



**VOSS**



**VOSS Automate  
Health Checks for Cluster Installations  
Guide**

Release 24.1-PB1

August 28, 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: 20240828152721

# Contents

- 1 Introduction** **1**
  
- 2 Notification/SNMP setup** **2**
  - 2.1 Purpose . . . . . 2
  - 2.2 Procedure . . . . . 2
  
- 3 Check General Cluster Health** **3**
  - 3.1 Services . . . . . 3
  - 3.2 Nodes in Cluster . . . . . 3
  - 3.3 Node Communication . . . . . 4
  - 3.4 NTP Connectivity . . . . . 4
  
- 4 Verify Web Proxy Sanity** **6**
  - 4.1 Purpose . . . . . 6
  - 4.2 Procedure . . . . . 6
  - 4.3 Step to Resolve . . . . . 6
  
- 5 Verify Database Status** **7**
  - 5.1 Database Health . . . . . 7
  - 5.2 Primary Database . . . . . 7
  - 5.3 Database Weights . . . . . 8
  
- 6 Resource Utilisation Checks** **9**
  - 6.1 Check Disk Space . . . . . 9
  - 6.2 Check Available RAM . . . . . 10
  
- 7 VMWare Checks** **11**
  - 7.1 Check VMWare Disk Space . . . . . 11
  - 7.2 VMWare Snapshots . . . . . 11
  - 7.3 VMWare recommendations . . . . . 12
  
- 8 Backup / Export** **13**
  - 8.1 Purpose . . . . . 13
  - 8.2 Procedure . . . . . 13
  
- 9 Latency** **14**
  - 9.1 Transaction Processing . . . . . 14
  - 9.2 UC Apps Latency . . . . . 14
  
- 10 Firewall configurations** **16**
  
- 11 Data Collection for Offline Analysis** **17**
  - 11.1 Database Data Information . . . . . 17
  - 11.2 Log Collection . . . . . 17

11.3 Export Installed Patch Information . . . . . 18

# 1. Introduction

This document serves as a brief checklist for VOSS Automate installation, deployment and maintenance.

For comprehensive details on the CLI commands, requirements, installation and maintenance procedures, please refer to the following guides:

- Architecture and Hardware Specification Guide
- Installation Guide
- Platform Guide
- Multi-Cluster Deployments Technical Guide

## 2. Notification/SNMP setup

### 2.1. Purpose

As good practice, we recommend that SNMP and/or notification be enabled and set up to proactively monitor the system.

### 2.2. Procedure

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology), and run:  
**notify list**
2. For the warn and error sections, please ensure that these have been set up with valid details (mail/snmp)
3. Then run:  
**snmp list**
4. The query field should have the valid credentials for the customers NMS system.

## 3. Check General Cluster Health

### 3.1. Services

#### 3.1.1. Purpose

If any services on the cluster are not running, it could indicate a problem in the system.

#### 3.1.2. Procedure

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology).
2. Run the following commands:  
**cluster status**  
and  
**cluster run all app status**
3. Check for any anomalous output, e.g. stopped services or unknown nodes or mismatched service versions.

#### 3.1.3. Step to Resolve

Start stopped services, resolve issues on non-responsive nodes. Escalate unresolvable issues to VOSS L2 helpdesk.

### 3.2. Nodes in Cluster

#### 3.2.1. Purpose

If all nodes in the cluster are not known to all other nodes, provisioning may fail.

### 3.2.2. Procedure

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology).
2. Run the following command:  
**cluster run database cluster list**
3. Ensure all nodes list the correct number of nodes.

### 3.2.3. Step to resolve

If one or more nodes do not list all nodes, the nodes may need to be deleted and re-added, possibly from a different unified node. Nodes can be added or deleted without any harm until all nodes show the same output of the cluster list command.

Escalate unresolvable issues to VOSS L2 helpdesk.

## 3.3. Node Communication

### 3.3.1. Purpose

Ensure the nodes in the cluster can freely communicate.

### 3.3.2. Procedure

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology).
2. Run a cluster command across all nodes, for example:  
**cluster run all network list**
3. Verify that all nodes respond with the expected output.

### 3.3.3. Step to resolve

Go back to checking the general health of the cluster.

## 3.4. NTP Connectivity

### 3.4.1. Purpose

Ensure NTP is accessible in order to prevent failures such as unexpected session timeout.



### 3.4.2. Procedure

For each node:

1. Log in as root.
2. Run the following command:

```
ntpq -p
```

3. The output will show a result for the **reach** metric. A value of 377 indicates that there has been no packet loss, while a value less than 377 shows that there was some packet loss. A value of zero will be a cause for concern.

### 3.4.3. Step to resolve

In the event that the **reach** parameter returns with a value of 0, restart the time service by running the following command:

```
app start services:time --force
```

Repeat the procedure above. If the problem persists, contact VOSS L2 Helpdesk.

