



VOSS Automate Installation Guide

Release 25.2

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: 20250821113000

Contents

1	What's New	1
1.1	Installation Guide: Release 25.2	1
2	Introduction	2
2.1	Overview	2
2.2	File transfer protocol and commands	2
2.3	Clustering commands	3
2.4	Determine the primary node role in a cluster	4
2.5	Geo-redundancy/redundancy and disaster recovery	4
2.6	Using the <code>tmux</code> command	5
3	Deployment Topologies	8
3.1	Automate deployment topologies	8
3.2	Multi data center deployments	25
3.3	Clustering	29
3.4	Network communications between nodes within the cluster	30
3.5	Network communications external to the cluster	31
4	Prepare to Install	34
4.1	Network Docker container range	34
4.2	Backup size considerations	35
4.3	Hardware Specifications	36
5	Automate Platform Install OVA	44
5.1	Installation overview	44
5.2	Platform install OVA on a VM	44
5.3	Inspect the logs to troubleshoot installation	52
5.4	Unified Node Topology Installation	54
5.5	Modular Cluster Topology Installation	69
6	Automate Azure Cloud Installation	77
6.1	Azure Cloud Deployment	77
7	Automate AWS Cloud Installation	82
7.1	AWS Deployment	82
	Index	88

1. What's New

1.1. Installation Guide: Release 25.2

- EKB-24140: Deprecate `screen` command and replace with `tmux`. See: [Using the *tmux* command](#)
Added details on the “`tmux`” command replacing the “`screen`” command.

2. Introduction

2.1. Overview

The Automate Installation Guide (this document) provides details and steps for fresh installs of the Automate system - three install options are provided:

- *Platform install OVA on a VM*
- *Azure Cloud Deployment* - Deploy into the Microsoft Azure cloud with terraform scripts
- *AWS Deployment* - Deploy onto the Amazon Web Services (AWS) cloud, using private Amazon Machine Image (AMI)

Note:

- Use the *Upgrade Guide with Delta or Patch Bundle* to upgrade to a maintenance release, for example, 24.2-PB2.
- Use the *Upgrade Guide with ISO and Template* to upgrade to a major version release, for example, 24.1, 24.2.
- For cluster installations, also refer to the *Health Checks for Cluster Installations Guide*.
- Before installing, review the release notes for the relevant version.

Automate 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 the Automate system. This guide describes the product in general and is not specific to a particular deployment/solution. Details may vary depending on the installation environment.

2.2. File transfer protocol and commands

For file transfers during install, you can use either Secure File Transfer Protocol (SFTP) or Secure Copy Protocol (SCP).

The table provides an example for each file transfer method:

Protocol	Example
SFTP	1. sftp platform@<node_hostname> 2. cd media 3. put <upgrade_template_file>
SCP	This guide shows file transfers using SCP. <pre>scp <upgrade_template_file> platform@<node_hostname>:~/media</pre>

2.3. Clustering 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 the following should be run on a standalone topology:

- cluster provision
- cluster list
- cluster status
- cluster maintenance-mode

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 <start stop status>	- Display the status, start or stop maintenance mode across the cluster
cluster prepnode	- Prepares the system so that it can be joined to a
→cluster	
cluster primary	- Check if the system is considered the primary by
→patch	
	or delta install scripts
cluster primary role <application database>	- Check if the system is considered the 'primary' in the given role
cluster provision [datacentre <location>] [role <role>]	- perform cluster-wide provisioning
cluster run <where> <command>	- run the command on a particular host. <where> can either be a name prefix, ip, role, or

(continues on next page)

(continued from previous page)

```

→ '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

```

2.4. Determine the primary node role in a cluster

To determine if the node is the primary database node or the application node, run the `cluster primary` command (with additional parameters) from the CLI on a node:

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

```

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:~$

```

2.5. Geo-redundancy/redundancy and disaster recovery

High Availability (HA) is an approach to IT system design and configuration that ensures 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 web proxy 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.

High availability 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 *Automate Platform Guide* for details.

2.6. Using the `tmux` command

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

Note: From release 25.2, the `screen` command is no longer available on the Automate CLI - typing in `screen` on the CLI will return a message to use `tmux` instead.

The following commands require the running of `tmux`:

- **backup create <location-name>**
- **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 `tmux` should be run first:

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

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

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

- List sessions:

```
tmux list-sessions
0: 1 windows (created Thu Apr 10 22:32:29 2025)
1: 1 windows (created Thu Apr 10 22:32:41 2025) (attached)
```

- Detach

`Ctrl+b + d`

- Exit (kill session)

`Ctrl+d`

- Attach to a session


```
tmux attach -t 0
```

- Scroll

```
Ctrl+b + [
```

- Scroll off

```
Esc
```

- Execute a command (example)

```
$ tmux -c 'voss export type license_audit'
Starting license_audit export, please wait...
Completed license_audit export, created vlf_[...]
```

To create a `tmux` log file in the `/var/log/platform/tmux/` directory:

```
tmux -Log
```

For details, refer to the on-line help - also referenced below.

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

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

2.6.1. On-line help

The platform CLI offers an enhanced online help to the `tmux` command that also includes details and examples, as well as custom parameters to output the session to a log file.

The enhanced, full on-line help information available from the CLI is shown below:

```
platform@VOSS:~$ tmux -h
usage: tmux [-2CDlNuvV] [-c shell-command] [-f file] [-L socket-name]
          [-S socket-path] [-T features] [command [flags]]

COMMON COMMANDS
  attach-session (alias: attach) [-t target-session]  target-session is the number
                                                         from 'tmux ls'.
  list-commands  (alias: lscm) [command]              List the syntax of command or
                                                         - if omitted - of all commands
                                                         supported by tmux.
  list-sessions  (alias: ls)                          List all sessions managed by
                                                         the server.
  man                                                    Displays tmux manual.

COMMON KEY BINDINGS
  tmux may be controlled from an attached client by using a key combination of a
  prefix key, 'Ctrl-b', followed by a command key
```

(continues on next page)

