



VOSS-4-UC Platform Guide

Release 21.1

Sep 02, 2021

Legal Information

Please take careful note of the following legal notices:

- Copyright © 2021 VisionOSS Limited.
All rights reserved.
- VOSS, VisionOSS and VOSS-4-UC are trademarks of VisionOSS Limited.
- No part of this document may be reproduced or transmitted in any form without the prior written permission of VOSS.
- VOSS does not guarantee that this document is technically correct, complete, or that the product is free from minor flaws. VOSS endeavors to ensure that the information contained in this document is correct, whilst every effort is made to ensure the accuracy of such information, VOSS accepts no liability for any loss (however caused) sustained as a result of any error or omission in the same.
- This document is used entirely at the users own risk. VOSS cannot be held responsible or liable for any damage to property, loss of income, and or business disruption arising from the use of this document.
- The product capabilities described in this document and the actual capabilities of the product provided by VOSS are subject to change without notice.
- VOSS reserves the right to publish corrections to this document whenever VOSS deems it necessary.
- All vendor/product names mentioned in this document are registered trademarks and belong to their respective owners. VOSS does not own, nor is related to, these products and vendors. These terms have been included to showcase the potential of the VOSS solution and to simplify the deployment of these products with VOSS should you select to utilize them.

Security Information

This product may contain cryptographic features that may be subject to state and local country laws that govern the import, export, transfer and use of such features. The provision of this software does not imply that third-party authorization to import, export, distribute or use encryption in your particular region has been obtained. By using this product, you agree to comply with all applicable laws and regulations within your region of operation. If you require further assistance, please contact your dedicated VOSS support person.

Contents

1	What's New	1
1.1	Platform Guide: Release 20.1.1	1
1.2	Platform Guide: Release 21.1	2
2	Overview	4
2.1	High-level Functions	4
3	The Command Line Interface	5
3.1	Overview	5
3.2	CLI Commands	5
3.3	Using the <code>screen</code> command	7
3.4	System Specific Commands	8
3.5	System Metrics	22
3.6	Data Export Commands	28
4	Node Deployment	32
4.1	Deployment	32
4.2	Modular Cluster Deployment	32
4.3	Migrate a Unified Node Cluster to a Modular Cluster	34
4.4	Primary Node Role in a Cluster	38
5	Provisioning	39
5.1	Provisioning	39
6	Networking	40
6.1	Network interfaces	40
6.2	Network services	40
6.3	Network URI specification	41
6.4	Network Docker Container Range	42
7	Application Control	43
7.1	Application Control and Status	43
7.2	Starting and Stopping	44
7.3	Installing Applications	45
7.4	Updating Applications	47
7.5	Summary Attribute Migration	48
7.6	Remote Execution in Clusters	48
7.7	List of Unused Cluster Commands	49
7.8	Self-Service Localization Management	50
7.9	Web Services	51
8	System Control	52
8.1	System Commands Overview	52

8.2	System restart	52
8.3	Passwords	52
8.4	System Boot Passwords	53
8.5	Federal Information Processing Standards (FIPS)	54
8.6	File Management	55
8.7	Drive control	56
8.8	Transaction Prioritization	58
8.9	Banner	59
8.10	Checksum	59
9	Diagnostics	61
9.1	Health Report	61
9.2	Cluster Check	62
9.3	Enable Health Monitoring	65
9.4	Enable Database Scheduling	66
9.5	Command History	66
9.6	Logs	67
9.7	Viewing Logs	68
9.8	Sending and Collecting Logs	69
9.9	Log Types	72
9.10	Audit Log Rule Sets	74
9.11	Log Type Commands	77
9.12	Audit Log Format and Details	79
9.13	Event Log Format and Details	85
9.14	Log Streaming	87
9.15	Remote Log Type Encryption	89
9.16	The Mail Command	90
9.17	Diagnostic Tools	91
9.18	Diagnostic Troubleshooting	92
10	Notifications	96
10.1	Warnings and Notifications	96
10.2	Events and SNMP Messages	97
10.3	SNMP Configuration and Queries	99
11	SNMP	102
11.1	Introduction to SNMP and MIB	102
11.2	SNMP Traps	103
11.3	Management Information Bases	104
11.4	MIB and Trap Details	105
11.5	VOSS-4-UC System Monitoring Traps	138
12	Scheduling	145
12.1	Scheduling	145
12.2	Internal Report Schedules	146
13	Backups	148
13.1	Backups	148
13.2	Backup and Restore in a Modular Architecture	148
13.3	Backup Destinations	149
13.4	Backup Passphrase	150
13.5	Backup Size Considerations	150
13.6	Create a Backup	151
13.7	Restore a Backup in a Clustered Environment	152
13.8	Create Space for a Backup or Restore	152

13.9	Maintaining Backups	153
13.10	Backup and Import to a New Environment	153
13.11	VMware Snapshot Maintenance	155
13.12	Restoring a Backup on a New Environment	156
14	System Security	160
14.1	Security Overview	160
14.2	Security Patches and Updates	160
14.3	Configuration Encrypted	161
14.4	Backup Encrypted	162
14.5	Application Install Files Encrypted	162
14.6	File Integrity	162
14.7	Protected Application Environments (Jails)	162
14.8	Restricted User Shell	163
14.9	User Security and Security Policy Management	163
14.10	Creating Additional Users	166
14.11	Creating and Managing SFTP Users	166
14.12	Granting and revoking user rights	167
14.13	Password Strength Rules	168
14.14	SSH Login Fail Limit	168
14.15	SSH Session Limit	170
14.16	SSH key management	171
14.17	SSH Algorithm Management	171
14.18	Adding a Key for Automatic User Login	172
14.19	Prevention of DOS Attacks	173
14.20	Memory Dumps and Security	174
14.21	Manage Read-Only Database Users	174
15	Network Security	175
15.1	Network Communications between Nodes within the Cluster	175
15.2	Network Communications External to the Cluster	176
15.3	Dynamic Firewall	177
15.4	Web Certificate Setup Options	178
15.5	VOSS-4-UC Setup a Web Certificate	180
15.6	Own Web Certificate Setup	181
15.7	Web Certificate Expiration Notice	182
15.8	Convert Web Certificates from P7B to PEM Format	183
15.9	Web Certificate Commands	184
15.10	Web TLS Protocol Configuration	184
15.11	Web TLS Cipher Management	186
15.12	Network URI specification	187
15.13	Web External Proxies	188
16	High Availability and Disaster Recovery (DR)	189
16.1	High Availability Overview	189
16.2	Default HA and DR scenario	189
16.3	HA and DR scenario with Cisco VMDC geo-redundancy architecture	190
16.4	Configuring a HA System Platform on VMware	190
16.5	DR Failover and Recovery Overview	191
16.6	DR Failover and Recovery in a Unified Node Cluster Topology	198
16.7	DR Failover and Recovery in a Modular Cluster Topology	225
17	Troubleshooting	246
17.1	Platform User Password Recovery Procedure	246
17.2	'No Space Left on Device' Error	247

17.3	Loss of the whole cluster and redeploying new servers	247
17.4	Memory (RAM) Increase for Large End User Capacity	250
17.5	Error Messages	250
18	Appendices	268
18.1	MIBs	268
18.2	Data Export Types	278
18.3	Command Examples	298
	Index	304

1. What's New

1.1. Platform Guide: Release 20.1.1

- EKB-4115: Hourly SNMP Trap WARNING: High CPU usage. See: [SNMP Trap: Excessive Load](#)
- EKB-4149: The system download command fails when URL is percent-encoded. See: [Network URI specification](#)
- EKB-4891: CLI command to install multiple patches in order. See: [Installing Applications](#)
- EKB-5146: Manual runs of license audit do not complete (without screen session). See: [Subscriber Data Export Command](#)
- EKB-5475: Update SDE to run off secondary unified nodes. See: [Subscriber Data Export Command](#)
- EKB-5610: Add support for tar.gz to app multi_install command. See: [Installing Applications](#)
- VOSS-612: Audit log filtering. See: [Log Types](#)
- VOSS-612: Audit log filtering. See: [Audit Log Rule Sets](#)
- VOSS-661: Upgrade platform OS to Ubuntu 18.04 (EKB-4494: TLSv1.3 does not allow disabling of ciphers, and supercedes v1.1 and 1.2 completely) . See: [Web TLS Protocol Configuration](#)
- VOSS-670: Small Enhancements 20.1.1 (EKB-4964: Include install logs in log collect) . See: [Sending and Collecting Logs](#)
- VOSS-670: Small Enhancements 20.1.1 (EKB-5806: Write cluster check output to log file) . See: [Sending and Collecting Logs](#)
- VOSS-670: Small Enhancements 20.1.1 (EKB-4720: Change backups to run off secondary unified nodes) . See: [Create a Backup](#)
- VOSS-722: VOSS Phone Server Solution. See: [Subscriber Data Export](#)
- VOSS-724: Upgrade Robustness and Improvements 20.1.1 (EKB-5299: Make run of upgrade commands in screen mandatory) . See: [Using the screen command](#)
- VOSS-724: Upgrade Robustness and Improvements 20.1.1 (EKB-4942: Limit the output of cluster check platform to only failed indications) . See: [Cluster Check](#)
- VOSS-724: Upgrade Robustness and Improvements 20.1.1 (EKB-5579: Incorporate security check in cluster check) . See: [Cluster Check](#)
- VOSS-724: Upgrade Robustness and Improvements 20.1.1 (EKB-5603: Incorporate cluster status into cluster check) . See: [Cluster Check](#)
- VOSS-724: Upgrade Robustness and Improvements 20.1.1 (EKB-5298: Check on database weights for the cluster and report on discrepancies) . See: [Cluster Check](#)

1.2. Platform Guide: Release 21.1

- EKB-8282: app multi_install requires delete-on-success. See: [Installing Applications](#)
- EKB-9459: Add Phone Expansion Module fields (Module 1, 2 and 3) to Phones SDE Extract. See: [Phones Data Export](#)
- VOSS-643: System Monitoring Enhancements (EKB-8451: Monitor lock/unlock events for platform CLI users) . See: [Audit Log Format and Details](#)
- VOSS-643: System Monitoring Enhancements (EKB-8452: Monitor lock/unlock events for Self-service/admin users) . See: [Audit Log Format and Details](#)
- VOSS-643: System Monitoring Enhancements (EKB-8744: Add capability to stream all platform log files to a remote syslog server) . See: [Log Streaming](#)
- VOSS-643: System Monitoring Enhancements (EKB-8453: Monitor nginx certificate expiry) . See: [SNMP Trap: Web Certificate Expiry](#)
- VOSS-669: Drop Support for Standalone Deployments. See: [Deployment](#)
- VOSS-728: Improve Log Collection and Logging (EKB-5111: Improve log collect) . See: [Sending and Collecting Logs](#)
- VOSS-772: Allow redirects from VOSS-4-UC proxy nodes. See: [Web External Proxies](#)
- VOSS-790: Pexip Support. See: [Subscriber Data Export](#)
- VOSS-794: Improve transaction log management (EKB-1173: Support change cap size for selected collections) . See: [Database Commands](#)
- VOSS-837: Support for Modular Architecture. See: [Modular Cluster Deployment](#)
- VOSS-837: Support for Modular Architecture. See: [Migrate a Unified Node Cluster to a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture (EKB-8170: Migrate existing unified node into a database node) . See: [Migrate a Unified Node Cluster to a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture (EKB-8171: Migrate existing unified node into an application node) . See: [Migrate a Unified Node Cluster to a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture (EKB-8463: Handle backup and restore in new topology) . See: [Backup and Restore in a Modular Architecture](#)
- VOSS-837: Support for Modular Architecture. See: [Create a New VM Using the Platform-Install OVA](#)
- VOSS-837: Support for Modular Architecture. See: [Modular Cluster: Election of a New Primary and Failover](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Power Off and On of a Node in a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of an app node: Modular Cluster](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of the Primary Database Server in a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of a non-primary database: Modular Cluster](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of a Primary Site in a Modular Cluster](#)

- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of a DR Site in a Modular Cluster](#)
- VOSS-837: Support for Modular Architecture. See: [Scenario: Loss of Full Cluster in a Modular Cluster](#)
- VOSS-847: Support for Microsoft O365 and MS Teams Management. See: [Subscriber Data Export](#)

2. Overview

2.1. High-level Functions

The VOSS-4-UC platform is an Infrastructure As A Service layer (IAAS) built on top of Ubuntu Linux.

This platform layer supports the following high-level functions:

- Installation, upgrades
- Application and process manipulation
- Clustering of multiple nodes with High Availability (HA) and Disaster Recovery (DR) capabilities
- Backup creation and restore
- Scheduling of tasks
- Security implementation
- System diagnostics

Both the platform and application are designed as a loose collection of processes which can be deployed in a wide range of topologies. Individual nodes can be clustered and provisioned together to provide High Availability and Disaster Recovery.

3. The Command Line Interface

3.1. Overview

Maintenance is carried out from a platform user login application command line, either by SSH or from the VM console command line. The password is configured during installation and can be changed using **system password**. On initial login, the system displays a banner indicating the general system health.

A local home directory is available to the user and must be managed by the user with standard Unix commands:

- **ls**
- **cp**
- **mv**
- **rm**
- **less**
- **grep**

The user is not permitted to view directories or run commands outside the home directory.

During system maintenance, a specially configured rbash shell enables a set of commands to be executed.

The exact list of commands users can run is determined by the user's specific privileges and the specific setup of the machine. Different installed applications can add their own additional commands. The list of commands are displayed on login and can be redisplayed by typing the **help** command.

3.2. CLI Commands

Enter **help** (or **health**) to display the following screen::

```
platform@development:~$ help

host: AS01, role: webproxy,application,database, LOAD: 3.85
date: 2014-08-28 11:24:22 +00:00, up: 6 days, 3:03
network: 172.29.42.100, ntp: 196.26.5.10
HEALTH: NOT MONITORED
database: 20Gb
application: up
    mail - local mail management          keys - ssh/sftp credentials
```

(continues on next page)

(continued from previous page)

network - network management	backup - manage backups
voss - voss management tools	log - manage system logs
database - database management	notify - notifications control
schedule - scheduling commands	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	template - template pack creator

Note: From release 21.1 onwards, **cluster** commands should be used in standalone (“cluster-of-one”) deployments.

On web proxy nodes, the only cluster command you can run is **cluster prenode**. Database commands are also not available on web proxy nodes.

Notices and critical warnings are shown in colors upon login.

```
Last login: Tue Jun 1 05:02:23 UTC 2021 from 172.29.41.2 on pts/1
host: VOSS-UN-1, role: webproxy,application, load: 0.59, USERS: 2
date: 2021-06-01 05:26:28 +00:00, up: 4 days, 19:57
network: 192.168.100.3, ntp: 172.29.1.15
SWAP USED: 103 MB
MEMORY FREE: 417MB
DATABASE TRANSACTION SIZE: 83.61GB
DATABASE TRANSACTION COUNT: 1267721
WARNING: TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED - SETUP SCHEDULE FOR REGULAR MAINTENANCE

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

platform@VOSS-UN-1:~$
```

Entering any valid command name displays the usage parameters of that command. The **system** command help display is shown below::

```
platform@development:~$ system
USAGE:
-----
system date                - Display the system date and time
system download <url>      - Download a specific URL to media directory
system history              - Display a history of all executed UI commands
system keyboard <kbd-type> - Change the keyboard type (e.g. dvorak, us)
system mount                - Mount all removable media
system password             - Change the platform password
system provision            - Provision all the applications
system reboot               - Reboot the system
system root                 - Support administration via one-time-password
```

(continues on next page)

(continued from previous page)

```
system shutdown      - Halt the system
system unmount       - Unmount all removable media
```

When commands are run on a cluster, a number of options are available to specify the nodes on which the commands can be run. In other words, there is a *<where>* clause: **cluster run <where>**. The clause can take:

- *role* - the role of the node: `application`, `database`, `webproxy`
- *all* - the entire cluster
- *notme* - all nodes except the one the command is run on

For example, **cluster run notme system shutdown** would issue the command to shut down all nodes except the one the command is run on.

Note: In a cluster, reboot and shutdown of the entire cluster should be done on each node and not with the **cluster run all** command - see: [Remote Execution in Clusters](#).

Tab completion is available from the CLI for commands, parameters and partial filenames, for example:

```
$ log <Tab>
audit      collect   follow   list     merge    purge    send     ↵
↪sendnewer view

$ log audit <Tab>
locallog   remotelog ssl      status

$ log view process/nginx <Tab>
$ log view process/nginx.proxy.log
```

See also [Using the screen command](#).

3.3. Using the `screen` command

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

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

- **cluster provision**
- **cluster upgrade**
- **app template**
- **voss export type <args>**
- **voss export group <args>**
- **voss subscriber_data_export**

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

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

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

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

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

The version of **screen** used in VOSS-4-UC also supports the creation of a log file. If long-running commands will be run, the log file captures screen console output up to the session timeout. A message shows:

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

To create a screen log file:

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

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

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

3.4. System Specific Commands

3.4.1. VOSS Management Tools

The CLI (Command Line Interface) menu provides access to a number of commands specifically related to VOSS-4-UC.

In addition to the description of the commands available from the CLI **voss - voss management tools** menu, further details are provided for a selection of the commands. Note that some of the commands are used by developers only.

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

The commands have been arranged into functional categories:

- Install commands: commands typically used during the install process.
- Database commands: commands that directly manage the database.
- Performance commands: commands to manage the system performance.
- System specific commands: general commands not specifically related to the categories above.

3.4.2. Install Commands

Important: On a Modular Architecture deployment, `vooss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

- **vooss cleardown** - the command reinitialises the VOSS-4-UC database. It is usually run on a fresh installation and care should be taken with its use, as it deletes all system data.

Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow vooss-deviceapi/app.log**.

- **vooss get_extra_functions_version <[-h] [-c] [-d] [-m] [-q]>** - Display details of the currently installed extra functions file (`extra_functions.py`). Details can be the checksum (md5), created date and modified date - with or without titles. The command **vooss get_extra_functions_version -h** displays information on these parameters.
- **vooss migrate_summary_attributes <model_type>** - Migrates the summary attribute schema for instances of the specified model.

3.4.3. Database Commands

Important: On a Modular Architecture deployment, `vooss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

- **vooss cleardown** - the command re-initializes the VOSS-4-UC database. It is usually run on a fresh installation and care should be taken with its use, as it deletes all system data.

Note that this step may take some time. You can follow the process by running **log follow upgrade_db.log** or **log follow vooss-deviceapi/app.log**.

- **vooss db_collection_stats [all | <collection,collection,...>]** - Display detailed statistics of all the VOSS-4-UC databases, or of only a list of collections.

Refer to the example snippets below.

```
$ vooss db_collection_stats
```

Collection	Cnt	% Size	DB Size	AvgObj Size	Idx	Idx Size
VOSS.DATA_ACCESS..	8	0.0%	92.88K	11.61K	2	15.97K
VOSS.DATA_APIVER..	1	0.0%	496.00b	496.00b	2	15.97K
VOSS.DATA_APPLIC..	1	0.0%	1.98K	1.98K	2	15.97K
...						
VOSS.WORKER_QUEUE	9508	9.4%	117.76M	12.68K	9	4.34M

```
Total Documents: 144624
Total Data Size: 1.22G
Total Index Size: 73.31M
```

```
$ voss db_collection_stats DATA_RULEMODELHIERARCHYTYPE,SCHEMA
```

Collection	Cnt	% Size	DB Size	AvOb Size	Idx	Idx Size
VOSS.DATA_RULEMO..	318	100.0%	97.61K	314.00b	2	96.00K
VOSS.SCHEMA	12	95.2%	1.89M	160.92K	2	72.00K

Total Documents: 330
Total Data Size: 1.98M
Total Index Size: 168.00K

- **voss db_index_stats** - Display detailed statistics of the VOSS-4-UC database indices, including the five largest.

Refer to the example below.

```
$ voss db_index_stats
```

Index Overview

Collection	Index	% Size	Index Size
VOSS.DATA_USER	_id_	0.1%	44.00K
VOSS.DATA_USER	username_1__hierarchy_1	0.1%	52.00K
VOSS.RESOURCE	_id_	0.6%	460.00K
VOSS.RESOURCE	_search._i.v_1	10.8%	7.83M
VOSS.RESOURCE	lock_1	0.3%	196.00K

...

VOSS.WORKER_QUEUE	parent_transaction_id_1_..	0.0%	12.00K
-------------------	----------------------------	------	--------

Top 5 Largest Indexes

Collection	Index	% Size	Index Size
VOSS.RESOURCE	meta.model_type_1__search._i..	16.5%	11.97M
VOSS.TRANSACTION_LOG	_id_	10.8%	7.85M
VOSS.RESOURCE	_search._i.v_1	10.8%	7.83M
VOSS.TRANSACTION_LOG	_id.t_id_1_time_1	8.6%	6.26M
VOSS.TRANSACTION	submitted_time_-1_parent_1_a..	6.3%	4.59M

Total Documents: 743332
Total Data Size: 2.86G
Total Index Size: 72.56M

- **voss db_collection_cap_check TRANSACTION_LOG** - Returns the TRANSACTION_LOG database collection cap size in GB.
- **voss db_collection_cap TRANSACTION_LOG <10-50GB>** - Updates the TRANSACTION_LOG database collection to a cap size in GB provided as parameter.

Important:

1. The resize operation will impact the usage on the disk size allocated for the database (typically, 250GB is reserved upon installation). Consider a larger initial disk size allocation if a larger cap size is set.
2. When resizing a collection, the existing collection is dropped and a new collection is created. Transactions in the existing collection can therefore not be re-run.
Therefore, if you resize the transaction log, all existing entries will be flushed and are unrecoverable.
3. The resize operation must be carried out during a maintenance window, since the `voss-deviceapi` service is restarted during this process.

– cap size in GB parameter values: from 10GB up to 50GB

This command to change the cap size is to enable larger providers to do this for operational reasons, for example for diagnostics and a longer transaction replay window.

3.4.4. Database Commands for Transaction Management

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

Commands are available to count, delete and export transactions from the database. All the commands take a `<days>` parameter that indicates transactions tasks are for transactions *older than* this number of days, counting from the current time.

For transaction **archive** and **delete** commands, the user is prompted on the command line to proceed or not.

If the transaction commands fail or are aborted:

- For *all* commands, a notification message is sent.
- For transaction **archive** and **delete** commands, the number of successful transaction deletes are applied.
- For transaction **export** commands, no export file is created.

If transactions are exported, the exported archive file will be in the `media/txn_archive/` directory. Available partition space is checked before any transaction export carried out and reported on. Estimated export sizes are based on average transaction sizes. If this directory contains files older than 30 days, an error notification is sent, with a message to remove these files.

The exported file is of the format `transaction_archive-YYYYMMDD_HHMMSS.gz`, where the UTC date stamp is the *current time*, for example:

```
media/txn_archive/transaction_archive-20190130_110122.gz
```

The exported `.gz` archive file contains a text file with lines of JSON formatted strings of the transactions.

- For suggestions on transaction archiving best practices, also refer to [Transaction Archiving](#).
- For more details on scheduling the transaction archiving to happen automatically, see [Enable Database Scheduling](#).

- **voss transaction count <days>** - Count the number of transaction entries in the database that are older than the number of days specified. When running the commands interactively, this command is typically used before deleting or exporting.

Example:

```
$ voss transaction count 11
167,582
```

- **voss transaction delete <days> [limit <number>]** - Delete transaction entries from the database that are older than the number of days specified, optionally limiting the number of *oldest* transactions to delete.

The optional **[limit <number>]** parameter limits the number of transactions deleted and is typically used when a large number of transactions are older than the specified number of days (using **voss transaction count <days>**), which would impact the time to delete transactions. The parameter can then be used to manage the delete transaction time.

The user is prompted to continue or not.

Example:

```
$ voss transaction delete 11
You are about to delete transactions from the system. Do you wish to continue?y
Available space:      27,219,492 KB
Estimated space:     25,737 KB
Deleting 167,582 transactions [#####]_
↪100%
```

- **voss transaction export <days>** - Create an archive file of the transaction entries in the database that are older than the number of days specified. No entries are deleted from the database. The export file has the format indicated above.

Example:

```
$ voss transaction export 11
Available space:      27,219,492 KB
Estimated space:     25,737 KB
Exporting 167,582 transactions [#####]_
↪100%
```

- **voss transaction archive <days>** - First create an archive file of and then delete transaction entries in the database that are older than the number of days specified. This command therefore combines two commands: **voss transaction export <days>** and **voss transaction delete <days>**. The export file has the format indicated above.

The user is prompted to continue or not.

Example:

```
$ voss transaction archive 11
You are about to delete transactions from the system. Do you wish to continue?y
Available space:      27,219,492 KB
Estimated space:     25,737 KB
Exporting 167,582 transactions [#####]_
↪100%
    Used space:        24,767 KB
Archive saved to media/txn_archive/transaction_archive-20190130_110122.gz
Deleting 167,582 transactions [#####]_
↪100%
```

3.4.5. Performance Commands

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

- **voss throttle-rates** - A command with parameters to set, show and disable the API request rate for an interface or for any user. The command may be used to manage API request overload.
 - Throttle rates apply to each unified node in a cluster.
 - Use **voss throttle-rates help** to see command parameters and options.
 - Please contact support before changing any settings for throttle rates. Great care should be taken when adjusting throttle rates, as a change can have a significant impact on system performance and behavior

By default, the following interface throttle rates apply:

- administration: disabled
- selfservice: 300 req/min
- per user: 20 req/sec

If the throttle rates are exceeded, the API returns the HTTP status code and message:

```
Error 429: Too Many Requests
```

- To set throttling:

voss throttle-rates type <administration|selfservice|user> requests <number of requests> unit <time unit>

- * The requests parameter is defined as an integer which is the number of requests per unit.
- * The time unit can be second (`sec, s`) or minute (`min, m`).
- * Command output will show the interface and configuration change and prompt for a service restart.

Examples:

```
$ voss throttle-rates type administration requests 10 unit min
Administration:
  Current Configuration: Disabled
  New Configuration: 10/min
Self Service:
  Current Configuration: Disabled
User:
  Current Configuration: Disabled

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi

$ voss throttle-rates type selfservice requests 20 unit sec
Administration:
  Current Configuration: 10/min
  Current Rates: 0/min
Self Service:
```

(continues on next page)

(continued from previous page)

```

Current Configuration: Disabled
New Configuration: 20/sec
User:
  Current Configuration: Disabled

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi

$ voss throttle-rates type user requests 30 unit min
Administration:
  Current Configuration: 10/min
  Current Rates: 0/min
Self Service:
  Current Configuration: 20/sec
  Current Rates: 0/sec
User:
  Current Configuration: Disabled
  New Configuration: 30/min

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi

```

Note:

- * If the command is used *without* parameters, the user will be prompted to enter them. Press **Ctrl-C** to exit this interactive mode.
 - * The user throttle rate can be limited by an interface throttle rate.
- To show current throttling continuously:

voss throttle-rates list-refresh

The current request values are updated until the command is canceled with **Ctrl-C**.

For example:

```

$ voss throttle-rates list-refresh
Administration:
  Current Configuration: 10/min
  Current Rates: 0/min
Self Service:
  Current Configuration: 20/sec
  Current Rates: 0/sec
User:
  Current Configuration: 30/min
Refreshing Ctrl-C to exit..

```

- To show current throttling and exit:

voss throttle-rates list

The current request values are updated and shown. The command then exits.

