the Platform Guide for more details.

The resize operation will impact the usage on the disk size allocated for the database (typically, 250GB is reserved upon installation). Consider a larger database disk size allocation upon installation if a larger cap size is set.

The backup disk should be Thick Provisioned and Eager Zeroed for better performance immediately after installation.

Web Proxies are optional, but if Web Proxies are used, then they form part of the cluster to allow sharing of static data and other content as needed (for example, themes).

To set up the disk requirements, the disk should be set up on the VMware GUI Resources tab where a disk can be created. This task should be done after the OVA import but prior to the boot of the system.

For VMware details, refer to the VMware Support topic in the Compatibility Matrix.

2 Node cluster hardware specification

Virtual machine requirements are specified in the table below.

Node type	Quantity	VM	Memory	CPU	Disk	Network
Unified	= 2	At least VMware 11, compatible with ESXi 6.0 and up	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	Total: 440 GB as allocated below. 370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40 GB for our apps • 50 GB for compressed backups • 250 GB for database 70 GB database disk to be added after upgrade or installation. Refer to the Upgrade Guide with ISO and Template or Installation Guide.	1 Gbit/s minimum
WebProxy	>= 0	At least VMware 11, compatible with ESXi 6.0 and up	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

For Memory and CPU, the Resource Allocation Reservation on VMware should correspond with these requirements.

Note:

- From release 24.1, allowance should be made for an additional 70GB database disk to be added after upgrade or installation. This disk will be used for Insights sync.

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Add disks in the AWS or MS Azure cloud hosted platform](#).

- If memory allocations are customized, ensure that the memory reservation remains equal to the allocated memory in order to prevent possible negative side-effects due to memory reclamation.

For VMware details, refer to the VMware Support topic in the Compatibility Matrix.

Modular cluster hardware specifications**Multinode modular cluster hardware specification**

Virtual machine requirements are specified in the table below.

Node type	Quantity	VM	Memory	CPU	Disk	Network
Application	3	At least VMware 11, compatible with ESXi 6.0 and up	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40 GB for our apps 	1 Gbit/s minimum
Database	3	At least VMware 11, compatible with ESXi 6.0 and up	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	Total: 440 GB as allocated below. 370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for compressed backups • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40 GB for our apps • 250 GB for database 70 GB database disk to be added after upgrade or installation. Refer to the Upgrade Guide with ISO and Template or Installation Guide.	1 Gbit/s minimum
WebProxy	2	At least VMware 11, compatible with ESXi 6.0 and up	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

For Memory and CPU, the Resource Allocation Reservation on VMware is indicated in the table. Hyper-threading is supported.

Note:

- From release 24.1, allowance should be made for an additional 250GB database disk to be added after upgrade or installation. This disk will be used for Insights sync.

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Add disks in the AWS or MS Azure cloud hosted platform](#).

- If memory allocations are customized, ensure that the memory reservation remains equal to the allocated memory in order to prevent possible negative side-effects due to memory reclamation.

The OS disk requirement is fixed and logs are rotated to ensure that 10 GB is sufficient. 40 GB for an applications role node is a generous allocation and the size will not have to be increased with the number of users.

The Database storage partition is sized to support 250 K users. Database backups are compressed and the partition is sized to ensure that sufficient space is available to support backup of 250 GB database.

Note: To change the TRANSACTION_LOG cap size to greater than 10GB at larger providers for operational reasons, for example for diagnostics, the **voss db_collection_cap TRANSACTION_LOG <10-50GB>** command can be used from the command line.

Refer to Database Commands in the Platform Guide for more details.

The resize operation will impact the usage on the size of the disk allocated for the database (typically, 250GB is reserved upon installation). Consider a larger database disk size allocation upon installation if a larger cap size is set.

The backup disk should be Thick Provisioned and Eager Zeroed for better performance immediately after installation.

Web Proxies are optional, but if Web Proxies are used, then they form part of the cluster to allow sharing of static data and other content as needed (for example, themes).

To set up the disk requirements, the disk should be set up on the VMware GUI Resources tab where a disk can be created. This task should be done after the OVA import but prior to the boot of the system.

For VMware details, refer to the VMware Support topic in the Compatibility Matrix.

Add disks in the AWS or MS Azure cloud hosted platform

The steps below are required to add a disk that provides for the Insights database in release 24.1 - that should then be assigned to the `insights-voss-sync:database` mount point (refer to the final step in the *Upgrade Guide with ISO and Template* for your topology).

AWS

1. Create the EBS Volumes for each DB node in the Amazon EC2 console.

Go to **EC2 > Volumes > Create volume**

For **Volume settings**, enter:

- Volume type: Provisioned IOPS SSD (io2)
- Size (GiB): 70GB
- IOPS: 750

For **Availability Zone**:

- Create 3 volumes in each of the zones (for example: us-east-1a, us-east-1b, us-east-1c)

2. Attach the newly created volumes to each of the database nodes.

Go to **EC2 > Volumes > volume_id > Attach volume**

- **Instance:** Select the database instance within the same corresponding az
- **Device Name:** /dev/sde (This will display as xvde in drives list)

Microsoft Azure

1. In the Microsoft Azure portal, search for Virtual Machines
 - Select each of the database nodes
 - Select **Disk** under **Properties**
2. Click **Create and attach a new disk**.
 - **LUN:** Next available
 - **Disk Name:** Label according to your recommended naming convention
 - **Storage Type:** Premium SSD LRS
 - **Size:** 70GB
 - **Encryption:** Set according to your requirements
 - **Host Caching:** Read/Write

4. Install VOSS Automate

4.1. Installation Process

The installation process is divided into:

1. The VMWare installation of a node (*Create a New VM Using the Platform-Install OVA*).
2. Node setup for single node or multinode installations (*Multinode Installation* and *Single-node cluster (cluster-of-one) Installation*).

The node setup stage requires one or more prerequisite VMWare node installations.

4.2. Create a New VM Using the Platform-Install OVA

Note: If an OVA file is not available for your current release, you need to obtain the most recent release OVA for which there is an upgrade path to your release.

The steps below show the common setup of a *single node* from the OVA file - either for the purposes of:

- a standalone installation

If an OVA file is not available for your current release:

1. Obtain and install the most recent release OVA for which there is an upgrade path to your release.
2. Apply the Delta Bundle Upgrade steps for the current release to the OVA to upgrade it.

- a node installation during multinode installation - see *Notes on Multi-Node Installation*

If an OVA file is not available for your current release:

1. Obtain and install the most recent release OVA for which there is an upgrade path to your release.
2. Apply the Delta Bundle Upgrade steps for the current release *to the cluster* to upgrade it. Refer to the *Upgrade Guide with Delta Bundle*.

- or during a failover recovery

If an OVA file is not available for your current release:

1. Obtain and install the most recent release OVA for which there is an upgrade path to your release.
2. Add it to your cluster. Use the same configure options in the table below as were applied to the lost node.

Note that the node version mismatch in the cluster can be ignored, since the next upgrade step aligns the versions.

3. Apply the Delta Bundle Upgrade steps for the current release *to the cluster* to upgrade it.

For details, refer to the *Upgrade Guide with Delta Bundle* and to the specific scenario Disaster Recovery steps in the *Platform Guide*.

The steps will therefore be followed either once or multiple times during installation - in accordance with the required topology.

The downloaded OVA file is imported into VMware vCenter Server. Only one OVA file is used to deploy all the functional roles. You choose the specific node `role` when the installation wizard is run.

1. Log in to vSphere to access the ESXi Host.
2. Choose **File > Deploy OVF Template**.
3. Choose Source, browse to the location of the .ova file, and click **Next**.
4. On the Name and Location page, enter a Name for this server.
5. On the Deployment Configuration page, select the appropriate node type. (Refer to the list at the `role` option below.)
6. Choose the resource pool in which to locate the VM.
7. Choose the data store you want to use to deploy the new VM.
8. Choose the disk format to use when deploying the new VM.
 - For non-SSD-based drives in production environments, “thick provisioning” is mandatory. Thick Provision Eager Zeroed is recommended.
 - For SSD-based drives, “thin provisioning” is supported.
9. On the Network Mapping, choose your network on which this VM will reside.
10. Do not select Power on after deployment.
11. On the Ready to Complete page, click **Finish** to start the deployment.
12. After the VM is created, select the CD ROM configuration and verify the **Connect at Power On** check box is enabled. Also, verify the memory, CPU, and disk settings against the requirements shown in either the Single-node cluster (cluster-of-one) System Hardware Specification or Multinode Cluster Hardware Specification section in the Install Guide.
13. Power on the VM.
14. Configure the options in the installation wizard:

Option	Option name	Description
1	network device	The network device name.
2	IP	The IP address of the server. The required format is with Classless Inter-Domain Routing (CIDR): ip/netmask. ¹
3	gateway	The IP address of the network gateway. ^{Page 37, 1}
4	DNS	The DNS server is optional. Ensure that the DNS server is capable of looking up all hostnames referred to, including NTP server and remote backup locations. ¹
5	NTP	The NTP server is mandatory to ensure that time keeping is accurate and synchronized among nodes in the same cluster. ¹
6	boot password	Enable boot loader configuration password. See the example below.
7	hostname	The hostname, not the fully qualified domain name (FQDN). The maximum character length for the hostname is 56.
8	role	<p>Note: only WebProxy, Application and Database nodes are used for a modular architecture installation.</p> <ul style="list-style-type: none"> • A WebProxy role installs only the front-end web server together with ability to distribute load among multiple middleware nodes. • An Application node is the main transaction processing engine and includes a web server which can operate by itself, or route transactions from a web node. • A Database node provides persistent storage of data. • A Standalone node consists of the Web, Application, and Database roles on one node. For Single-node cluster (cluster-of-one). • A Unified node consists of the Web, Application, and Database roles on one node. On installation, the system needs to be clustered with other nodes and the cluster provisioned. • A General node used for M2UC, NBI.
9	data center	The system's geographic location (data center name, city, country that a customer can use to identify the system location). You cannot change this setting once set.
10	platform password	Platform password must be at least eight characters long and must contain both uppercase and lowercase letters and at least one numeric or special character.