(continued from previous page)

```

d          Detach the current client.
:          Enter the tmux command prompt.

VOSS LOGGING ENHANCEMENTS:

-Log          Turn on output logging. Log folder /var/log/platform/tmux/
-LogPrefix [prefix] Prefix added to the log filename.

LOG NAME PATTERN:

<[PREFIX]->tmuxlog-[TIMESTAMP]-[USER_ID]-S[SESSION_NAME]-W[WINDOW_ID]-P[PROCESS_ID].log
I.e.
    tmuxlog-20250527120013-1401-0-W0-P1329199.log
    my-prefix-tmuxlog-20250527120013-1401-S0-W0-P1329199.log

Examples:
1. Start a logged session with default name pattern:
    tmux -Log
    I.e. tmuxlog-20250527120013-1401-S0-W0-P1329199.log
2. Run specific command with custom log prefix:
    tmux -LogPrefix my-log -c "my-long-running-command"
    I.e. my-log-tmuxlog-20250527120013-1401-S0-W0-P1329199.log

```

The `list-commands` parameter offers details on available commands, such as shown in the snippet below:

```

platform@VOSS:~$ tmux list-commands
attach-session (attach) [-dErX] [-c working-directory] [-f flags] [-t target-session]
bind-key (bind) [-nr] [-T key-table] [-N note] key [command [arguments]]
break-pane (breakp) [-abdp] [-F format] [-n window-name] [-s src-pane] [-t dst-window]
capture-pane (capturep) [-aCeJNpPq] [-b buffer-name] [-E end-line] [-S start-line] [-t
↪target-pane]
choose-buffer [-NrZ] [-F format] [-f filter] [-K key-format] [-O sort-order] [-t target-
↪pane] [template]
choose-client [-NrZ] [-F format] [-f filter] [-K key-format] [-O sort-order] [-t target-
↪pane] [template]
choose-tree [-GNrswZ] [-F format] [-f filter] [-K key-format] [-O sort-order] [-t target-
↪pane] [template]
clear-history (clearhist) [-t target-pane]
clock-mode [-t target-pane]
command-prompt [-lkiNTW] [-I inputs] [-p prompts] [-t target-client] [template]
[...]

```

3. Deployment Topologies

3.1. Automate deployment topologies

3.1.1. Overview

Automate offers two main deployment topologies:

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

Two additional deployment options are available:

- *Cloud deployments*
- *VOSS Automate Cloudv1 (SaaS)*

3.1.2. Node types

Automate deployment topologies are comprised of a configuration of the following types of nodes, each performing specific functions within the topologies:

- Web proxy node
- Unified/single node
- Application node
- Database node

Each node type is comprised of one or more of the following components (software subsystems):

Component	Description
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, for example, two 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 matrix outlined in the table describes the set of components in each node type:

Node type	Components				
	Operating system	Platform	Web server	Database	Application
Web proxy	X	X	X		
Unified/single node	X	X	X	X	X
Application	X	X			X
Database	X	X		X	

3.1.3. Unified node cluster topology

Automate's **unified node cluster** topology provides the following options:

- *Single-node cluster (cluster-of-one/standalone) (testing-only)*
- *Single-node cluster (cluster-of-one/standalone) with VMWare HA*
- Two node with web proxies
- *Four node with web proxies*
- *Six node with web proxies*

Important: Choose between a unified node deployment or a modular architecture deployment.

In a *unified node cluster* deployment, Automate is deployed as *one* of the following:

- A single unified node cluster
- Two unified nodes
- 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:

Functional role	Description
Web proxy	Load balances incoming HTTP requests across unified nodes.
Single unified node	Combines the Application and Database roles for use in a non-multi-clustered test environment.
Unified	Similar to the <i>Single unified node</i> role Application and Database roles, but clustered with other nodes to provide HA and DR capabilities.

The nginx web server is installed on the web proxy, *Single Unified Node*, and the *Unified Node Cluster*, but is configured differently for each role.

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

Automate supports deployment of either the web proxy node or a DNS load balancer. Consider the following when deciding whether to select a web proxy node or a DNS:

- The web proxy node takes load off the *unified node cluster* to deliver static content (HTML/JAVA scripts). When using DNS or a third-party load balancer, the *unified node cluster* must process this information.
- DNS is unaware of the state of the *unified node cluster*.
- The web proxy detects if a *unified node cluster* is down or corrupt. In this case, the web proxy will select the next *unified node cluster* in a round robin scheme.

Important: It is recommended that you run no more than two *unified node clusters* and one web proxy node on a physical (VMware) server.

Additionally, it is recommended that the disk sub-systems are unique for each *unified node cluster*.

The table describes the defined deployment topologies for test and production:

Deployment topology	Description
Test	A standalone, <i>single unified node</i>, with Application and Database roles combined. No high availability or disaster recovery (HA/DR) is available. Important: A test deployment must be used only for test purposes.
Production with unified nodes	In a clustered system, comprising: • Two, three, four, or six unified nodes (each with combined Application and Database roles) • Zero to four (maximum two if two unified nodes) web proxy nodes offering load balancing. The web proxy nodes can be omitted if an external load balancer is available.

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

Note: A *Single-node cluster (cluster-of-one/standalone)* deployment should be used *only* for test purposes.



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

Advantages	Disadvantages
• Smallest hardware footprint	• 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
• Smallest hardware footprint • Disaster recovery available	• Less throughput than clusters

Multi-node 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, ensuring that an Application role failure is gracefully handled. More may be added if web proxy nodes are required in a DMZ.

Important: 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. For this reason, 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, thus offering disaster recovery and high availability.
- For six unified nodes, all nodes in the cluster are active. For an eight node cluster (with latency between data centers greater than 10ms), the two nodes in the disaster recovery node are passive; that is, the `voss workers 0` command has been run on the disaster recovery nodes.


Note: Primary and fall-back secondary database servers can be configured manually. Refer to the *Automate Platform Guide* for details.

Example: Six node cluster

The diagram illustrates an example of a *six node cluster*:


**Example: Eight node cluster**

The diagram illustrates an example of an *eight node cluster*:




Example: Two web proxy nodes in a DMZ

The diagram illustrates an example of *two web proxy nodes in a DMZ*:



Example: Four web proxy nodes in a DMZ (two admin, two Self-service)

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



Two 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 in the following instances:

- 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 *Disaster recovery failover and recovery in a two node cluster* in the Platform Guide.