For example, to list when enabled:

```

$ voss throttle-rates list
Administration:
  Current Configuration: 10/min

```

(continues on next page)

(continued from previous page)

```

    Current Rates: 0/min
  Self Service:
    Current Configuration: 20/sec
    Current Rates: 0/sec
  User:
    Current Configuration: 30/min

```

For example, to list when throttling is disabled:

```

$ voss throttle-rates list
  Administration:
    Current Configuration: Disabled
  Self Service:
    Current Configuration: Disabled
  User:
    Current Configuration: Disabled

```

– To disable throttling:

voss throttle-rates disable

For example:

```

$ voss throttle-rates disable
  Administration:
    Current Configuration: 10/min
    Current Rates: 0/min
    New Configuration: Disabled
  Self Service:
    Current Configuration: 20/sec
    Current Rates: 0/sec
    New Configuration: Disabled
  User:
    Current Configuration: 30/min
    New Configuration: Disabled

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi

```

Throttle rates are disabled by default. To restore any rates that were disabled, the throttle rates need to be set again.

- **voss session-limits** - A command with parameters to set, show and disable the number of sessions - based on interface and customer. In other words, for each interface: Administration or Self Service, a Global and a Per Customer Hierarchy session limit can be set. The number of concurrent login sessions per user is determined by the user's Credential Policy. Highest level (above provider administrators) administrator logins do not affect and are not affected by session limits.

Use **voss session-limits help** to see command parameters and options. Please contact support before changing any settings for session limits.

The command line output of the command to adjust the limits show the current and new values and prompt to restart the `voss-deviceapi:voss-wsgi` service before the limit is changed. For clusters, the session limit is set per cluster, requiring a cluster wide service restart.

By default, the following limits apply:

- global administration: 200

- global selfservice: 20000
- per customer administration : 10
- per customer selfservice : 1000

The global session limit would always be set to a larger value than a customer hierarchy limit.

A session is active until it expires or the user logs out. If the session limits are exceeded, the API returns the HTTP status code and message:

```
Error 503: "Login is currently disabled due to a temporary overload.
Please try again later.
```

- To set session limits:

voss session-limits type <customer|global> interface <administration|selfservice> limit <number>

- * If the command is used *without* parameters, the user will be prompted to enter them. Press **Ctrl-C** to exit this interactive mode.
- * The type is the hierarchy to which the limit applies: `global` or `customer`
- * The interface is the user interface to which the limit applies: `Administration` or `Self Service`.
- * The limit number is an integer value for the number of sessions.

Examples:

```
$ voss session-limits type customer interface administration limit 100
Administration:
  Global:
    Current Limit: 200
    Current Sessions: 1
  Per Customer Hierarchy:
    Current Limit: 10
    New Limit: 100
  Self Service:
    Global:
    Current Limit: 20000
    Current Sessions: 0
  Per Customer Hierarchy:
    Current Limit: 1000

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi
```

- To list session limits:

voss session-limits <list|list-refresh>

The current session values are shown. if the `list-refresh` option is used, the values are updated every second until the command is canceled with **Ctrl-C**.

Examples for two customers “GenCorp” and “VS-Corp”:

```
$ voss session-limits list-refresh
Administration:
  Global:
    Current Limit: 200
    Current Sessions: 2
```

(continues on next page)

(continued from previous page)

```

Per Customer Hierarchy:
  Current Limit: 10
  GenCorp Current Sessions: 0
  VS-Corp Current Sessions: 1
Self Service:
  Global:
    Current Limit: 20000
    Current Sessions: 0
  Per Customer Hierarchy:
    Current Limit: 1000
    GenCorp Current Sessions: 0
    VS-Corp Current Sessions: 0

Refreshing, Ctrl-C to exit...

```

- To disable session limits:

voss session-limits disable

- * This setting *removes all* session limit settings.
- * Recall that for the command to take effect, *all* nodes in the cluster need the service to be restarted.
- * To restore any settings that were disabled, they need to be set again.

Example:

```

$ voss session-limits disable
Administration:
  Global:
    Current Limit: 200
    Current Sessions: 1
    New Limit: Disabled
  Per Customer Hierarchy:
    Current Limit: 10
    New Limit: Disabled
Self Service:
  Global:
    Current Limit: 20000
    Current Sessions: 0
    New Limit: Disabled
  Per Customer Hierarchy:
    Current Limit: 1000
    New Limit: Disabled

An application restart is required for this change to take effect, e.g.:
$ cluster run application app start voss-deviceapi:voss-wsgi

```

- **voss workers** - By default, 30 is the maximum number of parent transactions per unified node from the queue that will be processed at once. Use the command **voss workers <number>** to modify this value.

This command restarts the `voss-queue` services.

- **voss queues <number>** - Set the number of queue processes

Increasing the number of queues that can run in parallel. This command restarts the `voss-queue` services.

voss queues - Get the number of queue processes

Important: It is strongly advised that VOSS Support is consulted before making changes to the number of queue processes.

For details, see the Configurable Number of Queue Processes topic in the Best Practices Guide.

Transactions resulting from a number of system components are by default carried out with a low priority that is currently limited to 50% of the maximum allowed parent transactions to be processed at once (the default is 30 per unified node). The system components are:

- Bulk loaders
- Data Sync (CUCM, Unity, LDAP)

Since these transactions may place too much load on the system during business hours when users are using the system for other activities, commands are available to change the percentage of the maximum allowed number of transactions these low priority tasks can use at peak and off-peak times.

- Percentage options are: 20%, 50% and 80%
- Defaults are as follows:
 - Off-peak percentage is 50% of the maximum number of parent transactions per unified node from the queue that will be processed at once. This means if no off-peak percentage is set manually, this value will be the default.
 - Off-peak start and end times on weekdays (Mon - Fri) are 00:00 - 23:59
 - Off-peak start and end times on weekends (Sat - Sun) are 00:00 - 23:59

In other words, the default is 50% for all hours of the week.

The command that is available to change peak and off-peak times and percentages, is **voss worker low_priority_schedule** and takes a number of parameters.

Without parameters, the command shows the current times and percentages, for example:

```
$ voss worker low_priority_schedule

Off-peak Time:
  Weekday:
    Start Time: 00:00
    End Time: 23:59
  Weekend:
    Start Time: 00:00
    End Time: 23:59
Off-peak Percentage: 50
Peak Percentage: 50
```

The following list shows command parameters to set times and percentages for the low priority schedule. Note that if an end time is set to a value that falls into another schedule, the start time of the other schedule will apply.

- **voss worker low_priority_schedule off_peak_time weekday HH:MM HH:MM**
 - Set the weekday schedule off-peak start and end time. The command prompts for a queue service restart.

The command supports nodes in different time zones. When times are set on a node in a cluster, the equivalent schedule times are set on nodes in other time zones.

- **voss worker low_priority_schedule off_peak_time weekend HH:MM HH:MM**

- Set the weekend schedule off-peak start and end time. The command prompts for a queue service restart.

The command supports nodes in different time zones: when times are set on a node in a cluster, the equivalent schedule times are set on nodes in other time zones.

- **voss worker low_priority_schedule peak_percentage [20|50|80]**

- Set the percentage of maximum allowed workers for low priority tasks to either 20%, 50%, or 80% during the peak period - the period outside any defined off-peak period. The command prompts for a queue service restart.

- **voss worker low_priority_schedule off_peak_percentage [20|50|80]**

- Set the percentage of maximum allowed workers for low priority tasks to either 20%, 50%, or 80% at the defined off-peak period. The command prompts for a queue service restart.

For example, the sequence of commands:

```
voss worker low_priority_schedule off_peak_time weekday 23:00 04:00
voss worker low_priority_schedule off_peak_time weekend 13:00 10:00
voss worker low_priority_schedule off_peak_percentage 80
voss worker low_priority_schedule peak_percentage 20
```

will result in a low priority schedules that uses a higher percentage (80%) of the maximum workers on weekdays between 23:00 and 04:00 and on weekends between 13:00 and 10:00, and only 20% during the rest of the time - the peak time.

3.4.6. System Specific Commands

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

- **voss finalize_transaction <id>** - If a transaction status incorrectly shows as processing *after a service restart* - for example where a service has stopped - use this command to mark it as finalized.

If a transaction in a 'Processing' state has child transactions in a 'Processing' state, these also need to be finalized.

The finalize transaction command should never be run for transactions which are still processing. It should only be used for transactions that are still in a 'Processing' state after a service restart.

- **voss get_extra_functions_version <[-h] [-c] [-d] [-m] [-q]>** - Display details of the currently installed extra functions file (`extra_functions.py`). Details can be the checksum (md5), created date and modified date - with or without titles. The command **voss get_extra_functions_version -h** displays information on these parameters.

- **voss migrate_summary_attributes <model_type>** - Migrates the summary attribute schema for instances of the specified model.

For example:

```
voss migrate_summary_attributes data/InternalNumberInventory
```

- **voss reset_device_concurrency [all | device <device_name> | --show_locks]** - Call Manager can handle eight concurrent connections. VOSS-4-UC monitors the number of connections it continuously

makes and removes or adds to the number as connections are made and closed. For debugging purposes, this “tracking number” can be reset.

If the concurrency remains at the maximum for more than 10 minutes, it is automatically reset to zero and a log error message `Resetting stale Device Concurrency..` (containing details) is displayed. The parameters are as follows:

- **all** - this is the default parameter if no parameter is provided. Concurrent connections on *all* devices are reset.
- **device <device_name>** - reset the concurrent connections on a specific device with `<device_name>`, for example **voss reset_device_concurrency device 192.168.100.15**
- **–show_locks** - For all devices, show device type, device name , concurrent connections and max current concurrent connections. For example:

```
$ voss reset_device_concurrency --show_locks
INFO:__main__: [device/hcmf - 192.168.100.11:8443] [0/1]
INFO:__main__: [CUCM - 192.168.100.15:8443] [0/8]
INFO:__main__: [CUC - 192.168.100.12:8443] [0/8]
```

- **voss clear_device_pending_changes <device_name | all>** - Since Unified CM version 10.0, a Change Notification Feature is available that stores changes to device objects in a cache. The VOSS-4-UC application service called `voss-deviceapi::voss-cnf_collector` collects these changes from *all* the Unified CM devices as they are added, manages and stores the changes in a data collection.

This collection can then be used to update the system by means of a Data Sync option on the GUI. An additional tool on the GUI displays the status and manages the collection on a specified device and also allows for the polling interval between collections to be configured. Each device keeps the last time that a collection process ran on it, so that a new collection will only be run on it once its interval has expired.

The **voss clear_device_pending_changes** command clears all pending changes by the Change Notification Collector for a particular device or for all devices.

To clear the collection of pending changes from a single device, the command and output is for example:

```
$ voss reset_device_pending_changes 10.120.10.190
This will clear all pending changes.
Do you wish to continue? yes
```

The status of changes collected from a device can be checked from the GUI.

- **voss set_debug <level[0/1]>** - By default, the level is “level0”, so no debug information is present in the logs.
Setting the value to “level1” is not supported and will for example result in performance degradation.
- **voss unlock_sysadmin_account** - Unlocks and forces a password reset on the system administrator account. The system administrator is prompted to enter a new password and is then prompted to verify the new password.
- **voss update_device_schemas <schema models>** - Regenerate device model schemas in the database and in the device schemas fixtures file. No data is lost, so this can be done on a production system, although the system should not be used while this is happening. Parameters can be passed through (indicated as “<schema models>”). The available parameters are shown by using the “–help” parameter. Used by developers.

- **voss post-upgrade-migrations** - Schedule a transaction to execute long running data migrations after an upgrade.

Data migrations that are not critical to system operation can have significant execution time at scale. These need to be performed after the primary upgrade, allowing the migration to proceed whilst the system is in use - thereby limiting upgrade windows.

- This command is a mandatory step after an upgrade.
- The command only needs to be run on a single node of a cluster.
- The command will display progress information of the migration transaction until the transaction concludes, or until the user exits the command using Ctrl-c.
- Transaction progress can also be followed on the Transaction GUI.
- If the transaction is cancelled on the GUI or interrupted by a system or queue restart, the command can be run again to re-queue a migration transaction, which will resume the migration process.
- Console examples are shown below:

```
$ voss post-upgrade-migrations
Queueing post upgrade migrations, please wait...
Post upgrade migrations transaction queued, ID 6.
Transaction progress(ctrl-c to exit, transaction will continue to execute):
-Post upgrade migrations:  0%|                | 0/1 [00:00<?]
--Migrating objects in TRANSACTION colle...: 27%|##### | 53814/200000
↪[00:30<01:21]
^C
Aborting progress monitoring, transaction ID 6 is continuing execution,
↪exiting...
```

Subsequent re-execution with the same previously queued transaction still executing - progress display is resumed and displayed until the transaction concludes:

```
$ voss post-upgrade-migrations
Existing post upgrade migrations transaction found, not requeueing.
Transaction progress(ctrl-c to exit, transaction will continue to execute):
-Post upgrade migrations: 100%|#####| 1/1 [01:40<00:00]
--Migrating objects in TRANSACTION colle...: 100%|#####| 200000/200000
↪[01:40<00:00]

Post upgrade migrations complete, exiting...
```

3.4.7. Transaction Archiving

The following are considerations when determining the frequency of the transaction archiving schedule to set up on the system. If a schedule is not set up for transaction archiving, system Alerts will be raised as well as a warning on the platform CLI login:

```
TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED
```

- Run **voss transaction count <days>** on your system to inspect the number of transactions during a given period to determine your usage metrics.

Refer to the *Database Commands for Transaction Management* topic in the Platform Guide for details on transaction archive command use and scheduling:

- **voss transaction delete <days>**

- **voss transaction export <days>**
- **voss transaction archive <days>**
- Business policies - company policies may drive your choices: the immediate access to transaction logs for a period of time, security policy on data/audit retention, and so on.

Note: The transaction archive process does mean the logs are not lost, just that they are not immediately accessible in the administrator graphical interface for searching.

- You can also set up system monitoring thresholds so that you receive alerts via the GUI and SNMP if the threshold is exceeded - which might indicate you need to review the archive schedule to increase how frequently is runs.

See the *SNMP* and *VOSS-4-UC System Monitoring Traps* topics in the Platform Guide.

3.5. System Metrics

3.5.1. Report Transaction Commands

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

VOSS-4-UC provides a set of utilities available to provide transaction metrics. The commands are typically used to inspect and monitor transaction performance, for example in the case where transaction performance issues are encountered.

Note: For clusters, the number of workers are set to be the total default for 4 unified nodes ($4 \times 30 = 120$).

Profile

A command is available that aggregates the transactions performed during the given date-time range by grouping them by model type and sorting them by duration.

A **Pct** column is also available to indicate the percentage of the interval time used by the transaction.

The command and parameters are:

- **voss report transaction profile [OPTIONS] START_TIME END_TIME**
- `START_TIME` - The date-time value from which to start the sample collection.
- `END_TIME` - The date-time value from which to start the sample collection. Valid formats: `'%Y-%m-%d_%H:%M:%S'`, `'%Y-%m-%d_%H:%M'`, `'%Y-%m-%d'`

Example:

```
platform@VOSS:~$ voss report transaction profile 2019-05-06_16:39 2019-05-06_23:00
      Entity Operation      Calls  RespTime  QTime  SvcTime_
↔      Pct
```

(continues on next page)

(continued from previous page)

↔	99.95	tool/DataImport	execute	11.00	3128.24	1.11	3127.13
↔	0.03	data/User	update	2.00	0.86	0.09	0.77
↔	0.01	data/Schedule	update	1.00	0.28	0.18	0.10
↔	0.01	data/PasswordReset	create	2.00	0.17	0.10	0.08
↔	0.01	data/Schedule	Execute	1.00	0.15	0.00	0.14
↔	0.00						

Throughput

Given a number of output rows to display and duration in seconds between samples, a command is available to display the information about transaction throughput during the selected sample periods.

Throughput is the projected count per minute based on the sample interval and transaction count.

The command and parameters are:

- **voss report transaction throughput [OPTIONS] START_TIME INTERVALS DURATION**
- **START_TIME** - The date-time value from which to start the sample collection. Valid formats: '%Y-%m-%d_%H:%M:%S', '%Y-%m-%d_%H:%M', '%Y-%m-%d'
- **INTERVALS** - Specifies the number of samples to display.
- **DURATION** - The duration of the samples in seconds, for example every 30 seconds or every 60 seconds.

The command output shows parent and child transaction data in columns:

- **Count:** Transaction count
- **Utilisation:** ratio of total workers (120) used: $((\text{Throughput} * \text{RespTime}) / 120) * 100$

Example of parent transactions:

formula: $((\text{Throughput} * \text{RespTime}) / 120) * 100$

calculation: $(2.00 * 4.86 / 120) * 100 = 9.71$

- **Qtime:** ratio: (total queue time for transactions at time stamp / transaction count at the time stamp). Else, zero if no transactions.
- **SvcTime:** ratio: (total service time for transactions at time stamp / transaction count at the time stamp). Else, zero if no transactions.
- **RespTime:** calculated as total transaction time (**Qtime** + **SvcTime**). Else, zero if no transactions.
- **Throughput:** ratio: (transaction count at the time stamp / minutes (end_time - start_time)). Else, zero if no transactions.

Example of child transactions where **DURATION** is 30 seconds (0.5 minutes):

21 child transactions / 0.5 minutes = 42.00

- **Workers:** calculated as $(\text{Throughput} * \text{RespTime})$

Example command output:

```
platform@VOSS:~$ voss report transaction throughput 2019-05-06_16:39 2 30
                                Parent Transactions
↳
                                Child Transactions
TS                               Count  Utilisation  Workers  Qtime  SvcTime
↳RespTime Throughput  Count      Qtime      SvcTime  RespTime Throughput
2019-05-06_16:39:00            0      0.00        0.00    0.00   0.00
↳ 0.00      0.00   21      0.07        0.86    0.93   42.00
2019-05-06_16:39:30            0      0.00        0.00    0.00   0.00
↳ 0.00      0.00   80      0.07        0.14    0.22  160.00
```

```
platform@VOSS:~$ voss report transaction throughput 2019-05-06_16:42 2 30
                                Parent Transactions
↳
                                Child Transactions
TS                               Count  Utilisation  Workers  Qtime  SvcTime
↳RespTime Throughput  Count      Qtime      SvcTime  RespTime Throughput
2019-05-06_16:42:00            1      9.71        8.10    0.07   4.79
↳ 4.86      2.00   98      0.07        0.05    0.13  196.00
2019-05-06_16:42:30            2     3421.50    2851.25  0.12   855.26
↳ 855.37    4.00   34      0.07        0.21    0.28  68.00
```

Worker-Usage

A command is available to show the active and queued transactions at particular intervals from a given start time and then to provide a *utilisation* value for the transaction, relative to a worker value.

The command and parameters are:

- **voss report transaction worker-usage [OPTIONS] START_TIME INTERVALS DURATION**
- **START_TIME** - The date-time value from which to start the sample collection. Valid formats: '%Y-%m-%d_%H:%M:%S', '%Y-%m-%d_%H:%M', '%Y-%m-%d'
- **INTERVALS** - Number of report entries to list.
- **DURATION** - Number of seconds between each report entry.

The command output shows parent and child transaction data in columns:

- **Utilisation**: ratio: (active parents / total workers (=120)) * 100
- **Zombie**: number of transactions that have not been updated (changed from Queued to Active) for 60 minutes

Example output

```
platform@VOSS:~$ voss report transaction worker-usage 2019-05-06_16:39 50 30
REPORT:
Timestamp      Utilisation  QueuedParents  ActiveParents  QueuedChildren
↳ ActiveChildren  Zombies
2019-05-06_16:39:00    0.83         0              1              0
↳ 1              0
2019-05-06_16:39:30    0.83         0              1              0
↳ 1              0
2019-05-06_16:40:00    0.83         0              1              0
↳ 0              0
2019-05-06_16:40:30    0.83         0              1              0
↳ 0              0
```

(continues on next page)

(continued from previous page)

2019-05-06_16:41:00	0.83	0	1	0	↳
↳ 1	0				
2019-05-06_16:41:30	0.83	0	1	1	↳
↳ 0	0				
2019-05-06_16:42:00	0.83	0	1	0	↳
↳ 0	0				
2019-05-06_16:42:30	0.00	0	0	0	↳
↳ 0	0				
2019-05-06_16:43:00	0.83	0	1	0	↳
↳ 1	0				
2019-05-06_16:43:30	0.83	0	1	0	↳
↳ 0	0				
2019-05-06_16:44:00	0.83	0	1	0	↳
↳ 0	0				

Current-Usage

A command is available to show a table of transactions grouped by:

- Processing and Queued transactions
- Parent and child transactions
- Node name on which the transaction is running

Queued transactions will be shown with a node name of `default` because it is not known on which node they will be executed.

- Transaction priority (High, Medium, Low)

The report is useful to verify that all unified nodes are correctly processing transactions.

The command is:

- **voss report transaction current-usage**

Example output from a single node:

```
platform@VOSS:~$ voss report transaction current-usage

REPORT:
parent          | Processing transactions | Queued transactions  ↳
↳ |
|               | Priority                | Priority              ↳
↳ |
|               | High  Medium  Low | High  Medium        ↳
↳Low|
| node1-voss-queue| 1      0      0 | 0      0            ↳
↳ 0|

child          | Processing transactions | Queued transactions  ↳
↳ |
|               | Priority                | Priority              ↳
↳ |
|               | High  Medium  Low | High  Medium        ↳
↳Low|
| node1-voss-queue| 0      0      0 | 0      0            ↳
↳ 0|
```

3.5.2. Report API Commands

Important: On a Modular Architecture deployment, `voss` system specific commands can only be run on application nodes and for Unified and Single Node architectures, *not* on database or web proxy type nodes.

VOSS-4-UC provides a set of utilities available to provide API request metrics. The commands are typically used to inspect and monitor API request performance over a defined period.

Profile

A command is available that aggregates the API requests performed during a given date-time range by grouping them by model type and sorting them by duration.

The command and parameters are:

- **voss report api profile [OPTIONS] START_TIME END_TIME**
- `START_TIME` - The date-time value from which to start the sample collection.
- `END_TIME` - The date-time value from which to start the sample collection. Valid formats: `'%Y-%m-%d_%H:%M:%S'`, `'%Y-%m-%d_%H:%M'`, `'%Y-%m-%d'`
- Options:
 - `--path TEXT` API log file path (default is `nginx/access.log`)
 - `--limit INTEGER` Limits the number of profile lines to display. If limited, the last row is listed as Other

The command output shows Request_URI count data in columns.

Column headers:

- **90th-PCTL:** 90th percentile time value of the count
- **CV:** Coefficient of variation of the count

Example:

```
platform@VOSS:~$ voss report api profile --limit 6 2019-05-06_15:49 2019-05-06_16:00
  Count      TotalTime  AverageTime  Percentage  90th-PCTL
↔CV          Method  Request-URI
  435         199.97      0.46         72.11%     0.69      34.66
↔%           PUT   /api/data/DataModel
  69          64.92      0.94         23.41%     1.32      29.76
↔%           LIST  /api/data/DataModel
  44          10.94      0.25         3.95%     0.33      38.26
↔%           LIST  /api/data/Migration
  1           0.12      0.12         0.04%     0.00      0.00
↔%           LIST  /api/data/Cuccx
  1           0.09      0.09         0.03%     0.00      0.00
↔%           LIST  /api/data/Sntp
  1           0.07      0.07         0.02%     0.00      0.00
↔%           LIST  /api/data/Package
  26          1.19      0.00         0.43%     0.00      0.00
↔%           Other
```


Throughput

A command is available to show the number of API calls at particular intervals from a given start time and then to provide a *throughput* value for the transaction, in other words number of requests per time interval.

The command and parameters are:

- **voss report api throughput [OPTIONS] START_TIME INTERVALS DURATION**
- **START_TIME** - The date-time value from which to start the API call collection. Valid formats: '%Y-%m-%d_%H:%M:%S', '%Y-%m-%d_%H:%M', '%Y-%m-%d'
- **INTERVALS** - Number of report entries to list.
- **DURATION** - Number of seconds between each report entry.
- **Options:**
 - --path TEXT API log file path (default is nginx/access.log)

The command output shows API request count data and throughput in columns.

Column headers:

- **Throughput:** ratio: (Count / minutes (end_time - start_time))
- **SvcTime, RespTime:** time per request over timestamp period
- **Qtime:** defaults to 0 - not used
- Example **Throughput** of API requests:

DURATION is 30 seconds (0.5 minutes)

1 API request / 0.5 minutes = 2.00

```
platform@VOSS:~$ voss report api throughput 2019-05-06_15:49 10 30
```

Timestamp	Count	Qtime	SvcTime	RespTime	Throughput
2019-05-06_15:49:00	1	0.00	1.47	1.47	2.00
2019-05-06_15:49:30	27	0.00	0.56	0.56	54.00
2019-05-06_15:50:00	28	0.00	0.49	0.49	56.00
2019-05-06_15:50:30	32	0.00	0.42	0.42	64.00
2019-05-06_15:51:00	25	0.00	0.49	0.49	50.00
2019-05-06_15:51:30	21	0.00	0.59	0.59	42.00
2019-05-06_15:52:00	20	0.00	0.59	0.59	40.00
2019-05-06_15:52:30	28	0.00	0.51	0.51	56.00
2019-05-06_15:53:00	29	0.00	0.44	0.44	58.00
2019-05-06_15:53:30	21	0.00	0.63	0.63	42.00

3.6. Data Export Commands

3.6.1. Introduction

The Subscriber Data Extract is a capability to provide an adhoc or scheduled basic data feed via the generation of a set of files.

The content of the files is settings around Subscribers and other key services in order to support billing/expense management operations and is not intended as a general data feed.

Additional fields or files may be added to the collection over time via the roadmap, so any planned file consumption should take that into account to minimize impact.

Data Export Overview

The **voss export** command is used to carry out a bulk data export from the VOSS-4-UC system database. The exported data can for example be imported into a warehouse.

Important: Since a data export can take time, the **voss export** command can only be run in a `screen` session. First run **screen** and then **voss export** and its parameters.

Type **voss export help** for details.

The data extract schedule can be managed with the **schedule** command. For details on the use of the command, see: [Scheduling](#). Since bulk data exports can typically take more than an hour on a scale system, it is recommended to schedule this task instead of running it manually from the console.

The export file format is JSON as per RFC 7159. For details on the filename, format and contents of the export files, refer to the Data Export Types topic in the Appendices.

The **voss export** command takes a `type` or `group` parameter to indicate the type of data to export.

The following are values of the `group` parameter:

- `subscriber`
- `license`

For example:

voss export group subscriber

```
platform@VOSS:~$ voss export group subscriber
Starting subscriber group export consisting of analogue_line_mgcp,
analogue_line_sccp, call_pickup_group, contact_center_enterprise,
contact_center_express, customer, extension_mobility, fmc,
hunt_group, line, phones, site, subscriber, webex_teams, please wait...
Starting analogue_line_mgcp export, please wait...
Completed analogue_line_mgcp export,
created 2019-09-30_0859_analogue_line_mgcp.json.gz.
[...]
```

3.6.2. Subscriber Data Export Command

Note: The command **voss subscriber_data_export** is equivalent to **voss export group subscriber**.