¹ VOSS Automate supports IPv4 or IPv6
IPv6 allows the following input formats to be used:

- IPv6 Compressed, e.g.: ::ffff:c0a8:6403/64
- IPv6 Expanded (Shortened), e.g.: 0:0:0:0:ffff:c0a8:6403/64
- IPv6 Expanded, e.g.: 0000:0000:0000:0000:ffff:c0a8:6403/64

From Automate release 24.1 onwards, network addresses are in CIDR (Classless Inter-Domain Routing) format, for example: 192.168.100.3/27 or e00d::fafe:23/112. The use of a netmask in the 255.255.255.0 format is no longer supported.

Note: On a fresh installation, if you run the install on a network with a DHCP server and encounter an error: “Error: DNS server <DNS server> is either invalid or cannot be reached on the network” you can enter a valid DNS server address to continue the installation.

Once all details are entered, installation will commence. When installation is complete, the system will reboot. Since all services will be stopped, this takes some time.

4.2.1. Notes on Passwords and Security

The default security protocol for the web server is TLSv1.2.

Password protection can be enabled on the VOSS Automate boot loader configuration from the install wizard upon first install and also from the CLI - see the topic on System Boot Passwords in the Platform Guide for commands to enable, disable or reset the boot password.

Important: The boot password is non-recoverable.

The console example below shows the boot password configuration output:

```
(1)          ip      (199.29.21.89)
(2)          netmask (255.255.255.0)
(3)          gateway (199.29.21.1)
(4)          dns     (199.29.88.56)
(5)          ntp     (199.29.88.56)
(6)  boot password (disabled)
(7)          hostname (atlantic)
(8)          role     (UNDEFINED)
(9)          data centre (earth)
(10) platform password (UNDEFINED)
Select option ? 6
Valid passwords must contain:
    at least one lower- and one upper-case letter,
    at least one numeric digit
    and a special character eg. !#$%&^*
Password: Please enter platform user password:
Please re-enter password
Password:
NOTE: The system boot password is now set for user platform.
```

When the boot password is set, the wizard will show:

```
(6)  boot password  (*****)
```

4.2.2. Notes on Multi-Node Installation

According to the multi-node deployment topology and specification, the *role* of each VM installation is as indicated below.

- For each WebProxy instance, create a new VM using the platform-install OVA. For *role*, select **(3) WebProxy**. Specify the appropriate data center (Primary/DR site) for each WebProxy instance.
- *Standard Topology only*: For each Unified instance, create a new VM using the platform-install OVA. For *role*, select **(2) Unified**. Specify the appropriate data center (Primary/DR Site) for each Unified instance.

The following Unified nodes are required in the cluster:

- One Unified node as the Primary node at the Primary site
- One Unified node as the Secondary node at the Primary site
- Two Unified nodes as the Secondary nodes at the Disaster Recovery (DR) site

Note:

- For a six Node Multi Cluster deployment there are; two Unified nodes (one Primary and one Secondary) and one WebProxy node at the Primary site, and two Unified nodes (both Secondary) and one WebProxy node at the DR site.
- For an eight Node Multi Cluster deployment, there are four Unified nodes (one Primary and three Secondary) and one WebProxy node at the Primary site. Two Unified nodes (both Secondary) and one WebProxy node are at the DR site.
- *Modular Architecture Topology*

The following nodes are required in a typical Modular Architecture cluster:

- One Application node as the Primary node at the Primary site
- One additional Application node at the Primary site
- One Database node as the Primary Database node at the Primary site
- One additional Database node at the Primary site
- One Application node at the Disaster Recovery (DR) site
- One Database node at the Disaster Recovery (DR) site

Note: For a typical Modular Architecture cluster there is one WebProxy node at the Primary site and one WebProxy node at the DR site.

For each *Database* instance, create a new VM using the platform-install OVA. For *role*, select **(2) Database**. Specify the appropriate data center (Primary/DR Site) for each Database instance.

For each *Application* instance, create a new VM using the platform-install OVA. For *role*, select **(2) Application**. Specify the appropriate data center (Primary/DR Site) for each Application instance.

Also refer to Multinode Installation section in the Install Guide.

Detailed configuration can be applied from the Command Line Interface (CLI). Use **network help** or **network** for details. For example, domain can be configured using **network domain add <domain-name>**. For a geo-redundant deployment, the data center information entered in the wizard is equivalent to the location information.

4.2.3. Finalize the Installation

When the installation of the OVA is complete, a sign-in prompt for the platform user is displayed. The system is ready for use.

Connect to newly deployed server CLI as the platform user.

The login message would for example looks the same as below:

```
Last login: Wed Nov  2 11:12:45 UTC 2016 from thwh on pts/6
Last failed login: Wed Nov  2 11:19:53 UTC 2016 from iza on ssh:notty
There were 2 failed login attempts since the last successful login.

host: dev-test, role: webproxy,application,database, load: 0.21, USERS: 3
date: 2016-11-02 11:19:57 +00:00, up: 14:19
network: 172.29.253.14, ntp: 172.29.1.15
HEALTH: NOT MONITORED
database: 31Gb
Failed logins: 2 since Wed Nov 02 11:19:53 2016 from iza

    mail - local mail management          keys - ssh/sftp credentials
    network - network management          backup - manage backups
    voss - voss management tools           log - manage system logs
    database - database management         notify - notifications control
    schedule - scheduling commands         selfservice - selfservice management
    diag - system diagnostic tools         system - system administration
    snmp - snmp configuration             user - manage users
    cluster - cluster management           drives - manage disk drives
    web - web server management            app - manage applications
    security - security update tools
```

If the user failed to log in prior to a successful login, the count, date and origin of the attempts are shown as Failed logins. A successful login resets this login count.

Run **app status** on all application nodes and ensure the services are all running and reporting the correct version before continuing.

Note: Return to Multinode Installation, Standalone Installation or Failover step to complete the overall installation or failover recovery procedure.

4.3. Using the screen command

The **screen** command is available to execute long-running commands (for example, when upgrading) in the background.

The following commands require the running of **screen**:

- **cluster provision**
- **cluster upgrade**
- **app template**

- **voss export type <args>**
- **voss export group <args>**
- **voss subscriber_data_export**

A message is displayed to indicate that **screen** should be run first:

This is a potentially long-running command and should be executed in a screen session
Run ``screen`` and then execute the command again

The use of **screen** is *not affected* by the use of the `--force` parameter with any of these commands.

The commands then run in a screen session that can be reconnected. The standard screen command parameters are available, in particular:

- **screen** - start a new session
- **screen -ls** - show sessions already available
- **screen -r [screen PID]** - reconnect to a disconnected session

The version of **screen** used in VOSS Automate also supports the creation of a log file. If long-running commands will be run, the log file captures screen console output up to the session timeout. A message shows:

timed out waiting **for input**: auto-logout

To create a screen log file:

1. Run **screen** and wait for screen to open.
2. Press **<Ctrl>-a** then **:** (colon). This will enter screen command mode at the bottom of the console.
3. Create your screen logfile in the `media/` directory:
 - a. In screen command mode, type **logfile media/<screen-logfilename>.log**
 - b. Press **<Enter>**
 - c. Press **<Ctrl>-a** and then **H** to start writing to the log file
 - d. Run your commands.

If the **screen** session times out, you can obtain console output from the log file, for example:

```
$ sftp platform@<host>:media/<screen-logfilename>.log
```

4.4. Installation Logs

To troubleshoot an installation, log files can be inspected. For example, detailed platform commands show in the `execute.log` file. Here, log entries for the command execution have a `ui` column. Log entries that follow these show related commands.

To only see the commands in `execute.log` example, you can open a new console and run:

```
log follow execute.log | grep " ui "
```

Note: Logs are rotated and install commands may not show after log rotation.

The list below shows examples of installation commands and corresponding ui and following entries in execute.log.