Example: Two node cluster

The diagram illustrates a *two node cluster*:



Four node with web proxies

The table describes the advantages and disadvantages of a *four node with web proxies* deployment topology:

Advantages	Disadvantages
• More disaster recovery scenarios supported • More throughput than 3 Node	• More hardware than 3 Node

Six node with web proxies

The following are characteristics of a *six node with web proxies* deployment topology:

- Typically deployed for multi-data center deployments
- Supports Active/Standby

3.1.4. Modular node cluster deployment topology

Overview

A *modular node cluster* topology has separate Application and Database nodes:

- Three Database nodes
- One to eight Application nodes
- Web proxies

A *modular node 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 Node Cluster* deployment.

Automate is deployed as a *Modular Node 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:

Functional role	Description
Web proxy	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 web proxy and application role node, but is configured differently for each role.

Related topics

Modular Architecture Multi-node Installation in the Install Guide.

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

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

Automate supports deployment of either the web proxy node or a DNS load balancer. When choosing between a web proxy node and a DNS, consider the following:

- The web 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 is unaware of the state of the Application role node.
- The web proxy detects if an Application role node is down or corrupt. In this case, the web proxy will select the next Application role node in a round robin scheme.

Important: It is recommended that you run no more than one Application role node and one Database role node and one web proxy node on a physical (VMWare) server. When choosing disk infrastructure, high volume data access by database role replica sets must be considered where different disk sub-systems may be required depending on the performance of the disk infrastructure.

The following *Modular Node Cluster* topology is recommended (minimum):

Important: *Single Unified Node* topologies are not available for *Modular Node Cluster* deployments.

- Production with nodes (in a clustered system of two data centers):
 - DC1 = Data center 1, a primary data center containing primary database node (highest database weight)
 - DC2 = Data center 2, a data recovery data center

The system comprises of the following nodes:

- Three nodes with application roles (two in DC1; one in DC2)
- Three nodes with database roles (two in DC1; one in DC2)

- Maximum two web proxy nodes if two data centers; offering load balancing. The web proxy nodes can be omitted if an external load balancer is available.

Multi-node modular node cluster with application and database nodes

To achieve geo-redundancy using Application and Database nodes, consider the following:

- Six Application and Database nodes (three nodes with an application role and three nodes with a database role) are clustered and split over two geographically disparate locations.
- Two web proxy nodes to provide High Availability so that an Application role failure is gracefully handled. More may be added if web proxy nodes are required in a DMZ.

Important: 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. For this reason, 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 places between all database role nodes, thus offering disaster recovery and high availability.
- All nodes in the cluster are active.


Note: Primary and fall-back secondary database servers can be configured manually. Refer to the *Automate Platform Guide* for details.

Example: Six node cluster


The diagram illustrates an example of a *six node cluster*:

**Example: Two web proxy nodes in a DMZ**

The diagram illustrates an example of *two web proxy nodes in a DMZ*:

**Example: Four web proxy nodes in a DMZ**

The diagram illustrates an example of *four web proxy nodes in a DMZ* (two admin, two Self-service):



3.1.5. Cloud deployments

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.


Advantages of a Cloud deployment topology:

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

3.1.6. 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, such as ServiceNow and Identity Providers (IdPs) for Single Sign-On (SSO) authentication.



3.2. Multi data center deployments

3.2.1. Overview

Multi-node clusters can be deployed in both *Active/Active* or *Active/Standby* configurations:

Configuration	Description
Active/Active	Active/Active configurations have all the unified nodes enabled for transaction processing. In order to run an Active/Active configuration, the latency (RTT) between data centers must not exceed 20ms.
Active/Standby	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. Use an Active/Standby configuration for higher latencies.

3.2.2. Switch to an Active/Standby configuration

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; that is, this setting will still apply after upgrades and system restart.

3.2.3. Setting web weights for an Active/Standby configuration

For an *Active/Standby* configuration, the proxy server web weights should be set to the unified nodes on the primary data center, using 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 the following web weights configuration:

```
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.

Note: Refer to the *Best Practices Guide* for more information on deployment models and web weight settings.

Example:

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 follows:

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

To ensure that the secondary data center is configured for a *Standby* mode, change the web weights to show userweights, as displayed 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
```

(continues on next page)

(continued from previous page)

```

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

```

The outcome is that the load balancing web weights have been changed to the unified nodes on the *primary* data center.

3.3. Clustering

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 (described in more detail in the *Automate Platform Guide*).

All communication between nodes is encrypted.

Node type	Ports
Web proxy	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)

Note:

- **22/ssh:** Used for remote administration
- **80 and 443:** Used for the web server
- **161 and 162:** Used for sending and receiving SNMP
- **8443:** Used for inter-cluster communication
- **27020 and 27030:** Used for database queries and replication

3.4. Network communications between nodes within the cluster

The cluster contains multiple nodes, which can be contained in separate, firewalled networks. Network ports must 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. Review and implement according to the deployment model in use.

Note: *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, refer to 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, only where the database node is separate from the application node (that is, 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, 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, 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

3.5. Network communications external to the cluster

Details in this section are all based on the default settings, which can vary depending on the application setup and network design (such as NAT) of the solution. Adjust 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. Details are provided for the following:

- Outbound communications to devices from the application/unified nodes
- Outbound to external systems from the proxy node
- Outbound to external systems from all nodes
- Inbound communications from external systems to the proxy node
- Inbound communications to all nodes
- On-line help links to external documentation

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

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

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

Inbound communications from external systems to the proxy node

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

Inbound communications to all nodes

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

On-line Help links to external documentation

To have access to the online help website URL, you may need to request that your network administrator provides access to the website.

4. Prepare to Install

Note: From release 21.4 onwards, 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.

4.1. Network Docker container range

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

If it is in use, use the following command to modify the Private Address Space, as described below: `network container range add <private IP>`

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.