Important:

- To optimise performance, run and schedule the data export command from the *secondary* database server if possible.
- Since a data export can take time, the **voss subscriber_data_export** and **voss export** commands can only be run in a `screen` session. First run **screen** and then **voss export** and its parameters. See also: *Using the screen command*.
- Since the data export command runs database queries, it is recommended that the data exports be scheduled. Refer to the topic on scheduling for details and syntax.

for example:

```
schedule add subscriber_export voss export group subscriber
```

```
schedule time subscriber_export weekly 1
```

Best practices for scheduling to consider, are:

- Individual report exports should be scheduled in a serial manner so that they do not overlap and result in a high database load.
- For resilience:
 - * Stagger the schedule based on how long it is expected to run - in accordance with the number of subscribers in the database.
 - * For better failover support, schedules can be created on all active Unified Nodes. This requires a more complex schedule staggering and collection management.
 - * For simplified schedule staggering and the export collection management, schedules can be created and staggered on a single Unified Node. This option but requires a manual re-schedule in the case of node failover.

More than one `type` parameter can be specified for the command by using the `type` parameter for each. For example:

```
voss export type line type site.
```

The `type` parameter values by `subscriber` group are listed below, as well as a reference to the content details:

- `analogue_line_mgcp` (*Analogue line MGCP Data Export*)
- `analogue_line_sccp` (*Analogue Line SCCP Data Export*)
- `call_pickup_group` (*Call Pickup Group Data Export*)
- `contact_center_enterprise` (*Contact Center Enterprise Data Export*)
- `contact_center_express` (*Contact Center Express Data Export*)
- `customer` (*Customer Data Export*)
- `extension_mobility` (*Extension Mobility Data Export*)

- `fmc` (*FMC Data Export*)
- `hunt_group` (*Hunt Group Data Export*)
- `line` (*Line Data Export*)
- `phones` (*Phones Data Export*)
- `site` (*Site Data Export*)
- `subscriber` (*Subscriber Data Export*)
- `webex_teams` (*Webex Teams Data Export*)

The export file directory and file format of the `subscriber` group is:

- **directory:** `media/data_export/<YYYY-MM-DD>`
- **file naming format:** `<YYYY-MM-DD_HHMM>_<type>.json.gz`

For `subscriber` group files:

- A retention policy of 30 days is in place. After each successful extraction of the data, any extract files 31 days old or older will be removed.
- If an export contains no data, a JSON file will contain an empty JSON list: `[]`.

Example:

```
media/data_export/2018-10-11/2018-10-11_1236_analogue_line_sccp.json.gz
```

Command examples:

- **Single type**

```
$ voss export type line
Starting line export, please wait...
Completed line export, created 2018-10-11_1236_line.json.gz.
```

- **Multiple types**

```
$ voss export type line type site
Starting line export, please wait...
Completed line export, created 2018-10-11_1236_line.json.gz.
Starting site export, please wait...
Completed site export, created 2018-10-11_1236_site.json.gz.
```

- **Group**

All types in a group are exported.

```
$ voss export group subscriber
Starting subscriber group export consisting of analogue_line_mgcp, analogue_line_sccp,
↪ [...]
Starting analogue_line_mgcp export, please wait...
Completed analogue_line_mgcp export, created 2018-10-11_1236_analogue_line_mgcp.json.
↪gz.
Starting analogue_line_sccp export, please wait...
Completed analogue_line_sccp export, created 2018-10-11_1236_analogue_line_sccp.json.
↪gz.
[...]
Completed subscriber group export.
```

The export files can then be copied to a remote system. For example, from the VOSS-4-UC system, list out the data export files:

```
$ ls media/data_export/2018-10-11  
2018-10-11_1236_analogue_line_sccp.json.gz
```

The exported files can be copied to a remote system using SCP or SFTP on port 22. For example:

```
remote_system:~$ scp <platform_user>@<voss_system>:media/data_export/2018-10-11/2018-  
↪10-11_1236_analogue_line_sccp.json.gz .
```

4. Node Deployment

4.1. Deployment

Single and Clustered node deployment is described in detail in the Install Guide.

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

In a clustered topology, a number of nodes with different roles are clustered together and provisioned to form a networked system. When nodes are clustered together, High Availability and Disaster Recovery can be achieved by ensuring that there are redundant services. Nodes can be deployed in any order.

Once two or more nodes have been deployed, the nodes can be grouped into a cluster by executing **cluster add <ip>**. Note that a node already in one cluster cannot be added to another cluster. Likewise, nodes can be removed from a cluster with the **cluster del <ip>** command. The nodes in the cluster can be displayed using **cluster list**.

Cluster roles cannot be changed after installation because it is dependent on installed software and other configuration at time of deployment.

The status of the cluster can be viewed using **cluster status**.

If the node topology needs to be changed, the following procedure can be followed:

- A node can be removed from the cluster with **cluster del <ip>**
- The node can be redeployed with the correct parameters
- Add the new node to the cluster with **cluster add <ip>**
- Provision the cluster with **cluster provision**. This command should only be run on one node in the cluster, usually an application server.

4.2. Modular Cluster Deployment

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

VOSS-4-UC 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.

Note: Also refer to the following topics:

Multinode Modular Cluster with Application and Database Nodes in the Install Guide.

Modular Cluster Deployment in the Install Guide.

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-4-UC 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 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.

4.3. Migrate a Unified Node Cluster to a Modular Cluster

From release 21.1 onwards, a new Modular Architecture deployment is available.

Important: Since individual requirements and topologies differ, contact VOSS for guidance and assistance *before* migrating.

Note:

- A topology migration from a unified node cluster topology to a modular cluster is recommended under certain circumstances to improve the system behaviour at scale and under load.
- Since the *minimum* number of nodes in a modular architecture is 8 (2 web proxy and 6 application/database), *only* a migration of a unified node deployment of 8 nodes is supported.
- The migration to a modular architecture deployment is a one-way process. A migrated system cannot be migrated back to a unified node cluster topology.
- The migration process requires a maintenance window to complete, since a VM power down is needed for the migration of some nodes.

- For an overview of this topology, see: [Modular Cluster Deployment](#).
- For details on *installing* with this topology, see the Installation Guide topics:
 - Multinode Modular Cluster with Application and Database Nodes
 - Multinode Modular Cluster Hardware Specification
 - Modular Architecture Multinode Installation

Commands are available to migrate unified nodes from the unified node cluster topology. Since unified nodes combine database and application roles, the commands remove one of these roles. The commands are:

- **app remove-role db**: remove the database role from a unified node, leaving the application role.
- **app remove-role app**: remove the application role from a unified node, leaving the database role.

4.3.1. Considerations and Prerequisites

- VOSS has been contacted and is assisting.
- The migration requires a maintenance window.
- The new topology should be decided, in other words, which nodes take the db and app roles in the database centers.
- Backups: create a VM snapshot before starting.

Database backups from the Primary Unified node cannot be used to restore the database on the new primary database node after migration.

Note: If you need to roll back from the migration using the snapshots, first revert *all* the cluster nodes to the correct snapshot, then power the nodes on. Do not revert the snapshot and power on each node

in sequence.

- Hardware impact:
 - Conversion from unified to app node - has disk space impact (db space unused and can be removed in the VM)
 - Conversion from unified to database node - unified node requirement on VM of 16GB RAM needs to be changed to 32GB RAM for database node.
- Schedules:
 - system
 - admin created
 - application related: NBI

4.3.2. Steps

The example and steps below shows a migration from a unified node cluster topology with 6 unified nodes in 2 data centers to modular topology with 3 database nodes and 3 application nodes:

Data Center 1		
Unified Node Topology	Migration Command	Modular Topology
UN1	app remove_role application	db node (primary db)
UN2	app remove_role database	app node
UN3	app remove_role application	db node
UN4	app remove_role database	app node
WP1	N/A	WP1

Data Center 2 (Data Recovery site)		
Unified Node Topology	Migration Command	Modular Topology
UN5	app remove_role application	db node
UN6	app remove_role database	app node
WP2	N/A	WP2

Before you start, verify that the nodes on which you run the migrate commands are *unified* nodes.

1. Check the memory of the nodes that will have application roles removed (the nodes: UN1, UN3 and UN5) to identify those with memory *less* than 32GB.

You can run **diag mem** on the nodes to verify this and add Mb used and Mb free.

2. Shut down the system:
 - a. From the primary database node: **cluster run notme system shutdown**
 - b. Then **system shutdown**
3. Take VM snapshots.

4. For the VMs of unified nodes that will become db nodes and that you checked to have less than 32GB memory (from the nodes: UN1, UN3 and UN5), resize the VM memory to 32GB:
 - a. Log into VMWare GUI.
 - b. Adjust the **Memory** assigned to the system to 32GB.
5. Power up the VMs.
6. Run **app remove_role database** on UN2, UN4 and UN6.
7. Run **app remove_role application** on UN1, UN3 and UN5.
8. Do not modify the web proxy nodes WP1, WP2.
9. Run **cluster provision** on UN1.
10. Run **cluster status** to ensure the cluster roles are as expected.

Note: When the deployment is migrated to a modular system, the **cluster provision** command will run a check to verify that all nodes have been migrated, in other words, that no “hybrid” system is present.

An error message is displayed otherwise:

```
Cluster provisioning not allowed on cluster with mixed unified and db or app nodes
```

11. Run **database config** on the first new *database* node and ensure the database nodes are in the correct state. In this example it would be UN1.
12. Verify the primary application node (UN2) with the **cluster primary role application** command run on the node. The output should be *true*, for example:

```
platform@UN2:~$ cluster primary role application
is_primary: true
```

See: [Primary Node Role in a Cluster](#).

13. For the new *application* nodes (nodes where the database role has been removed, in other words UN2, UN4 and UN6)
 - a. Check the number of queues using **voss queues** and if the number is *less than 2*, set the queues to 2 with **voss queues 2**.

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

- b. Remove the database drives and volumes. Use the following commands to check and carry out the tasks. See the output example: [Removing Database Drives and Volumes](#).
 - `drives list`
 - `drives remove_logical voss dbroot`
 - `drives remove_volume voss`

Important: Ensure the nodes are application nodes: you can run **app status** and verify no `mongodb` service is running on them.

14. You can remove disks from VMs of nodes with *application* roles (UN2, UN4 and UN6).

Important: Ensure the nodes you remove VM disks are application nodes: run **app status** and verify no `mongodb` service is running on them.

If disks are hot-mounted, a system restart may be needed.

- a. Run **drive list** and confirm that LVM volume is shown as `available`. Make note of VG size.
- b. Log into VMWare GUI.
- c. Look at disks attached to system, and specifically ones that are the size noted above.
- d. Remove that disk from the system.
- e. On the CLI, rerun **drive list** to confirm that LVM and VG are removed.

4.3.3. Removing Database Drives and Volumes

Example of commands and output when removing database drives and volumes.

```
platform@UN2:~$ drives list
Used disks and mountpoints:
  sdc1 - services:backups

Unused disks:
dm-0

Unused mountpoints:
  services:SWAPSPACE

Volume Groups
  voss - 25.0 GB free, 250.0 GB total
  Physical volumes:
    sdd1
  Logical volumes:
    dbroot/dm-0 - 225.0 GB
platform@UN2:~$
```

```
platform@UN2:~$ drives remove_logical voss dbroot
Logical volume dbroot has been removed from voss
platform@UN2:~$ drives list
Used disks and mountpoints:
  sdc1 - services:backups

Unused disks:
  none - if disks have been hot-mounted, it may be necessary to reboot the system

Unused mountpoints:
  services:SWAPSPACE

Volume Groups
  voss - 250.0 GB free, 250.0 GB total
  Physical volumes:
    sdd1
```

(continues on next page)

(continued from previous page)

```
Logical volumes:
platform@UN2:~$
```

```
platform@UN2:~$ drives remove_volume voss
Volume group voss has been removed
platform@UN2:~$ drives list
Used disks and mountpoints:
    sdc1 - services:backups

Unused disks:
sdd

Unused mountpoints:
    services:SWAPSPACE

platform@UN2:~$
```

4.4. Primary Node Role in a Cluster

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

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

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

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

5. Provisioning

5.1. Provisioning

The system is installed as a loosely bundled set of applications. In order for the applications to be coupled, a process called ‘provision’ must take place.

Standalone systems are single-node (“cluster-of-one”) clusters and are provisioned automatically since there is only one node in the system. This can be performed manually afterward with **cluster provision**.

When the topology of the cluster changes, e.g. additional nodes or applications are added; or to reprovision the system to bypass a faulty node, the cluster must be reprovisioned using **cluster provision**.

Note that for multi-node clusters, the cluster provisioning needs to reconfigure and restart services across the cluster in a complex arrangement and the provisioning duration is dependent on the number of nodes - it can take a number of hours for large installations.

Provision the cluster from the primary node with **cluster provision** For backwards compatibility, this command is the same as for example **cluster provision fast**.

Use the **cluster provision serial** command if the VMware host is under load.

The provisioning step may take some time, because all applications must be cross-configured to work with one another and the database is also provisioned. If the system discovers that no primary database server exists (or multiple database servers exist), the **cluster provision** command prompts the user to select a primary server manually.

See also: *Using the screen command*.

6. Networking

6.1. Network interfaces

The command **network interfaces** will display the available network interfaces and their configuration. The hostname can be set or changed with **network name <hostname>**.

A network interface can be configured or changed as follows:

network <interface-name> <ip> <netmask> <gateway>

For example: **network eth0 172.29.89.100 255.255.255.0 172.29.89.1**

The IP address can be changed without affecting the netmask and gateway using:

network <interface-name> ip <ip>

For example, **network eth0 ip 172.29.89.100**.

The system should be rebooted after a network interface configuration or change. In the case of a standalone single-node cluster topology, the system should be provisioned again with the **cluster provision** command as the final step of the change.

Note: In a clustered environment an IP address change may show the following message:

```
$ network eth0 178.29.21.253 255.255.255.0 176.29.22.1
This change will require a reprovision. Do you wish to continue [y/n] ?y
Unable to change IP address while clustered.
Please remove host from cluster before changing IP address
```

To remove the host from the cluster, run the command below *on the primary unified node*:

cluster del <IP address of node to be changed>

Network routes can be displayed with **network routes**.

- A new network route can be configured with **network routes <network-address> <netmask> <gateway>**.
- Network routes can be deleted with **network routes del <network-address>**.

6.2. Network services

Network security is described in detail under the Security section, including detail regarding firewall ports.

NTP servers can be configured using the following commands:

- **network ntp** will display the list of configured NTP servers
- **network ntp set <ntp-server1> [<ntp-server2> ...]** will set up one or more NTP servers.

Note: This command will overwrite any existing list of configured NTP servers.

- **network ntp add <ntp-server1> [<ntp-server2> ...]** will add one or more NTP servers to the existing list
- **network ntp del <ntp-server>** will delete a NTP server

DNS servers can be configured using the following commands:

- **network dns** will display the configured DNS servers
- **network dns set <dns-server1> [<dns-server2> ...]** set up one or more DNS servers

Note: This command will overwrite any existing list of configured DNS servers.

- **network dns add <dns-server1> [<dns-server2> ...]** will add one or more DNS servers to the existing list
- **network dns del <dns-server>** will delete a DNS server
- **network domain <domain-name>** sets the default DNS domain
- Alternate DNS search domains can be configured with **network search add <domain>** and **network search del <domain>**

6.3. Network URI specification

All network locations are specified as a URI, for example download locations, backup destinations, notification destinations, and so on.

The following list shows the URI syntax:

- **ftp:** ftp://user[:password]@host[:port] [/path]
- **http:** http(s)://user[:password]@host[:port]/path
- **file:** file://{/path}+{/filename}
- **sftp:** sftp://user[:password]@host[:port] [/path]
- **scp:** scp://[user@]host[:port] : [/path]
- **email:** mailto:user@host
- **snmpv2:** snmp://community@host[:port]
- **snmpv3:** snmp://user:auth:password@host[:port] ... minimum auth/password

The [password] in the URI is optional when authentication is set up. Refer to [SSH key management](#).

Note:

- If a password contains special characters, it should not be added to the URI, but typed in at the password prompt.
- If necessary, the URI should be URL-encoded: reserved characters should be encoded with percent-encoding.

6.4. Network Docker Container Range

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

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

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

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

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

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

For example:

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

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

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

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

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

For example:

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

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


7. Application Control

7.1. Application Control and Status

The functioning system is comprised of applications. Each application has a name and a version number. An application may have multiple processes running within and each process has its own state.

7.1.1. Application Status

The command **app status** is used to display the status of the system. When the command is executed, it requests an up-to-date status of every application and process, and hence may take a few seconds to return.

A typical **app status** screen from the command line interface:

```
selfservice v19.3.1 (2019-11-13 14:38)
|-node                running
voss-deviceapi v19.3.1 (2019-11-13 14:40)
|-voss-cnf_collector  running
|-voss-monitoring     running
|-voss-queue          running
|-voss-risapi_collector running
|-voss-wsgi           running
cluster v19.3.1 (2019-11-13 14:39)
template_runner v19.3.1 (2019-11-13 14:42)
mongodb v19.3.1 (2019-11-13 14:39)
|-arbiter             running
|-database            running
support v19.3.1 (2019-11-13 14:42)
selenium v19.3.1 (2019-11-13 14:46)
|-gui_orchestration   running
platform v19.3.1 (2019-11-13 14:40)
nginx v19.3.1 (2019-11-13 14:40)
|-proxy              running
services v19.3.1 (2019-11-13 14:41)
|-wsgi               running
|-logs              running
|-firewall           running
|-mount             running
|-scheduler         running
|-syslog            running (completed)
|-time              running (completed)
voss-portal v19.3.1 (2019-11-13 14:48)
|-gui                running
```

(continues on next page)

(continued from previous page)

```

security v19.3.1 (2019-11-13 14:43)
snmp v19.3.1 (2019-11-13 14:42)
  |-daemon                running (completed)
  |-traps                  running (completed)
vmware v19.3.1 (2019-11-13 14:42)

```

Applications can have multiple processes running within it - each with its state displayed. The following states are defined:

- `running` indicates that the process is running correctly.
- `completed` indicates that the process ran to completion successfully.
- `suspended` indicates that the process is suspended while waiting for another process.
- `stopped` indicates that the process is not running. An error message indicates that the process stopped for an unexpected reason.
- `disabled` indicates that the application is not licensed.

7.2. Starting and Stopping

The system application may be stopped with **app stop** and restarted with **app start**.

Important: While services can be started across a cluster, they should not be stopped using the **cluster** command.

In other words, do not run **cluster run <where> app stop <no arg>**

For details on the cluster command and the `<where>` clause, see: [CLI Commands](#).

By default this is a non-blocking command, which means that the console prompt will be available after running this command while processes that are a part of it are running.

It is possible to start or stop individual applications and/or processes by appending the `<application-name>[:<process-name>]`. The list of applications can be seen by using the command **app status**.

For example, to start the process `voss-queue`:

app start voss-queue

or

app start voss-deviceapi:voss-queue

It is possible to perform a blocking start by including `blocking` after start but before the `<application-name>[:<process-name>]`. For example:

app start blocking

app start blocking voss-queue

This will ensure that all background processes that are started by **app start** will be completed before the console prompt is available.

7.3. Installing Applications

In general, it is not necessary to install single applications on a running system. Instead, the system is upgraded using **cluster upgrade** as described in the Upgrading Applications section.

The system collects all visible application versions for display using **app list**.

A screen for the **app list** command from the command line interface:

```
platform@cscluster1:~$ app list
selfservice - selfservice install script
  latest version 1.5.0 (2015-11-23 17:28) is installed

voss-deviceapi - voss-deviceapi install script
  latest version 1.5.0 (2015-11-23 17:26) is installed

cluster - cluster management
  latest version 1.5.0 (2015-11-23 17:25) is installed

template_runner - Template Runner
  not installed
  version available: 1.5.0 (2015-11-23 17:36)

mongodb - no-sql database server
  latest version 1.5.0 (2015-11-23 17:24) is installed

support - Diagnostic tools for tech support
  latest version 1.5.0 (2015-11-23 17:27) is installed

snmp - snmp management client and server
  latest version 1.5.0 (2015-11-23 17:31) is installed

nginx - lightweight web server
  latest version 1.5.0 (2015-11-23 17:24) is installed

services - Platform core services
  latest version 1.5.0 (2015-11-23 17:26) is installed

platform - Platform install/upgrade
  latest version 1.5.0 (2015-11-23 17:26) is installed

nrs - NRS install script
  not installed
  version available: 1.2.0 (2015-11-23 17:36)

security - Latest system security updates
  latest version 1.5.0 (2015-11-23 17:34) is installed

vmware - VMware tools
  not installed
  version available: 1.5.0 (2015-11-23 17:24)
```

Each application will display a short description of the application, the version installed, and whether other versions are available.

Additional applications can be downloaded using the instructions detailed under the System Control:Download section.

7.3.1. Single Application Installation

An application can be installed with **app install <application-name>[version:<version>] [delete-on-success <[yes|y]||[no|n]>]**

For example,

- **app install snmp**
- **app install platform version 0.8.1 2014-01-09 00:46**

For release 19.1.2 and later; when using the **app install** command to install patch scripts that have been added to the `media/` directory, the **delete-on-success** parameter and `yes|no` value can be added to remove or keep the the patch file in the directory after successful installation.

Example: **app install media/patch.script delete-on-success yes**

By default, in other words without this parameter, the user is prompted whether to delete or keep the patch script:

```
Do you want to remove patch after successful completion?
```

If the `--force` parameter is appended to the parameter, no prompt is shown.

Note: This script file removal using the `delete-on-success` parameter only applies to *patch installation* and not to the installation of delta bundles or other bundles.

The application will automatically start its processes on install. In isolated rare cases, it may be necessary to manually provision the system afterward with **cluster provision** so that other applications are configured to work with the new application. This is taken care of automatically during the upgrade process described below.

7.3.2. Multiple Application Installation

More than one application can be installed by adding them as `multi_install` parameters:

A file `media/<manifest_file>.manifest` can be created to list the applications per line or space-separated, in the install sequence.

The `delete-on-success` parameter is optional.

- If the applications are listed on the command line, the sequence of the application parameters in the command is the sequence of the installation.

app multi_install [<manifest_file.manifest>|<app-name_1> <app-name_2> ...] [delete-on-success <[yes|y]||[no|n]>]

Note:

- It is advisable to make a system backup prior to installation so that a backup restore can be carried out in the case of a failure. Contact VOSS support if needed.
- If an application URI is specified as a parameter, it is downloaded.
- If an application is already installed, it is skipped with a message.
- If an application with more than one version is found in `media/`, the latest version is installed.

- If an application in the list fails to install, the installation is terminated. Refer to the install logs.
- If the `delete-on-success` parameter is `yes`, the manifest file and application files are removed after successful installation.

- The applications and manifest file can also be bundled into a single archive with extension: `.tar.gz`, `.tar`.

The command will then be of the format:

app multi_install <archive_file> [delete-on-success <[yes|y]||[no|n]>]

The archive is extracted to the `media/<archive_basename>` directory.

Note:

- If no manifest file is found, the installation exists with an error message
 - If a file in the manifest is not available, an error message for the entry is displayed.
 - The same considerations apply as noted with a multiple installation that lists the files on the command line.
 - If the `delete-on-success` parameter is `yes`, the `media/<archive_basename>` directory is removed after successful installation.
-

7.4. Updating Applications

The entire cluster can be upgraded from a single node with the command:

cluster upgrade <ISO>|<URL>

Some issues to note:

- By default, the cluster upgrade is carried in parallel on all nodes and without any backup in order to provide a fast upgrade. For backwards compatibility, this command is the same as for example **cluster upgrade <ISO> backup none fast**.
- Use the **cluster upgrade <ISO>|<URL> serial** command if the VMware host is under load.
- If a backup is required, use the **backup <location>** parameter with a location as it was added with the **backup add** command. The command parameter can for example be:

cluster upgrade <ISO>|<URL> backup <location>

A downloadable URL can be specified using **cluster upgrade <URL>** which will first be downloaded before upgrading. For example:

cluster upgrade http://myserver/path/myfile.iso

If a downloadable URL is not available, use the instructions detailed under the Download section to copy the application upgrade package to the local server. Once complete, use **cluster upgrade <ISO path>** to upgrade the application, or for example: **cluster upgrade myfile.iso**.

A single, cluster-of-one node can be upgraded using the local variant:

cluster upgrade <ISO>|<URL>

By default, the **upgrade** will upgrade the system using the latest ISO upgrade package in the `platform` user's directory.

The system will automatically reprovision itself after upgrading if necessary.

Note that the system will automatically perform a full backup before the upgrade so that the system can be rolled back if necessary. The destination for this backup can be specified using the following syntax:

cluster upgrade <URL>|<path> backup <backup-destination>

Valid backup destinations can be listed with **backup list**. If necessary, it is also possible to instruct the system not to perform a backup by specifying the backup destination as `None`.

See the backup restore section on how to restore a system to a former snapshot in order to revert to the snapshot prior to upgrade.

7.5. Summary Attribute Migration

If template updates that modify the summary attributes of existing models are installed, then the summary attributes in the data of existing instances of the models need to be migrated.

If the summary attributes are not migrated, the list view representation of these model types will not contain the correct columns or values for display.

The command to carry out this migration for a specific model is:

voss migrate_summary_attributes <model_type>

For example:

```
voss migrate_summary_attributes data/InternalNumberInventory
```

7.6. Remote Execution in Clusters

When commands are run on a cluster, a number of options are available to specify the nodes on which the commands can be run.

There is a *<where>* clause: **cluster run <where>**.

The clause can take:

- *role* - the role of the node: `application`, `database`, `webproxy`
- `all` - all the roles, in other words, the entire cluster
- `notme` - all nodes except the one the command is run on
- `<data center name>` - all nodes in the data center
- `<nodename | IP>` - the hostname or IP address of the node.

For example:

- **cluster run database app start mongod** will restart the mongod service on all database nodes.
- **cluster run all app status** will display the app status of all nodes on the cluster.
- **cluster run notme system shutdown** would issue the command to shut down all nodes except the one the command is run on.

Important: In a cluster, reboot and shutdown of the entire cluster should be done on each node and not with the **cluster run all** command.

Sometimes there are long-running processes running on a server. To display such jobs, use the **cluster job list** command.

Note: The **cluster job list** command is not available on a web proxy node.

It is also possible to re-attach to those jobs to see the output, using **cluster job reconnect <pid>**.

7.7. List of Unused Cluster Commands

The following table shows **cluster** commands that should not be used:

command
cluster run all system shutdown
cluster run all system reboot
cluster run <any> backup create XXX
cluster run <any> cluster upgrade
cluster run <any> cluster run XXX
cluster run <any> cluster provision
cluster run all cluster add <ip>
cluster run all cluster del <ip>
cluster run all cluster job kill <pid>
cluster run all cluster job list
cluster run all cluster job reconnect <pid>
cluster run all cluster list
cluster run all cluster status

Only one **cluster** command can be run on web proxy nodes. A message is shown if running the other **cluster** commands, for example running the **cluster list** and **cluster status** commands *on a web proxy node*:

```
$ cluster list
Invalid command syntax - refer to help descriptions

USAGE:
-----
cluster prenode - Prepares the system so that it can be joined to a cluster

$ cluster status
cluster status not available on webproxy nodes.
```

To exclude these **cluster** commands from running on web proxy nodes, use the `<where>` parameter to specify the node types, for example, with the `database` node type as specified below:

```
$ cluster run database cluster list
You are about to a run a command across the cluster, which could affect service_
↪availability.
Do you wish to continue? y

----- VOSS-UN-1, ip=164.168.151.3, role=webproxy,application,database, loc=cpt

Cluster has 8 nodes:

    application : 164.168.151.3, 164.168.151.5, 164.168.151.4, 164.168.151.7, 164.
↪168.151.6, 164.168.151.8
    webproxy : 164.168.151.3, 164.168.151.5, 164.168.151.4, 164.168.151.7, 164.
↪168.151.6, 164.168.151.9, 164.168.151.8, 164.168.151.10
    database : 164.168.151.3, 164.168.151.5, 164.168.151.4, 164.168.151.7, 164.
↪168.151.6, 164.168.151.8

----- VOSS-UN-3, ip=164.168.151.5, role=webproxy,application,database, loc=cpt

[...]
```

7.8. Self-Service Localization Management

Translation template files for the Self-Service application can be exported for translation and can be added or imported.