- **app install vmware.**

execute.log:

```
<timestamp><user><execnum>: ui - /opt/platform/bin/scripts.py install 'vmware'
```

- **app list.**

execute.log:

```
<timestamp><user><execnum>: ui - /opt/platform/bin/execute get /scripts/
<timestamp><user><execnum>: run: /opt/platform/bin/execute get /scripts/
<timestamp><user><execnum>: run: /opt/platform/bin/scripts.py list
```

- **database config**

execute.log:

```
<timestamp><user><execnum>: ui - /opt/platform/apps/mongodb/bin/database_
↳helper.py config
<timestamp><user><execnum>: run: /opt/platform/apps/mongodb/bin/database_
↳helper.py config
<timestamp><user><execnum>: run: /opt/platform/apps/mongodb/bin/database_
↳helper.py config returned 0
```

- **cluster list.**

execute.log:

```
<timestamp><user><execnum>: ui - /opt/platform/bin/execute get /apps/cluster/engine/
↳list
<timestamp><user><execnum>: run: /opt/platform/bin/execute get /apps/cluster/engine/
↳list
<timestamp><user><execnum>: run: /opt/platform/bin/execute --app=cluster get /list
<timestamp><user><execnum>: run: /opt/platform/apps/cluster/cluster.py list
<timestamp><user><execnum>: run: /opt/platform/apps/cluster/cluster.py list returned
↳0
```

- **cluster status.**

execute.log:

```
<timestamp><user><execnum>: ui - /opt/platform/bin/execute get /apps/cluster/
↳engine/status
<timestamp><user><execnum>: run: /opt/platform/bin/execute get /apps/cluster/
↳engine/status
<timestamp><user><execnum>: run: /opt/platform/bin/execute --app=cluster
↳get /status
<timestamp><user><execnum>: run: /opt/platform/apps/cluster/cluster.py status
```

- **web service list**

execute.log:

```

<timestamp><user><execnum>: ui - /opt/platform/bin/execute get /apps/nginx/
↳ engine/disable
<timestamp><user><execnum>: run: /opt/platform/bin/execute get /apps/nginx/
↳ engine/disable
<timestamp><user><execnum>: run: /opt/platform/bin/execute --app=nginx get
↳ /disable
<timestamp><user><execnum>: run: /opt/platform/bin/config.py get --app=nginx
↳ disable
<timestamp><user><execnum>: run: /opt/platform/bin/config.py get --app=nginx
↳ disable returned 0

```

- **log follow upgrade_db.log**

execute.log:

```

<timestamp><user><execnum>: ui - /opt/platform/bin/execute post /apps/services/
↳ process/log/engine/log/follow '{"follow":"upgrade_db.log"}'
<timestamp><user><execnum>: run: /opt/platform/bin/execute post /apps/services/
↳ process/log/engine/log/follow '{"follow":"upgrade_db.log"}' --method=os.system
<timestamp><user><execnum>: run: /opt/platform/bin/execute --app=services:log post
↳ /log/follow '{"follow":"upgrade_db.log"}' --method=os.system
<timestamp><user><execnum>: run: /opt/platform/apps/services/logviewer.sh follow
↳ upgrade_db.log

```

- **app template media/<VOSS Automate_template_file>**

execute.log:

```

<timestamp><user><execnum> ui - /opt/platform/apps/template_runner/template media/
↳ install.template platform
[...]
<timestamp><user><execnum> ui - /opt/platform/bin/execute --app=template_runner post
↳ /methods/import
  '{"filename":"DummyTestImport.json","import":"DummyTestImport.json"}'

Please enter a password for ...

<timestamp><user><execnum> ui - /usr/bin/docker exec -it voss-wsgi /opt/voss-
↳ deviceapi/bin/python
  /opt/voss-deviceapi/src/deviceapi/utils/get_user_password.py set_details
↳ sysadmin@sys

[...]

<timestamp><user><execnum>: ui - /opt/platform/bin/execute --app=template_runner
↳ post /methods/import
  '{"filename":"UpgradeChecks.json","import":"UpgradeChecks.json -p sys"}'

[...]

  '{"filename":"EndToEnd.application.json","import":"EndToEnd.application.json -p sys"}'
↳ '

```

(continues on next page)

(continued from previous page)

```
'{"filename":"SYS.json","import":"SYS.json -p sys"}'
'{"filename":"SYSnoPKG.json","import":"SYSnoPKG.json -p sys"}'

[...]
```

4.5. Commands to Determine the Primary Node Role in a Cluster

Commands are available to determine if a node is the primary database or application node.

From the CLI on a node, the **cluster primary** command can be run with additional parameters to determine if the node is the primary database or application node.

```
platform@VOSS-UN-1:~$ cluster primary
is_primary: true
platform@VOSS-UN-1:~$ cluster primary role application
is_primary: true
platform@VOSS-UN-1:~$ cluster primary role database
is_primary: false
platform@VOSS-UN-1:~$
```

This command should be used to establish and confirm the primary application node during patch, patch bundle or delta bundle installation.

4.6. Unified Node Topology Installation

4.6.1. Multinode Installation

Before You Begin

Before continuing, you should have followed the OVA installation on each node according to the steps and preliminary requirements specified in: [Create a New VM Using the Platform-Install OVA](#) and according to the node roles as indicated in [Notes on Multi-Node Installation](#).

Optionally download or extract language pack template files to support languages other than English.

Note:

- For geo-redundant Multinode Cluster deployment with six Unified Nodes, there are four Unified nodes in the Primary Site and two Unified nodes in the Disaster Recovery (DR) Site in active-standby setup. The worker count (**voss workers** command) needs to be set on the DR nodes. Refer to:
 - [Multinode cluster with unified nodes](#)
 - [Multi Data Center Deployments](#)
- For 2 node cluster deployment there are 2 unified nodes.
- Template installation and upgrade takes approximately two hours. You can follow the progress on the Admin Portal transaction list.

- It is strongly recommended *not* to allow customer end-users the same level of administrator access as the restricted groups of provider- and customer administrators. This is why Self-service web proxies as well as Administrator web proxies should be used.

Systems with Self-service only web proxies are *only* recommended where the system is customer facing, but where the customer does not administer the system themselves.

- For cluster installations, also refer to the Health Checks for Cluster Installations Guide.
- If it is necessary to change an IP address of a node in a cluster, first remove it from the cluster by running the command below *on the node to be changed*:

cluster del <IP address of node to be changed>

- Refer to [Installation Logs](#) for troubleshooting logs during an installation.
- Before installing release 24.2, ensure that an additional 70 GB disk has been made available for the Insights database.

See the Adding Hard Disk Space topic in the Platform Guide and VOSS Automate Hardware Specifications in this guide.

This disk is needed to assign to the `insights-voss-sync:database` mount point. See the final installation step below.

The standard **screen** command should be used where indicated. See: [Using the screen command](#).

Installation

1. Install VMware tools on each node.
 - a. Log in to each node and run **app install vmware**.
 - b. Verify that vmware is running: **app list**.
2. Prepare each node to be added to the cluster:
 - a. Select a primary Unified node that will become the primary database node. The designation of primary unified node is arbitrary. The deploying administrator can pick any unified node that they see fit.
 - b. On each WebProxy and Unified node, *excluding the primary node*, run **cluster prenode**.
3. Add nodes to the cluster.
 - a. Log in to the selected primary Unified node.
 - b. Add the Unified and WebProxy nodes to the cluster: **cluster add <ip_addr>**.

Note that you do not have to add the selected primary node to the cluster. It will automatically be added to the cluster.
 - c. Verify the list of nodes in the cluster: **cluster list**.
4. Add the network domain (optional if a domain name is needed). From the selected primary Unified node:
 - a. Configure the domain: **cluster run all network domain <domain_name>**.
 - b. Verify the configured network domain: **cluster run all network domain**. Each node shows the domain that you configured.
5. Check the network:

- a. From the selected primary Unified node, run **cluster check** to verify the status of the cluster, network connectivity, disk status and NTP.

Note:

- Since database weights have not been added yet, database: not configured errors can be ignored. Verification of database weights should be done when this command is run during the step following the provisioning step.
- If a cluster has not been provisioned, port 443 errors from web proxies are raised upon this check. These errors can be ignored.
- If errors on other ports (for example 27020) are raised by the check, verify that the firewall service has started. The following command can be run:

```
cluster run database app start services:firewall --force
```

- b. Verify the DNS configuration: **cluster run all network dns**. Each node responds with the DNS server address.
6. Create a restore point. As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
7. Configure the cluster.

- a. Provide a weight for each database server with the **database weight add <database_ip> <priority>** command.
 - Weights of 40, 30 are recommended for two Unified nodes
 - Weights of 40, 30, 20, and 10 are recommended for four Unified nodes
 - Weights of 60, 50, 40, 30, 20, and 10 are recommended for six Unified nodes

The higher the value, the more priority.

For Multinode Cluster deployment with four Unified Nodes in a geo-redundant system containing two data center infrastructures in two physical locations the following weights are used:

- Specify a weight of 40 for the Primary node at the Primary site
- Specify a weight of 30 for the Secondary node at the Primary site
- Specify weights of 20 and 10 for the Secondary nodes at the DR site

For Multinode Cluster deployment with six Unified Nodes in a geo-redundant system containing two data center infrastructures in two physical locations the following weights are used:

- Specify a weight of 60 for the Primary node at the Primary site
- Specify a weight of 50 for the Secondary node at the Primary site
- Specify a weight of 40 for the Secondary node at the Primary site
- Specify a weight of 30 for the Secondary node at the Primary site
- Specify weights of 20 and 10 for the Secondary nodes at the DR site

- b. From the selected primary Unified node, now set it up as the primary Unified node. It is recommended that this step is run in a terminal opened with the **screen** command.
 - i. Run **screen**.

ii. Run **cluster provision**

Allow approximately 2 hours for the operation to complete for two WebProxy and four Unified nodes.

