



VOSS



**VOSS Automate
Architecture and Hardware Specification
Guide**

Release 24.1

Sep 12, 2024

Legal Information

- Copyright © 2024 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: 20240912150015

Contents

- 1 What's New** **1**
- 1.1 Architecture and Hardware Specification Guide: Release 24.1 1
- 2 VOSS Automate** **2**
- 2.1 VOSS Automate Deployment Topologies 2
- 2.2 VOSS Automate Hardware Specifications 17
- 2.3 Scale and Performance 23
- Index** **26**

1. What's New

1.1. Architecture and Hardware Specification Guide: Release 24.1

- VOSS-1187: Add Visualization Options and Configurable Dashboards to the Automate Portal. See: [VOSS Automate Hardware Specifications](#)
Docs added for the new Analytics dashboards in Automate.

2. VOSS Automate

2.1. VOSS Automate Deployment Topologies

2.1.1. Overview

Architecturally, VOSS Automate offers two main deployment topologies:

- *Unified Node Cluster Topology*
- *Modular Cluster Deployment Topology*

In addition, the following options are available:

- Cloud deployments
- VOSS Automate Cloudv1

2.1.2. Unified Node Cluster Topology

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

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

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

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

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

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

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

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

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

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

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

The following deployment topologies are defined:

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

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

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

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

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



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

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

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

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

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

Multinode Cluster with Unified Nodes

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

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

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


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

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

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


Example: 6-Node Cluster

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




Example: 8 Node Cluster

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



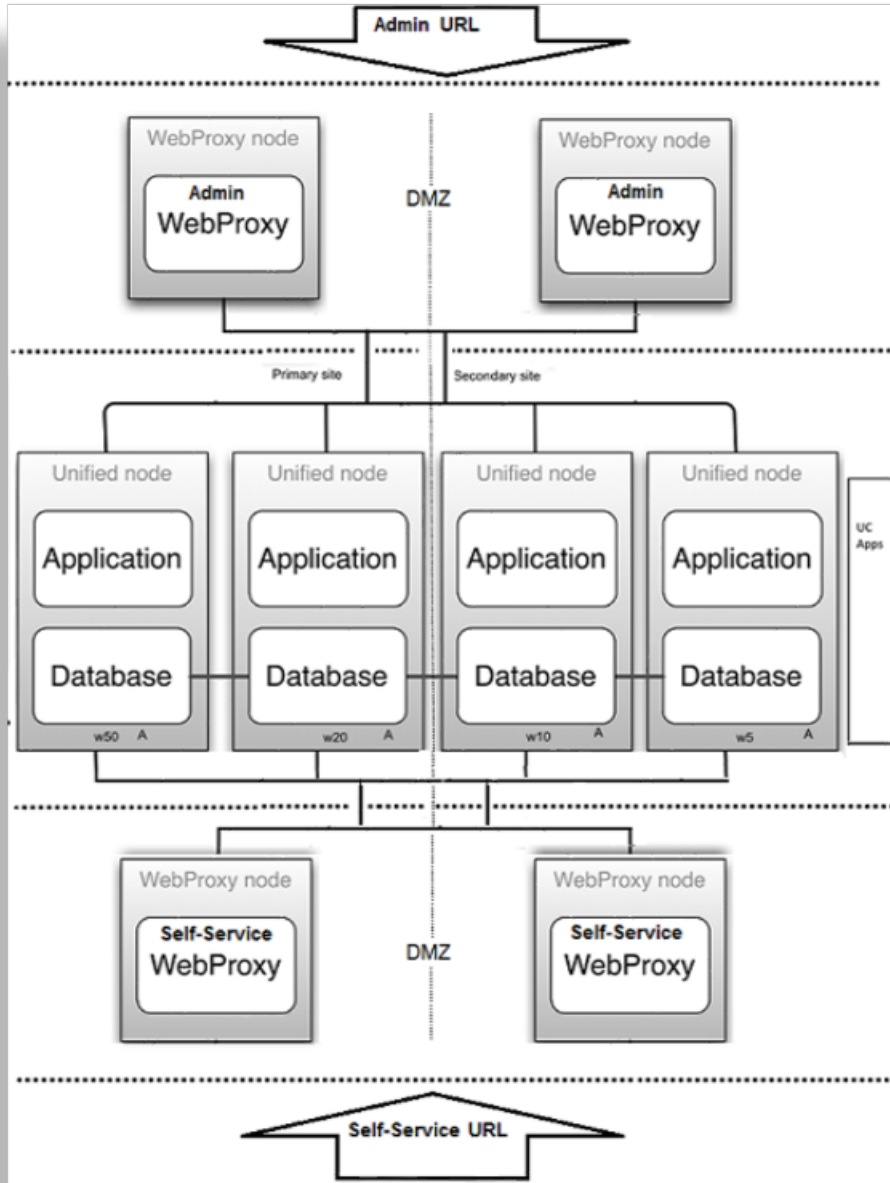
Example: 2 Web Proxy Nodes in a DMZ

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



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

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



2 Node Cluster with Unified Nodes

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

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

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

Important:

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


Only redundancy is available.

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

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

Example: 2 Node Cluster

The diagram illustrates a 2 Node Cluster:



4 Node with Web Proxies

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

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

6 Node with Web Proxies

A 6 Node with Web proxies deployment topology:

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

2.1.3. Modular Cluster Deployment Topology

Overview

A Modular Cluster topology has separate Application and Database nodes:

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

A Modular Cluster topology has the following advantages:

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

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

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

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

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

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

Related Topics

Modular Architecture Multinode Installation in the Install Guide.