To export the Self-Service translation template:

```
selfservice get_translation_template
```

To add or import a Self-Service Translation template:

```
scp <local.json file> platform@<host>:media/
```

Log into system as platform user:

```
selfservice import_translation media/<local.json file>
```

Translation template files need to follow the naming convention:

```
locale-<languagecode>.json
```

For example:

- locale-en-us.json
- locale-es-es.json
- locale-de-de.json

7.9. Web Services

On a web proxy node only, Self-service and admin Web services can be disabled, re-enabled and listed if required. The task should be carried out after provisioning and if the Admin Portal or Self-Service GUI for example needs to be disabled for security purposes.

Note: 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.

The commands should be run on the relevant web proxy node. It is not recommended that the commands be run on a standalone single-node cluster system, but only on a multi-node cluster.

In particular, the commands to disable or enable web services 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 admin or self-service web services on a web proxy node, run the command on the relevant node:

web service disable <selfservice|admin>

- To enable admin or self-service web services on a web proxy node, run the command on the relevant node:

web service enable <selfservice|admin>

- To list disabled web services on an admin or self-service web services web proxy node:

web service list

For example:

```
platform@cscluster1:~$ web service list
disable: admin
```

8. System Control

8.1. System Commands Overview

This section covers commands available from the CLI that are started with the **system** command prefix. The commands are generic, operating system type commands. System commands that start with the prefix **voss**, are VOSS-4-UC specific system commands.

Note that the system commands that start with **system ssh_session_limit** are covered in the security topic on SSH session limits.

8.2. System restart

The system can be restarted with **system reboot** and shutdown with **system shutdown**.

Note: In a cluster, reboot and shutdown of the entire cluster should be done on each node and not with the **cluster run all** command - see: *Remote Execution in Clusters*.

8.3. Passwords

The password for the platform user is chosen at install time, but can be changed using **system password** which will then prompt for the old password, the new password and confirmation.

Passwords be least 8 characters long and must contain:

- at least one upper case letter
- at least one lower case letter
- at least one number
- at least one symbol

User passwords must be changed at least every 60 days.

Additional users can be created with **user add <username>**. Refer to the System Security : Creating additional users section.

Each user can be granted access to specific commands offering role-based access control - Refer to System Security : Granting and revoking user rights.

8.4. System Boot Passwords

Password protection can be enabled on the VOSS-4-UC boot loader configuration from the CLI. Commands are available to check, enable and disable the bootloader password.

- **system boot password** - Check if a bootloader password is enabled or disabled.
- **system boot password enable** - Prompts for and sets the platform user boot loader password. Refer to the topic on Passwords for password text requirements.
- **system boot password disable** - Disable the bootloader password if it is enabled.
- **system boot password reset** - If a password has been set, reset the bootloader password and enter a new password. If the system boot password is disabled when the command is run, a message will show this.

For example:

```
$ system boot password
System boot password disabled.

$ system boot password reset
You are about to reset the boot password. Do you wish to continue? y
System boot password is disabled. Enable the system boot password first.

$ system boot password enable
You are about to enable the boot password. Do you wish to continue? y
Valid passwords must contain:
    at least one lower- and one upper-case letter,
    at least one numeric digit
    and a special character eg. !#$%&^*
Please enter platform user boot password:
Password:
Please re-enter password
Password:
System boot password enabled.

$ system boot password reset
You are about to reset the boot password. Do you wish to continue? y
status true
Valid passwords must contain:
    at least one lower- and one upper-case letter,
    at least one numeric digit
    and a special character eg. !#$%&^*
Please enter platform user boot password:
Password:
Please re-enter password
Password:
System boot password enabled.
```

System boot passwords can also be enabled and set upon installation. Refer to the topic on Installation Details in the Installation Guide.

8.5. Federal Information Processing Standards (FIPS)

An administrator can check and enable the system for adherence to Federal Information Processing Standards (FIPS).

To check the system FIPS status, use **system fips**.

If FIPS is not enabled, the command output look as follows:

```
platform@nicodel:~$ system fips
FIPS mode is disabled
```

To enable FIPS on the system, use **system fips enable**.

Important: The use of FIPS on the system requires a subscription to the Ubuntu Advantage service package from Canonical in order to obtain the necessary cryptographic modules.

Internet access will be required from your system to the necessary Ubuntu Advantage service package URLs.

You will be prompted to:

- input the base and update system URLs as given for the program
- indicate if you wish to use a proxy and to provide its URL

Contact your VOSS account manager or VOSS support for detailed information on using the Ubuntu Advantage service package in the system.

Console output will be similar to the example below:

```
platform@nic-fips-un1:~$ system fips enable
Please enter the URL as given by Canonical for the base Ubuntu Advantage program
eg. deb https://<user>:<password>@private-ppa.launchpad.net/ubuntu-advantage/fips/
↵ubuntu <ubuntu version> main

URL: <URL>

Please enter the URL as given by Canonical for the Ubuntu Advantage update program
eg. deb https://<user>:<password>@private-ppa.launchpad.net/ubuntu-advantage/fips-
↵updates/ubuntu <ubuntu version> main

URL: <URL>

Do you want to use an apt proxy? y

What is the proxy URL?
<URL>

Installing required packages
```

If FIPS is enabled, the **system fips** command output is:

```
platform@nic-fips-un1:~$ system fips
FIPS mode is enabled
```

It is important to note:

- After running **system fips enable**, run **system reboot** to apply the FIPS enable changes.
- If fips mode is to be enabled on a cluster, it should be enabled on all nodes.
- *If FIPS is enabled on a system, it cannot be disabled.*
- All system passwords are stored using FIPS 140-2 complaint encryption algorithms, when FIPS mode is enabled or not.
- If FIPS is enabled on a system, all install scripts and templates are encrypted and decrypted using FIPS 140-2 complaint encryption algorithms.

8.6. File Management

Each user has a unique home directory in which local files can be stored. It is the user's responsibility to manage the disk space used by these files.

The command **diag disk** displays the disk usage. Files in the user's directory are displayed using the standard **ls** command, and deleted with **rm**.

New applications or upgrade packages are uploaded to the platform user using **scp** or **sftp**, for example **scp <filename> platform@192.168.0.1:** on the remote Unix file server. Refer to Network URI Specification for usage. If **sftp** or **ftp** is used, the following FTP servers are supported:

- OpenSSH for Unix or Linux based systems
- Titan SFTP and Cygwin (which is OpenSSH based) for Windows based systems.

A **sftp** or **scp** of files to VOSS-4-UC must be done in the `media` directory (`/opt/platform/admin/home/media`), which is a writable directory.

Alternatively a downloadable URL can be downloaded directly on the VOSS-4-UC system using **system download <URL>** and the downloaded file is placed in the platform user's directory, For example: **system download http://myserver/path/myfile.iso**

Individual applications are installed using **app install <filename>.script**. For details, see [Installing Applications](#). A list of available applications and versions is displayed using the command **app list**.

ISO packages include all the individual packages required for upgrading. Upgrade the system using **cluster upgrade <filename>.iso**. Alternatively, the ISO package file system can be mounted with the system **mount** command, and the individual applications are visible under the `media` directory, and visible via the **app list** command.

On a running system, the **app list** command shows versions and details of applications and patches installed or not installed on the node where the command is run, as for example in this output snippet on a web proxy node:

```
...
mongodb - NoSQL database server
  not installed
  version available: 11.5.2 (2017-09-14 16:16)

nginx - Nginx web server
  latest version 21.1.0 (2021-07-11 16:01) is installed
  other version available: 11.5.2 (2017-09-14 16:15)
```

(continues on next page)

(continued from previous page)

```

services - Platform core services
  latest version 21.1.0 (2021-07-11 16:07) is installed
  other version available: 11.5.2 (2017-09-14 16:15)

selfservice - selfservice install script
  not installed
  version available: 11.5.2 (2017-09-14 16:16)

voss-deviceapi - voss-deviceapi install script
  not installed
  version available: 11.5.2 (2017-09-14 16:16)

...

```

8.7. Drive control

In order to reduce the risk of *disk full* errors, the platform divides the file system over several disks keeping areas liable to grow outside the main root filesystem. The areas with the highest growth such as logs and database storage are kept on their own private file systems.

Note: The database mount point is stored in a logical volume.

These disk mounts can be migrated onto new, larger disks and some other locations can optionally be moved onto their own disks. This is managed through the **drives** command.

The current mounted filesystems and mount points can be displayed using **drives list mounted** and **drives list mountpoints** respectively.

A screen showing drives list mounted and drives list mountpoints:

```

platform@development:~$ drives list mountpoints
Used disks and mountpoints:
  sdc1 - services:backups
  dm-0 - mongodb:dbroot

Unused disks:
sde

Unused mountpoints:
  services:SWAPSPACE

Volume Groups
  voss - 25.0 GB free, 250.0 GB total
  Physical volumes:
    sdd1
  Logical volumes:
    dbroot/dm-0 - 225.0 GB

platform@development:~$ drives list mounted

```

(continues on next page)

(continued from previous page)

```
Used disks and mountpoints:
  sdc1 - services:backups
  dm-0 - mongodb:dbroot
platform@development:~$
```

The mount points are as follows:

- `mongodb:dbroot` is the volume used for database storage
- `services:backups` is used for default backup storage
- `services:appdata` is the main system volume used for application data in the users account
- `services:SWAPSPACE` is the swap volume used by the system

Note: While the system is carrying out a backup, additional *Unused disks*, for example `dm-1`, `dm-2`, may show when the **drives list** command is run. These disks are used for snapshots and will not display once the backup is completed.

In order to add or extend an existing disk volume, follow the following steps:

- Under VMware, add an additional disk volume to the VM
- **drives list** displays any unused available volumes
- A free mount point can be linked to a new disk using **drives add <disk> <mountpoint>**.

Note: The **drive add** command on a Generic NBI node (for Billing Data Extract) is not in use.

- An existing used mountpoint (i.e. currently linked to a disk volume) can be linked to a new disk volume of greater size using **drives reassign <disk> <mountpoint> <new disk size in GB>**.

The *<new disk size in GB>* specification means the volume need not be the size of the entire disk. However, a specified size must *not* be more than 90% of the disk size (or more than *<disk size less 10GB>* if the disk size is 100GB or smaller).

Existing data on the current disk will be copied to the new disk volume, and once successful, the new disk volume will be linked.

For example, the following steps can be followed to add a 250GB hard disk to the system:

1. Log into the VMware console and select Server.
2. Right-click and select Edit settings.
3. Click **Add...** and select Hard Disk.
4. Step through the rest of the wizard and edit parameters - in this case 250GB, thick provisioned.
5. Once done, log into system as the platform user.
6. Carry out a disk listing with the command **drives list**.
7. Reassign the disks with the command:
 - For the *database mount point (mongodb:dbroot)*, a Volume Group must be reassigned:
 - a. **drives create_volume <volume_name> <new disk name>**.
 - b. Carry out a disk listing to check the Volume Group with the command **drives list**.

- c. **drives reassign** <volume_name> **mongodb:dbroot** <size in GB>. The <size in GB> specification means the volume need not be the size of the entire disk. However, a specified size must *not* be more than 90% of the disk size (or more than <disk size less 10GB> if the disk size is 100GB or smaller).
 - d. Old volumes can be removed with:
 - **drives remove_logical** <volume_name> <logical_volume_name>
 - **drives remove_volume** <volume_name>
 - For *other mount points*, a disk must be reassigned: **drives reassign** <disk> <mountpoint>
8. Start the application with **app start**.
 9. Verify the new reassignment with the command **drives list**.

Note: Volume Groups for database mount points reserve a 10% or 10GB space - whichever is the largest - which is used for and then released during database backups.

SAN alignment is implemented using the offset value in **drives offset**. This value can be changed if necessary; however the default should be sufficient for most SAN hardware.

For swap partitions, use **drives checkswap** to check their alignment. Use **drives alignswap** to fix a misaligned swap partition.

8.8. Transaction Prioritization

There are three buckets for transactions in the VOSS-4-UC system Priority Queue, namely high, medium and low priority:

1. High Priority:
 - Self-service transactions: carried out by end users on the Self-Service interface
2. Medium Priority:
 - MACD operations on the VOSS-4-UC Admin Portal by Administrator users
 - API-based provisioning (HIL)
 - Any other transaction not in the Low Priority bucket
3. Low Priority:
 - Data Sync (LDAP, any device import transaction)
 - Bulk load transactions
 - Data import (import of data in JSON format)

From the Command Line Interface (CLI), a command is available to modify the default number of queue workers:

- Use **voss workers** to show the current number of queue workers. The default is 30.
- Use **voss workers <number>** set the number of queue workers.

An adjustment of the number of queue workers will impact on the number of parallel transactions that will run, which is a factor of this value as well the queue priority bucket to which the transaction belongs.

For example, Data Sync transactions may execute asynchronous workflows which are executed in parallel, or a Bulk Load transaction may have the Parallel flag set to True.

8.9. Banner

An administrator can manage a custom banner on the system from the CLI to display before login. The banner needs to be configured on a per node basis, in other words on each node in a cluster.

Banners are maintained on system update.

Banner text:

- Must be ASCII text in a UTF-8 file
- Can be up to 1600 characters. This includes spaces, tabs and other non-printing characters.
- Displays before login (SSH, SFTP, SCP)

An error message will display if the banner text is longer than the required length.

The following is an example:

```
$ ssh username@host

#####
# WARNING: Unauthorized access to this system is forbidden and will be #
# prosecuted by law. By accessing this system, you agree that your   #
# actions may be monitored if unauthorized usage is suspected.        #
#####

username@host's password:
```

The banner can be created in a file, uploaded to the system and then enabled.

1. Create the banner in a file, for example `banner`.
2. Upload the banner file, for example:

```
scp banner username@host:media/
```

3. When logged in on the system, remove any previous banner and add the uploaded banner:

```
system banner remove
system banner add media/banner
```

4. The banner will then be shown as in the example above.
 - To show the current banner, use: **system banner read** If no banner is available, a message will show.
 - To remove the current banner, use: **system banner remove**

8.10. Checksum

An administrator can generate the SHA256 checksum of a file such as an `.iso` image by using the **system checksum <path-to-file>** command.

The checksum of the file can then be compared to the one originally provided with the file, to verify its integrity.

9. Diagnostics

9.1. Health Report

On login, the system displays a health report indicating the status of the system before displaying the CLI user prompt. This health report shows the following:

```
Last login: Fri Aug 23 07:26:25 UTC 2019 from 172.29.41.201 on pts/0
host: VOSS-UN-1, role: webproxy,application,database, load: 0.27
date: 2019-08-23 07:28:14 +00:00, up: 1 day, 18:09
network: 192.168.100.3, ntp: 172.29.1.15
SECURITY: There are security updates available for your system. Please
  run 'security check' for more information.
HEALTH: NOT MONITORED
database: 25Gb
DATABASE TRANSACTION SIZE: 21.75GB
DATABASE TRANSACTION COUNT: 500003
Failed logins: 1 since Thu Aug 22 13:44:47 2019 from atlantic.biz

    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          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           template - template pack creator
security - security update tools

platform@development:~$
```

The report explanation is shown below:

Name	Description
Last login	Last console login and IP address source. This is only shown if there has been a previous login.
load	The load average of the system.
USERS	The number of CLI users currently logged in. This is only shown if more than one user is logged in.
up	The system uptime.
services	The status of the system services.
SECURITY	Whether security updates are available - shown if updates are available. Refer to the Security Patches section. Security updates are installed using security update .
HEALTH	A Health notification, for example a scheduled mail message, is set up or not.
database	Current database size.
DATABASE TRANSACTION SIZE	If the size of the TRANSACTION database collection exceeds 20GB, the current size is reported.
DATABASE TRANSACTION COUNT	If the number of entries in the TRANSACTION database collection exceeds 500,000, the current number is reported.
TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED	A schedule with either the command <code>voss transaction archive</code> or <code>voss transaction delete</code> is not set up.
Failed logins	If the user failed to log in prior to a successful login, the count, date and origin of the attempts are shown. A successful login resets this login count.

- disk, CPU and memory warnings are shown if applicable
- warnings are displayed in upper-case to draw attention
- DATABASE TRANSACTION warnings are shown for the following thresholds:
 - Transaction collection exceeds 500,000 documents
 - Transaction collection exceeds 20GB

A list of diagnostic tools is available in the topic on Diagnostic Tools.

- For schedule setups if the `HEALTH: NOT MONITORED` message is shown, see [Enable Health Monitoring](#).
- For schedule setups if the `TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED` message shows, see [Enable Database Scheduling](#).

9.2. Cluster Check

On a cluster, the **cluster check [verbose]** command is available to check:

- `network`: test and validate connectivity from each node to every other node, for each port required, as well as the time taken to connect to each node.
 - Checks for access to port 27020 on database hosts is not required from web proxy nodes.

- Checks for access to port 443 is only required from web proxy nodes to unified nodes.
- `database`: carry out a check of database configuration
 - `info`: displays database weights and whether the node state is primary, secondary or arbiter
 - `error`:
 - * if there is no connection to the database IP on a port
 - * if the current database weight does not match the configured weight
 - * if a node is marked as an arbiter but is not in the list of arbiters
 - `warn`: if the primary database node does not have the highest weight
- `disk`: carry out a drive space percentage check
- `ntp`: at NTP is functioning
- `packages`: Check status of packages installed by the system package manager. If an error occurs for a package, a message next to the package name shows: `package in an undesired state`.
- `security`: Check for security updates. Error status:
 - `info`: zero or one security update missed
 - `error`: more than one security update missed
- `cluster status`: also check the cluster status and
 - `info`: show status as OK
 - `error`: display a message to run **cluster status** for details
 - `warn`: It is advisable that these be resolved prior to upgrading where possible. Some warnings may be resolved by upgrading.

Note: If *only* node versions mismatch or some nodes are missing components, a `warning` status is displayed. This status will allow for an upgrade of a node during failover recovery.

This caters for scenarios during repair/recovery of nodes. The **cluster check** will warn about version mismatches and not prevent upgrade commands. The cluster check cannot distinguish between whether a recovery process is ongoing or a general fault exists. When no node recovery process is ongoing, then the warning should be treated as an error and resolved before upgrade commences.

This command should also be run *before* carrying out a system upgrade.

Note: Without the `verbose` parameter, the **cluster check** command will *only show warnings and errors*. Otherwise it would only show the message `No issues found with host checks`.

Use the `verbose` parameter to see detailed output.

Example output (abbreviated):

```
$ cluster check
warn
  192.168.322.3:
    drives
      /: 47 % utilised
  192.168.322.5:
```

(continues on next page)

(continued from previous page)

```

    drives
      /: 47 % utilised
192.168.322.6:
  drives
    /: 47 % utilised
error
  192.168.322.3
    network
      => 192.168.322.4:27020: Failed
192.168.322.4: Failed to connect to host
192.168.322.5
  network
    => 192.168.322.4:27020: Failed
192.168.322.6
  database
    arbiter: not configured
    weight: mismatched

[...]
  cluster
    status
      Error, please run `cluster status` for more information

```

Using the `verbose` parameter to see detailed output Any warnings and errors are then shown at the end of the verbose output.

Abbreviated example, `info` only; no issues:

```

$ cluster check verbose
info
  192.168.100.3
    database
      arbiter: ok
      state: ok
      weight: ok
    disk
      /: 28%
      /opt/platform: 27%
      /opt/platform/apps/mongodb/dbroot: 1%
      /tmp: 1%
      /var/log: 3%
    network
      => 192.168.100.4:8443: 0.223ms
      => 192.168.100.4:27020: 0.205ms
      => 192.168.100.5:8443: 0.246ms
      => 192.168.100.5:27020: 0.405ms
      => 192.168.100.6:8443: 0.169ms
      => 192.168.100.6:27020: 0.218ms
      => 192.168.100.7:8443: 0.225ms
      => 192.168.100.8:8443: 0.208ms
    ntp
      172.29.88.56: 18.313ms
  packages
    package database: ok
  security
    updates: 0 missed

```

(continues on next page)

(continued from previous page)

```

192.168.100.4
  database
    arbiter: ok
    state: ok
    weight: ok
  disk
    /: 28%
    /opt/platform: 27%
    /opt/platform/apps/mongodb/dbroot: 1%
    /tmp: 1%
    /var/log: 2%
  network

[...]
cluster
  status
    OK

```

9.3. Enable Health Monitoring

The steps below are to enable health monitoring if the system status displays `HEALTH: NOT MONITORED` upon login or when typing `help`.

1. Add an email relay address for outgoing email: add the SMTP IP address:

notify emailrelay <smtp ip address>

For example:

```
platform@host:~$ notify emailrelay 192.29.42.30
emailrelay: 192.29.42.30
```

The email relay can be verified with:

notify emailrelay

To disable or remove an external email relay or set the email relay back to default, use the following command:

notify emailrelay 127.0.0.1

2. Add a schedule instance for health reporting, for example with a schedule name `reports`:

schedule add reports log send output mailto:user@server.com diag health

For example:

```
$ schedule add reports log send output mailto:user@server.com diag health

Automatically setting time to midnight and enabling

reports:
active: true
command: log send output mailto:user@server.com diag health --force
hour: 0
min: 0
```

Typing **help** at the command line will now *not* show the `HEALTH: NOT MONITORED` message.

The schedule instance can be modified, for example weekly on Sunday:

```
schedule time reports weekly 0
```

9.4. Enable Database Scheduling

If the following message is shown upon login or when typing **health login**, the steps below should be carried out:

- `TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED`

For general information on scheduling, see: [Scheduling](#).

For considerations and guidance on frequency of the schedule to set up for your system. see the System Maintenance section of the Best Practices Guide.

Important: Schedules must be set up on *all* unified nodes. This ensures that they still run in the event of a failover and DR scenario.

- For *transaction* database maintenance:

With for example a schedule name `dbtxn`:

```
schedule add dbtxn voss transaction archive 7
```

or

```
schedule add dbtxn voss transaction delete 7
```

Note that transaction archiving also deletes transactions. For further details, see [Database Commands for Transaction Management](#)

Typing **health login** at the command line will now *not* show the `TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED` message.

The schedule instance can be modified, for example weekly on Sunday:

```
schedule time dbtxn weekly 0
```

9.5. Command History

A history of CLI commands issued can be displayed using **system history**. This command is the same as **log view platform/ui.log**, for example:

```
platform@development:~$ system history
Aug 15 10:55:44 node00 ui[6348]: app install support_install.script
Aug 15 10:56:26 node00 ui[11345]: diag
Aug 15 10:56:27 node00 ui[11351]: app install security_install.script
Aug 15 10:59:57 node00 ui[11609]: security
Aug 15 11:00:23 node00 ui[11704]: notify list
Aug 15 11:00:28 node00 ui[11892]: backup list
Aug 15 11:01:23 node00 ui[12483]: ssl
Aug 15 11:03:12 node00 ui[13562]: drives add sdd mongodb:dbroot
```

(continues on next page)

(continued from previous page)

```
Aug 15 11:07:21 node00 ui[16568]: voss
Aug 15 11:09:34 node00 ui[20740]: snmp
/var/log/platform/ui.log
```

Press **q** to exit the file.

9.6. Logs

The system maintains a comprehensive list of logs under `/var/log`:

- The `platform/` directory has logs pertaining to the general platform or to specific log types.
 - `apps.log` contains application and process control logging
 - `audit.log` a log type available if enabled with **log audit locallog on** - contains audit log information
 - `database.log` contains database logs spawned by VOSS-4-UC transactions
 - `dockerd.log` contains logs spawned by the Docker container management daemon
 - `backup.log` contains all logging pertaining to backups - available after the first backup
 - `cluster.log` contains all control level management of the cluster
 - `cluster_check.log` contains output of **cluster check** command. Run **cluster check** to have the latest cluster check information logged to `cluster_check.log` - see: [Cluster Check](#).
 - `config.log` contains information relating to the platform-level configuration
 - `event.log` a log type available if enabled with **log event locallog on** - contains event log information
 - `execute.log` contains low-level information about command execution
 - `notifications.log` contains information relating to SNMP notifications
 - `reports.log` contains information relating to system reports. Refer to the Scheduling section on how reports can be created.
 - `security.log` contains information relating to user security
 - `security_install_GUI.log` contains information relating to user graphical interface security
 - `security_update.log` contains information relating to security updates
 - `ui.log` contains higher-level information relating to UI commands being executed.
 - `wsgi.log` contains information relating to API-level commands via the WSGI server
- The `provision/` directory contains logs relating to provisioning. Every module provision is logged to component log files.
- The `health/` directory contains health logs. These are stored automatically every half hour, or whenever health is run, and are of the format `health/summary_report-<date>-<time>`.
- The `process/` directory contains process logs instrumental in debugging particular processes. All of the output from each process is logged to an individual file `process/<application>.<process>.log`
- The `install/` directory contains logs detailing the install process.

- The `mongodb/` directory contains logs relating to the Database function.
- The `nginx/` directory contains logs relating to the WebProxy function.
- The `voss-deviceapi/` directory contains logs relating to the Application function. For example, `voss-deviceapi/cnf_collector.log` is the Change Notification Collector log.
- The `billing-data-extract/` directory contains service logs relating to the Billing Data Extract module, if installed. For example, `callback.log`, `extractor.log`, `notifier.log`, `scheduler.log`.
- The `nbi-api/` directory contains logs relating to the Billing Data Extract module API, if installed.
- The `voss-portal/` directory contains nginx logs relating to the portal interface.

log list [`<search_string>`] is used to display a list of logs, optionally matching `search_string`. For example:

```
platform@clusternode:~$ log list alternatives.log
selfservice/alternatives.log
voss-deviceapi/alternatives.log
nginx/alternatives.log
mongodb/alternatives.log
alternatives.log
```

The main log rotation scripts will rotate log files only when files exceed 100M or if the disk containing `/var/log/` is over 80% full. This is checked once per hour. The system will attempt to keep 5 historic zipped files of each log. If the disk containing `/var/log` is over 90% full, files will be purged to ensure that the system continues to function.

100M size logs:

- `mongodb/*.log`
- `nginx/*.log`
- `selfservice/*.log`
- `voss-deviceapi/*.log`

10M size logs:

- other logs in `/var/log/` and sub-directories not specified above

All rotated log files and log files exceeding 1GB can be purged manually with **log purge**.

9.7. Viewing Logs

Once a file name is known, the particular log can be viewed with **log view** `<logfile>`. For example: **log view process/mongodb.router.log**.

When the log file is viewed, it can be searched for a particular regular expression using the forward slash `/`, as when using the **less** command.

A log file can also be watched or followed with **log follow** `<logfile>` The Unix command equivalent is: **tail -f**, so quit with **Ctrl-C**.

Log entries in the `voss-deviceapi/` directory have key-value pairs. The keys are as follows:

- `hostname` - the hostname of the server
- `level` - debug level

- `message` - the actual log message
- `name` - module where log occurred
- `parent process id` - Linux process parent id
- `process id` - Linux process id
- `request uuid` - identifier to group all logs generated in a request.
- `user` - user that generated the log
- `user hierarchy` - user that generated the log's hierarchy.
- `txn_id` - in logs that generate transactions (e.g. not in `request.log`) - the transaction uuid
- `txn_seq_id` - in logs that generate transactions (e.g. not in `request.log`) - transaction ID as seen in the Admin Portal
- `toplevel_txn_seq_id` - in logs that generate transactions (e.g. not in `request.log`) - toplevel transaction ID

Note that the system will attempt to auto-complete the prefix if it uniquely identifies a file, for example:

log view process/nginx

9.8. Sending and Collecting Logs

• log send

Single or multiple log files can be sent to a URI destination using **log send <URI> <logfile>** and **log send <URI> <prefix>** respectively.