- c. When provisioning is complete, check that each node is contactable and that the time server is running on each with **cluster check**.

If a service is down, run **cluster run <node_ip> app start** to restart the service.

If provisioning is successful, the screen session can be closed by typing **exit** in the screen terminal. If errors occurred, keep the screen terminal open for troubleshooting purposes and contact VOSS support.

- d. (Optional) If required, set the web weights configurations (Active-Active, Active-Standby, Single node cluster). From the primary Unified node, run the required **web weight** commands for the Web Proxy nodes. For details, refer to [Multi Data Center Deployments](#) and the VOSS Automate Best Practices Guide.
- e. (Optional) If required, enable or disable Self-service or admin web services on the web proxy nodes. This may for example be needed for security purposes.

The commands must be run on the relevant web proxy node. It is not recommended that the commands be run on a single node cluster system, but only on a cluster. The commands will automatically reconfigure and restart the nginx process, so some downtime will result. Request URLs to a disabled service will redirect the user to the active service.

- To disable or enable admin or Self-service web services on the web proxy node:

web service disable <selfservice|admin>

web service enable <selfservice|admin>

- To list web services on the web proxy node:

web service list

8. Create a restore point. As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
9. Initialize the database and clear all data. On the primary Unified node, run **voss cleardown**.

Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow voss-deviceapi/app.log** in a separate console on the primary Unified node.

10. Import the templates.

- a. Copy the VOSS Automate template file to the primary Unified node with the command:

scp <VOSS Automate_template_file> platform@<unified_node_ip_address>:~/media

- b. Log in to the primary Unified node and install the template. It is recommended that this step is run in a terminal opened with the **screen** command.

- i. Run **screen**.

- ii. Run **app template media/<VOSS Automate_template_file>**

- The console will display a message:

Deploying the template-product **for** VOSS Automate <<RELEASE_VERSION>> ...

- c. When prompted to select the product deployment type, provide and confirm the deployment type:

- Enterprise
- Provider

(For information on the “Insights Netflow” deployment type when installing release 24.2, contact VOSS.)

In accordance with the selected deployment type, you are prompted to enter and verify:

- a top-level administrator password:
Please enter a password for "sysadmin"
- and one administrator password - depending on the deployment:
 - Enterprise : Please enter a password for "entadmin"
 - Provider : Please enter a password for "hcsadmin"

Upon installation, the password length should be at least 8 characters.

Deployment-specific artifacts are installed according to the selected type of product deployment. A message displays according to the selected deployment type - one of:

```
"Importing EnterpriseOverlay.json"
"Importing ProviderOverlay.json"
```

Deployment specific system artifacts are imported and a message is displayed:

```
Deployment-specific Overlay artifacts successfully imported.
```

- Python functions are deployed
- System artifacts are imported.
- You are prompted to provide administrator passwords.

The template install automatically restarts necessary applications. If a cluster the installation propagates changes throughout the cluster.

11. Review the output from the **app template** commands and confirm that the install message appears:

```
Deployment summary of UPDATED template solution (i.e. current values after
↪ installation):
-----
↪ ----

Product: [PRODUCT]
Version: [UPDATED PRODUCT RELEASE]
Iteration-version: [UPDATED ITERATION]
Platform-version: [UPDATED PLATFORM VERSION]
```

You can also monitor the template installation from the Admin Portal transaction list.

- If there are no errors, as part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
- If there was an error, the install script stops with a failure message listing the problem. Contact VOSS Support.

12. Check for needed security updates by running the **cluster run all security check** command on the primary node. If at least one update is required for any node, run the **cluster run all security update** command on the primary Unified node.

After the security update is successful, reboot the cluster:

- a. From the selected primary Unified node, run **cluster run notme system reboot**. Since all services will be stopped, this takes some time.
- b. From the selected primary Unified node, run **system reboot**. Since all services will be stopped, this takes some time.

If a node does not properly reboot but the console shows that all processes have terminated, you can manually reboot the node without any system corruption.

13. (Optional) Install language templates for languages other than English.

- a. Copy the language template file to the primary Unified node with the command:

```
scp <language_template_file> platform@<unified_node_ip_address>:~/media
```

- b. Log in to the primary Unified node and install the template with the command:

```
app template media/<language_template_file>
```

For example, to install French:

```
app template media/VOSS AutomateLanguagePack_fr-fr.template
```

There is no need to run this command on all nodes.

14. (Optional) If the VOSS Automate Phone Based Registration Add-on is required, follow the installation instructions in the Appendix of your Core Feature Guide:

“Install the Phone Based Registration Web Service”

15. Run the following command:

```
voss migrate_summary_attributes device/cucm/HuntPilot
```

16. License the installation:

From release 21.4 onwards, the deployment needs to be licensed. After installation, a 7-day grace period is available to license the product.

- a. Obtain the required license token from VOSS.
- b. License:
 - i. To license through the GUI, follow steps indicated in Product License Management in the Core Feature Guide.
 - ii. To license through the CLI, follow steps indicated in Product Licensing in the Platform Guide.

17. Mount the Insights database drive

On each unified node, assign the `insights-voss-sync:database` mount point to the drive added for the Insights database prior to installation.

For example, if `drives list` shows the added disk as:

```
Unused disks:
sde
```

then run the command

```
drives add sde insights-voss-sync:database
```

on each unified node where the drive has been added.

Sample output (the message below can be ignored on release 24.1:

```
WARNING: Failed to connect to lvmetad. Falling back to device scanning.)
```

```
$ drives add sde insights-voss-sync:database
Configuration setting "devices/scan_lvs" unknown.
Configuration setting "devices/allow_mixed_block_sizes" unknown.
WARNING: Failed to connect to lvmetad. Falling back to device scanning.
71ad98e0-7622-49ad-9fg9-db04055e82bc
Application insights-voss-sync processes stopped.
Migrating data to new drive - this can take several minutes
Data migration complete - reassigning drive
Checking that /dev/sde1 is mounted
Checking that /dev/dm-0 is mounted
/opt/platform/apps/mongodb/dbroot
Checking that /dev/sdc1 is mounted
/backups

Application services:firewall processes stopped.
Reconfiguring applications...
Application insights-voss-sync processes started.
```

4.6.2. Single-node cluster (cluster-of-one) Installation

Before You Begin

Before continuing, you should have followed the OVA installation according to the steps and preliminary requirements specified in: [Create a New VM Using the Platform-Install OVA](#)

Note:

- Template installation and upgrade takes approximately two hours. You can follow the progress on the GUI transaction list.
- Before installing release 24.2, ensure that an additional 70 GB disk has been made available for the Insights database.

See the Adding Hard Disk Space topic in the Platform Guide and VOSS Automate Hardware Specifications in this guide.

This disk is needed to assign to the insights-voss-sync:database mount point. See the final installation step below.

The standard **screen** command should be used where indicated. See: [Using the screen command](#).

Procedure

1. Install VMware tools:
 - a. Log in on the node as the platform user and run **app install vmware**.
 - b. Verify that vmware is running: **app list**.
2. Prepare the node to be added as a single-node cluster ("cluster-of-one").
Run **cluster prepnode**.
3. Add the node to the cluster: **cluster add <ip_addr>**.
 - a. Verify the node is listed in the cluster: **cluster list**. Example:


```
platform@VOSS:~$ cluster list
Cluster has 1 nodes:
  application : 192.168.100.3
  webproxy   : 192.168.100.3
  database    : 192.168.100.3
```
4. Initialize the database and clear all data with the **voss cleardown** command.
Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow voss-deviceapi/app.log**.
5. Add the network domain (optional if a domain name is needed).
 - a. Configure the domain: **cluster run all network domain <domain_name>**.
 - b. Verify the configured network domain: **cluster run all network domain**. The node shows the domain that you configured.
6. Check the network:
 - a. Run **cluster check** to verify the status of the cluster, network connectivity, disk status and NTP.
 - b. Verify the DNS configuration: **cluster run all network dns**. The node responds with the DNS server address.
7. As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
8. Import the templates.
 - a. Copy the VOSS Automate template file to the node with the command:
scp <VOSS Automate_template_file> platform@<ip_address>:~/media
 - b. Log in to the node and install the template. It is recommended that this step is run in a terminal opened with the **screen** command.
 - i. Run **screen**.
 - ii. Run **app template media/<VOSS Automate_template_file>**
 - The console will display a message:


```
Deploying the template-product for VOSS Automate <<RELEASE_VERSION>> ...
```
 - c. When prompted to select the product deployment type, provide and confirm the deployment type, either "Enterprise" or "Provider".

(For information on the “Insights Netflow” deployment type when installing release 24.2, contact VOSS.)