View the current Private Address Space

Use the following command to display the current Private Address Space: `network container range list`

For example:

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

Modify the Private Address Space

Use the following command to modify the Private Address Space: `network container range add <private IP>`

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 the following command: `cluster run all network container range add <private IP>`

Setting a different Private Address Space on each node

If required, you can set the Private Address Space 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)
```

4.2. Backup size considerations

4.2.1. Default backup partition size

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 `tmux` session. For details, refer to [Using the `tmux` command](#).

4.2.2. Backup space requirements

Consider the following:

- A local backup requires free space of at least twice the size of the database. Preferably add an additional 30%.

At any given time, it is 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 an additional 30%.

- 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.

4.2.3. Run a size check

To determine the required space needed to perform a backup for a specific backup partition, from the CLI, run the following command:

```
backup create <location-name>
```

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.

4.2.4. Increase size

To increase the size of the backup partition, use the following command: `drives add`

Refer to 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.

4.2.5. Display size of current database

To display the size of the current database, run: `voss db_collection_stats all`

This command validates the size of the database. This total will be smaller than the suggested backup size.

4.3. Hardware Specifications

4.3.1. 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.

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 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: • 20 GB for OS • 50 GB for application: – 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: • 20 GB for OS • 50 GB for application: – 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: • 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: • 20 GB for OS • 50 GB for application: – 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: • 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: • 20 GB for OS • 50 GB for application: – 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: • 20 GB for OS • 50 GB for compressed backups • 50 GB for application: – 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: • 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

5. Automate Platform Install OVA

5.1. Installation overview

The installation process is divided into:

1. The VMWare installation of a node
 - *Platform install OVA on a VM*
2. Node setup for single node or multi-node installations
 - *Multinode installation*
 - *Single-node cluster (cluster-of-one) installation*

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

5.2. Platform install OVA on a VM

5.2.1. Overview

This topic describes the steps for a fresh install, using the latest OVA file.

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. Refer to [What to do if the OVA file is not available for your current version](#)

The steps for creating a VM and running the wizard should be completed for *each* node in your topology. Therefore, these tasks will need to be performed either once or multiple times during installation, depending on the topology you're installing.

The install tasks (creating the VM and running the install wizard) describe the steps for a common setup of a *single node* from the OVA file for each of the following fresh install scenarios:

- Standalone installation
- A node install during multi-node installation

Note: For the *node install during multi-node installation* option, refer to [Role of each VM installation for multi-node installation](#).

- Failover recovery

5.2.2. Step 1: Download the OVA file

Download the **OVA** file for your release from the **New Installation** folder on the client portal.

The downloaded OVA file is imported into VMware vCenter Server (when creating the VM). Only one OVA file is used to deploy all the functional roles. Each time you run the install wizard on the VM, you will select the relevant role for the node type you're installing on.

5.2.3. Step 2: Create the VM

Prerequisites:

- [Step 1: Download the OVA file](#)

To create the VM:

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, fill out a name for this server.
5. On the **Deployment Configuration** page, select the appropriate node type.

Note: Refer to the description for **role** in the install wizard step, 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** page, 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** checkbox 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 Multi-node Cluster Hardware Specification section in the Install Guide.
13. Next steps: [Step 3: Run the installation wizard on the VM](#)

5.2.4. Step 3: Run the installation wizard on the VM

Prerequisites:

- [Step 2: Create the VM](#)

To run the install wizard on the VM:

1. Power on the VM.
2. Configure the options in the installation wizard:

Step	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 47, 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	Choose a role for the node you're installing on. Note: only WebProxy, Application and Database nodes are used for a modular architecture installation. • A WebProxy role installs only the front-end web server together with ability to distribute load among multiple middle-ware 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.: 0:0:0:0:0:ffff:c0a8:6403/64

3. Once all details are entered, installation proceeds. Monitor install progress.

When the installation completes, the system reboots. Since all services will be stopped, this takes some time.

Related topics

- [Boot password and security](#)

5.2.5. Step 4: Finalize the installation

1. Once the OVA installation completes, a sign-in prompt for the platform user displays. This confirms that the system is ready for use.
2. Connect to newly deployed server CLI as the platform user.

A login message such as the following displays:

```
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
```

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.

3. Run `app status` on all application nodes and ensure the services are all running and reporting the correct version before continuing.
4. Return to Multi-node Installation, Standalone Installation or Failover step to complete the overall installation or failover recovery procedure.

- 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.

On a fresh install, if you run the install on a network with a DHCP server and encounter the following error, you can enter a valid DNS server address to continue the installation:

“Error: DNS server <DNS server> is either invalid or cannot be reached on the network”

5.2.6. Boot password 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  (*****)
```

5.2.7. Role of each VM installation for multi-node installation

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

- **For each web proxy 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 web proxy instance.
- **For each unified instance** (*Standard Topology only*):
 - 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)
- One web proxy node at the Primary site
- Two unified nodes (both Secondary)
- One web proxy node at the DR site

For an eight node multi cluster deployment, there are:

- Four unified nodes (one Primary and three Secondary)
 - One web proxy node at the Primary site
 - Two unified nodes (both Secondary)
 - One web proxy node 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 web proxy 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 Multi-node Installation section in the Install Guide.

Detailed configuration can be applied from the Command Line Interface (CLI). Use the following commands for details: `network help` or `network`

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.

5.2.8. What to do if the OVA file is not available for your current version

If the OVA file is not available for your current release, you will need to obtain and install the most recent release OVA for which there is an upgrade path to your current release. In this case however, you will be performing an upgrade of your system.

The table describes the steps for obtaining the relevant OVA file and applying the upgrade:

Scenario	If the OVA file is not available for your current release ...
Standalone installation	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 install during multi-node installation	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 <i>to the cluster</i> to upgrade it. Refer to the <i>Upgrade Guide with Delta Bundle</i>.
Failover recovery	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: 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 <i>to the cluster</i> to upgrade it. For details, refer to the <i>Upgrade Guide with Delta Bundle</i> and to the specific scenario Disaster Recovery steps in the <i>Platform Guide</i>.

5.3. Inspect the logs to troubleshoot installation

You can inspect the logs to troubleshoot an installation. For example, detailed platform commands display in the **execute.log** file.

In the **execute.log** file, log entries for the command execution have a **ui** column. Log entries that follow these show related commands.

To view only the commands in the `execute.log` file, open a new console and run the following command:

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

Note: Logs are rotated and install commands may no longer display after log rotation.