The URI must match the URI specification detailed under the Network URI Specification topic.

An example of an email URI is `mailto:user@server.com`. All email communication requires **notify emailrelay** to be configured with the IP address of your mail relay.

When using **log send** to a `scp` and `sftp` destination, no port should be specified.

For example:

```
$ log send sftp://usr:pass@172.21.21.122/ install/voss-deviceapi_install.script-
↪20150819.log
```

• log send output

The output of a VOSS-4-UC CLI command can be sent to a URI destination using

log send output <URI> <CLI command>

For example:

```
$ log send output sftp://usr:pass@172.21.21.122/ app status
File transfer successful - 172.21.21.122:None/VOSS_1558945096-combined_logs.tar.gz
```

The transferred file (archive file format example: `var/log/users/platform/<command>.<yyyymmdd>`) then contains the output of the CLI command.

Note that *only VOSS-4UC CLI commands* will generate a file with command output. For example, while the command was **ls media** can be run from the VOSS-4-UC CLI prompt, it is not a VOSS-4-UC command and the contents of the transferred `lsmedia.<yyyymmdd>` file will be `command not found: ls.`

- **log sendnewer**

Log files newer than a certain date can be sent using **log sendnewer <yyyy-mm-dd> <URI>**. If the remote URI destination requires a password, it will prompt for the password.

A passwordless **scp** session can be enabled by generating keys locally with **keys create** and then sending the local keyset to the remote destination with **keys send user@<hostname>**.

- **log collect**

Use **log collect** to collect logs into an archive file. Both system and transaction logs can be collected. Mandatory and optional parameters are available to refine the log collection.

The syntax is:

log collect start <start-time> [include <logs|db|all>] [end <end-time>] [limit]

Note:

- Run **cluster check** before **log collect** to have the latest cluster check information logged to `cluster_check.log`.
- If the command is run from a web proxy node, only the system logs can be collected.
- The `start` and `end` parameters do not affect date range of system logs - they only apply to the date range of transaction collection logs.

-
- The `start` parameter is mandatory.
 - The `end` parameter is optional.

If omitted, the transaction collection logs are collected up to the current time.

- The `<start_time>` and `<end_time>` date value format can be:
 - * `+%Y-%m-%d_%H:%M:%S`, for example `2016-01-10_00:00:00`
 - * `+%Y-%m-%d`, for example `2016-01-10`
 - For a `<start_time>`, if the value is entered as `+%Y-%m-%d`, then the value used will be the start of that day, in other words: `+%Y-%m-%d_00:00:00`.
 - For an `<end_time>`, if the value is entered as `+%Y-%m-%d`, then the value used will be the end of that day, in other words: `+%Y-%m-%d_23:59:59`.
 - * Year-Month-Day format without leading zeroes, for example: `2016-1-10`
- The `include` parameter is optional. If *not* used, both transaction collection and system logs are collected.

If used, valid options are:

- * `logs`: collect database server, system, application and install log files:
 - `db_server_status.log` - database server status: the output of `mongo db.serverStatus()`
 - `disk_usage.log` - disk usage: output of command `df -h`

- `top.log` - current CPU consumption per process: output of command `top -b -n 1 -o +%CPU`
- `sar<DD>` - today's system activity by the system activity reporter `sar`: output of `sa2 -A`
- `db_collection_stats.log` - database collection stats: output of CLI command **voss db_collection_stats all**
- `db_index_stats.log` - database index stats: output of CLI command **voss db_index_stats**
- from the `/var/log` directory:

Log files with the following file wild cards are collected:

```
/var/log/syslog*
/var/log/dmesg*
/var/log/psmem*
```

Log files in the following directories are collected:

```
/var/log/nginx
/var/log/sysstat
/var/log/platform
/var/log/process
/var/log/voss-deviceapi
/var/log/provision
/var/log/install
/var/log/mongodb
/var/log/voss-portal
/var/log/billing-data-extract (if available)
/var/log/nbi-api (if available)
/var/log/m2uc-nx (if available)
```

You can inspect the list of collected files with the **log list** command and a search parameter, for example **log list install** to see all the install log files, including those for patches, for example:

```
platform@VOSS:~$ log list install
install/voss-deviceapi_install.script-20210709.log
install/disable-platform-password-expiry_patch.script-20210709.log
install/selenium_install.script-20210709.log
install/support_install.script-20210709.log
install/mongodb_install.script-20210709.log
install/nginx_install.script-20210709.log
install/snmp_install.script-20210709.log
...
```

This option *excludes* the transaction collection. The `start` and `end` parameters do not affect this collection.

- * `db`: collect transaction collection logs and exclude the system logs. By default, this includes:
 - transaction activity log records (`TRANSACTION.json.gz`)
 - the detailed content of transactions as seen on the GUI in the Log transaction list (`TRANSACTION_LOG.json.gz`).
 - worker transaction queue activity log records (`WORKER_QUEUE.json.gz`)

This option can only be added when the command is run on a unified or database node.

- * all: *both* transaction collection (db) and system logs (logs) are collected
 - The `limit` option *only* affects the transaction collection logs. It specifies that the detailed transaction collection logs (`TRANSACTION_LOG.json.gz`) are *not* included.
- This parameter is usually used if the logs are required for a task such as performance analysis, but not for debugging. Typically, all collection logs are needed for debugging.

Important:

- The `log collect` command with default parameters should not be run in parallel on all nodes.
- Transaction database data should not be collected from secondary nodes. So the commands issued on a primary would be different from that on secondaries, e.g.:

Primary

```
log collect start 2020-04-10 end 2020-04-14
```

Secondaries

```
log collect start 2020-04-10 include logs end 2020-04-14
```

An example of the console input and output of the command is shown below:

```
$ log collect start 2019-08-27 include db limit
2019-08-27T08:16:02.140+0000      connected to: localhost:27020
2019-08-27T08:16:02.148+0000      exported 3 records
Output saved to media/logs.VOSS.2019-08-27_08-16-01.tar.gz
```

The log file archive is of the format: `logs.<hostname>.<timestamp>.tar.gz`, where `<timestamp>` is the time that the collection was requested, in the format: `%Y-%m-%d_%H-%M-%S`. This file is created in the `media/` directory.

The log file archive can then for example be fetched with `scp`, for example:

```
scp platform@VOSS:media/logs.VOSS.2019-08-27_08-16-01.tar.gz
```

9.9. Log Types

The VOSS-4-UC system can log records of certain types, that can be logged locally or remotely. The log types contain events or transactions that originate from the:

- User interface
- API
- Command Line Interface (CLI)
- System activity (for example database connections)

The minimum specifications of the remote system for audit and event logs are:

- 2 VCPU's
- 80 GB HDD
- 2GB RAM

Log types:

1. Audit

Important: The available types of audit logs are determined by the audit log rule set that is active - see: [Audit Log Rule Sets](#)

The details below show details on the contents of types of audit logs.

- On the Admin Portal and Self-Service Portal GUI and API:
 - login and logout attempts (successful and unsuccessful) and session login time, logout time and expire events using any of the authentication methods:
 - * SSO
 - * LDAP
 - * VOSS-4-UC

Expired sessions will only be logged at 5 minute intervals.

- User account creations, modifications, disabling, and termination events. This means all create, update, delete operations on the `data/User` data model.

User modifications include user move operations from one hierarchy to another.

In particular, operations on the list of VOSS-4-UC models or attributes below, for: add, modify and delete.

- * `data/Role`
- * `data/AccessProfile`
- * `data/User.role`
- * `data/CredentialPolicy`

Note that these operations on any created models that refer to these core models are not logged.

- On the Command Line Interface (CLI):
 - login and logout attempts (successful and unsuccessful) and session login time, logout time and expire events; and also including:
 - * root shell login and logout using the `nrs` script
 - * `ssh`
 - * `scp`
 - * `sftp`
 - All root shell CLI commands are logged.
 - All CLI commands are logged. The audit log will show “CLI” or “Cluster” depending on how command was run.

For the creation of schedules (using **schedule**), these are logged, but the scheduled commands are not logged when they execute.

This includes for example user account creations, modifications, disabling, and termination events commands from the CLI:

- * **user add**
- * **user del**

- * **user grant**
- * **user revoke**
- * **user list**
- * **voss unlock_sysadmin_account**
- * See: [Audit Log Format and Details](#)

2. Event

- All transactions, sub-transactions as well as their details as seen when viewing the Transaction Log in the GUI.

Note that the detailed logs are not recorded. In other words, the rows of entries under the Logs table of a transaction as seen in the GUI under Administration Tools > Transaction are not shown in the event log, since the primary purpose of the log is auditing: “who did what”.

- See [Event Log Format and Details](#)

3. Stream

- Refers to the method of distribution of selected log files to a syslog server.
- See: [Log Streaming](#)

9.10. Audit Log Rule Sets

Audit log rule sets are available to manage the level of detail in audit logs. Types of logs can be added to or removed from rule sets by means of **log audit ruleset** command line parameters.

The following table shows rule sets and their default state:

Table 1: Rule Sets

Option	Name	Enabled
1	Default Rules	true
2	CLI Commands	true
3	Users and Groups	false
4	Network Events	false
5	Security	false
6	Software Management	false
7	Root Commands	false
8	File Access	false

For details on the logs associated with the rules, see:

- [Types of command and change logs in audit rules](#)
- the audit log description under [Log Types](#)

This means that by default, the audit log only shows logs associated with the default audit rules (1) and any VOSS-4-UC platform CLI commands (2).

The following parameters are available for the command **log audit ruleset**:

- **log audit ruleset list**

Show the current ruleset, in other words the enabled and disabled rules.

For example, consider the following (non-default option 7 has been enabled):

```
$ log audit ruleset list

          Option   Name
          -
Rules Enabled
          1   Default Rules
          2   CLI Commands
          7   Root Commands

Rules Disabled
          3   Users and Groups
          4   Network Events
          5   Security
          6   Software Management
          8   File Access
```

- **log audit ruleset disable 1,2**

Disable rules 1 and 2 from the rule set.

Note: The parameter syntax is a comma separated list of option numbers *without* spaces.

For example:

```
$ log audit ruleset disable 1,2

$ log audit ruleset list

          Option   Name
          -
Rules Enabled
          7   Root Commands

Rules Disabled
          1   Default Rules
          2   CLI Commands
          3   Users and Groups
          4   Network Events
          5   Security
          6   Software Management
          8   File Access
```

- **log audit ruleset enable 2**

Enable rule 2 from the rule set.

For example:

```

$ log audit ruleset enable 2

$ log audit ruleset list

          Option   Name
          -
Rules Enabled
          2      CLI Commands
          7      Root Commands

Rules Disabled
          1      Default Rules
          3      Users and Groups
          4      Network Events
          5      Security
          6      Software Management
          8      File Access

```

- **log audit ruleset enable all**

Enable all the rules.

For example:

```

$ log audit ruleset enable all

$ log audit ruleset list

          Option   Name
          -
Rules Enabled
          1      Default Rules
          2      CLI Commands
          3      Users and Groups
          4      Network Events
          5      Security
          6      Software Management
          8      File Access
          7      Root Commands

Rules Disabled

```

- **log audit ruleset default**

Reset the rules to the default set.

For example:

```

$ log audit ruleset default

$ log audit ruleset list

```

(continues on next page)

(continued from previous page)

Option	Name
-----	----
Rules Enabled	
1	Default Rules
2	CLI Commands
Rules Disabled	
3	Users and Groups
4	Network Events
5	Security
6	Software Management
8	File Access
7	Root Commands

9.10.1. Types of command and change logs in audit rules

Option #	Name	Purpose
1	Default Rules	Audit mgt tool, kernel , mount, swap, stunnel, cron events
2	CLI Commands	All Voss CLI commands logged in a clear text format
3	Users and Groups	User, group, sudo, password, login/logout events
4	Network Events	Hostname, pam, ssh, systemd, access failures, power state, session initiation, access control, etc
5	Security	Suspicious activity, reconnaissance, code injection, and privilege abuse
6	Software Management	Package management (dpkg, apt, aptitude)
7	Root Commands	Commands executed as root (high volume)
8	File Access	File access failures and deletion

9.11. Log Type Commands

The **log** command takes a log type parameter, as can be seen from the command syntax [audit|event|stream]:

```
$ log
USAGE:
-----
log [audit|event] locallog on|off          - Enable or disable audit/event
↳logging
log [audit|event|stream] remotelog        - Get the config for remote system
↳logging
log [audit|event|stream] remotelog <IP:port>|off - Configure a remote system for
↳sending logs
log [audit|event|stream|ssl] status       - Get the status for audit/event/
↳stream/ssl logging
```

For an overview of the log types and formats, see:

- [Log Types](#)
- [Audit Log Format and Details](#)
- [Event Log Format and Details](#)
- [Log Streaming](#)

Note: Audit log details are determined by the audit log ruleset - see: [Audit Log Rule Sets](#).

To enable or disable local audit and event logging, use the command and its respective option:

- **log audit locallog on|off**
- **log event locallog on|off**

Important: In a clustered environment, logging should be enabled or disabled on all application nodes in order to generate or stop logs completely, since a single transaction queue is utilized in the cluster and transactions can run on all application nodes. For commands on a cluster, see the **cluster run** command: [Remote Execution in Clusters](#).

If local logs are enabled, local log files of the type are available:

- Audit log files can be viewed as with all logs: **log view platform/audit.log**
- Event logs: **log view platform/event.log**

To enable remote logs of a type requires a remote system IP address and port as input parameters. The location and format of the logged data on the remote system would depend on the syslog application being used and the configuration of that application.

For remote system requirements, see: [Log Types](#).

Note:

- When audit or event logging is enabled or disabled locally or remotely, the syslog service restarts.
 - When stream logging is enabled or disabled, the syslog service restarts.
-

Remote log type disable CLI output example:

```
$ log audit remotelog off
You are about to restart syslog. Do you wish to continue? yes
You have new mail in /var/mail/<username>
```

The log type status for *both* local and remote logging can be checked with: **log audit status** or **log event status**, for example:

```
$ log audit status
audit:
  ip: 112.19.42.249:10514
  locallog: true
```

To check *only the remote* logging status of a log type: **log audit remotelog** or **log event remotelog**, for example:

```
$ log audit remotelog
   ip: 112.19.42.249:10514
```

Note:

- The internal rsyslog statistics are checked every 60 seconds to detect failed actions. If a failure is detected, the failure notification is retransmitted every 10 minutes.
- If the remote syslog server stops receiving logs, an email message or SNMP trap is generated, with the email message:

```
Subject: Log processing failure

Message: System unable to send <event type> messages to <IP>
```

In the case of an SNMP trap:

```
mteHotTrigger: Log processing failure

mteHotContextName: System unable to send <event type> messages to <IP>
```

- If the remote syslog server stops receiving logs, the local disk space of the queue of logs can grow to a maximum of 1GB before logs are not queued and log messages are discarded.

See [Warnings and Notifications](#) to set up the notification.

9.12. Audit Log Format and Details

The following is the format of an audit log entry. Line breaks have been added here for readability.

```
%b %d %Y %H:%M:%S.%f %Z|
UserID : %s
ClientAddress : %s
Severity : %s
EventType : %s
ResourceAccessed: %s
EventStatus : %s
CompulsoryEvent : No
AuditCategory : %s
ComponentID : CUCDM
AuditDetails : %s
App ID: %s
```

The first entry is the string format of the timestamp, while the %s is a variable for a value.

An example of the timestamp would be:

```
Oct 23 2015 10:54:28.615377 UTC
```

- Audit logs include logs for `auditd` and `audispd` which include system events. If system events are not required, they must be filtered by the client.
- All remote syslog streaming from VOSS-4-UC is via TCP. UDP is not supported.

The tables below show key and example descriptions in the audit log.

UserID	Username
"johnB"	Username on CLI or database
"johnB prov1.cust1"	GUI username and hierarchy
"ProviderUser@Provider.com"	User email address from GUI login
hidden	Invalid username

ClientAddress	IP address / pseudo terminal
"102.29.232.50:/dev/pts/1"	From IP: 102.29.232.50 and pseudo terminal /dev/pts/1
127.0.0.1	Internal API user
102.29.232.50	IP of GUI or API. Also Bulk Load, JSON import.

Severity	0-2. Higher is more severe
0	Basic log activity on the CLI. All log activity on the GUI or API.
1	All Rootshell activity
2	CLI: AuditCategory : Privileged, AuditDetails : user list and App ID: CLI - user may not run user list command

EventType	Type of event
UserLogging	Login, logout, expiry activity
FileDetection	File checksum activity
<AuditCategory>	GUI or API event type is the AuditCategory

ResourceAccessed	Resource accessed
CLI	CLI transaction
DB	Database logging
Application REST API	GUI or API resource

EventStatus	Status of the event
Success	Successful transaction
Failed	Failed transaction
Unknown	Note: Mongo successful login has this status

CompulsoryEvent	Not in use
No	Currently always No

AuditCategory	Activity category
AdministrativeEvent	non-privileged CLI command
Privileged	CLI transactions as root user, and commands by any user from the list below.
SecurityEvent	Login or logout to CLI, database,
PrivilegedDataModelAdd	e.g. GUI or API system user, including the type and operation. Type can also be <code>Mod</code> and <code>Del</code> . Details in <code>AuditDetails</code> .
DataModelAdd	e.g. GUI or API ordinary user, including the type and operation. Type can also be <code>Mod</code> and <code>Del</code> . Details in <code>AuditDetails</code> .
UserRoleChange	Transactions on the GUI, API flagged as privileged, including the type and operation. Details in <code>AuditDetails</code> .
UserLogin	Login on the GUI, API.
UserLogout	Logout on the GUI, API.
MultipleSourceLogin	Simultaneous login on GUI, API. Multiple sources in <code>AuditDetails</code> .

The CLI commands that are flagged as `Privileged`, are:

- **user** (and any parameters, such as **user del**)
- **voss unlock_sysadmin_account**
- **voss cleardown**
- **system password**
- **system reboot**
- **system shutdown**

The GUI and API commands flagged as privileged, are:

- carried out by a system user
- operations on the models:
 - `data/Role`
 - `data/AccessProfile`
 - `data/User.role`
 - `data/CredentialPolicy`

Audit Category for GUI and API transaction on a data model can be: *[Privileged]DataModel(Add/Delete/Update)*

Component ID	Identifier
CUCDM	The value is always CUCDM

App ID	Application
CUCDM	The application GUI and API interface
CLI	CLI command
CUCDM CLI	Rootshell login
CUCDM SSH	SSH login
CUCDM DB	Database, for example Mongo connect, login, logout

Audit Details	Details of transaction
Login	CLI or database login
“Login from 172.29.232.88”	GUI or API login also shows IP address
Logout	CLI or database logout
Login Invalid User	CLI or database login
Login Invalid Password	CLI or database login
User account locked - {} / {}	CLI or database login. Account locked after failed_login_attempts / allowed_attempts
User account expired	CLI or database login. Account expired
RootShell login	Root shell login
RootShell logout	Root shell logout
File checksum initialized	File checksum process initialized. The EventType is FileDetection.
<CLI command>	The CLI command that is run
“Resource type data/User named User Name: Joe”	Example of a create transaction on the data/User model.
“User Joe role updated to admin”	Example of a role update on a user.
“Login failed with Unknown from 172.29.232.88”	
[Basic NonInteractive SSO LDAP] Authentication on Log [in out]	Login or log out by a user using the indicated credentials (Basic, NonInteractive, SSO, LDAP). The log entry includes Client Address for source of the login.
Session Expired	Session timeout
Permission Error	Access control error: the user has no permission for an operation on a resource type from a hierarchy.
Invalid Request	If the request URL is not found (HTTP response is 400, 404)
Password retry limit reached. Locking account with username ..	When an account is locked due to failed password attempts
Unlocking account with username ..	When an account is unlocked
Locking account with username ..	When an account is locked

9.12.1. Example Syslog Messages

The following are example audit log entries.

Note: Line breaks have been added for readability.

```
API,Login,2019-10-29T21:11:20+00:00 VOSS audit: Oct 29 2019 21:11:20.042962 UTC|
UserID : CS-PAdmin
ClientAddress : 172.29.90.25
Severity : 0
EventType : UserLogin
ResourceAccessed : Application REST API
EventStatus : Success
CompulsoryEvent : No
AuditCategory : UserLogin
ComponentID : CUCDM
AuditDetails : Login with Mongo from 172.29.90.25 using interface None
App ID: CUCDM

API,Logout,2019-10-29T21:11:11+00:00 VOSS audit: Oct 29 2019 21:11:11.449544 UTC|
UserID : CS-PAdmin
ClientAddress : 172.29.90.25
Severity : 0
EventType : AuthLogout
ResourceAccessed : Application REST API
EventStatus : Success
CompulsoryEvent : No
AuditCategory : AuthLogout
ComponentID : CUCDM
AuditDetails : Logged out from 172.29.90.25
App ID: CUCDM

API,Access Control Bypass,2019-10-29T21:14:36+00:00 VOSS audit: Oct 29 2019 21:14:36.
→016777 UTC|
UserID : CS-PAdmin sys.hcs.CS-P
ClientAddress : 172.29.90.25
Severity : 0
EventType : PermissionError
ResourceAccessed : Application REST API
EventStatus : Failed
CompulsoryEvent : No
AuditCategory : PermissionError
ComponentID : CUCDM
AuditDetails : Read operation on model type data/Countries
App ID: CUCDM

API,Data Model Add,2019-10-29T21:31:33+00:00 VOSS audit: Oct 29 2019 21:31:33.872904.
→UTC|
UserID : CS-PAdmin sys.hcs.CS-P
ClientAddress : 172.31.252.1
Severity : 0
EventType : DataModelAdd
ResourceAccessed : Application REST API
EventStatus : Success
CompulsoryEvent : No
```

(continues on next page)

(continued from previous page)

```

AuditCategory : DataModelAdd
ComponentID : CUCDM
  AuditDetails : Resource type data/Role named
Name: Test
App ID: CUCDM

CLI, User Add,
"2019-10-29T21:45:42+00:00
VOSS audispd:
  node=VOSS
  type=ADD_GROUP
msg=audit (1572385542.608:242353) :
  pid=421859
  uid=0
  auid=1401
  ses=4
  msg='op=adding group acct=""testuser"" exe=""/usr/sbin/useradd"" hostname=? addr=?
↪ terminal=pts/0 res=success'

2019-10-29T21:45:42+00:00
VOSS audispd:
  node=VOSS
  type=USER_CHAUTHTOK
msg=audit (1572385542.736:242401) :
  pid=421872
  uid=0
  auid=1401
  ses=4
  msg='op=PAM:chauthtok acct=""testuser"" exe=""/usr/sbin/chpasswd"" hostname=? addr=?
↪ terminal=? res=success'

2019-10-29T21:45:42+00:00
VOSS audispd:
  node=VOSS
  type=PATH
msg=audit (1572385542.764:242413) :
  item=0
  name=""/opt/platform/users/testuser""
  inode=1654786
  dev=08:12
  mode=040700
  ouid=0
  ogid=0
  rdev=00:00
  nametype=NORMAL

2019-10-29T21:45:42+00:00
VOSS audispd:
  node=VOSS
  type=PATH
msg=audit (1572385542.768:242417) :
  item=0
  name=""/opt/platform/users/testuser/media""
  inode=1654788
  dev=08:12
  mode=040500

```

(continues on next page)

(continued from previous page)

```

oid=0
ogid=0
rdev=00:00
nametype=NORMAL

2021-05-26T15:27:33.715215+00:00 VOSS audit: May 26 2021 15:27:33.714993 UTC|
UserID : system
ClientAddress : 172.29.90.57
Severity : 0
EventType : SecurityEvent
ResourceAccessed : Application REST API
EventStatus : Failed
CompulsoryEvent : No
AuditCategory : SecurityEvent
ComponentID : CUCDM
AuditDetails : Password retry limit reached. Locking account with username john_smith.
App ID: CUCDM

...

```

9.13. Event Log Format and Details

Event log streaming sends all transaction data via `syslog`. It should be noted that the data is simply raw transaction data with no hierarchical grouping of parent and associated child transactions. If required, the remote `syslog` server must recreate a transaction tree hierarchy as part of log processing.

The following describes the format of an event log entry. The event log file contains single lines of data in JSON format, with meta data and data elements.

- Meta data has `event_` - attributes, and describes the type of event log.

For example, for `"event_type": "transaction.finalise"`: when a transaction is finalized in the system:

```

{
  "event_level": "INFO",
  "event_type": "transaction.finalise",
  "event_source": "voss-un1",
  "event_id": "abc08383-5adb-48cb-8181-ef6adc546791",
  "event_timestamp": "2017-12-04T12:18:07.025595Z",
  "event_message": "Transaction 1267 finalised.",
  "event_data": {
    [...]
  }
}

```

- `event_id`: Unique ID associated with the log entry.
- `event_message`: Message specified during log entry creation.
- `event_level`: Log level with which log entry was created.
- `event_timestamp`: Datetime string of timestamp when the log entry was created.
- `event_type`: Unique type associated with event described by log entry.
- `event_source`: Hostname of host from which log was created.

- `event_data`: Additional data associated with log entry, containing a `transaction` object.
- Data is recorded in the `"event_data"` element of the event type, with each event type determining its own `event_data` JSON structure.

For example, for `"event_type": "transaction.finalise"`, the `event_data` is `"transaction"`, the start of the structure is for example:

```
[...]
"event_data": {
  "transaction": {
    "status": "Success",
    "username": "sysadmin",
    "rolled_back": false,
    "resource": {
      "hierarchy": "1c0nfmo2c0deab90da595101",
      [...]
    }
  }
}
```

- `transaction`: Transaction specific data.
 - * `action`: Transaction's `action` field, which is displayed by the Admin Portal in its Action field.
 - * `completed_time`: Datetime string of the transaction's `completed_time` field, which is displayed by the Admin Portal in its Completed Time field.
 - * `detail`: Transaction's `detail` field, which is displayed by the Admin Portal in its Detail field.
 - * `duration`: Transaction's `duration` field (in seconds), which is displayed by the Admin Portal in its Duration field.
 - * `hierarchy`: Transaction's `execution_hierarchy` field.
 - * `message`: Transaction's `exception_message` field (if any), which is displayed by the Admin Portal in its Message field.
 - * `operation`: Transaction's `operation` field.
 - * `parent_pkid`: Transaction's `parent` field (if present, can be used to identify a parent transaction `pkid` value and any sub-transactions).
 - * `pkid`: Transaction's `_id` field (this value will be the `parent_pkid` of sub-transactions if there are any).
 - * `priority`: Transaction's `config['priority']` field, which is displayed by the Admin Portal in its Priority field.
 - * `processor_host_name`: Transaction's `processor_host_name` field, which is displayed by the Admin Portal in its Processor Host Name field.
 - * `resource`: resource object associated with transaction
 - `hierarchy`: Transaction's `resource hierarchy` field.
 - `model_type`: Transaction's `resource model_type` field, which is displayed by the Admin Portal in its Model Type field.
 - `pkid`: Transaction's `resource pkid` field.
 - * `rolled_back`: Transaction's `rollback` field, which is displayed by the Admin Portal in its Rolled Back field.
 - * `started_time`: Datetime string of the transaction's `started_time` field, which is displayed by the Admin Portal in its Started Time field.