In accordance with the selected deployment type, you are prompted to enter and verify:

- a top-level administrator password:
 - Please enter a password for "sysadmin"
- and one administrator password - depending on the deployment:
 - Enterprise : Please enter a password for "entadmin"
 - Provider : Please enter a password for "hcsadmin"

Upon installation, the password length should be at least 8 characters.

Deployment-specific artifacts are installed according to the selected type of product deployment. A message displays according to the selected deployment type - one of:

```
"Importing EnterpriseOverlay.json"
```

```
"Importing ProviderOverlay.json"
```

Deployment specific system artifacts are imported and a message is displayed:

```
Deployment-specific Overlay artifacts successfully imported.
```

- i. Python functions are deployed
- ii. System artifacts are imported.
- iii. You are prompted to provide administrator passwords.

The template install automatically restarts necessary applications. The installation propagates changes throughout the cluster.

9. Review the output from the **app template** commands and confirm that the install message appears:

```
Deployment summary of UPDATED template solution (i.e. current values after ↵
↵installation):
```

```
-----
↵-----
```

```
Product: [PRODUCT]
Version: [UPDATED PRODUCT RELEASE]
Iteration-version: [UPDATED ITERATION]
Platform-version: [UPDATED PLATFORM VERSION]
```

You can also monitor the template installation from the Admin Portal transaction list.

- If there are no errors indicated, we recommend that you make a restore point.

As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.

- If there was an error, the install script stops with a failure message listing the problem. Contact VOSS Support.

10. Check for needed security updates by running the **security check** command. If at least one update is required for any node, run the **security update** command.

After the security update is successful, reboot: run **system reboot**.

Since all services will be stopped, this takes some time.

If the node does not properly reboot but the console shows that all processes have terminated, you can manually reboot without any system corruption.

11. (Optional) Install language templates for languages other than English.
 - a. Copy the language template file to the primary Unified node with the command:

scp <language_template_file> platform@<ip_address>:~/media

- b. Log in to the primary Unified node and install the template with the command:

app template media/<language_template_file>

For example, to install French:

app template media/VOSS AutomateLanguagePack_fr-fr.template

12. (Optional) If the VOSS Automate Phone Based Registration Add-on is required, follow the installation instructions in the Appendix of your Core Feature Guide:

“Install the Phone Based Registration Web Service”

13. Run the following command:

voss migrate_summary_attributes device/cucm/HuntPilot

14. License the installation:

From release 21.4 onwards, the deployment needs to be licensed. After installation, a 7-day grace period is available to license the product.

- a. Obtain the required license token from VOSS.
- b. License:
 - i. To license through the GUI, follow steps indicated in Product License Management in the Core Feature Guide.
 - ii. To license through the CLI, follow steps indicated in Product Licensing in the Platform Guide.

15. Mount the Insights database drive

On the primary unified node, assign the `insights-voss-sync:database` mount point to the drive added for the Insights database prior to installation.

For example, if `drives list` shows the added disk as:

```
Unused disks:
sde
```

then run the command

```
drives add sde insights-voss-sync:database
```

on each unified node where the drive has been added.

Sample output (the message below can be ignored on release 24.1:

```
WARNING: Failed to connect to lvmetad. Falling back to device scanning.)
```



```

$ drives add sde insights-voss-sync:database
Configuration setting "devices/scan_lvs" unknown.
Configuration setting "devices/allow_mixed_block_sizes" unknown.
WARNING: Failed to connect to lvmetad. Falling back to device scanning.
71ad98e0-7622-49ad-9fg9-db04055e82bc
Application insights-voss-sync processes stopped.
Migrating data to new drive - this can take several minutes
Data migration complete - reassigning drive
Checking that /dev/sde1 is mounted
Checking that /dev/dm-0 is mounted
/opt/platform/apps/mongodb/dbroot
Checking that /dev/sdc1 is mounted
/backups

Application services:firewall processes stopped.
Reconfiguring applications...
Application insights-voss-sync processes started.

```

4.6.3. View Installation and Upgrade Transactions

Use this procedure to view transactions from a VOSS Automate installation or upgrade.

Procedure

1. Log in as sysadmin administrator.
2. Select **Administration Menu > Transactions**.
3. To view details on a transaction, click the transaction.

4.6.4. Installation Quick Reference

Note:

- These steps are described in depth in the VOSS Automate Install Guide.
- The standard **screen** command should be used where indicated, and the *reconnect* parameter is available if needed:
 - **screen** - start a new session
 - **screen -ls** - show sessions already available
 - **screen -r [screen PID]** - reconnect to a disconnected session

VOSS recommends to use the **screen** command to avoid failures if the connection is interrupted whilst running the command. If the connection is interrupted whilst running the command in **screen** then the session can be retrieved by first listing the sessions PID currently running in **screen**: **screen -ls**, and then reconnecting to the session using **screen -r [screen PID]**.

General Steps

1. Download VOSS Automate install and patch media from:
<https://voss.portalshape.com> > Downloads > VOSS Automate > XXX > New Installation
 where XXX is the release number.
2. Review sizing requirements and define the deployment model:
 - Single node cluster/cluster-of-one (Lab Only)
 - MicroCluster (Two Unified nodes clustered)
 - Cluster (4 Unified Nodes and 2 Web Proxies)
 - DR Cluster (6 Unified Nodes and 2 Web Proxies)
3. Define VMHost space on VMWare servers
4. Deploy the VOSS Automate OVA to VMHost(s)
5. After the VM is created, select the CD ROM configuration and verify the **Connect at Power On** check box is enabled.
6. Power on the VM.
7. Configure the options in the installation wizard.
8. Install VMWare Tools from VOSS Automate CLI as platform user:
 command: **app install vmware**
9. Continue below to chosen deployment model.

Single Node Cluster Deployment

1. Prepare the node for cluster command via the SSH CLI:
 command: **cluster prepnode**
2. Add the node to the cluster:
 command: **cluster add <node_ip-address>**
3. Verify the node is a member of the cluster:
 command: **cluster list**
4. Initialize the database and clear all data with the **voss cleardown** command.
 Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow voss-deviceapi/app.log**.
5. (Optional) Set VOSS Automate Network Domain:
 command: **cluster run all network domain <yourdomain>**
6. Check application status:
 command: **cluster status**
 - Should any services be in a down state, restart all services on that affected node:
 command: **cluster run <node_ip> app start**
7. Run the command: **voss cleardown**.

8. SFTP the install templates to the VOSS Automate server `media` directory
9. Install VOSS Automate Templates via VMWare Console CLI. VOSS recommends that this step is run in a terminal opened with the **screen** command, or on the VMWare console.

- a. Run **screen**.

- b. Run **app template media/<VOSS Automate_Template_Name>**.

- A Deployment Type choice must be made during the template install execution. Choose either "Enterprise" or "Provider"

In accordance with the selected deployment type, you are prompted to enter and verify:

- a top-level administrator password:

Please enter a password for "sysadmin"

- and one administrator password - depending on the deployment:

* Enterprise : Please enter a password for "entadmin"

* Provider : Please enter a password for "hcsadmin"

Upon installation, the password length should be at least 8 characters.

Multi Node Deployment

All of the following commands will be run on the primary node via the SSH CLI until specified to use ESX Console CLI. The designation of primary unified node is arbitrary. The deploying administrator can pick any unified node that they see fit.

1. On each node that is not the designated primary unified node. prepare the servers for cluster command via the SSH CLI:

command: **cluster prenode**

2. Add all of the other nodes to the cluster:

command: **cluster add <non-primary-node_ip-address>**

Repeat this command for each other node - binding each individual node IP Address to the cluster. This command does not need to be run for the primary unified node.

3. Verify all nodes are members of the cluster:

command: **cluster list**

4. (Optional) Set VOSS Automate Network Domain:

command: **cluster run all network domain <yourdomain>**

5. Set each unified node's database weight:

command: **database weight <un-node_ip-address> <priority_weight>**

- This command must be run for all unified nodes primary and secondary.
- Priority weights of 40, 30 are recommended for *two* Unified nodes.
- Priority weights of 40, 30, 20, and 10 are recommended for *four* Unified nodes
- Priority weights of 60, 50, 40, 30, 20, and 10 are recommended for *six* Unified nodes.

6. Provision the VOSS Automate cluster database. VOSS recommends that this step is run in a terminal opened with the **screen** command, or on the VMWare console.

- a. Run **screen**.
 - b. command: **cluster provision**
7. Check cluster application status:

command: **cluster status**

 - Should any services be in a down state, restart all services on that affected node:

command: **cluster run <node_ip> app start**
8. Run the command: **voss cleardown**.
9. SFTP the install templates to the VOSS Automate server media directory of the primary unified node.
10. Install VOSS Automate Templates via VMWare Console CLI of primary unified node. VOSS recommends that this step is run in a terminal opened with the **screen** command, or on the VMWare console.
 - a. Run **screen**.
 - b. Run **app template media/<VOSS Automate_Template_Name>**
11. A Deployment Type choice must be made during the template install execution. Choose one of:
 - Enterprise
 - Provider

In accordance with the selected deployment type, you are prompted to enter and verify:

 - a top-level administrator password:

Please enter a password for "sysadmin"
 - and one administrator password - depending on the deployment:
 - Enterprise : Please enter a password for "entadmin"
 - Provider : Please enter a password for "hcsadmin"

Upon installation, the password length should be at least 8 characters.