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

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

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

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

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

The following modular cluster topology is recommended (minimum):

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

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

The system comprises of the following nodes:

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

Multinode Modular Cluster with Application and Database Nodes

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

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

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

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


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

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

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


Example: 6 Node Cluster

The diagram illustrates an example of a 6 Node Cluster:




Example: 2 Web Proxy Nodes in a DMZ

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

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

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



2.1.4. Cloud Deployments

VOSS Automate supports the following Cloud deployments:

- Azure
- Google Cloud Platform (GCP)
- Support all Standalone, Unified and Modular cluster topologies


The advantages of a Cloud deployment topology:

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

2.1.5. VOSS Automate Cloudv1 (SaaS)

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

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



2.2. VOSS Automate Hardware Specifications

2.2.1. Overview

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

Virtualized Hardware and Resource Oversubscription

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

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

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

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

2.2.2. Unified Node Hardware Specifications

Single-node Cluster (cluster-of-one) Hardware Specification

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

Node type	Quantity	VM	Memory	CPU	Disk	Network
Single node cluster	1	>= VMware 5.1	16 GB with 16 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: 10 GB for logs, 40GB 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 [Steps to 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	>= VMware 5.1	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: 10 GB for logs, 40GB 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	>= VMware 5.1	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

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

Note:

- From release 24.1, allowance should be made for an additional 70GB database disk to be added after upgrade or installation. This disk will be used for Insights sync.
Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Steps to 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	>= VMware 5.1	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	370 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: 10 GB for logs, 40GB 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	>= VMware 5.1	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

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

Note:

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

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Steps to 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.

2.2.3. 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	>= VMware 5.1	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	80 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application: 10 GB for logs, 40GB for our apps 	1 Gbit/s minimum
Database	3	>= VMware 5.1	32 GB with 32 GB reservation	4 vCPU @ 2 GHz with 4000 MHz reservation	380 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for compressed backups • 50 GB for application: <ul style="list-style-type: none"> – 10 GB for logs – 40GB 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	>= VMware 5.1	4 GB with 4 GB reservation	2 vCPU @ 2 GHz with no reservation	70 GB partitioned: <ul style="list-style-type: none"> • 20 GB for OS • 50 GB for application 	1 Gbit/s minimum

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

Note:

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

Refer to the Upgrade Guide with ISO and Template or Installation Guide and [Steps to 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.

2.2.4. Steps to 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).

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)

1. In the MS 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

2.3. Scale and Performance

2.3.1. Overview

This section describes the supported configurations and corresponding scale for unified deployment topologies and modular deployment topologies, including their Geo-Redundancy and Round Trip Time (RTT) requirements.

The supported configurations and scale are based on a standard profile for each subscriber, which includes the following:

- Two endpoints (physical phone and soft client) per subscriber.
- Voice Mail
- Extension Mobility (EM)
- Single Number Reach (SNR)
- IM and Presence

Scale limits are not enforced by the system. However, exceeding the limits can result in significant performance degradation.

With regards Geo-Redundancy options, the table describes the difference between Active and Standby nodes:

Active Node	A Node that can process transactions
Standby Node	A Node that is powered on and running the VOSS Automate software, but cannot process transactions.

For RTT across different UC components, this should not exceed 400ms. RTT will impact the duration of a transaction, but only for the remote leg, in other words, performing an action on the remote device.

2.3.2. Unified Node Deployments

Configuration	Number of Unified Nodes	Number of Web Proxy Nodes	Supported Scale (# Subscribers)	Geo-Redundancy (Y/N)
Single node cluster VOSS Automate	1	0	50,000	NA
Two-Node VOSS Automate Across Data Centers	2	1 ¹	50,000	Yes ² (Active-Standby) If RTT <= 10ms
Multi-Node VOSS Automate Across Data Centers	4	2 ¹	500,000	Yes (Active-Active) If RTT <= 10ms
	6	2 ¹	500,000	Yes (Active-Standby)
Multi-Node VOSS Automate One Data Center	4	2 ¹	500,000	No

¹ Two web proxy nodes can be used. It may be omitted if an external load balancer is available.

² With only two Unified nodes, with or without Web proxies, there is no High Availability. The database on the primary node is read/write, while the database on the secondary is read only.

Only redundancy is available.

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

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

2.3.3. Modular Node Deployments

Configuration	Number of Modular Nodes	Number of Web Proxy Nodes	Supported Scale (# Subscribers)	Geo-Redundancy (Y/N)
Application VOSS Automate Across Data Centers	1-8 ³	2 ¹	500,000	Yes (Active-Active) If RTT <= 10ms
	1	2 ^{Page 24, 1}	500,000	Yes (Active-Standby)
Modular Database VOSS Automate Across Data Centers	2	2 ^{Page 24, 1}	500,000	Yes If RTT <= 10ms
	1	2 ^{Page 24, 1}	500,000	Yes
Multi-Node VOSS Automate One Data Center	4	2 ^{Page 24, 1}	500,000	No
Modular Application VOSS Automate One Data Center	1-8 ³	2 ^{Page 24, 1}	500,000	No
Modular Database VOSS Automate One Data Center	3	2 ^{Page 24, 1}	500,000	No

³ Horizontal scaling of Application nodes allows for increased API performance and transaction throughput.

Index

V

voss

voss db_collection_cap, 19

voss workers, 5, 13

W

web

web service, 5, 13