The following list displays 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 ...
```

(continues on next page)

(continued from previous page)

```

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

[...]

```

5.4. Unified Node Topology Installation

5.4.1. Installation Quick Reference

Note:

- These steps are described in detail in the Automate Install Guide.
- The standard `tmux` command should be used where indicated, and the `reconnect` parameter is available if needed:
 - `tmux` - start a new session
 - `tmux list-sessions` - show sessions already available
 - `tmux attach -t [session number]` - reconnect to a disconnected session

It is recommended that you use the `tmux` command to avoid failures if the connection is interrupted while running the command. If the connection is interrupted while running the command in `tmux` then the session can be retrieved by first listing the sessions PID currently running in `tmux`: `tmux list-sessions`, and then reconnecting to the session using `tmux attach -t [session number]`.

General Steps

1. Download 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 (2 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 Automate OVA to VMHost(s)
5. After the VM is created, select the CD ROM configuration and verify the **Connect at Power On** checkbox is enabled.
6. Power on the VM.
7. Configure options in the installation wizard.
8. Install VMWare Tools from Automate CLI as platform user using the following command: `app install vmware`
9. See the sections below for your chosen deployment model.

Single node cluster deployment

1. Prepare the node for cluster command via the SSH CLI: `cluster prepnode`
2. Add the node to the cluster: `cluster add <node_ip-address>`
3. Verify the node is a member of the cluster: `cluster list`
4. Initialize the database and clear all data: `voss cleardown`
 This step may take some time. Monitor progress using either of the following commands:
`Run log follow upgrade_db.log` or `log follow voss-deviceapi/app.log`
5. (Optional) Set Automate network domain: `cluster run all network domain <yourdomain>`
6. Check application status: `cluster status`
 - If any services are down, use the following command on the affected node to restart all services on that node: `cluster run <node_ip> app start`
7. Run the following command: `voss cleardown`
8. SFTP the install templates to the Automate server **media** directory.
9. Install Automate templates via VMWare Console CLI. It is recommended that you run this step in a terminal opened with the `tmux` command, or on the VMWare console:
 - a. Run `tmux`
 - b. Run `app template media/<VOSS Automate_Template_Name>`

- At the prompt, choose your deployment type during the template install execution, either “Enterprise” or “Provider”.

Depending on the selected deployment type, at the prompt, fill out and verify passwords:

- A top-level administrator password - fill out the password for “sysadmin”
- One administrator password - depending on the deployment:
 - * For Enterprise, fill out a password for “entadmin”
 - * For Provider, fill out 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 they prefer.

1. On each node that is not the designated primary unified node, prepare the servers for cluster command via the SSH CLI: `cluster prepnode`
2. Add all of the other nodes to the cluster: `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: `cluster list`
4. (Optional) Set Automate network domain: `cluster run all network domain <yourdomain>`
5. Set each unified node's database weight: `database weight <un-node_ip-address> <priority_weight>`
 - Run this command for all unified nodes primary and secondary.
 - For *two* unified node, priority weights of 40, 30 are recommended.
 - For *four* unified nodes, priority weights of 40, 30, 20, and 10 are recommended.
 - For *six* unified nodes, priority weights of 60, 50, 40, 30, 20, and 10 are recommended.
6. Provision the Automate cluster database. It is recommended that you run this step in a terminal opened with the `tmux` command, or on the VMWare console.
 - a. Run `tmux`
 - b. Run `cluster provision`
7. Check cluster application status: `cluster status`
 - If any services are down, use the following command to restart all services on the affected node:
`cluster run <node_ip> app start`
8. Run `voss cleardown`
9. SFTP the install templates to the Automate server **media** directory of the primary unified node.
10. Install Automate templates via VMWare Console CLI of primary unified node. It is recommended that you run this step in a terminal opened with the `tmux` command, or on the VMWare console.
 - a. Run `tmux`
 - b. Run `app template media/<VOSS Automate_Template_Name>`

11. During the template install execution, choose the deployment type, either “Enterprise” or “Provider”.

At the prompt, depending on your selected deployment type, enter and verify passwords:

- A top-level administrator password - fill out the password for “sysadmin”
- And one administrator password - depending on the deployment:
 - For Enterprise, enter a password for “entadmin”
 - For Provider, enter a password for “hcsadmin”

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

Post deployment

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

`https://<ip_address_or_dns_name_of_VOSS_Automate_PrimaryUN_or_WebProxy>`

2. Run the following security command: `cluster check`

Then inspect entries under **security**.

3. Run the following command: `voss migrate_summary_attributes device/cucm/HuntPilot`

5.4.2. 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: [Platform install OVA on a VM](#) and according to the node roles as indicated in [Role of each VM installation for multi-node installation](#).

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

Note:

- For a 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:

- [Multi-node 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 [Inspect the logs to troubleshoot installation](#) for troubleshooting logs during an installation.
- Before installing from release 24.2 onwards, ensure that an additional 70 GB disk has been made available for the Insights database.

See topics around adding hard disk space in the Automate Platform Guide and 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 `tmux` command should be used where indicated. See: [Using the tmux command](#).

Installation

Step 1: Install VMWare tools

Install VMware tools on each node:

1. Log in to each node, then run the following command: `app install vmware`
2. Verify that vmware is running: `app list`

Step 2: Add nodes to the cluster

1. 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 web proxy and unified node, *excluding the primary node*, run `cluster prepnode`.
2. Add nodes to the cluster.
 - a. Log in to the selected primary unified node.
 - b. Add the unified and web proxy nodes to the cluster: `cluster add <ip_addr>`

Note: You don't need 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`

Step 3: Add network domain and check network

1. 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.

2. 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.

Since database weights are not yet added, you can ignore the following errors:

```
database: not configured
```

Verification of database weights should be done when the `cluster check` command is run during the step following the provisioning step.

If a cluster is not yet provisioned, you can ignore port 443 errors from web proxies triggered by the `cluster check` command.

If the `cluster check` command triggers errors on other ports, for example, port 27020, you can run the following command to verify that the firewall service has started: `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.

Step 4: 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 Automate platform is deployed.

Step 5: Configure the cluster

To configure the cluster:

1. Use the following command to provide a weight for each database server: `database weight add <database_ip> <priority>`

Recommended weights:

- For two unified nodes, weights of 40, 30 are recommended
- For four unified nodes, weights of 40, 30, 20, and 10 are recommended
- For six unified nodes, weights of 60, 50, 40, 30, 20, and 10 are recommended

Higher values are prioritized.

Weights used for multinode cluster deployment with four unified nodes in a geo-redundant system containing two data center infrastructures in two physical locations:

- For the primary 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 nodes at the DR site, specify weights of 20 and 10

Weights used for multinode cluster deployment with six unified nodes in a geo-redundant system containing two data center infrastructures in two physical locations:

- For the primary node at the primary site, specify a weight of 60
- For the secondary 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 nodes at the DR site, specify weights of 20 and 10

2. From the selected primary unified node, now set it up as the primary unified node.

Run this step in a terminal opened with the `tmux` command.

- a. Run `tmux`
- b. Run `cluster provision`

For two web proxies and four unified nodes, allow approximately two hours for the operation to complete.

3. When provisioning is complete, use the following command to check that each node is contactable and that the time server is running on each: `cluster check`

- Restart any stopped services: `cluster run <node_ip> app start`
- **Is provisioning successful?**
 - **Yes.** Type `exit` in the terminal to close the `tmux` session.
 - **No.** If there are errors, keep the `tmux` terminal open for troubleshooting purposes, and then contact VOSS support.

4. (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 Automate Best Practices Guide.

5. (Optional) If required, enable or disable Self-service or admin web services on the web proxy nodes.

Note:

- This may be needed for security purposes.
 - Commands must be run on the relevant web proxy node, and *only* on a cluster (not on a single node cluster system).
 - 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`

Step 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 Automate platform is deployed.

Step 7: Initialize the database

To initialize the database and clear all data, run the following command on the primary unified node:

```
voss cleardown
```

This step may take some time.

To monitor progress, in a separate console on the primary unified node, you can run either `log follow upgrade_db.log` or `log follow voss-deviceapi/app.log`.

Step 8: Import the templates

To import the templates:

1. Copy the Automate template file to the primary unified node: `scp <VOSS Automate_template_file> platform@<unified_node_ip_address>:~/media`
2. Log in to the primary unified node and install the template.

Run this step in a terminal opened with the `tmux` command.

- a. Run `tmux`
- b. Run `app template media/<VOSS Automate_template_file>`

View console message:

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

3. At the prompt to select the product deployment type, choose an option:

- Enterprise
- Provider

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

4. Depending on the deployment type selected, at the prompt, fill out and verify the following:

- A top-level administrator password:

```
Please enter a password for "sysadmin"
```

- And one administrator password - depending on the deployment:

- For Enterprise deployment:

```
Please enter a password for "entadmin"
```

- For Provider deployment:

```
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 product deployment type. View console message based on the deployment type, either of the following:

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

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

Deployment-specific system artifacts are imported, and the following message displays:

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

- Python functions are deployed.
- System artifacts are imported.
- At the prompt, provide administrator passwords.

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

5. Review the output from the `app template` commands, and confirm that the install message displays:

```
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.

6. Are there errors?

- **No.** 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 Automate platform is deployed.
- **Yes.** The install script stops with a failure message describing the problem. Contact VOSS Support.

Step 9: (Optional) Install language template

This step is optional and required only for installing language templates for languages other than English.

1. Copy the language template file to the primary unified node: `scp <language_template_file> platform@<unified_node_ip_address>:~/media`
2. Log in to the primary unified node and install the template: `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.

Step 10: (Optional) Install Automate Phone-based registration

This step is optional and required only if the Automate Phone Based Registration Add-on is required.

If required, follow installation instructions in the Appendix of the Automate Core Feature Guide: *Install the Phone Based Registration Web Service*

Step 11: device/cucm/HuntPilot

1. Run the following command: `voss migrate_summary_attributes device/cucm/HuntPilot`

Step 12: 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.

1. Obtain the required license token from VOSS.
2. License through the GUI or CLI:
 - To license through the GUI, follow steps indicated in *Product License Management* in the Core Feature Guide.
 - To license through the CLI, follow steps indicated in *Product Licensing* in the Automate Platform Guide.

Step 13: 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 following command on each unified node where the drive has been added: `drives add sde insights-voss-sync:database`

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
```

(continues on next page)

(continued from previous page)

```
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.
```

5.4.3. 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: [Platform install OVA on a VM](#)

Note:

- Template installation and upgrade takes approximately two hours. You can follow the progress on the GUI transaction list.
- Before installing from release 24.2 onwards, 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.

It is recommended that you use the standard `tmux` command where indicated. See: [Using the tmux command](#).

Install single-node cluster (cluster-of-one)

Step 1: Install VMWare tools

1. Log in on the node as the platform user and run `app install vmware`.
2. Verify that vmware is running: `app list`

Step 2: Add nodes to the cluster

1. Prepare the node to be added as a single-node cluster ("cluster-of-one"): `cluster prenode`
2. Add the node to the cluster: `cluster add <ip_addr>`
3. 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. For Azure and AWS install installations *only*, first run: `cluster provision`
 Initialize the database and clear all data: `voss cleardown`
 This step may take some time. Monitor progress using either of the following commands:
 Run either `log follow upgrade_db.log` or `log follow voss-deviceapi/app.log`

Step 3: Add network domain and check network

1. 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.
2. 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.
3. Create 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 Automate platform is deployed.

Step 4: Import templates

1. Copy the Automate template file to the node: `scp <VOSS Automate_template_file> platform@<ip_address>:~/media`
2. Log in to the node and install the template.
 It is recommended that you run this step in a terminal opened with the `tmux` command.
 - i. Run `tmux 11. Run app template media/<VOSS Automate_template_file>`

View console message:

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

- At the prompt, select the product deployment type, either “Enterprise” or “Provider”.

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

Depending on the deployment type you selected, at the prompt, provide and verify passwords:

- A top-level administrator password - provide a password for “sysadmin”
- And one administrator password - depending on the deployment:
 - For Enterprise, enter a password for “entadmin”
 - For Provider, 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 product deployment type. View console message, based on the deployment type, either of the following:

"Importing EnterpriseOverlay.json"

"Importing ProviderOverlay.json"

Deployment-specific system artifacts are imported and a message displays:

Deployment-specific Overlay artifacts successfully imported.