Post Deployment

1. Access the VOSS Automate web interface via any web browser:

`https://<ip_address_or_dns_name_of_VOSS_Automate_PrimaryUN_or_WebProxy>`
2. Run the following security commands:
 - **cluster check** - inspect entries under security.
 - If needed, **security update**.
3. Run the following command:

voss migrate_summary_attributes device/cucm/HuntPilot

4.6.5. Migrating from a 6 Node to 8 Node System

On a standard topology, to migrate a clustered 6 node system (4 unified nodes and 2 WebProxy nodes) to a clustered 8 node system (6 unified nodes and 2 WebProxy nodes), the considerations and steps below are required.

1. Check and make a restore point of the clustered 6 node system *before* adding the nodes:
 - a. Run **cluster list** to ensure the node count is correct.
 - b. Run **cluster status** to check all nodes are online and services reported as running.
 - c. Run **cluster run database cluster list** to make sure all unified nodes are aware of the current cluster nodes.
 - d. Run **cluster run all app status** to make sure all services are running on all nodes.
 - e. Make a restore point of the entire 6 node cluster.

As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
2. Add the 2 unified nodes:
 - a. Create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
 - b. Run **cluster prepnode** on all nodes, including new nodes.
 - c. From the primary unified node, run **cluster add <ip>** for new nodes, excluding itself.
3. Reset the cluster database weights. When nodes are removed from and added to a cluster, remove all database weights completely and add them back in *before provisioning* to reset the configuration.
 - a. Delete all database weights in the cluster. For each IP, run **database weight del <IP>**.
 - b. To add database weights in, you set the weight of the intended primary, but must always specify the current primary (using **database primary**), regardless of whether the new intended primary is the same node or not. During the provision process, the role of primary will then be transferred from the existing primary to the node with the highest weight.

Determine the current primary database node with **database primary**.
4. Run the command **database weight add <IP> <numeric>** on the primary database node for each IP, making the value of the intended primary database node the highest value.
5. Check the cluster before provisioning:
 - a. Run **cluster list** to ensure the node count is correct.
 - b. Run **cluster status** to check all nodes are online and services reported as running.
 - c. Run **cluster run database cluster list** to make sure all unified nodes are aware of the current cluster nodes.
 - d. Run **cluster run all app status** to make sure all services are running on all nodes. Fresh nodes that have not been provisioned will show a message: *suspended waiting for mongo*.
6. Run **cluster provision** to provision the cluster.
7. After a successful migration, the restore point made in step 1. can be removed.

4.7. Modular Cluster Topology Installation

4.7.1. Modular Architecture Multinode Installation

Note:

- A modular architecture installation is not supported for a single node cluster (“cluster of one”) topology.
- Before installing release 24.2, ensure that an additional 70 GB disk has been made available for the Insights database.

See the Adding Hard Disk Space topic in the Platform Guide and VOSS Automate Hardware Specifications in this guide.

This disk is needed to assign to the `insights-voss-sync:database` mount point. See the final installation step below.

Before You Begin

Before continuing, you should have followed the OVA installation on each node according to the steps and preliminary requirements specified in: [Create a New VM Using the Platform-Install OVA](#) and according to the node roles as indicated in [Notes on Multi-Node Installation](#). Data center names are also selected at this stage.

For example, for an 8-node modular cluster in 2 data centers:

- DC1 = primary site or data center containing primary database node (highest database weight)
- DC2 = data recovery (DR) data center

Install:

- 3 nodes with Database roles (2 in DC1, 1 in DC2)
- 3 nodes with Application roles (2 in DC1, 1 in DC2)
- 2 nodes with WebProxy roles (1 in DC1, 1 in DC2)

Optionally download or extract language pack template files to support languages other than English.

Note:

- For typical modular geo-redundant multinode cluster deployment with 3 database and 3 application nodes, there are:
 - two application nodes in the primary Site
 - two database nodes in the primary Site
 - one application node in the Disaster Recovery (DR) Site
 - one database node in the Disaster Recovery (DR) Site

The worker count (**voss workers** command) needs to be set on the DR nodes. Refer to:

- [Multinode cluster with unified nodes](#)
- [Multi Data Center Deployments](#)

- NAT between nodes is not allowed
 - If there is a firewall between nodes, then specific ports must be configured. For port configuration, refer to:
 - [Clustering Considerations](#)
 - [Network Communications between Nodes within the Cluster](#)
 - [Network Communications External to the Cluster](#)
 - Template installation and upgrade takes approximately two hours. You can follow the progress on the Admin Portal transaction list.
 - It is strongly recommended *not* to allow customer end-users the same level of administrator access as the restricted groups of provider- and customer administrators. This is why Self-service web proxies as well as Administrator web proxies should be used.

Systems with Self-service only web proxies are *only* recommended where the system is customer facing, but where the customer does not administer the system themselves.
 - For cluster installations, also refer to the Health Checks for Cluster Installations Guide.
 - If it is necessary to change an IP address of a node in a cluster, first remove it from the cluster by running the command below *on the node to be changed*:
cluster del <IP address of node to be changed>
 - Refer to [Installation Logs](#) for troubleshooting logs during an installation.
-

The standard **screen** command should be used where indicated. See: [Using the screen command](#).

Procedure

1. Install VMware tools on each node.
 - a. Log in to each node and run **app install vmware**.
 - b. Verify that vmware is running: **app list**.
2. Prepare each node to be added to the cluster:
 - a. Select a database node that will become the primary database node. The *primary site* or data center will contain the primary database node. The deploying administrator can pick any database node that they see fit.
 - b. On each node, run **cluster prepnode**.
3. Add nodes to the cluster.
 - a. Log in to the selected primary database node.
 - b. Add the other database, application and WebProxy nodes to the cluster: **cluster add <ip_addr>**.

Note that you do not have to add the selected primary database node to the cluster. It will automatically be added to the cluster.
 - c. Verify the list of nodes in the cluster: **cluster list**.
4. Add the network domain (optional if a domain name is needed). From the selected primary database node:
 - a. Configure the domain: **cluster run all network domain <domain_name>**.

- b. Verify the configured network domain: **cluster run all network domain**. Each node shows the domain that you configured.
5. Check the network:
 - a. From the selected primary database node, run **cluster check** to verify the status of the cluster, network connectivity, disk status and NTP.
 - b. Verify the DNS configuration: **cluster run all network dns**. Each node responds with the DNS server address.
6. Create a restore point. As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
7. Configure the cluster.
 - a. From the selected primary database node, provide a weight for each database server with the **database weight add <database_ip> <priority>** command.
 The higher the value, the more priority.
 - A weight of 40 for the primary database node at the primary site (DC1)
 - A weight of 30 for the secondary database node at the primary site (DC1)
 - A weight of 10 for the secondary database node at the DR (Data Recovery) site (DC2)
 - b. From the selected primary database node:

It is recommended that this step is run in a terminal opened with the **screen** command.

 - i. Run **screen**.
 - ii. Run **cluster provision**
 Allow approximately 2 hours for the operation to complete for two WebProxy and four database, application nodes.
 - c. When provisioning is complete, check that each node is contactable and that the time server is running on each with **cluster check**.
 If a service is down, run **cluster run <node_ip> app start** to restart the service.
 If provisioning is successful, the screen session can be closed by typing **exit** in the screen terminal. If errors occurred, keep the screen terminal open for troubleshooting purposes and contact VOSS support.
 - d. On each of the new *application* nodes, set the queues to 2 with the command **voss queues 2**.

Note: Applications are reconfigured and the voss-queue process is restarted.

- e. (Optional) If required, set the web weights configurations (Active-Active, Active-Standby, Single node cluster). From the primary database node, run the required **web weight** commands for the Web Proxy nodes. For details, refer to [Multi Data Center Deployments](#) and the VOSS Automate Best Practices Guide.
 - f. (Optional) If required, enable or disable Self-service or admin web services on the web proxy nodes. This may for example be needed for security purposes.
 The commands must be run on the relevant web proxy node. The commands will automatically reconfigure and restart the *nginx* process, so some downtime will result. Request URLs to a disabled service will redirect the user to the active service.

- To disable or enable admin or Self-service web services on the web proxy node:


```
web service disable <selfservice|admin>
```

```
web service enable <selfservice|admin>
```
 - To list web services on the web proxy node:


```
web service list
```
8. Create a restore point. As part of the rollback procedure, ensure that a suitable restore point is obtained prior to the start of the activity, as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
 9. Initialize the database and clear all data. On an *application node*, run **voss cleardown**.
 Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow voss-deviceapi/app.log** in a separate console on the application node.
 10. Import the templates.
 - a. Copy the VOSS Automate template file to an *application node* with the command:


```
scp <VOSS Automate_template_file> platform@<app_node_ip_address>:~/media
```
 - b. Log in to *this application* node and install the template. It is recommended that this step is run in a terminal opened with the **screen** command.
 - i. Run **screen**.
 - ii. Run **app template media/<VOSS Automate_template_file>**
 - The console will display a message:

```
Deploying the template-product for VOSS Automate <<RELEASE_VERSION>> ...
```

- c. When prompted to select the product deployment type, provide and confirm the deployment type:
 - Enterprise
 - Provider

(For information on the “Insights Netflow” deployment type when installing release 24.2, contact VOSS.)