- * `status`: Transaction's `status` field, which is displayed by the Admin Portal in its Status field.
- * `submitted_time`: Datetime string of the transaction's `submitted_time` field, which is displayed by the Admin Portal in its Submitted Time field.
- * `submitter_host_name`: Transaction's `submitter_host_name` field, which is displayed by the Admin Portal in its Submitter Host Name field.
- * `txn_seq_id`: Transaction's `txn_seq_id` field, which is displayed by the Admin Portal in its Transaction ID field.
- * `username`: Transaction's `username` field, which is displayed by the Admin Portal in its Username field.

A full example of a `transaction.finalise` type event log entry is shown below (formatted multiline):

```
{
  "event_level": "INFO",
  "event_type": "transaction.finalise",
  "event_source": "voss-un1",
  "event_id": "abc08383-5adb-48cb-8181-ef6adc546791",
  "event_timestamp": "2017-12-04T12:18:07.025595Z",
  "event_message": "Transaction 1267 finalised.",
  "event_data": {
    "transaction": {
      "status": "Success",
      "username": "sysadmin",
      "rolled_back": false,
      "resource": {
        "hierarchy": "1c0nfmo2c0deab90da595101",
        "model_type": "data\\~/Countries",
        "pkid": "5a203da004222e1c67f93c83"
      },
    },
    "processor_host_name": "voss-un1",
    "pkid": "233cd3b1-8acc-4702-bd64-90653c02cd81",
    "hierarchy": "1c0nfmo2c0deab90da595101",
    "started_time": "2017-12-04T12:18:04.946000Z",
    "detail": "Australia",
    "completed_time": "2017-12-04T12:18:07.022000Z",
    "priority": "Normal",
    "duration": 2.076404,
    "submitted_time": "2017-12-04T12:18:04.461000Z",
    "submitter_host_name": "voss-un2",
    "txn_seq_id": "1267",
    "parent_pkid": null,
    "action": "Update Countries",
    "message": null,
    "operation": "update"
  }
}
```

9.14. Log Streaming

VOSS-4-UC platform CLI commands are available to enable and manage the streaming of a configurable list of log files to a remote syslog server that can then for example generate alarms based on log messages. The

VOSS Assurance product can for example be configured to be the remote server that monitors the system.

The following parameters are available:

- Remote IP
- Remote port
- Protocol
- Files to stream

The **log stream** command parameters can therefore be used to:

- set up and remove up the server
- add and remove files to stream to a server

Note: The **log stream** commands that make changes will restart the `syslog` service.

Refer to the command examples for parameter usage and CLI output.

Command Examples

- **log stream remotelog**

```
platform@VOSS:~$ log stream remotelog
ip: Unset
```

- **log stream remotelog <IP:port>|off**

```
platform@VOSS:~$ log stream remotelog 192.168.100.25:514
You are about to restart syslog. Do you wish to continue? y
System updated to use 192.168.100.25 as remote syslog server
You have new mail in /var/mail/platform
```

- **log stream status**

```
platform@VOSS:~$ log stream status
stream:
  ip: 192.168.100.25
  port: 514
  protocol: udp
```

- **log stream remotelog status**

```
platform@VOSS:~$ log stream remotelog status
ip: 192.168.100.25:514
```

- **log stream enable <filename>[,<filename>...]**

```
platform@VOSS:~$ log stream enable syslog,auth.log
You are about to restart syslog. Do you wish to continue? y
syslog already in list of files to send to remote
Updating list of files to stream to remote
You have new mail in /var/mail/platform

platform@VOSS:~$ log stream status
```

(continues on next page)

(continued from previous page)

```

stream:
  filenames:
    syslog
    auth.log
  ip: 192.168.100.25
  port: 514
  protocol: udp

platform@VOSS:~$ log stream enable dmesg,dpkg.log
You are about to restart syslog. Do you wish to continue? y
Updating list of files to stream to remote
platform@VOSS:~$ log stream status
stream:
  filenames:
    syslog
    auth.log
    dmesg
    dpkg.log
  ip: 192.168.100.25
  port: 514
  protocol: udp

```

If a file is already in the list, a message shows: “already in list of files to send to remote”.

- **log stream disable <filename>**

```

platform@VOSS:~$ log stream disable syslog
You are about to restart syslog. Do you wish to continue? y
syslog not currently in list of files to send to remote
You have mail in /var/mail/platform

```

- **log stream protocol <tcp|udp>**: udp is default if not specified.

See also:

- [Log Types](#)
- [Log Type Commands](#)

9.15. Remote Log Type Encryption

The VOSS-4-UC system can encrypt remote log types: `audit`, `event` or `stream`.

The steps and commands to follow for remote log type encryption are set out below:

1. Edit SSL details on the system. (The user is prompted for C,ST,O,OU,FQDN):

log cert details edit

Inspect the edited SSL details:

log cert details

2. Generate a Certificate Signing Request (CSR) file and submit it to the certificate authority (CA).

log cert gen_csr

The CSR file can also be printed out:

log cert print_csr

3. Receive the signed certificate. Then upload it to the system (using for example **scp**) and add your signed certificate with:

log cert add <filename>

For example:

```
$ log cert add media/cert.pem
```

Add the CA certificate to the system with:

log cert addca <filename>

For example:

```
$ log cert addca media/ca-chain.cert.pem
```

Inspecting the SSL details at this stage, using **log cert details**, shows the SSL details for:

- Issuer data
- Key data
- User set details

4. Enable remote logging of the log type. This will restart the syslog server.

log [audit|event|stream] remotelog <IP:port>

5. Enable SSL on log type logging. This will restart the syslog server.

log ssl enable

SSL logging of log type can be disabled by the command **log ssl disable**. This will restart the syslog server.

To see SSL logging details and if it is enabled or not, run **log ssl status**.

For example, the output below shows `enabled: false`:

```
user@host:~$ log ssl status
ssl:
  C: ZA
  CN: VOSS.visionoss.int
  L: Cape Town
  O: Voss-Solutions
  OU: Platform
  ST: WP
  email: user@host.com
  enabled: false
```

9.16. The Mail Command

The system monitors a number of events – these are described in more detail in the topic on Warnings and Notifications. The events can be signalled externally using email and snmp. However, a local copy of all events is maintained in the platform user's mailbox.

Command	Description
mail list	Display a list of events stored in the mailbox.
mail read all	Read all mail.
mail read <number>	Read a specific mail message.
mail del <number>	Delete a specific mail message.
mail del <from> <to>	Delete a range of mail messages.
mail del all	Delete all mail messages.

Mail events may accumulate over time. The system will purge old events automatically if the mailbox becomes too full (more than 500 messages).

9.17. Diagnostic Tools

There is an extensive list of diagnostic tools available under the **diag** menu.

For example output of **diag** commands, see: [diag Command Examples](#).

```
platform@VOSS:~$ diag
USAGE:
-----
diag config [<config-item>] - Display configuration, optionally
↳specifying the configuration prefix, diag config /
↳platform/network/name
diag config app <app-config> [<config-item>] - Display application configuration,
↳optionally specifying the configuration prefix, eg diag config
↳app snmp /community
diag date <epoch timestamp> - Convert epoch timestamp to datetime
diag disk - Display diagnostics for disk usage
diag epoch <yyyy-mm-dd hh:mm> - Convert datetime to epoch timestamp
diag filehash - Display the file system hash integrity
↳check
diag filehash init - Initialise the file system hash
↳integrity check
diag free - Display diagnostics relating to free
↳memory
diag health - Display a health report
diag health report - Save a health report as a logfile
diag iostat - IO subsystem statistics
diag largefiles - Find the largest files on your system
↳no more
than the top 10 items are display
diag mem - Display memory diagnostics
diag monitor - Display the system resource analysis
diag monitor list - Display the long-term average and
↳maximum system resource usage
diag nicstat - Network interface statistics
diag ping <host> - Ping a remote host to test network
↳reachability
```

(continues on next page)

(continued from previous page)

<code>diag proc</code>	- Display a list of system processes
<code>diag resolve <host></code>	- Resolve a hostname to IP address
<code>diag stats</code>	- Display the container statistics
<code>diag tasks</code>	- Display constant task listing
<code>diag test_connection <host> <port></code> →to a remote port	- Test if system can open a connection.
<code>diag top</code>	- Process resource statistics
<code>diag traceroute <host></code>	- Discover the network path to <host>
<code>diag vmstat</code>	- Virtual memory subsystem statistics

Command Descriptions

Some of the commands are provided with details below:

Command	Description
<code>diag ping <host></code>	Test network reachability to a network host.
<code>diag resolve <hostname></code>	Test DNS resolution of a hostname.
<code>diag test_connection <host> <port></code>	Given a host IP and port number, return a message: “Successfully connected to <host>:<port>” or “Failed to connect to <host>:<port>”.
<code>diag free</code>	Display the memory usage.
<code>diag disk</code>	Display the disk usage. Logical volumes for the database have the format <code>/dev/mapper/voss-dbroot</code> .
<code>diag mem</code>	Display a more detailed memory usage by process.
<code>diag health</code>	Display a comprehensive health summary. This includes status on the following: CPU, Memory, Disk, Security Update, Application, Cluster, Cluster Failover and Health email scheduling. Logical volumes used by the database have the format: <code>/dev/mapper/voss-dbroot</code> .
<code>diag top</code>	Display a single Unix top summary.
<code>diag unittests</code>	Utility for developers only. Note that services will be restarted by this utility.
<code>diag filehash</code>	Although a checksum of system and configuration files is carried out regularly, a manual check for changes since the previous check can be carried out. If any files have changed, these will be listed in the command output.

For `diag health`, see also *Diagnostic Troubleshooting*.

9.18. Diagnostic Troubleshooting

The health displayed on login will normally include sufficient information to determine that the system is either working, or experiencing a fault. More detailed health reports can be displayed with **diag health**.

Important: Since the `diag health` command output is paged on the console, you can scroll up or down to

see all the output.

Type `q` at the `:` prompt to quit the console pager and output (*not Ctrl-C*).

```
platform@atlantic:~$ diag health

Health summary report for date:
  Mon Aug 16 12:56:51 UTC 2021
CPU Status:      12:56:51 up 12:13, 1 user, load average: 1.75, 1.72, 1.75
Platform version:
  platform v21.1.0 (2021-08-15 13:37) Network Status:      System name:  VOSS
  Device:  ens160 Ip:      Netmask:      Gateway: 192.168.100.1
Memory Status:
  total used free shared buff/cache available
Mem: 8152816 5820696 392336 164500 1939784 1872452
Swap: 2096124 112640 1983484
Disk Status:
  Filesystem Size Used Avail Use% Mounted on
  /dev/sdal 18G 4.7G 13G 28% /
  /dev/sdb1 9.9G 154M 9.2G 2% /var/log
  /dev/sdb2 40G 8.4G 30G 23% /opt/platform
  /dev/sdc1 49G 53M 47G 1% /backups
  /dev/mapper/voss-dbroot 225G 5.1G 220G 3% /opt/platform/apps/mongodb/dbroot
Security Update Status:
  There are 0 security updates available for the base system.
Application Status:
  selfservice v21.1.0 (2021-08-15 13:36)
    |-node running
  voss-deviceapi v21.1.0 (2021-08-15 13:36)
    |-voss-cnf_collector running
    |-voss-queue running
    |-voss-wsgi running
    |-voss-risapi_collector running
    |-voss-monitoring running
  cluster v21.1.0 (2021-08-15 13:36)
  template_runner v21.1.0 (2021-08-15 13:43)
  mongodb v21.1.0 (2021-08-15 13:36)
    |-arbiter running
    |-database running
  support v21.1.0 (2021-08-15 13:43)
  selenium v21.1.0 (2021-08-15 13:42)
  ...
```

A rich set of SNMP and SMTP traps are described in the Notifications section which can be used to automate fault discovery.

Determine if all processes are running using **app status**. If a process is not running, investigate its log file with:

log view process/<application>.<process>

For example, checking processes:

```
platform@development:~$ app status
development v0.8.0 (2013-08-12 12:41)
voss-deviceapi v0.6.0 (2013-11-19 07:37)
  |-voss-celerycam          running
  |-voss-queue_high_priority  running
```

(continues on next page)

(continued from previous page)

```

...
core_services v0.8.0 (2013-08-27 10:46)
  |-wsgi                running
  |-logsizeimon         running
  |-firewall            running
  |-mountall            running
  |-syslog              running (completed)
  |-timesync            stopped (failed with error 1)
nginx v0.8.0 (2013-08-27 10:53)
  |-nginx               running
security v0.8.0 (2013-08-27 11:02)

```

Followed by a log investigation for a stopped process:

```

platform@development:~$ log view process/core_services.timesync
2013-08-15 10:55:20.234932 is stopping from basic_stop
2013-08-15 10:55:20:   core_services:timesync killed
  successfully
2013-08-15 10:55:20: Apps.StatusGenerator core_services:timesync
  returned 1 after 1 loops
App core_services:timesync is not running with status stopped

...

+ /usr/sbin/ntpdate 172.29.1.15
2014-02-04 09:27:31: Apps.StatusGenerator core_services:timesync
  returned 0 after 1 loops
2014-02-04 09:27:31: WaitRunning core_services:timesync is reporting
  return code 0
core_services:timesync:/opt/platform/apps/core_services/timesync
  started
4 Feb 09:27:38 ntpdate[2766]: no server suitable for
  synchronization found
+ echo 'Failed to contact server: 172.29.1.15 - retrying'
Failed to contact server: 172.29.1.15 - retrying
+ COUNTER=2
+ sleep 1
+ test 2 -lt 3
+ /usr/sbin/ntpdate 172.29.1.15
4 Feb 09:27:48 ntpdate[3197]: no server suitable for
  synchronization found
+ echo 'Failed to contact server: 172.29.1.15 - retrying'
Failed to contact server: 172.29.1.15 - retrying
+ COUNTER=3
+ sleep 1
+ test 3 -lt 3
+ test 3 -eq 3
+ echo 'Timesync - could not contact server 172.29.1.15 after
  three tries. Giving up'
Timesync - could not contact server 172.29.1.15 after
  three tries. Giving up
+ exit 1

```

The error message and return code being displayed in the browser is also invaluable in determining the cause of the problem.

The system resources can be inspected as follows:

- **diag disk** will display the disk status
- **diag free** and **diag mem** will display the memory status
- **diag top** will display the CPU status

10. Notifications

10.1. Warnings and Notifications

On console sign-in, a health report indicates the system status. This health report shows this data:

```
Last login: Tue Sep 3 10:19:07 2013 from 172.29.232.68
host: alan, role: standalone, load: 0.35, USERS: 3
date: 2013-09-03 10:20:02 +00:00, up: 2:05
network: 172.29.89.182, ntp: 172.29.1.15
SECURITY UPDATES: 136 security updates available
database: 8.0Gb
services: ok
```

The report values mean:

- last console sign-in and IP address source
- the load average of the system
- the number of users currently signed in
- the system uptime
- the status of the system services
- whether security updates are available
- disk, CPU, and memory warnings if applicable
- warnings are displayed in uppercase to draw attention

The report can be redisplayed by typing the command:

health

The system can be configured to forward warnings and notifications to various destinations, including:

- local email
- remote email addresses
- remote SNMP destinations

Local email allows the administrator to view a list of warnings, and delete them as necessary.

The notification destinations can be displayed with **notify list**. The destinations for each event level can be set with **notify add info|warn|error <destination-URI>** Refer to the Network URI Specification topic for a detailed description of URIs. Note that email notifications require the mail relay to be set with **notify emailrelay <relayhost>**. A test event can be generated with **notify test info|warn|error** to test the notification delivery mechanism.

Examples:

- **notify add info mailto:sysadmin@mycompany.com**
- **notify add error snmp://public@mynmpserver.com**

```
$ notify add error snmp://public@mynmpserver.com
notifications:
  emailrelay: 172.1.1.1
  level:
    error:
      snmp://public@mynmpserver.com
      mailto:platform@localhost
    info:
      mailto:platform@localhost
    warn:
      mailto:platform@localhost
```

In addition to external email and SNMP alerts, the system also records various events to a local mailbox. Refer to the Mail Command section for details.

SNMP CPU load notifications are set using:

snmp load <1min load> <5min load> <15min load>

This results in notifications being sent should the threshold be exceeded. For a server with 2 CPUs, it is recommended that this setting be:

snmp load 8 4 2

This means that notifications are sent if the 2-CPU system load averages over the last 1, 5, and 15 minutes reach these values.

10.2. Events and SNMP Messages

The following conditions are monitored for which SNMP traps can be sent. The trap levels and message strings are shown for the condition:

- Script install failures
 - ‘error’, ‘upgrade failed’, ‘upgrade failed as other activity is in progress’
- Backup Success/Failure
 - ‘error’, ‘Backup failed’,
 - ‘info’, ‘Backup completed’
 - ‘error’, ‘ERROR: The last backup was more than 2 days ago’, ‘Backup list:
 - ‘info’, ‘INFO: Backups now runs regularly’, ‘Backup list:
- Restore Success/Failure
 - ‘info’, ‘Backup restored
 - ‘error’, ‘Backup restoration failed
- Nginx reconfiguration - If a webproxy is unable to contact one of the upstream systems
 - “error”, “nginx upstream failure”, “upstream %s server %s failed: %s”
 - “info”, “nginx upstreams OK”, “nginx upstream servers returned to normal”

- Disk full/cleared (if a monitored disk is above 80%, it will send a trap, and also when this is cleared)
 - level, 'DISK ALMOST FULL: Disk <disk name> is more than 80 percent full'
 - level, 'DISK STATUS: Disk <disk name> is now running below 80 percent'
 - 'error', 'DISK ALMOST FULL: Disk /var/log is more than 80 percent full','Use log purge to purge all rotated logsnnCurrent disk status:
- Email failure (if a monitoring email was set up, and system cannot reach it)
 - 'error', 'ERROR: Trouble sending health email', 'Trouble sending health email'
 - 'info', 'INFO: Health emails is now being sent', 'Health emails is now being sent'
 - 'info', 'INFO: Messages for <username> auto archived as it reached more than 500' % user, 'Use the following command to view archived messages:nnlog view <username>' %
 - 'info', 'INFO: The total local messages for <username> is now under 200'
 - 'warn', 'WARNING: Not all notify levels is configured with an external email address '
 - 'info', 'INFO: All notify levels is now configured with an external email address'
- Database usage
 - 'warn', reason='WARN: Database <database name> exceeded threshold'
 - 'info', reason='INFO: Database <database name> returned to normal'
- High disk latency
 - 'error', 'ERROR: Disk slow ', 'Disk latency info:
 - 'info', 'INFO: The disk latency returned to normal', 'The disk latency returned to normal.'
- Database failover (if the DB fails over from one node to another twice in a 5 minute period)
 - 'error', 'ERROR: The db is failing over constantly within 5 min', 'Cluster failover status:
 - 'info', 'INFO: The db failover status returned to normal', 'Cluster failover status:
- Large log file warning
 - 'error', 'ERROR: Log files larger than 1Gig found in /var/log ', 'Logrotation was executed to rotate the following logs: <log filename>'
 - 'info', 'INFO: /var/log rotated', '/var/log rotated'
- Network
 - 'error', 'ERROR: Network Failures', 'The following network failures occured: <network errors>'
 - 'info', 'INFO: Network failures resolved', 'Network failures resolved',
- Service failures
 - 'error', 'ERROR: Service Failures'
 - 'info', 'INFO: Services started successfully'
- Security updates available
 - 'warn', 'WARNING: Security Updates available', '<number> updates available'
 - 'info', 'INFO: Security Updates applied', 'There are 0 security updates available'
- High memory and CPU usage
 - 'error', 'ERROR: High memory usage', 'Memory activity:

- ‘info’, ‘INFO: Memory usage returned to normal’, ‘Memory more than 1024MB’
- ‘warn’, ‘WARNING: High CPU usage’, ‘CPU activity:
- ‘error’, ‘ERROR: Extremely high CPU usage’, ‘CPU activity:
- ‘info’, ‘INFO: CPU usage returned to normal’
- High swap usage
- NTP configuration issues
 - ‘warn’, ‘WARNING: The ntp daemon has stopped on <server name>’, ‘Run ‘app start services:time’ to restart ntpd’,
 - ‘warn’, ‘WARNING: The ntp offset exceeds 1 second on %s’ % system_info, ‘ntp offset exceeds 1 secondnCurrent ntp offset value: <ntp offset>’
 - ‘info’, ‘INFO: ntp is now configured for <server name>’, ‘NTP cleared’, value=0
 - ‘info’, ‘INFO: The ntp offset restored to normal on <server name>’, ‘ntp offset clearednCurrent ntp offset value: <ntp offset>’
- DNS configuration issues
 - ‘warn’, ‘WARNING: No dns configured for <server name>’, ‘It is recommended that the dns is configured.nnTo configure dns use the following command:nnetwork dns <server1> <server2>’
 - ‘info’, ‘INFO: dns is now configured for <server name>’, ‘DNS cleared’, value=0
- Domain configuration issues
 - ‘info’, ‘INFO: domain is now configured for <server name>, ‘Domain cleared’, value=0
 - ‘warn’, ‘WARNING: No domain configured for <server name>, ‘It is recommended that the domain is configured.nnTo configure the domain use the following command:nnetwork domain <server1> <server2>’,
- Log processing failure
 - ‘error’, ‘System unable to send <event type> messages to <IP>’
- Monthly report export failure: (example at [Internal Report Schedules](#))
 - ERROR: License file generation failed. The license audit report scheduled for <month> <year> was not successful. Please contact your VOSS account manager.

For details, refer to the topic on the individual SNMP trap.

10.3. SNMP Configuration and Queries

This topic covers configurations for various versions.

10.3.1. SNMP Configuration

SNMP must be configured under the SNMP menu and the SNMP URI needs to be configured for all the notify severity levels(info|warn|error).

Note: If special characters are used in the SNMP URI, these should be URL encoded.

SNMP URI usage:

- snmpv2: **snmp://community@host[:port]**
- snmpv3: **snmp://user:auth:password@host[:port]** ... minimum auth/password length is 8 characters.

For example:

- snmpv2: **notify add info snmp://public@1.2.3.4**
- snmpv3: **notify add error snmp://public:publicauth:password@1.2.3.4**

The following options can be configured under the SNMP menu in the CLI.

- Enabled - Enable or disable SNMP Queries
- Community - SNMP v2c Community String used to query this server
- Authorized Username - SNMP v3 Username to query this server
- Password - SNMP v3 Password to query this server
- Query - IP address that is allowed to query this server
- Sysname - Name of this server, as it will appear when queried via SNMP
- Syslocation - Location of this server
- Syscontact - Contact person(s) for this server (email address)
- Load1 - 1 Minute load average alarm value
- Load5 - 5 Minute load average alarm value
- Load15 - 15 Minute load average alarm value

The following options can be configured in the CLI:

- Hostname - Server name to send SNMP traps to.
- Version - Version of SNMP to use for sending trap, version 2c or 3.
- Community - refer to the SNMP-URI command usage.

10.3.2. SNMP Queries

The VOSS-4-UC server permits multiple remote query sources to perform SNMP queries against.

The following commands are available to set SNMP details:

- **snmp contact <system contact>**
- **snmp name <system name>**
- **snmp location <system location>**

SNMP query sources can be added with

snmp query add <uri>

SNMP v2 can be set with:

snmp query add snmp://<community string>@<ip>

SNMP v3 username and password can be set with:

snmp query add snmp://user:auth:password@<ip>:<port>

Where:

- **user:** the username for the SNMPv3 server
- **auth:** the SNMPv3 authKey, with a minimum length of 8 characters (SHA authentication)
- **password:** the SNMPv3 privKey, with a minimum length of 8 characters (AES encryption)
- To see the list of added query sources, run **snmp query list**.
- To remove a query source, run **snmp query del <uri>**.

The screen console output below are example of the use of **add**, **list**, and **del** parameters with SNMPv2:

```
platform@host:~$ snmp query add snmp://private@192.29.21.3
You are about to restart the SNMP service. Do you wish to continue? y
Please update notify to reflect your latest changes.
You have new mail in /var/mail/platform
platform@host:~$

platform@host:~$ snmp list
load1: 4
load15: 1
load5: 2
query:
    snmp://public@192.29.21.2
    snmp://private@192.29.21.3
syscontact: Robert Frame
syslocation: Dublin
sysname: host

platform@host:~$ snmp query del snmp://private@192.29.21.3
You are about to restart the SNMP service. Do you wish to continue? y

Application services:firewall processes stopped.
Application snmp processes stopped.
Application snmp processes started.
Please update notify to reflect your latest changes.
You have new mail in /var/mail/platform
platform@host:~$
```

SNMP CPU load notifications are set using:

snmp load <1min load> <5min load> <15min load>

This results in notifications being sent should the threshold be exceeded. For a server with 2 CPUs, it is recommended that this setting be:

snmp load 8 4 2

This means that notifications are sent if the 2-CPU system load averages over the last 1, 5, and 15 minutes reach these values.

11. SNMP

11.1. Introduction to SNMP and MIB

Simple Network Management Protocol(SNMP) is a UDP-based network protocol used mostly in network management systems to monitor network-attached devices. SNMP is a component of the Internet Protocol Suite as defined by the Internet Engineering Task Force(IETF) and consists of a set of standards for network management, including an application layer protocol, a database schema and a set of data objects.

SNMP exposes management data in the form of variables on the managed systems that describe the system configuration. These variables can be queried using SNMP management applications.

SNMP allows a Network Management Station to do the following:

- Poll a device for info or to trend data i.e. VOSS-4-UC server load graph via HOST-SYSTEMS-MIB
- Receive notifications in the form of traps or informs in response to events, threshold violations, whatever the trap definitions in the loaded MIBs are. We enable process monitoring and disk space checks - when triggered, these send out a trap.

A management information base (MIB) is a form of virtual database used for managing the entities in a communications network. Working closely with SNMP, the hierarchical data structure describes all of the objects that a device can report the status of.

The MIB is structured based on the RFC 1155 standard. This standard defines how the MIB information is organized, what data types are allowed and how resources within the MIB are named. Each MIB contains the name, object identifier (a numeral), data type and the permissions relating to whether the value can be read or written to. The top hierarchies of the MIB are fixed; however, certain sub trees can be defined by product vendors and other organizations.

The variables within MIB are named using the Abstract Syntax Notation 1 (ASN.1). This is an international standard for representing data.

SNMP Terminology:

- MIB: The term MIB is used to refer to the complete collection of management information available on an entity, while MIB subsets are referred to as MIB-modules.
- NMS: A Network Management System is a combination of hardware and software used to monitor and administer a network and the devices associated with that network.

SNMP on VOSS-4-UC Platform is configured after initial system setup. The following SNMP parameters can be configured. Refer to [SNMP Configuration and Queries](#) and the index for commands.

- The SNMP system name
Identifies the system being monitored on the NMS (Network Management System). Defaults to nodename.domainname.

- The SNMP system location
Describes the location of the system. Defaults to Unknown.
- The SNMP system contact
Defines the email address of administrator responsible for the system. Defaults to None.
- The SNMP query source
URI from which the system accepts SNMP queries. Formatted as `snmp://community@host[:port]` for version 2 and `snmp://user:auth:password@host[:port]` for version 3
CIDR-style IP (e.g. 196.0.0.0/8) network allowed to query SNMP from this host. This is used to limit the hosts allowed to manage the system via SNMP. Defaults to all hosts.
- The SNMP load triggers
The 1, 5 and 15 minute load averages that will trigger warnings via SNMP. Defaults to values dynamically calculated from the number of CPUs in the system. This should be formatted as `8n/4n/2n` (where `n` represents the number of processors available) when entered into the configuration wizard during setup.
- The SNMP trap destination
Formatting identical to query source.