- Python functions are deployed.
- System artifacts are imported.
- At the prompt, provide administrator passwords.

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

- 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.

5. Are there errors?

- **No.** If there are no errors indicated, it is recommended 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 Automate platform is deployed.

- **Yes.** If there are errors, the install script stops with a failure message listing the problem. Contact VOSS Support.

Step 5: (Optional) Install language templates

This step is required only for installing language templates for languages other than English.

1. Copy the language template file to the primary unified node: `scp <language_template_file> platform@<ip_address>:~/media`
2. Log in to the primary unified node and install the template: `app template media/<language_template_file>`

For example, to install French: `app template media/VOSS AutomateLanguagePack_fr-fr.template`

Step 6: (Optional) Install Automate Phone-based Registration

This step is optional and required only if the Automate Phone Based Registration Add-on is required. If required, follow the installation instructions in the Appendix of your Core Feature Guide: “Install the Phone Based Registration Web Service”

Step 7: device/cucm/HuntPilot

1. Run the following command: `voss migrate_summary_attributes device/cucm/HuntPilot`

Step 8: 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.

1. Obtain the required license token from VOSS.
2. License through the GUI or CLI:
 - License through the GUI? Follow steps indicated in Product License Management in the Core Feature Guide.
 - License through the CLI? Follow steps indicated in Product Licensing in the Platform Guide.

Step 9: 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 following command on each unified node where the drive has been added:

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

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.
```

5.4.4. 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.

5.4.5. 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: [Platform install OVA on a VM](#).
 - 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.

5.5. Modular Cluster Topology Installation

5.5.1. Modular architecture multinode installation

Overview

Note:

- A modular architecture installation is not supported for a single node cluster ("cluster of one") topology.

- Before installing from release 24.2 onwards, 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 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 start

Before continuing, you should have followed the OVA installation on each node according to the steps and preliminary requirements specified in: [Platform install OVA on a VM](#) and according to the node roles as indicated in [Role of each VM installation for 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:

- [Multi-node 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](#)
 - [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 [Inspect the logs to troubleshoot installation](#) for troubleshooting logs during an installation.
-

The standard `tmux` command should be used where indicated. Refer to [Using the tmux command](#).

Install modular multinode

Step 1: Install VMWare tools and add nodes to the cluster

Start by installing VMware tools on each node, then prepare each node to be added to the cluster, then add nodes to the cluster.

1. Install VMware tools on each node.
 - a. Log in to each node, then 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.

Note: The *primary site* or data center will contain the primary database node. The deploying admin user can choose any database node they prefer.

- b. On each node, run `cluster prepnode`.
3. Add nodes to the cluster.
 - a. Log in to the selected primary database node.
 - b. Run the following command to add the other database, application, and web proxy nodes to the cluster: `cluster add <ip_addr>`

Note: You won't need to add the selected primary database node to the cluster as it is added to the cluster automatically.

- c. Run the following command to verify the list of nodes in the cluster: `cluster list`

Step 2: Add network domain and check network

1. 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 you configured.

2. Check the network:

- a. From the selected primary database node, run the following command to verify the status of the cluster, network connectivity, disk status, and NTP: `cluster check`
- b. Verify the DNS configuration: `cluster run all network dns`

Each node responds with the DNS server address.

3. 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 Automate platform is deployed.

Step 3: Configure the cluster

1. From the selected primary database node, use the following command to provide a weight for each database server: `database weight add <database_ip> <priority>`

Recommended weights are as follows (the higher the value, the higher the priority):

- For the primary database node at the primary site (DC1), a weight of *40* is recommended
- For the secondary database node at the primary site (DC1), a weight of *30* is recommended
- For the secondary database node at the DR (Data Recovery) site (DC2), a weight of *10* is recommended

2. From the selected primary database node, run the provisioning step in a terminal opened with the `tmux` command:

Run `tmux`, then run `cluster provision`.

For two web proxies and four database, application nodes, allow approximately two hours for the operation to complete.

3. When provisioning is complete, check that each node is contactable and that the time server is running on each: `cluster check`

- If a service is down, to restart the service, run: `cluster run <node_ip> app start`
- If provisioning is successful, type `exit` in the terminal to close the `tmux` session.
- If there are errors, keep the `tmux` terminal open for troubleshooting purposes, and contact VOSS support.

4. On each of the new *application* nodes, use the following command to set the queues to "2": `voss queues 2`

Once you run the command, applications are reconfigured and the *voss-queue* process is restarted.

5. (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 Automate Best Practices Guide.

6. (Optional) If required, enable or disable Self-service or admin web services on the web proxy nodes. This may be needed, for example, for security purposes.

The commands must be run on the relevant web proxy node, and will automatically reconfigure and restart the *nginx* process. Some downtime will therefore 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`

7. 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.

Step 4: Initialize the database

To initialize the database and clear all data:

1. On an application node, run the following command: `voss cleardown`

This step may take some time. To monitor progress, in a separate console on the application node, run either: `log follow upgrade_db.log` or `log follow voss-deviceapi/app.log`

Step 5: Import the templates

To import the templates:

1. Run the following command to copy the Automate template file to an *application node*: `scp <VOSS Automate_template_file> platform@<app_node_ip_address>:~/media`
2. Log in to this application node, and install the template. Run this step in a terminal opened with the `tmux` command.