In accordance with the selected deployment type, you are prompted to enter and verify:

- a top-level administrator password:


```
Please enter a password for "sysadmin"
```
- and one administrator password - depending on the deployment:
 - Enterprise : Please enter a password for "entadmin"
 - Provider : Please enter a password for "hcsadmin"

Upon installation, the password length should be at least 8 characters.

Deployment-specific artifacts are installed according to the selected type of product deployment. A message displays according to the selected deployment type - one of:

```
"Importing EnterpriseOverlay.json"
```

```
"Importing ProviderOverlay.json"
```

Deployment specific system artifacts are imported and a message is displayed:

Deployment-specific Overlay artifacts successfully imported.

- i. Python functions are deployed
- ii. System artifacts are imported.
- iii. You are prompted to provide administrator passwords.

The template install automatically restarts necessary applications. If a cluster the installation propagates changes throughout the cluster.

11. Review the output from the **app template** commands and confirm that the install message appears:

Deployment summary of UPDATED template solution (i.e. current values after ↪ installation):

↪ -----

Product: [PRODUCT]
Version: [UPDATED PRODUCT RELEASE]
Iteration-version: [UPDATED ITERATION]
Platform-version: [UPDATED PLATFORM VERSION]

You can also monitor the template installation from the Admin Portal transaction list.

- If there are no errors indicated, we recommend a suitable restore point is created as per the guidelines for the infrastructure on which the VOSS Automate platform is deployed.
- If there was an error, the install script stops with a failure message listing the problem. Contact VOSS Support.

12. Check for needed security updates by running the **cluster run all security check** command on the primary database node. If at least one update is required for any node, run the **cluster run all security update** command on the primary database node.

After the security update is successful, reboot the cluster:

- a. From the selected primary database node, run **cluster run notme system reboot**. Since all services will be stopped, this takes some time.
- b. From the selected primary database node, run **system reboot**. Since all services will be stopped, this takes some time.

If a node does not properly reboot but the console shows that all processes have terminated, you can manually reboot the node without any system corruption.

13. (Optional) Install language templates for languages other than English.

- a. Copy the language template file to the *selected application node* with the command:

scp <language_template_file> platform@<app_node_ip_address>:~/media

- b. Log in to the *application* node and install the template with the command:

app template media/<language_template_file>

For example, to install French:

app template media/VOSS AutomateLanguagePack_fr-fr.template

There is no need to run this command on all nodes.

14. (Optional) If the VOSS Automate Phone Based Registration Add-on is required, follow the installation instructions in the Appendix of your Core Feature Guide:

“Install the Phone Based Registration Web Service”

15. Run the following command:

voss migrate_summary_attributes device/cucm/HuntPilot

16. License the installation:

From release 21.4 onwards, the deployment needs to be licensed. After installation, a 7-day grace period is available to license the product.

- a. Obtain the required license token from VOSS.
- b. License:
 - i. To license through the GUI, follow steps indicated in Product License Management in the Core Feature Guide.
 - ii. To license through the CLI, follow steps indicated in Product Licensing in the Platform Guide.

17. Mount the Insights database drive

On each database node, assign the `insights-voss-sync:database` mount point to the drive added for the Insights database prior to installation.

For example, if `drives list` shows the added disk as:

```
Unused disks:
sde
```

then run the command

```
drives add sde insights-voss-sync:database
```

on each unified node where the drive has been added.

Sample output (the message below can be ignored on release 24.1:

```
WARNING: Failed to connect to lvmetad. Falling back to device scanning.)
```

```
$ drives add sde insights-voss-sync:database
Configuration setting "devices/scan_lvs" unknown.
Configuration setting "devices/allow_mixed_block_sizes" unknown.
WARNING: Failed to connect to lvmetad. Falling back to device scanning.
71ad98e0-7622-49ad-9fg9-db04055e82bc
Application insights-voss-sync processes stopped.
Migrating data to new drive - this can take several minutes
Data migration complete - reassigning drive
Checking that /dev/sde1 is mounted
Checking that /dev/dm-0 is mounted
/opt/platform/apps/mongodb/dbroot
Checking that /dev/sdc1 is mounted
/backups

Application services:firewall processes stopped.
Reconfiguring applications...
Application insights-voss-sync processes started.
```

5. VOSS Automate Azure Installation

5.1. Azure Cloud Deployment

VOSS Automate can be deployed into the Azure cloud by means of terraform scripts.

5.1.1. Hardware Requirements

For details on Standard and Modular Topologies, refer to the VOSS Automate Architecture and Hardware Specification Guide and Platform Guide.

Unified or Database Nodes:

- VM Size: E4ds_v4 Standard
- CPU: 4
- RAM: 32
- OS disk: 30GB, Premium_LRS
- application disk: 50GB, Standard_LRS
- backup disk: 55 GB, Standard_LRS
- DB disk: 250 GB, Premium_LRS
- Total disk size: 385GB

Application Nodes:

- VM Size: E4ds_v4 Standard
- CPU: 4
- RAM: 32
- OS disk: 30GB, Premium_LRS
- application disk: 50GB, Standard_LRS
- Total disk size: 80GB

Web Proxies:

- Web Proxies are replaced by an Azure Load Balancer or Application Gateway

5.1.2. Network Communications External to the Cluster

The following details are all based on the default settings. These can vary depending on the application setup and network design (such as NAT) of the solution, so may need adjustment accordingly. Where a dependent is noted, this is fully dependent on the configuration with no default.

These communications are all related to communications with devices external to the cluster.

- Outbound Communications to Devices from the Application/Unified nodes:

Communication	Protocol	Port
Cisco Unified Communications Manager (UCM)	HTTPS	TCP 8443
Cisco Unity Connection (CUXN)	HTTPS	TCP 443
Webex	HTTPS	TCP 443
LDAP directory	LDAP	TCP/UDP 389 and/or 636(TLS/SSL)

- VOSS Automate Communications

The cluster contains multiple nodes which can be contained in separate secured networks. Network ports need to be opened on firewalls and/or network security groups to allow inter-node communication – these are described in more detail in the Platform Guide.

All communication between nodes are encrypted.

Communication	Protocol	Port
Database access	Database	TCP 27020 and 27030 bi-directional
Cluster Communications	HTTPS	TCP 8443
Remote Administration	SSH	TCP 22
Web Server Communication	HTTP/HTTPS	TCP 80/443
Simple Network Management Protocol	SNMP	UDP 161 and 162
Network Time Protocol	NTP	UDP 123
Domain Name System	DNS	UDP 53

5.1.3. VOSS Automate Azure Deployment Procedure

1. The supplied terraform deployment configuration requires the following:

- Existing resources:
 - Resource group
 - Virtual Network
 - Virtual Network Subnet
 - Disk Encryption Set (Optional if required to have Customer-Managed Keys)

- App Registration with a role assignment of “Contributor” on the Resource Group to allow terraform to deploy
 - * Record “client_id”
 - * Record “client_secret”
 - * Record “tenant_id”
 - * Record “subscription_id”
 - Bastion Server (Required for the initial deployment and Administration Access)
 - * Deploy the Bastion Server within the same Virtual Network that Automate will be deployed in.
 - VM Size: 1vCPU, 2GB RAM, 30GB Storage is sufficient
 - OS Type: Ubuntu or Windows
 - * Install AZ CLI
 - Windows: <https://learn.microsoft.com/en-us/cli/azure/install-azure-cli-windows?tabs=azure-cli>
 - Ubuntu: <https://learn.microsoft.com/en-us/cli/azure/install-azure-cli-linux?pivot=apt>
 - * Install Terraform
 - <https://developer.hashicorp.com/terraform/install>
 - * Upload the **voss-automate-terraform.zip** file to the bastion server and extract it.
2. Update **values.yaml** with appropriate information
 3. Ensure that the local file sources are correct in modules/voss-node/provision.tf on line 18 and 47.

```
provisioner "file" {
  # Windows source path example
  # source      = "C:\\path\\to\\your\\platform-install.iso"
  source       = "~/voss-automate-terraform/files/platform-install.iso"
  destination = "/home/install/platform-install.iso"
}
```

4. Deployment Type

The defaults are configured as follows:

```
Unified:
  Availability Zone 1
  1x Unified Node

Modular: 2x Application, 3x Database
  Availability Zone 1
    App Node 1, Database Node 1
  Availability Zone 2
    App Node 2, Database Node 2
  Availability Zone 3
    Database Node 3
```

If the requirement is to deploy more than the defaulted Unified or Application(Modular) Nodes, edit the **voss-automate-terraform/modules/voss-deployment/main.tf** file.

```

module "voss-unified-node" {
  count = var.deployment_type == "Unified" ? 1 : 0
}

module "voss-app-node" {
  count = var.deployment_type == "Modular" ? 2 : 0
}

```

5. Configure the state backend in main.tf (defaults to a local file) if required
6. For the terraform deployment, it is recommended to use a screen session.

```

screen - start a new session
screen -ls - show sessions already available
screen -r [screen PID] - reconnect to a disconnected session

```

7. Run terraform

- From the root terraform directory
 - Run: **terraform init**
 - Run: **terraform plan** (Validate plan)
 - * Default total resources to be created (Example):

```

Plan: 11 to add, 0 to change, 0 to destroy.
Modular (2 App 3 DB): Plan: 43 to add, 0 to change, 0 to destroy.