11.2. SNMP Traps

When the managed system generates certain events, it will forward a SNMP trap. The reason for the event trap is contained in the SNMP MIB string.

For a list of notification messages, see: [Notification Messages](#).

Note that if the corresponding SNMP MIB is not loaded on the NMS, a numerical representation of the SNMP entry is provided.

The list of monitored events is described in the SNMP Trap section below. A detailed breakdown of each SNMP trap type is provided in the appendix.

The SNMP will send traps to the trap destination configured. If the trap destination is incorrect or not configured, the NMS will not receive the traps.

SNMP configuration settings can be managed from the CLI. Refer to the CLI **notify** command:

```
platform@development:~$ notify
USAGE:
-----
notify add [info|warn|error]          - Add the email or snmp URI to a
    <email/snmp-uri> ...              specified notification level.
```

The following system parameters are monitored by default

- Disk Space: warnings are issued if the file system breaches the following thresholds:

```
disk / 30% free
disk /opt/platform 30% free
disk /var/log 10% free
```

- System Load Monitoring: warnings are issued if the system load is excessive (the system load parameters can be defined during configuration)
- SNMP: standard SNMP System Events, for example, Cold Start
- Process state changes: Informative messages are sent to the NMS indicating that processes have been restarted.

In general, the originator of the SNMP traps is determined by originating hostname / IP address. Many Network Management Systems provide trap management and escalation per system being managed, including identification based on system name, location and contact details.

Those events monitored directly by VOSS-4-UC (e.g. disk space, system load and process warnings) include the system name as part of the variable bindings to assist identification of the originating system.

The state of the VOSS-4-UC system can be monitored either on the NMS or via the command line interface using the **diag** command.

It is recommended to monitor SNMP traps that can have a *critical* severity level. Refer to the examples in this document.

11.3. Management Information Bases

Important:

- The VOSS-4-UC system uses standard MIBs that are usually deployed as part of a Network Management System (NMS).

No VOSS-4-UC specific MIBs are available.

The standard MIBs can for example be inspected from on-line resources, such as <http://www.mibdepot.com>.

- String values in the trap descriptions and examples shown here is illustrative purposes only.
- Multi-line display in the trap examples shown here is done for formatting purposes only.

SNMP information is grouped together in Management Information Bases (MIBs). The MIBs loaded on the VOSS-4-UC system represent all the configuration/data items that can be queried or be used to generate traps (notifications) when certain events occur. A list of all MIBs loaded on the system is provided below.

In order to manage the system, a Network Management System (NMS) should be installed at the customer site (e.g. HP OpenView, iReasoningMib Browser). The NMS should be loaded with the same set of MIBs as those installed on the system. The NMS should be configured to send SNMP queries to the managed host (i.e. correct IP address, port number (default 161), community string (default public), and version (default version 2c).

Further, the NMS should be configured to receive traps from the managed host - the correct IP, port number (default 162), version (default version 2), and community strings (default public) should be provided).

SNMP items can be selected in the MIBs and the item queried on the remote managed system. The remote system will return a response to the MIB entry being queried.

For example, if the following entry is queried

```
.1.3.6.1.2.1.1.5.0 alias '.iso.org.dod.internet.mgmt.mib-2.system.sysName.0'
```

the system will return the system name that was assigned during setup (e.g. sysName.0 'Voss Node00').

Note that if any of the configured details on the NMS are incorrect, it is likely that the query will never reach the managed host and no response will be received.

Please ensure that version 2 is selected with the correct community string (default public).

When the managed system generates certain events, it will forward a SNMP trap. The reason for the event trap is contained in the SNMP MIB string.

Note that if the corresponding SNMP MIB is not loaded on the NMS, a numerical representation of the SNMP entry is provided.

The list of monitored events is described in the SNMP Trap section below.

Refer to the MIB List at the end of this document for the list of net-SNMP packages that ship with VOSS-4-UC.

11.4. MIB and Trap Details

11.4.1. SNMPv2-MIB - RFC 3418 - Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)

Basic information about SNMP on the entity. Includes:

- sysDescr: A text description of the entity
- sysObjectID: The vendor's authoritative identification of the network management subsystem contained in the entity.

Note: sysUpTime indicates how long the SNMP software has been running on the box, and not how long the box itself has been up (this is a common misconception).

- sysUpTime: The time since the network management portion of the system was last re-initialised.
- Counters for SNMP requests and responses

11.4.2. IF-MIB - RFC 2863 - The Interfaces Group MIB

Describes the network interfaces on the entity. For each interface the following information is given:

- ifType: The type of interface
- ifMtu: Size of the largest packet which can be sent/received on the interface
- ifSpeed: An estimate of the interface's current bandwidth
- ifPhysAddress: The interface's address at its protocol sub-layer. For 802.x interfaces, this is the MAC address
- The administrative and operational state of the interface
- The number of octets and packets sent and received on the interface

11.4.3. MIB-II - RFC 1213 - Management Information Base for Network Management of TCP/IP- based internets

TCP/IP network information not covered by the other MIBs, split into a number of groups:

- Address translation group:
- atPhysAddress: The media-dependent physical address
- atNetAddress: The network address (IP address) corresponding to the physical address
- IP group:
- ipRouteTable: IP routing table, contains an entry for each route presently known to this entity

11.4.4. IP-MIB - RFC 4293 - Management Information Base for the Internet Protocol (IP)

Internet Protocol information:

- Counters for IP packets sent and received
- For each IP address:
- The IP address
- Index of the physical interface (in the IF-MIB)
- Netmask
- ICMP counters

11.4.5. TCP-MIB - RFC 4022 - Management Information Base for the Transmission Control Protocol (TCP)

TCP information:

- Retransmission timeout information
- Overall counters for number of inbound and outbound connections
- For each current connection:
- Connection state
- Local and remote IP addresses and TCP port numbers

11.4.6. UDP-MIB - RFC 4113 - Management Information Base for the User Datagram Protocol (UDP)

UDP information:

- Counters for datagrams sent and received
- Local IP addresses and UDP port numbers

11.4.7. HOST-RESOURCES-MIB - RFC 2790 - Management Information Base for Host Resources

Objects useful for the management of host computers. These are split into a number of groups:

- System Group
 - hrSystemUptime: Amount of time since the host was last initialised (note this is different from sysUpTime).
 - hrSystemDate: The host's notion of the local date and time of day
 - hrSystemProcesses: The number of process contexts currently loaded or running on this system
- Storage Group
 - hrMemorySize: The amount of physical read-write main memory, typically RAM, contained by the host
- For each storage device:
 - hrStorageType: The type of storage (RAM, fixed disk etc.)
 - hrStorageDescr: A description of the storage (Swap Space, mount point etc.)
- Size of storage units, number available and number used
- Device Group
 - For each device:
 - * Type (processor, network, disk, printer etc.)
 - * Description
- For each disk storage device:
 - Access (read-write, read-only)
 - Fixed/removable
 - Capacity
- For each disk partition:
 - Label
- For each file system:
 - Mount point
 - Type
 - Access (read-write, read-only)
 - Bootable
- Running Software Group
 - For each running process:
 - * Name
 - * Path
 - * Parameters
 - * Status

- * Running Software Performance Group for each running process:
- * CPU resources consumed by this process
- * Amount of real system memory allocated to this process

11.4.8. SNMP Traps: System Startup

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID can be used to identify the cause of the SNMP trap
.iso.org.dod.internet.snmpV2.snmpModules.snmpMIB.snmpMIBObjects.snmpTraps.coldStart

Trap OID

.iso.org.dod.internet.snmpV2.snmpModules.snmpMIB.snmpMIBObjects.snmpTraps.coldStart

Variable Bindings

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 190 milliseconds (19)
- snmpTrapOID = coldStart
- .iso.org.dod.internet.snmpV2.snmpModules.snmpMIB.snmpMIBObjects.snmpTrap.snmpTrapEnterprise.0 = linux

Severity: Info

Example: coldStart

```
Mar 28 10:54:57 robot-sl snmptrapd[1214]:
2019-03-28 10:54:57 <UNKNOWN>
[UDP: [192.168.100.3]:50638->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (98820) 0:16:28.20
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.6.3.1.1.5.1
#011iso.3.6.1.6.3.1.1.4.3.0 = OID: iso.3.6.1.4.1.8072.3.2.10
```

11.4.9. SNMP Traps: Service Startup Changes Made

The following traps are generated at service startup and indicate the various services changing state:

SNMP 1.3.6.1.2.1.88.2.0.1

```
2014-07-04 15:40:30 <server_IP> [UDP: [<server_IP>]:56005->[<snmp_server_IP>]]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (8785393) 1 day, 0:24:13.93
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ProcessRestart"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: <resource>
iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "<hostname> "
```

The <resource> is any of the services listed when running **app status** for any node.

Severity levels, messages and resolution:

- Info : ProcessRestart, ProcessStart

Resolution: If this is an unexpected event, call Support should be called for further investigation. This trap can also be triggered as expected, when **app start** or **system reboot** is run.

- Urgent : ProcessWarning

Resolution: This trap should be seen when a process or service is being restarted or stopped. If this is an unexpected event, call Support should be called for further investigation.

- Critical : ProcessStop, ProcessError

Resolution: If this is an unexpected event, call Support should be called for further investigation. This trap can also be triggered as expected, when **app stop** or **system reboot** is run.

Example: ProcessRestart

- Severity: Info
- Message: ProcessRestart

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 26069):
Var-binds:
1.3.6.1.2.1.1.3.0 = 6619
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ProcessRestart
1.3.6.1.2.1.88.2.1.3.0 = Applications are restarting
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Example: ProcessWarning

- Severity: Urgent
- Message: ProcessWarning

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 50336):
Var-binds:
1.3.6.1.2.1.1.3.0 = 16212
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ProcessWarning
1.3.6.1.2.1.88.2.1.3.0 = mongodb:database has changed its state from \
                        stopped (not reachable) to running (not reachable)
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Example: ProcessStop

- Severity: Critical
- Message: ProcessStop

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 43961):
Var-binds:
1.3.6.1.2.1.1.3.0 = 6286121
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ProcessStop
1.3.6.1.2.1.88.2.1.3.0 = Applications are stopping nginx
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Example: ProcessStart

- Severity: Info
- Message: ProcessStart

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 43961):
Var-binds:
1.3.6.1.2.1.1.3.0 = 6286121
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ProcessStop
1.3.6.1.2.1.88.2.1.3.0 = Applications are started nginx
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

11.4.10. SNMP Traps: Service Monitoring - Changes Made

For each of the services listed above, the system will monitor the process and restart as necessary.

When the service shuts down, it sends a trap indicating a resource stopped in the following format:

```
2014-07-04 15:40:30 <server_IP>
[UDP: [<server_IP>]:56005->[<snmp_server_IP>]]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (8785393) 1 day, 0:24:13.93
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ProcessStop"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: <resource>
iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "<hostname> "
```

Service restart is indicated by the following:

```
2014-07-04 15:40:30 <server_IP>
[UDP: [<server_IP>]:56005->[<snmp_server_IP>]]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (8785393) 1 day, 0:24:13.93
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ProcessRestart"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: <resource>
iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "<hostname> "
```

11.4.11. SNMP Traps: Service Status

- Info: INFO: Services started successfully

- Critical: ERROR: Service Failures

Example: Info

```
Mar 19 15:21:45 robot-sl snmptrapd[1214]:
2019-03-19 15:21:45 <UNKNOWN>
[UDP: [192.168.100.3]:5245->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (578148) 1:36:21.48
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "INFO: Services started successfully"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "0"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Example: Critical

Note: Multi-line display is for example formatting purposes only.

```
Mar 19 15:13:46 robot-sl snmptrapd
[1214]: 2019-03-19 15:13:46 <UNKNOWN>
[UDP: [192.168.100.3]:38997->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (530243) 1:28:22.43
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Service Failures"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "
    selfservice v11.5.2 (2017-08-30 07:40)
    |-node                running
    voss-deviceapi v11.5.2 (2017-08-30 07:40)
    |-voss-cnf_collector  stopped
    |-voss-wsgi           stopped
    |-voss-queue
    Message is truncated"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

11.4.12. SNMP Traps: System Shutdown

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator.

Severity: Critical

Example: nsNotifyShutdown

```
Mar 28 10:54:46 robot-sl snmptrapd[1214]:
2019-03-28 10:54:46 <UNKNOWN>
[UDP: [192.168.100.3]:31384->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (97806) 0:16:18.06
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.4.1.8072.4.0.2
#011iso.3.6.1.6.3.1.1.4.3.0 = OID: iso.3.6.1.4.1.8072.4
```

11.4.13. SNMP Trap: Disk Status

For ERROR: Disk full:

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification.

Severity Messages

- Info : DISK STATUS: Disk <disk> is now running below 80 percent
- Critical : ERROR: Disk full, DISK ALMOST FULL: Disk <disk> is more than 80 percent full

Example: ERROR: Disk full

```
Mar 19 09:09:34 robot-sl snmptrapd[1234]:
2019-03-19 09:09:34 <UNKNOWN>
[UDP: [192.168.100.3]:52717->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (7163878) 19:53:58.78
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Disk full"
#011iso.3.6.1.2.1.88.2.1.2.0 = ""
#011iso.3.6.1.2.1.88.2.1.3.0 = ""
#011iso.3.6.1.2.1.88.2.1.4.0 = OID: iso.3.6.1.4.1.2021.9.1.100.1
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
#011iso.3.6.1.4.1.2021.9.1.2.1 = STRING: "/"
#011iso.3.6.1.4.1.2021.9.1.101.1 = STRING: "/: less than 30% free (= 23%)"
```

Resolution:

This trap depends on which disk is full:

- If it is the media or backup disks, then clean up the disk space.
- If it is any other disk then contact Support immediately.

Example: DISK STATUS

Note: Multi-line display is for example formatting purposes only.

```
Mar 19 08:12:14 robot-sl snmptrapd[1234]:
2019-03-19 08:12:14 <UNKNOWN>
[UDP: [192.168.100.3]:18540->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (6819861) 18:56:38.61
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "DISK STATUS: Disk /backup is now running_
↳below 80 percent"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Disk /backup cleared Disk status after it was_
↳cleared:
  Filesystem:
    /dev/sdc1      Size: 50G      Used: 857M      Avail: 46G      Use%: 2%
  Mounted on:
    /backups"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Resolution:

None needed.

Example: ERROR: DISK ALMOST FULL

Note: Multi-line display is for example formatting purposes only.

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 62210):
Var-binds:
1.3.6.1.2.1.1.3.0 = 25163513
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = DISK ALMOST FULL: Disk /opt/platform is more than 80 percent_
↳full
1.3.6.1.2.1.88.2.1.3.0 = Filesystem: /dev/sdb2
  Contact support to free space.
  Current disk status:
    Size: 40G
  Contact support to free space.
  Current disk status:
    Used: 38G
  Contact support to free space.
  Current disk s
    Message is truncated
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

Resolution:

This trap depends on which disk is full:

- If it is the media or backup disks, then clean up the disk space.
- If it is any other disk then contact Support immediately.

11.4.14. SNMP Trap: Database Usage

A trap is generated when the transaction and cache collections exceed usage thresholds.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the VOSS-4-UC system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Severity Messages:

- Info : WARN: Database transaction count returned to normal,
INFO: Database transaction size returned to normal,
INFO: Database cache size returned to normal
- Critical : WARN: Database cache size exceeded threshold,
WARN: Database transaction size exceeded threshold,
WARN: Database transaction count exceeded threshold

Variable Bindings - DB cache size

WARN

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARN: Database cache size exceeded threshold'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'size: 14960011111'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

INFO

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'INFO: Database cache size returned to normal'

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'The cache size returned to normal'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 0
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Variable Bindings - DB transaction size

WARN

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARN: Database transaction size exceeded threshold'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'size: 23353681111'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

INFO:

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'INFO: Database transaction size returned to normal'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'The transaction size returned to normal'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 0
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Variable Bindings - DB transaction count

WARN

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARN: Database transaction count exceeded threshold'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'count: 500001'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

INFO

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARN: Database transaction count returned to normal'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 = 'The transaction count returned to normal'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 0
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Example: WARN: Database transaction size exceeded threshold

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 11776):
Var-binds:
1.3.6.1.2.1.1.3.0 = 27396043
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = WARN: Database transaction size exceeded threshold
1.3.6.1.2.1.88.2.1.3.0 = size: 22574836480
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

Resolution:

The following commands can be used to resolve this:

- **voss transaction delete <days> [limit <number>]**
- **voss transaction archive <days>**

11.4.15. SNMP Trap: Database Maintenance

A trap is generated when the database maintenance schedules are not set up.

Severity Messages:

- ERROR: Database maintenance not scheduled

Example hourly notifications generated by the data/Alert instance in the Database

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('10.120.1.203', 60578):
Var-binds:
1.3.6.1.2.1.1.3.0 = 16128700
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = Database Maintenance
1.3.6.1.2.1.88.2.1.3.0 = ID: TRANSACTION_DATABASE_MAINTENANCE-VOSS-UN-1, Code: -1,
↪Occurrences: 27, Latest Occurrence: 2019-08-23T10:06:35.405Z
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = VOSS-UN-1
```

Example warnings raised once through platform monitoring

TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('10.120.1.203', 5194):
Var-binds:
1.3.6.1.2.1.1.3.0 = 16128774
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: Database maintenance not scheduled
1.3.6.1.2.1.88.2.1.3.0 = TRANSACTION DATABASE MAINTENANCE NOT SCHEDULED - SETUP
↳SCHEDULE FOR REGULAR MAINTENANCE
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = VOSS-UN-1
```

Resolution

Create schedule and time:

With for example a schedule name dbtxn:

schedule add dbtxn voss transaction archive or **schedule add dbtxn voss transaction delete**
schedule time dbtxn weekly 0

See: [Enable Database Scheduling](#).

11.4.16. SNMP Trap: Excessive Load

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

- The following variable binding can be used to determine that the load average threshold has been exceeded.

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix.
↳dismanEventMIBNotificationObjects.mteHotTrigger.0 = ERROR: Excessive load.
```

- The following variable binding can be used to further diagnose which time interval threshold has been exceeded

```
.iso.org.dod.internet.private.enterprises.ucdavis.laTable.laEntry.laNames.
↳<LoadIdx> = <LoadError>
.iso.org.dod.internet.private.enterprises.ucdavis.laTable.laEntry.laErrMessage.
↳<LoadIdx> = <LoadMessage>
```

Load average interval	<LoadIdx>	<LoadError>	<LoadMessage>
1 minute	1	Load-1	1 min Load Average too high (= 2.52)
5 minute	2	Load-5	5 min Load Average too high (= 1.27)
15 minute	3	Load-15	15 min Load Average too high (= 1.27)

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = ERROR: Excessive load.
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTargetName.0 =
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotContextName.0 =
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotOID.0 = laErrorFlag.1
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
- .iso.org.dod.internet.private.enterprises.ucdavis.laTable.laEntry.laNames.1 = Load-1
- .iso.org.dod.internet.private.enterprises.ucdavis.laTable.laEntry.laErrorMessage.1 = 1 min Load Average too high (= 1.36)

Severity:

- Critical:
 - ERROR: Excessive load

Example: Critical

```
Mar 19 08:08:34 robot-sl snmptrapd[1234]:
2019-03-19 08:08:34 <UNKNOWN>
[UDP: [192.168.100.3]:20997->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (6797884) 18:52:58.84
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Excessive load"
#011iso.3.6.1.2.1.88.2.1.2.0 = ""
#011iso.3.6.1.2.1.88.2.1.3.0 = ""
#011iso.3.6.1.2.1.88.2.1.4.0 = OID: iso.3.6.1.4.1.2021.10.1.100.1
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
#011iso.3.6.1.4.1.2021.10.1.2.1 = STRING: "Load-1"
```

(continues on next page)

(continued from previous page)

```
#011iso.3.6.1.4.1.2021.10.1.101.1 = STRING: "1 min Load Average too high (= 3.45)"

Mar 19 08:10:34 robot-sl snmptrapd[1234]:
2019-03-19 08:10:34 <UNKNOWN>
[UDP: [192.168.100.3]:49080->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (6809885) 18:54:58.85
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Excessive load"
#011iso.3.6.1.2.1.88.2.1.2.0 = ""
#011iso.3.6.1.2.1.88.2.1.3.0 = ""
#011iso.3.6.1.2.1.88.2.1.4.0 = OID: iso.3.6.1.4.1.2021.10.1.100.2
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
#011iso.3.6.1.4.1.2021.10.1.2.2 = STRING: "Load-5"
#011iso.3.6.1.4.1.2021.10.1.101.2 = STRING: "5 min Load Average too high (= 2.24)"

Mar 19 08:11:34 robot-sl snmptrapd[1234]:
2019-03-19 08:11:34 <UNKNOWN>
[UDP: [192.168.100.3]:47676->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (6815886) 18:55:58.86
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Excessive load"
#011iso.3.6.1.2.1.88.2.1.2.0 = ""
#011iso.3.6.1.2.1.88.2.1.3.0 = ""
#011iso.3.6.1.2.1.88.2.1.4.0 = OID: iso.3.6.1.4.1.2021.10.1.100.3
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
#011iso.3.6.1.4.1.2021.10.1.2.3 = STRING: "Load-15"
#011iso.3.6.1.4.1.2021.10.1.101.3 = STRING: "15 min Load Average too high (= 1.16)"
```

11.4.17. SNMP Trap: Backup and Restore

A trap is generated on every backup.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-
Notifications.mteTriggerFired
```

Variable Bindings - successful backup

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = "backup completed"
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 0
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Variable Bindings - failed backup

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = "backup failed"
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 5
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Messages:

- Info : Backup completed, INFO: Backup now runs regularly
- Critical : Backup failed, ERROR: The last backup was more than 2 days ago,

Example: Critical: Backup failed

Note: Multi-line display is for example formatting purposes only.

```
Notification message from (1, 3, 6, 1, 6, 1, 1): ('192.29.22.122', 29632):
Var-binds:
1.3.6.1.2.1.1.3.0 = 1365111
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = Backup failed
1.3.6.1.2.1.88.2.1.3.0 = Creating/updating backup at
    file:///backups/473bf9ae1027e3f2af6edbcca43d2c64f1ed4b87b
    Using archive dir:
        /root/.cache/duplicity/5e41da3a0b2fd68dbef152b9358f2c40
    Using backup name: 5e41da3a0b2fd68dbef15
    Message is truncated
1.3.6.1.2.1.88.2.1.5.0 = 0
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

Resolution:

There may have been some form of corruption during the time that the backup was run.

Run backup manually and see whether the same failure is detected. If it is the case, investigate. Otherwise, contact L2 Support to investigate.

Example: Critical: ERROR: The last backup was more than 2 days ago

Note: Multi-line display is for example formatting purposes only.

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 44687):
Var-binds:
1.3.6.1.2.1.1.3.0 = 24804734
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: The last backup was more than 2 days ago
1.3.6.1.2.1.88.2.1.3.0 = Backup list: \
    localbackup: \
    URI: file:///backups \
    Backups: \
    No backups created yet
1.3.6.1.2.1.88.2.1.5.0 = 1
```

Resolution:

Investigate if the scheduler is set correctly and that the backup is set up correctly. Run a manual backup to test if it is working as it should. Otherwise, call Support to investigate.

Example: INFO: Backup now runs regularly

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 10865):
Var-binds:
1.3.6.1.2.1.1.3.0 = 25108817
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = INFO: Backup now runs regularly
1.3.6.1.2.1.88.2.1.3.0 = Backup list: \
    localbackup: \
    URI: file:///backups \
    Backups: \
    1 backups have been created - most recently 2019-03-01 11:25
1.3.6.1.2.1.88.2.1.5.0 = 0
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

11.4.18. SNMP Trap: Health Emails

A trap is generated if health email send fail to be generated.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-
Notifications.mteTriggerFired
```

Variable Bindings

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)

- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: Trouble sending health email'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Messages:

- Info : INFO: Health emails is now being sent, INFO: All notify levels is now configured with an external email address
- Minor : ERROR: Trouble sending health email, WARNING: Not all notify levels is configured with an external email address

Example: ERROR: Trouble sending health email

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 44151):
Var-binds:
1.3.6.1.2.1.1.3.0 = 58347940
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: Trouble sending health email
1.3.6.1.2.1.88.2.1.3.0 = Trouble sending health email
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Example: WARNING: Some notify levels are configured with a local email address

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 25067):
Var-binds:
1.3.6.1.2.1.1.3.0 = 7420086
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = WARNING: Some notify levels are configured with a local_
↪email address
1.3.6.1.2.1.88.2.1.3.0 = Notify list:
  notifications:
  emailrelay: 172.29.42.30
  level:
  audit:
  snmp://rrako@192.29.21.225
  error:
  mailto:xlatom@loc Message is truncated
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

Resolution:

Ensure that all notifications are non-local.

11.4.19. SNMP Trap: Disk Latency

A trap is generated when the disk appear to be slow.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings - Disk slow

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: Disk slow'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Message

- Info : INFO: The disk latency returned to normal
- Critical : ERROR: Disk slow

INFO Example

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 38416):
Var-binds:
1.3.6.1.2.1.1.3.0 = 24804469
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = INFO: The disk latency is normal
1.3.6.1.2.1.88.2.1.3.0 = The disk latency is normal.
1.3.6.1.2.1.88.2.1.5.0 = 0
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Resolution:

None. This trap will only appear to state that the heavy disk activity has receded and that the disk is operating as per normal.

Critical Example

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 53997):
Var-binds:
1.3.6.1.2.1.1.3.0 = 52188064
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: Disk slow
1.3.6.1.2.1.88.2.1.3.0 = Disk latency info: True
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Resolution:

This trap indicates that there is heavy disk activity. This will be normal in many situations and should not always require immediate action. The one reason for this is that there are many transactions spawned at the same time. The disk utilization can be relieved by cancelling a few current transactions or by rescheduling some others.

However, this should be monitored closely and should it persist over several hours then Support should be contacted to investigate.