## 4. Verify Web Proxy Sanity

### 4.1. Purpose

To prevent cluster processes from failing silently due to a misconfigured cluster, a web proxy should not have any database or VOSS application services present or running.

### 4.2. Procedure

1. Log in on all webproxy nodes.
2. Run the following command:  
**app status | grep "voss-deviceapi|selfservice|mongodb"**
3. Check for any output from the command. A healthy node should simply return to the command prompt with no output from the command.

### 4.3. Step to Resolve

If there are problems identified on any of the webproxy nodes, contact VOSS L2 helpdesk.

## 5. Verify Database Status

### 5.1. Database Health

#### 5.1.1. Purpose

To ensure that the database is in healthy state

#### 5.1.2. Procedure

1. Log in on any unified node (multinode unified topology) / database node (modular cluster topology).
2. Run the following command:  
**database config**
3. Verify that the stateStr of each node is one of the following values:
  - stateStr: PRIMARY
  - stateStr: ARBITER
  - stateStr: SECONDARY

#### 5.1.3. Step to Resolve

If any node has a stateStr not listed above, contact VOSS L2 helpdesk. Provisioning must not take place if any of the database nodes are in STARTUP, STARTUP2 or RECOVERING state.

### 5.2. Primary Database

#### 5.2.1. Purpose

To ensure the primary database is the correct node

### 5.2.2. Procedure

1. Log in on any unified node (multinode unified topology) / database node (modular cluster topology).
2. Run the following commands:  
**database primary**  
and  
**cluster run database database primary**
3. Ensure the IP address matches the intended primary database expected.

### 5.2.3. Steps to Resolve

If a failover has occurred for any reason and the primary database has changed from what is expected, refer to the “Check General Cluster Health” section.

## 5.3. Database Weights

### 5.3.1. Purpose

Database weights are used to determine how a new primary node is elected in the event of database primary node failover. Although any values can be used, for 4 database nodes the weights: 40/30/20/10 is recommended and for 6 database nodes, 60/50/40/30/20/10. These numbers ensure that if a reprovision happens (when the primary data center goes offline for an indeterminate time), the remaining nodes have weights that will allow a new primary to be chosen.

### 5.3.2. Procedure

1. Log in on any unified node (multinode unified topology) / database node (modular cluster topology).
2. Run the following commands:  
**database config**
3. Verify that weights are set with highest numbers in primary Data Center (DC), and lesser weights in secondary DC.

### 5.3.3. Steps to Resolve

Fix database weights to have the highest numbers at primary DC.

## 6. Resource Utilisation Checks

### 6.1. Check Disk Space

#### 6.1.1. Purpose

To ensure that there is enough disk space.

#### 6.1.2. Procedure

1. Login to all nodes.
2. Run the following command:  
**diag disk**
3. Verify that the following disks are not over 85%:

```
/
/opt/platform
/tmp
/var/log
/opt/platform/apps/mongodb/chroot/dbroot
  (this is only on:

    unified nodes for multinode unified topology
    database nodes for modular cluster topology)
```

#### 6.1.3. Step to Resolve

If any node is above the threshold, please clean if possible, else contact VOSS L2 support.

## 6.2. Check Available RAM

### 6.2.1. Purpose

To ensure that the RAM available per node is aligns with the scale of the platform.

- Multinode unified topology: unified nodes must be allocated a minimum of 16GB.
- Modular cluster topology: application nodes must be allocated a minimum of 16GB.
- Modular cluster topology: database nodes must be allocated a minimum of 32GB.
- Each WebProxy should have a minimum of 4GB RAM.

### 6.2.2. Procedure

1. Login to one node as platform.
2. Run the following command:  
**cluster run all diag free**
3. Verify that total RAM aligns with the scale of the platform. Note: the command above shows RAM in kilobytes.

### 6.2.3. Step to Resolve

If any node is below the threshold, please allocate more RAM to the virtual machine.

## 7. VMWare Checks

### 7.1. Check VMWare Disk Space

#### 7.1.1. Purpose

To ensure the datastore(s) the cluster is using will be sufficient for the cluster.

#### 7.1.2. Procedure

1. Log into the VMWare server(s) hosting the cluster nodes.
2. Identify the Datastore each node is using.
3. Ensure there is sufficient space for all systems.

#### 7.1.3. Step to resolve

If there is any doubt that datastore space may not be sufficient, contact your VMWare administrator.

### 7.2. VMWare Snapshots

#### 7.2.1. Purpose

VMWare snapshot management is vital to ensure optimal performance.

### 7.2.2. Procedure

1. Log in on one VMWare host(s)
2. Ensure that there is no more than one snapshot, if possible. Delete any old snapshots.

## 7.3. VMWare recommendations

### 7.3.1. Purpose

VOSS recommends that there be a 1 to 1 mapping for memory and CPU in VMWare.