```

- Run: **terraform apply**

5.1.4. VOSS Automate Platform Config and Template Install Procedure

1. *On each of the newly deployed node(s)*, log in as the platform user, using the password configured in values.yaml (**automate_cli_password**)
 - 1.1. Run: **system reboot**. Alternatively you can reset the Virtual Machine from the Azure Portal. This is to ensure all services are running before proceeding with the cluster configuration.
2. You are now ready to configure the platform and install the template.

Note: For both Azure and AWS, deploy using Automate 21.4, then upgrade to the latest version.

- Single Node Cluster Guide:

https://documentation.voss-solutions.com/release_21.4-PB5/html/src/user/install/standalone-installation.html
- Modular Cluster Guide:

https://documentation.voss-solutions.com/release_21.4-PB5/html/src/user/install/modular-multinode-installation.html

3. *On each of the newly deployed node(s)*, log in as the platform user.

```
username: platform  
password: automate_cli_password (Specified in the values.yaml file)
```


6. VOSS Automate AWS Installation

6.1. AWS Deployment

6.1.1. Overview

VOSS Automate can be deployed onto the Amazon Web Services (AWS) cloud using private Amazon Machine Image (AMI). Two private AMIs are provided - one for deploying the application node, and the other for deploying the database node. Both AMIs are built as appliances that contain a self-contained operating system, and the required application or database.

Before you Start

The customer should supply the following to enable VOSS to create a private AMI:

- a. AWS Account ID
- b. Deployment Country
- c. Deployment Region

6.1.2. Hardware Requirements

Note: The AMI's storage is pre-configured as per below specifications. Max EBS IOPS and Throughput is dependent on the Instance Type, increase if required.

Example:

```
r6a.xlarge
- Max Throughput (MB/s) = 1250.0
- Max I/O Operations/second (IOPS) = 40000
```

Application Nodes:

- Instance Type: r6a.xlarge or equivalent
- Instance CPU Architecture: x86_64
- CPU: 4
- RAM: 32
- OS Disk: 30GB, GP3 @ 3000 IOPS / 125 throughput
- Application Disk: 50GB, GP3 @ 3000 IOPS / 125 throughput
- Total Disk size: 80GB

Database Nodes:

- Instance Type: r6a.xlarge or equivalent
- Instance CPU Architecture: x86_64
- CPU: 4
- RAM: 32
- OS Disk: 30GB, gp3 @ 3000 IOPS / 125 throughput
- Application Disk: 50GB, gp3 @ 3000 IOPS / 125 throughput
- Backup Disk: 125 GB, sc1
- DB Disk: 250 GB, io2 @ 750 IOPS
- Total Disk size: 455GB

Web Proxies:

- Web Proxies are replaced by an Application Load Balancer

6.1.3. Network Communications External to the Cluster

The following details are all based on the default settings. These can vary depending on the application setup and network design (such as NAT) of the solution, so may need adjustment accordingly. Where a dependent is noted, this is fully dependent on the configuration with no default.

These communications are all related to communications with devices external to the cluster.

- Outbound Communications to Devices from the Application/Unified nodes:

Communication	Protocol	Port
Cisco Unified Communications Manager (CUCM)	HTTPS	TCP 8443
Cisco Unity Connection (CUXN)	HTTPS	TCP 443
Webex	HTTPS	TCP 443
LDAP directory	LDAP	TCP/UDP 389 and/or 636(TLS/SSL)


- VOSS Automate Communications

The cluster contains multiple nodes which can be contained in separate secured networks. Network ports need to be opened on firewalls and/or network security groups to allow inter-node communication – these are described in more detail in the Platform Guide.

All communication between nodes are encrypted.

Communication	Protocol	Port
Database access	Database	TCP 27020 and 27030 bi-directional
Cluster Communications	HTTPS	TCP 8443
Remote Administration	SSH	TCP 22
Web Server Communication	HTTP/HTTPS	TCP 80/443
Simple Network Management Protocol	SNMP	UDP 161 and 162
Network Time Protocol	NTP	UDP 123
Domain Name System	DNS	UDP 53

6.1.4. VOSS Automate AWS Deployment Topology



6.1.5. AWS Application Load Balancer Configuration

Basic configuration

- Scheme: Internal
- IP address type: IPv4

Network mapping

- VPC: The VPC where the VOSS Automate Application Nodes reside
- Mappings: The Availability Zones where the VOSS Automate Application Nodes reside

Security groups

- Create a new security group or select an existing one.

Target group

- Basic configuration
 - Target type: Application Load Balancer
 - Protocol: TCP/443
 - VPC: The VPC where the VOSS Automate Application Nodes reside
- Health checks
 - Health check protocol: 443
 - Health check path: /portal
 - Advanced health check settings
 - * Health check port: Traffic port
 - * Success codes: 200, 202

Listeners and routing

- Protocol: HTTPS
- Port: 443
- Default Action: Forward to (above target group)

6.1.6. AWS Management Console Deployment Procedure

Prerequisites

- VPC with 3 subnets across 3 availability zones.

Login to your Account

- Navigate to EC2 Services.
- Click on **AMIs** under **Images** in the left pane.
- Select the correct Region from the drop-down at the top right.
- Select **Private Images** from the drop-down filter.
- The VOSS Automate AMIs should appear in the list.

- Select the VOSS Automate Application/Database AMI in the list, Launch Instance from AMI.

Configuration

- Name: Give the instance a descriptive name e.g automate-app-node-1
- Key pair (login): Proceed without a key pair. This is managed by VOSS Automate.
- Instance Type: As per hardware requirements.
- Network Settings: Configure the subnet based on the availability zone.

Example:

```
Subnet 1 - us-east-2a
Application Node 1
Database Node 1

Subnet 2 - us-east-2b
Application Node 2
Database Node 2

Subnet 3 - us-east-2c
Database Node 3
```

- Configure storage: As per hardware requirements.

6.1.7. Terraform Deployment Procedure

Terraform scripts have been provided as a starting point to deploy as per above topology diagram.

This will deploy the following:

- VOSS Automate Modular Cluster
- AWS Application Load Balancer
- Bastion Server (For Automate Remote Administration Access)

6.1.8. VOSS Automate Platform Config and Template Install Procedure

1. On each of the newly deployed node(s), log in as the platform user.

```
username: platform
password: platform
```

User will be prompted **for** a password change.
Enter the current password, new password **and** confirm the new password.

2. You are now ready to configure the platform and install the template.

Note: For both Azure and AWS, deploy using Automate 21.4, then upgrade to the latest version.

- Single Node Cluster Guide:

https://documentation.voss-solutions.com/release_21.4-PB5/html/src/user/install/standalone-installation.html

- Modular Cluster Guide:

https://documentation.voss-solutions.com/release_21.4-PB5/html/src/user/install/modular-multinode-installation.html

3. *On each of the newly deployed node(s)*, log in as the platform user.

```
username: platform
password: automate_cli_password (Specified in the values.yaml file)
```

Index

A

app

- app install, [41, 44, 50, 59](#)
- app template, [40, 41, 44, 59](#)

B

backup

- backup create, [27](#)

C

cluster, [24](#)

- cluster add, [41, 44, 58, 59](#)
- cluster del, [41, 44, 59](#)
- cluster list, [58](#)
- cluster prenode, [41, 44, 58, 59](#)
- cluster primary, [44](#)
- cluster primary role application, [44](#)
- cluster primary role database, [44](#)
- cluster provision, [40, 41, 44, 58, 59](#)
- cluster run, [26, 41, 44, 58, 59](#)
- cluster status, [41, 44, 58, 59](#)
- cluster upgrade, [40](#)

D

database

- database weight, [41, 44, 59](#)

drives

- drives add, [27](#)

L

log

- log follow, [35, 41, 44, 50, 59](#)

N

network

- network container range, [26](#)
- network domain, [50](#)

S

screen, [40, 41, 44, 54, 59](#)

security

- security update, [35, 50](#)

system

- system reboot, [50](#)

V

voss

- voss cleardown, [35, 41, 44, 50, 59](#)
- voss db_collection_cap, [29](#)
- voss db_collection_stats, [27](#)
- voss queues, [59](#)
- voss upgrade_db, [35](#)
- voss workers, [5, 13](#)

voss export

- voss export group, [40](#)
- voss export type, [40](#)

voss subscriber_data_export, [40](#)

W

web

- web service, [5, 13, 35](#)
- web weight, [17, 41, 44, 59](#)
- web weight add, [17](#)
- web weight list, [17](#)