```
tmux
```

```
app template media/<VOSS Automate_template_file>
```

View the console message:

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

3. At the prompt, select the product deployment type, either of the following:
 - Enterprise
 - Provider

Note: Contact VOSS for details on the *Insights Netflow* deployment type when installing release 24.2.

4. At the prompt, provide your password, depending on the deployment type you chose:

- Top-level admin password - fill out a password for *sysadmin*
- Fill out a deployment-specific admin password:
 - For Enterprise, fill out the *entadmin* password
 - For Provider, fill out the *hcsadmin* password

Note: At install, password length must be at least 8 characters.

Deployment-specific artifacts are installed. View the system message that displays, either of the following:

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

Deployment specific system artifacts are imported. View the system message that displays, either of the following:

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

- Python functions are deployed
- System artifacts are imported.
- At the prompt, provide administrator passwords.

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

5. 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.

6. **Are there errors?**

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

Step 6: (Optional) Install language templates

This step is required only for installing language templates for languages other than English.

1. Copy the language template file to the *selected application node*: `scp <language_template_file> platform@<app_node_ip_address>:~/media`
2. Log in to the *application* node and install the template:
`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.

Step 7: (Optional) Install Automate Phone-based Registration

This step is optional and required only if the Automate Phone Based Registration Add-on is required. If required, follow the installation instructions in the Appendix of your Core Feature Guide: "Install the Phone Based Registration Web Service"

Step 8: device/cucm/HuntPilot

1. Run the following command: `voss migrate_summary_attributes device/cucm/HuntPilot`

Step 9: 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.

1. Obtain the required license token from VOSS.
2. License through the GUI or CLI:
 - License through the GUI? Follow steps indicated in Product License Management in the Core Feature Guide.
 - License through the CLI? Follow steps indicated in Product Licensing in the Platform Guide.

Step 10: 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 following command on each unified node where the drive has been added:

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

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.
```

6. Automate Azure Cloud Installation

6.1. Azure Cloud Deployment

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

6.1.1. 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: Standard_E4as_v5
- CPU: 4
- RAM: 32
- OS Disk: 30 GB, Premium_LRS
- Application Disk: 50 GB, Standard_LRS
- Backup Disk: 55 GB, Standard_LRS
- DB Disk: 250 GB, Premium_LRS
- Insights Disk: 70 GB, Premium_LRS
- Total Disk size: 455 GB

Application Nodes:

- VM Size: Standard_E4as_v5
- CPU: 4
- RAM: 32
- OS Disk: 30 GB, Premium_LRS
- Application Disk: 50 GB, Standard_LRS
- Total Disk size: 80 GB

Web Proxies:

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

6.1.2. 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

6.1.3. 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 for each provisioner “file” block are correct in modules/voss-node/provision.tf.

```
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 `tmux` session.

```
tmux - start a new session
tmux list-sessions - show sessions already available
tmux attach -t [session number] - 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**

8. An install log file will be created to monitor the installation progress.

*On each of the newly deployed node(s), log in as the platform user, using the password configured in values.yaml (**automate_cli_password**).*

8.1. Run: **log list platform_install.run**

8.2. Run: **log follow install/platform_install.run-tty-<current_date>.log**

Once the installation is complete, the log file will return Platform installed successfully.

6.1.4. VOSS Automate Platform Config and Template Install Procedure

1. You are now ready to configure the platform and install the template.
2. *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)
```

Note:

- For both Azure and AWS, deploy using Automate 25.1 and later, the security check and security update commands are not available, since security updates are included during the release upgrade process.
-

- Single Node Cluster Guide:

https://documentation.voss-solutions.com/release_25.2-PB0/html/src/user/install/standalone-installation.html

- Modular Cluster Guide:

https://documentation.voss-solutions.com/release_25.2-PB0/html/src/user/install/modular-multinode-installation.html

7. Automate AWS Cloud Installation

7.1. AWS Deployment

7.1.1. 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

7.1.2. 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: 30 GB, GP3 @ 3000 IOPS / 125 throughput
- Application Disk: 50 GB, GP3 @ 3000 IOPS / 125 throughput
- Total Disk size: 80 GB

Database Nodes:

- Instance Type: r6a.xlarge or equivalent
- Instance CPU Architecture: x86_64
- CPU: 4
- RAM: 32
- OS Disk: 30 GB, gp3 @ 3000 IOPS / 125 throughput
- Application Disk: 50 GB, gp3 @ 3000 IOPS / 125 throughput
- Backup Disk: 125 GB, sc1
- DB Disk: 250 GB, io2 @ 750 IOPS
- Insights Disk: 70 GB, io2 @ 750 IOPS
- Total Disk size: 525 GB

Web Proxies:

- Web Proxies are replaced by an Application Load Balancer

7.1.3. 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

7.1.4. VOSS Automate AWS Deployment Topology



7.1.5. 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)

7.1.6. 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.

7.1.7. 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)

7.1.8. 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 25.1 and later, the security check and security update commands are not available, since security updates are included during the release upgrade process.

-
- Single Node Cluster Guide:

https://documentation.voss-solutions.com/release_25.2-PB0/html/src/user/install/standalone-installation.html

- Modular Cluster Guide:

https://documentation.voss-solutions.com/release_25.2-PB0/html/src/user/install/modular-multinode-installation.html

Index

A

app

- app install, [52](#), [57](#), [64](#), [69](#)
- app template, [5](#), [52](#), [57](#), [69](#)

B

backup

- backup create, [35](#)

C

cluster, [3](#)

- cluster add, [52](#), [57](#), [68](#), [69](#)
- cluster del, [52](#), [57](#), [69](#)
- cluster list, [68](#)
- cluster prenode, [52](#), [57](#), [68](#), [69](#)
- cluster primary, [4](#)
- cluster primary role application, [4](#)
- cluster primary role database, [4](#)
- cluster provision, [5](#), [52](#), [57](#), [68](#), [69](#)
- cluster run, [34](#), [52](#), [57](#), [68](#), [69](#)
- cluster status, [52](#), [57](#), [68](#), [69](#)
- cluster upgrade, [5](#)

D

database

- database weight, [52](#), [57](#), [69](#)

drives

- drives add, [35](#)

L

log

- log follow, [44](#), [52](#), [57](#), [64](#), [69](#)

N

network

- network container range, [34](#)
- network domain, [64](#)

S

system

- system reboot, [64](#)

T

- tmux, [5](#), [54](#), [57](#), [69](#)

V

voss

- voss cleardown, [44](#), [52](#), [57](#), [64](#), [69](#)
- voss db_collection_cap, [38](#)
- voss db_collection_stats, [35](#)
- voss queues, [69](#)
- voss upgrade_db, [44](#)
- voss workers, [12](#), [21](#)

voss export

- voss export group, [5](#)
- voss export type, [5](#)

- voss subscriber_data_export, [5](#)

W

web

- web service, [12](#), [21](#), [44](#)
- web weight, [25](#), [52](#), [57](#), [69](#)
- web weight add, [25](#)
- web weight list, [25](#)