### 7.3.2. Procedure

1. Log into the VMWare console, select the VOSS nodes one by one.
2. Ensure that the system has a 1 to 1 mapping for memory and CPU. For example, if the system is set to 16GB, there is 16GB RAM reserved, and if CPU is set to 4 cores, that the VMWare host has 4 cores available.



## 8. Backup / Export

### 8.1. Purpose

As good practice, a customer should run a backup on a regular basis.

---

**Note:** To carry out these tasks on a multinode unified node topology or on a modular cluster topology, log in on the database node with the second highest weight.

---

### 8.2. Procedure

1. Create a localbackup with:

**backup create localbackup**

2. Add a remote location for the export with:

**backup add <remotename> sftp://<sftpusername>:<sftppassword>@<IP address>**

3. Export the local backup with:

**backup export localbackup <remotename> <timestamp>**

(<timestamp> seen with the **backup list** command)

4. Ensure that the backup gets run regularly, using the **schedule** commands.

For more details and examples on backup and restore, refer to the Backup and Import topic in the Platform Guide.

## 9. Latency

### 9.1. Transaction Processing

#### 9.1.1. Purpose

Transactions should not be run by nodes over a high latency network, if possible.

#### 9.1.2. Procedure

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology).
2. Run the following for all the other nodes:  
**diag ping <IP>**
3. If any of the average round trip time is higher than 10ms, it is recommended that the nodes in the secondary DC should have `voss workers` set to 0 on version 11.5.3 and later, or the command **app stop voss-queue** is run on older versions.

---

**Note:** For software versions *pre* 11.5.3, the **app stop voss-queue** command needs to be run after every service restart or server reboot.

---

### 9.2. UC Apps Latency

#### 9.2.1. Purpose

Latency between unified nodes and UC apps impacts overall provisioning times against the UC apps. This check is done as part of diagnosis of provisioning performance problems.

## 9.2.2. Procedure

For each data center:

1. Log in on any unified node (multinode unified topology) / application node (modular cluster topology).
2. Run the following for all the UC apps under investigation:  
**diag ping <IP>**
3. Record the output and share results with VOSS L2.

## 10. Firewall configurations

Incorrect firewall rules can cause outages and make it difficult to resolve issues. These need to be verified by the customer's network/firewall team.

1. Ensure that the connectivity between all VOSS nodes allows bidirectional traffic for ports 80, 443 and 8443. For example, to test platform API connectivity on port 8443 from all other hosts back to a node with an IP address of `10.0.0.10`:
  - a. SSH to `10.0.0.10`
  - b. Run `cluster run all diag test_connection 10.0.0.10 8443 --force` to test connectivity **from** the other hosts in the cluster.
2. Ensure that ports 27020 and 27030 are bidirectionally open between:
  - unified nodes (multinode unified topology)
  - database nodes (modular cluster topology)
  - database and application nodes (modular cluster topology)

For example, to test connectivity from all unified to the arbiter running on a primary node with IP address `10.0.0.10`:

- a. SSH to `10.0.0.10`
  - b. Run `cluster run database diag test_connection 10.0.0.10 27030 --force` to test connectivity from the unified hosts (multinode unified topology) or database nodes (modular cluster topology) in the cluster.
3. From VOSS unified nodes (multinode unified topology) / application and database nodes (modular cluster topology), ensure that all Cisco equipment managed by VOSS is accessible on the relevant ports. For example, to test connectivity from a VOSS Automate cluster to a CUC on `172.16.0.10`:
    - a. SSH to the primary unified node (multinode unified topology) / application node (modular cluster topology).
    - b. Run `cluster run application diag test_connection 172.16.0.10 443` to test HTTPS connectivity to a remote host.

# 11. Data Collection for Offline Analysis

## 11.1. Database Data Information

### 11.1.1. Purpose

Collect information relating to actual database data size usage and index configuration.

### 11.1.2. Procedure

1. Log in on the primary unified node.
2. Run the following commands and save the output in a file:  
**voss db\_collection\_stats**  
and  
**voss db\_index\_stats**
3. Send the file with the output to VOSS L2.

## 11.2. Log Collection

### 11.2.1. Purpose

Logs must be extracted to enable offline performance analysis of a platform.

### 11.2.2. Procedure

For each node:

1. Log in as platform.
2. Run the following command:  
**log collect start YYYY-MM-DD end YYYY-MM-DD** Note: **start** and **end** must cover the period over which performance analysis is required.

3. Send the output files from each node to VOSS L2.

## 11.3. Export Installed Patch Information

### 11.3.1. Purpose

Ensures platform has all recommended patches installed.

### 11.3.2. Procedure

1. Log into VOSS Automate Admin Portal.
2. Navigate to **About > Extended Version** menu.
3. Open Patches tab.
4. Export to json (via actions button top right).
5. Send the output files from each node to VOSS L2.