11.4.20. SNMP Trap: Mailbox Status

A trap is generated when the local mailbox reaches 200 plus emails.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotifications.mteTriggerFired

Variable Bindings - Mailbox email messages reach 200

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARNING: The total messages in the local mailbox for %s has reached in excess of 200'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Variable Bindings - Mailbox email messages reach 500

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'INFO: Messages for <server info> auto archived as it reached more than 500'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Messages:

- Info : INFO: Messages for <server> auto archived as it reached more than 500, INFO: The total local messages for <server> is now under 200
- Warning: WARNING: The total messages in the local mailbox for platform has reached in excess of 200

Example: WARNING

```
Mar 22 09:33:46 robot-sl snmptrapd[1214]:
2019-03-22 09:33:46 <UNKNOWN>
[UDP: [192.168.100.3]:62862->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (15555245) 1 day, 19:12:32.45
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "WARNING: The total messages in the local
↳ mailbox for platform has reached in excess of 200"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Use the following mail commands to manage the
↳ local mailbox:
mail del <number> - delete the selected mail
mail del <from#> <to#> - deletes the selected range of mail message
mail del all Message is truncated"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

11.4.21. SNMP Trap: Cluster Status

A trap is generated when one or more nodes are down in a cluster.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotifications.mteTriggerFired

Variable Bindings - One or more nodes are down in the cluster

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)

- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: One or more nodes are down in the cluster'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

11.4.22. SNMP Trap: Database Failover Status

A trap is generated when one or more nodes are down in a cluster.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the VOSS-4-UC system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired Variable Bindings - db constantly fails over

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: The db is failing over constantly within 5 min'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Messages:

- Info : INFO: The db failover status returned to normal
- Critical : ERROR: The db is failing over constantly within 5 min

Example: INFO

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 31127):
Var-binds:
1.3.6.1.2.1.1.3.0 = 34697935
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = INFO: The db failover status returned to normal
1.3.6.1.2.1.88.2.1.3.0 = Cluster failover status: 4th last database failover uncured_
↪3099 seconds ago
```

(continues on next page)

(continued from previous page)

```
1.3.6.1.2.1.88.2.1.5.0 = 0
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

Example: ERROR

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 4884):
Var-binds:
1.3.6.1.2.1.1.3.0 = 34615003
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: The db is failing over constantly within 5 min
1.3.6.1.2.1.88.2.1.3.0 = Cluster failover status: 4th last database failover occurred_
↪2320 seconds ago
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

11.4.23. SNMP Trap: Large Log Files

A trap is generated when large log files are detected in `/var/log/`.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:
 - .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings - large log files detected.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: Log files larger than 1Gig found in /var/log'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity Messages:

- Info : INFO: /var/log rotated
- Urgent : ERROR: Log files larger than 1Gig found in /var/log

Severity: Info Trap Example

Message: INFO: /var/log rotated

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 25035):
Var-binds:
1.3.6.1.2.1.1.3.0 = 24804740
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = INFO: /var/log rotated
1.3.6.1.2.1.88.2.1.3.0 = /var/log rotated
1.3.6.1.2.1.88.2.1.5.0 = 0
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

Severity: Urgent Trap Example

Message: ERROR: Log files larger than 1Gig found in /var/log

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.22.21.124', 51928):
Var-binds:
1.3.6.1.2.1.1.3.0 = 52324087
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: Log files larger than 1Gig found in /var/log
1.3.6.1.2.1.88.2.1.3.0 = Logrotation was executed to rotate the \
following logs: /var/log/some.log: 7.3G
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.22.21.124
```

11.4.24. SNMP Trap: Network Status

A trap is generated when a network failures occur.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-
Notifications.mteTriggerFired
```

Variable Bindings - Network failures

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEvent-
MIBNotificationObjects.mteHotTrigger.0 = 'ERROR: Network Failures'

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity:

- Critical: ERROR: Network Failures
- Info: INFO: Network failures resolved

Example: Critical

```
Mar 22 14:07:58 robot-sl snmptrapd[1214]:
2019-03-22 14:07:58 <UNKNOWN>
[UDP: [192.168.100.3]:18751->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (1155411) 3:12:34.11
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Network Failures"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "The following network failures occured:
↪netntp: 172.29.1.15"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Example: Info

```
Mar 29 13:57:11 robot-sl snmptrapd
[1234]: 2019-03-29 13:57:11 <UNKNOWN>
[UDP: [192.168.100.3]:32794->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (165816) 0:27:38.16
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "INFO: Network failures resolved"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Network failures resolved"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

11.4.25. SNMP Trap: NGINX Status**Severity:**

- Critical: nginx upstream failure, upstream <node> server <server> failed
- Info: nginx upstreams OK

Example: Critical

```
Mar 25 13:30:12 robot-sl snmptrapd[1214]:
2019-03-25 13:30:12 <UNKNOWN>
[UDP: [172.29.21.129]:63573->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (293333) 0:48:53.33
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "nginx upstream failure"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "upstream selfservice server 192.29.22.122:443
↪failed: <urlopen error timed out>"
```

(continues on next page)

(continued from previous page)

```
#011iso.3.6.1.2.1.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "WP2-172.29.21.129"
```

Example: Info

```
Mar 25 13:34:02 robot-sl snmptrapd[1214]:
2019-03-25 13:34:02 <UNKNOWN>
[UDP: [172.29.21.129]:31265->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (316311) 0:52:43.11
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "nginx upstreams OK"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "nginx upstream servers returned to normal"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "WP2-172.29.21.129"
```

11.4.26. SNMP Trap: Web Certificate Expiry

A trap is generated when the nginx web certificate is about to expire.

Three intervals are checked which will return a result:

- 30 days prior to expiry
- 14 days prior to expiry
- 1 day prior to expiry

After expiry, an alert is sent every day.

Severity Messages:

- ERROR: Web certificate about to expire
- ERROR: WEB CERT EXPIRED AT:

SNMP Trap examples:

```
2021-07-09 02:56:31 <server_IP>
[UDP: [<server_IP>]:58567->[<snmp_server_IP>]:162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (795178) 2:12:31.78
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "Certificate_Maintenance"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: "ID: WEB_CERTIFICATE_VOSS,
Code: -1,
Occurences: 1,
Latest Occurrence: 2021-07-09T00:56:30.610Z"
iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"

2021-07-09 02:56:32 <server_IP>
[UDP: [<server_IP>]:45737->[<snmp_server_IP>]:162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (795286) 2:12:32.86
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: Web certificate about to expire"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Cert will expire in less than 30 day(s)"
```

(continues on next page)

(continued from previous page)

```

iso.3.6.1.2.1.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"

2021-07-09 03:00:40 <server_IP>
[UDP: [<server_IP>]:37982->[<snmp_server_IP>]:162]:
iso.3.6.1.2.1.1.3.0 = Timeticks: (820135) 2:16:41.35
iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
iso.3.6.1.2.1.88.2.1.1.0 = STRING: "INFO: Web certificate has been renewed"
iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Web certificate has been renewed"
iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"

```

Example GUI alert shown in the data/Alert instance in the Database

- AlertCode: "25005"
- Category: "Certificate_Maintenance"
- AlertId: "WEB_CERTIFICATE_<certificate-name>"
- Message: "The certificate will expire in less than {} day(s)." (number of days - alert will result)
- Severity: WARNING
- alert_timestamp: <timestamp>

Example error raised through platform monitoring

```

Notification message from (1, 3, 6, 1, 6, 1, 1):('10.120.1.203', 5194):
Var-binds:
1.3.6.1.2.1.1.3.0 = 16128774
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = ERROR: Web certificate about to expire
1.3.6.1.2.1.88.2.1.3.0 = Cert will expire in less than {} day(s)
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = VOSS-UN-1

```

Resolution

Renew the web certificate - see: [Web Certificate Setup Options](#).

11.4.27. SNMP Trap: Security Updates

A trap is generated when security updates are available.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system

- The SNMP system name is included as part of the variable binding to assist identification:
- `.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone`

Trap OID

`.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired`

Variable Bindings - Security updates available.

- `.iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)`
- `snmpTrapOID = mteTriggerFired`
- `.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARNING: Security Updates available'`
- `.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1`
- `.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone`

11.4.28. SNMP Trap: Memory Usage

A trap is generated for high memory usage.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

`.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired`

Variable Bindings - High memory usage.

- `.iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)`
- `snmpTrapOID = mteTriggerFired`
- `.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: High memory usage'`
- `.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1`
- `.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone`

Variable Bindings - Extremely high CPU usage.

- `.iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)`
- `snmpTrapOID = mteTriggerFired`

- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: Extremely high CPU usage'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity:

- Critical: ERROR: Memory swap error
- Info: INFO: Memory usage returned to normal

Example: Critical

```
Mar 28 22:32:24 robot-sl snmptrapd[1214]:
2019-03-28 22:32:24 <UNKNOWN>
[UDP: [192.168.100.3]:25747->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (997093) 2:46:10.93
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "INFO: Memory usage returned to normal"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Memory more than 1024MB"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Example: Info

```
Mar 28 22:32:24 robot-sl snmptrapd[1214]:
2019-03-28 22:32:24 <UNKNOWN>
[UDP: [192.168.100.3]:25747->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (997093) 2:46:10.93
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "INFO: Memory usage returned to normal"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Memory more than 1024MB"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 0
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

11.4.29. SNMP Trap: NTP Status

A trap is generated if NTP is not configured.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings - NTP not configured.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'ERROR: No ntp configured for <server info>'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity:

- Critical: WARNING: The ntp daemon has stopped on <server>
- Urgent:
 - ERROR: No ntp configured for <server>
 - WARNING: The ntp offset exceeds 1 second on <server>

Example: Critical

```
Mar 19 15:25:21 robot-sl snmptrapd[1214]:
2019-03-19 15:25:21 <UNKNOWN>
[UDP: [192.168.100.3]:7564->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (599782) 1:39:57.82
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "WARNING: The ntp daemon has stopped on VOSS-
↪192.168.100.3"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "Run
'app start services:time' to restart ntpd"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Example: Urgent

```
Mar 20 12:21:24 robot-sl snmptrapd[1214]:
2019-03-20 12:21:24 <UNKNOWN>
[UDP: [192.168.100.3]:11256->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (421144) 1:10:11.44
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "ERROR: No ntp configured for VOSS-192.168.100.
↪3"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "It is mandatory that the ntp is configured.
To configure ntp use the following command: network ntp <server1> <server2>"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

Example: Urgent

```
Notification message from (1, 3, 6, 1, 6, 1, 1):('192.29.22.122', 51983):
Var-binds:
1.3.6.1.2.1.1.3.0 = 242825
1.3.6.1.6.3.1.1.4.1.0 = 1.3.6.1.2.1.88.2.0.1
1.3.6.1.2.1.88.2.1.1.0 = WARNING: The ntp offset exceeds 1 second on UN1-192.29.22.
↪122-192.29.22.122
1.3.6.1.2.1.88.2.1.3.0 = ntp offset exceeds 1 second Current ntp offset value: 2
1.3.6.1.2.1.88.2.1.5.0 = 1
1.3.6.1.2.1.1.5.0 = UN1-192.29.22.122
```

11.4.30. SNMP Trap: DNS status

A trap is generated if DNS is not configured.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-
Notifications.mteTriggerFired
```

Variable Bindings - DNS not configured.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEvent-
MIBNotificationObjects.mteHotTrigger.0 = 'WARNING: No dns configured for <server info>'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEvent-
MIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity:

- Urgent: WARNING: No dns configured for <server>

Resolution:

Configure a DNS for the system with **network dns add <server>**

Example: Urgent

```

Mar 20 12:13:50 robot-sl snmptrapd[1214]:
2019-03-20 12:13:50 <UNKNOWN>
[UDP: [192.168.100.3]:37298->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (375782) 1:02:37.82
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "WARNING: No dns configured for VOSS-192.168.
↪100.3"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "It is recommended that the dns is configured.
  To configure dns use the following command: network dns <server1> <server2>"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"

```

11.4.31. SNMP Trap: Domain Status

A trap is generated if the domain is not configured.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:

```
.iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone
```

Trap OID

```
.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-
Notifications.mteTriggerFired
```

Variable Bindings - Domain not configured.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEvent-
MIBNotificationObjects.mteHotTrigger.0 = 'WARNING: No domain configured for <server info>'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEvent-
MIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Severity:

- Urgent: WARNING: No domain configured for <server>

Resolution:

Configure domain with **network domain <name>**

Example: Urgent

```

Mar 19 08:12:14 robot-sl snmptrapd[1234]:
2019-03-19 08:12:14 <UNKNOWN>
[UDP: [192.168.100.3]:58537->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (6819891) 18:56:38.91
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING: "WARNING: No domain configured for VOSS-192.168.
↪100.3"
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING: "It is recommended that the domain is_
↪configured.
  To configure the domain use the following command: network domain <server1> <server2>
↪"
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"

```

11.4.32. SNMP Trap: NTP Offset

A trap is generated when the NTP offset exceeds 1 second.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:
 - .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings - NTP exceeds 1 second.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = 'WARNING: The ntp offset exceeds 1 second on <server info>'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

11.4.33. SNMP Trap: Process Memory Threshold Status

A trap is generated when the a process memory exceeds its current threshold.

Identification

- The originating IP / hostname is used to identify the system generating the traps
- The NMS is responsible for associating traps with each managed system, along with clearing of alarms and escalation to the relevant system operator
- The trap OID is generic for various SNMP events monitored by the system
- The SNMP system name is included as part of the variable binding to assist identification:
 - .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

Trap OID

.iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIB-Notifications.mteTriggerFired

Variable Bindings - Process exceeds memory threshold.

- .iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 = 2 minutes (12065)
- snmpTrapOID = mteTriggerFired
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotTrigger.0 = '<process name: mem_<name> exceeded maximum value of current_threshold with current_reading>'
- .iso.org.dod.internet.mgmt.mib-2.dismanEventMIB.dismanEventMIBNotificationPrefix. dismanEventMIBNotificationObjects.mteHotValue.0 = 1
- .iso.org.dod.internet.mgmt.mib-2.system.sysName.0 = standalone

11.5. VOSS-4-UC System Monitoring Traps

11.5.1. SNMP Traps: System Monitoring

Administrators at `sysadmin` level can configure additional SNMP traps for alerts from the **System Monitoring > Configuration** menu on the GUI (menu model: `data/SystemMonitoringConfig`). Note: some traps are not configurable.

Refer to the topic on System Monitoring Configuration in the Advanced Configuration Guide.

The following alerts are configured:

Notification	Interval	Level	Configurable
Txn Queue Size	Hourly	warn	Yes
Failed Txn	Immediate	warn	Yes
data/Alert (CNF atm)	Immediate	Alert defined	No
Session Exceeded	Immediate	warn	Yes (via platform CLI command)
API Request Throttled	Immediate	warn	Yes (via platform CLI command)
Total DB Index Size	Daily	warn	Yes
Total DB Size	Daily	warn	Yes
Device Comms. Concurrency Limit	Immediate	warn	No

For platform CLI commands for session limits and throttle rates, see: [Performance Commands](#).

Transaction Queue Size

In accordance with the configurable threshold (default 500)

Identifying strings and example context:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Transaction Queue Size Exceeded Threshold
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING: Current Size: 520 Threshold: 500
```

Transactions: Model Operations Alerts

- Alerts on transactions failure
 - per model (wild cards allowed, default is data/*)
 - model operations (default is **Import**)

Identifying string:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Transaction Completed with Fail
```

Transaction trap context information (200 chars):

- ID: transaction ID (same as on GUI - further transaction details available on GUI)
- Action: transaction message (same as on GUI)
- Detail: source of resource (source host for import)
- Hierarchy: friendly path of the resource, else the execution hierarchy of transaction

Example: Import Fail

```
2019-03-28 10:54:46 <UNKNOWN>
[UDP: [192.168.100.3]:31384->[192.168.100.25]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (170158257) 19 days, 16:39:42.57
SNMPv2-MIB::snmpTrapOID.0 = OID: DISMAN-EVENT-MIB::mteTriggerFired
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING:
Transaction Completed with Fail
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING:
ID: 44967,
Action: Import Call Manager,
Detail: 192.168.100.15,
Hierarchy: sys
DISMAN-EVENT-MIB::mteHotValue.0 = INTEGER: 1
SNMPv2-MIB::sysName.0 = STRING: VOSS
```

Change Notification Feature (CNF)

CNF traps are triggered when Change Notification Sync transactions add or update instances on the data/Alerts model.

The identifying alert string is:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Device Change Notification
```

The data/Alerts attribute values of the model are provided in the traps details:

```

alert_severity
alert_category
alert_timestamp
alert_count
alert_id
alert_message
alert_code

```

For example , the trap Context information (200 chars) is:

- ID: Device Host business key (`alert_id`)
- Code: CNF Alert code (`alert_code`)
- Occurrences: number of occurrences
- Latest Occurrence:: time stamp (`alert_timestamp`)

Warning and Error Alert Codes

The following table shows alert codes and details

Code	Details	Resolution
72051	ERROR. Device connectivity failure	For connectivity checks, see <i>UC Apps Reachability</i> in the Advanced Configuration Guide and also the Platform Guide and Health Checks for Cluster Installations Guide.
72052	WARNING. Slow device connection: the roundtrip time (RTT) is greater than 400ms.	For latency checks, see <i>UC Apps Reachability</i> in the Advanced Configuration Guide and the Health Checks for Cluster Installations Guide.
72053	ERROR. Utilization approaching limit: if the maximum number of sessions during the interval exceeds 80% of the configured limit.	The threshold can be reached if many users are using the system at the same time, or by not logging out. For Utilization % , see <i>Login Sessions</i> in the Advanced Configuration Guide.

Example: CNF alert

```

2019-03-28 10:54:46 <UNKNOWN>
[UDP: [192.168.100.3]:31384->[192.168.100.25]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (170158257) 19 days, 16:39:42.57
SNMPv2-MIB::snmpTrapOID.0 = OID: DISMAN-EVENT-MIB::mteTriggerFired
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Device Change Notification
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING:
  ID: 44967,
  Code: 100034,
  Occurrences: 1,
  Latest Occurrence: 2019-03-28 10:54:44Z
DISMAN-EVENT-MIB::mteHotValue.0 = INTEGER: 1
SNMPv2-MIB::sysName.0 = STRING: VOSS

```

Session Limits

SNMP traps are triggered when session limits are reached.

Example:

For example, the customer administrator session limit default is 10 and a trap is triggered if it is exceeded. (The default can be configured with the **voss session-limits** command).

Note: Global session limits do not show a `Hierarchy` value in the message string.

```
2019-03-28 10:54:46 <UNKNOWN>
[UDP: [192.168.100.3]:31384->[192.168.100.25]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (170158257) 19 days, 16:39:42.57
SNMPv2-MIB::snmpTrapOID.0 = OID: DISMAN-EVENT-MIB::mteTriggerFired
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Customer Administration Session Limit_
↔Exceeded
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING:
  Limit: 10,
  Hierarchy: sys.hcs.Varidion.GSCorp
DISMAN-EVENT-MIB::mteHotValue.0 = INTEGER: 1
SNMPv2-MIB::sysName.0 = STRING: VOSS
```

API Request Throttle

SNMP traps are triggered when throttle rates are reached.

Throttle rates are configured with:

```
voss throttle-rates type <administration|selfservice|user> requests <number of requests>
unit <min|sec>
```

In other words, the SNMP trap would be triggered for request limits for any of:

- Administration
- Self-service
- User-specific

Identifying strings and Self-service as example:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: Selfservice Api Request Limit Exceeded
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING: Rate 20/min
```

Total DB Index Size

In accordance with the configurable threshold (default 50)

Identifying strings and example:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: DB Index Size Exceeded Threshold
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING: DB Index Size (60.00GB) exceeded_
↔threshold (50GB)
```

Total DB Size

In accordance with the configurable threshold (default 200)

Identifying strings and example:

```
DISMAN-EVENT-MIB::mteHotTrigger.0 = STRING: DB Size Exceeded Threshold
DISMAN-EVENT-MIB::mteHotContextName.0 = STRING: DB Size (210.30GB) exceeded threshold_
↪ (200GB)
```

Device Communications Concurrency Limit

SNMP traps are sent if there is a timeout failure while connecting to a device and waiting for the concurrency limit.

Current concurrency support:

- 8 concurrent requests to Unified CM
- 8 concurrent requests to Unify Connection
- 1 concurrent requests to HCM-F

11.5.2. Notification Messages

There are 3 levels of notifications:

- info - these are informational messages intended to communicate a positive event.
- warn - these are non-critical notifications indicating that there is a temporary error
- error - these are critical messages indicating a system fault that is affecting system functionality

Category	Level	Message
Maintenance mode	warn	WARN: System going into maintenance mode
Maintenance mode	info	INFO: System taken out of maintenance mode
Disk status	error	DISK ALMOST FULL: Disk {} is more than {} percent full ...
Disk status	warn	DISK STATUS: Disk {} is now running below {} percent. Disk %s cleared Disk status after it was cleared:
Disk status	info	DISK STATUS: Disk {} is now running below {} percent. Disk %s cleared Disk status after it was cleared:
Health email	error	ERROR: Trouble sending health email
Health email	info	INFO: Health emails is now being sent
Database usage	warn	WARN: Database %s exceeded threshold ...
Database usage	info	INFO: Database %s returned to normal ...
Disk latency	error	ERROR: Disk slow ', 'Disk latency info:..
Disk latency	info	INFO: The disk latency returned to normal
Local mail monitoring	warn	WARNING: The total messages in the local mailbox for %s has reached in excess of 200...
Local mail monitoring	info	INFO: Messages for %s auto archived as it reached more than 500...
Local mail monitoring	info	INFO: The total local messages for %s is now under 200...

Continued on next page

Table 1 -- continued from previous page

Category	Level	Message
External mail monitoring	warn	WARNING: Not all notify levels is configured with an external email address ...
External mail monitoring	info	INFO: All notify levels is now configured with an external email address ...
Backup monitoring	error	ERROR: The last backup was more than 2 days ago ...
Backup monitoring	info	INFO: Backups now runs regularly ...
Cluster failover	error	ERROR: The db is failing over constantly within 5 min ...
Cluster failover	info	INFO: The db failover status returned to normal ...
Logging	error	ERROR: Log files larger than 1Gig found in /var/log ...
Logging	info	INFO: /var/log rotated ...
Network Monitoring	error	ERROR: Network Failures. The following network failures occurred: ...
Network Monitoring	info	INFO: Network failures resolved ...
Service Monitoring	error	ERROR: Service Failures ...
Service Monitoring	info	INFO: Services started successfully ...
Security Update	warn	WARNING: Security Updates available ...
Security Update	info	INFO: Security Updates applied ...
Memory Monitoring	error	ERROR: High memory usage Memory activity: ...
Memory Monitoring	info	INFO: Memory usage returned to normal Memory more than 1024MB
Disk Status	error	DISK_USAGE_LIMITS /
Disk Status	error	DISK ALMOST FULL: Disk /var/log is more than {} percent full ...
Disk Status	info	DISK_USAGE_LIMITS /
Disk Status	info	DISK_USAGE_LIMITS /var/log
Disk Status	audit	DISK ALMOST FULL: Disk /var/log is more than {} percent full ...
Backups	error	DISK_USAGE_LIMITS /backups
Backups	info	DISK_USAGE_LIMITS /backups
Platform	error	DISK_USAGE_LIMITS /opt/platform
Platform	info	DISK_USAGE_LIMITS /opt/platform
NTP	error	ERROR: No ntp configured for %s ...
NTP	info	INFO: ntp is now configured for %s ...
DNS	warn	WARNING: No dns configured for %s...
DNS	info	INFO: dns is now configured for %s ...
Domain	warn	WARNING: No domain configured for %s...
Domain	info	INFO: domain is now configured for %s...
NTP Daemon	warn	WARNING: The ntp daemon has stopped on %s ...
NTP Daemon	info	INFO: The ntp daemon is now running on %s ...
NTP Offset	warn	WARNING: The ntp offset exceeds 1 second on %s ...

Continued on next page

Table 1 -- continued from previous page

Category	Level	Message
NTP Offset	info	INFO: The ntp offset restored to normal on %s . . .
License	error	ERROR: License file generation failed The license audit report scheduled for %s was not successful . . .
License	info	INFO: License file generated The license audit report scheduled for %s was successful. . .
Transaction Archive	error	ERROR: Transaction archive files older than 30 days. . .
Transaction Archive	info	INFO: Transaction archive files OK. . .
Database Maintenance	error	ERROR: Database maintenance not scheduled. . .
Database Maintenance	info	INFO: Database maintenance is scheduled . . .
Web Certificate Maintenance	error	ERROR: Web certificate about to expire. . .
Web Certificate Maintenance	info	INFO: Web certificate has been renewed . . .
Logs rotate	warn	Forcing log rotation as /var/log is at \$USAGE usage: \$message
Logs rotate	warn	Autopurging /var/log due to extremely high disk usage. Usage \$USAGE: \$message
Log processing	error	System unable to send {} messages to {}. . .
Backups	info	Backup completed. . .
Backups	info	Restore successful. . .
Backups	error	Restore failed . . .
Backups	error	Backup failed . . .
Upgrade	error	upgrade failed upgrade failed as other activity is in progress . . .

12. Scheduling

12.1. Scheduling

Any CLI command can be scheduled to run automatically, including but not restricted to backups and security upgrades.

By default there is no backup maintenance scheduled. Backup maintenance can be scheduled with the number of copies to be kept - refer to the backup maintenance topic.

The automated job schedule format is as follows:

- **schedule add <job-name> <user-command>**
- **schedule time <job-name> <hour> <minute>**
- **schedule time <job-name> every <N> hours**
- Alternatively the job can be scheduled to run every week on Monday with **schedule time <job-name> weekly 1**; where 0 is Sunday, 1 is Monday, 2 is Tuesday, 3 is Wednesday, 4 is Thursday, 5 is Friday and 6 is Saturday
- **schedule enable <job-name>**

Example:

```
schedule add mybackups backup create localbackup
```

```
schedule time mybackups 2 0
```

```
schedule time mybackups weekly 0
```

```
schedule enable mybackups
```

Among the tasks that can be scheduled are:

- Backup creation, e.g. **schedule add backupme backup create localbackup**
- Backup maintenance, e.g. **schedule add backupclean backup clean localbackup keep 5**
- Health reports, e.g. **schedule add reports diag report**

The example below shows the console output for some commands:

```
platform@host:~$ schedule add myexport voss export type license_initial_audit
Automatically setting time to midnight and enabling
myexport:
  active: true
  command: voss export type license_initial_audit --force
  hour: 0
  min: 0
```

(continues on next page)

(continued from previous page)

```
platform@host:~$ schedule time myexport weekly 0
myexport:
  active: true
  command: voss export type license_initial_audit --force
  hour: 0
  min: 0
  week: 0

platform@host:~$ schedule disable myexport
myexport:
  active: false
  command: voss export type license_initial_audit --force
  hour: 0
  min: 0
  week: 0
```

12.2. Internal Report Schedules

For v2 of the feature (VOSS-4-UC 18.1 patch and CUCDM 11.5.3 patch and later), the system runs an internal schedule to generate monthly license reports. For details on license reports and how to generate these manually, refer to the Licensing Guide.

This internal schedule cannot be disabled. The schedule is configured to run at 3AM UTC on the first day of the month. The date cannot be changed, but the time can. Please contact your VOSS account manager if a schedule time change is required.

After the monthly schedule is run, a check is carried out for the generated report. If the report was generated successfully, no messages are sent and no notifications are generated. If the report was not generated successfully, a message shows on the CLI console when logging in or when typing the **health** command:

```
LICENSE REPORT: FAILED - Please run 'voss export type license_initial_audit'
```

This message will continue to show until the report is generated successfully by running the command shown in the message.

An e-mail notification is also sent after the check fails:

```
ERROR: License file generation failed
The license audit report scheduled for <month> <year> was not successful.
Please contact your VOSS account manager.
```

An example SNMP trap that is generated when the report fails to run is show below - <month> <year> are variables in the example:

```
May 23 02:01:00 robot-slave snmptrapd[18891]: 2018-05-23 02:01:00 <UNKNOWN>
[UDP: [192.168.100.3]:11814->[192.168.100.25]:162]:
#012iso.3.6.1.2.1.1.3.0 = Timeticks: (207758) 0:34:37.58
#011iso.3.6.1.6.3.1.1.4.1.0 = OID: iso.3.6.1.2.1.88.2.0.1
#011iso.3.6.1.2.1.88.2.1.1.0 = STRING:
"ERROR: License file generation failed"
```

(continues on next page)

(continued from previous page)

```
#011iso.3.6.1.2.1.88.2.1.3.0 = STRING:  
"The license audit report scheduled for <month> <year> was not successful.  
#012Please contact your VOSS account manager. "  
#011iso.3.6.1.2.1.88.2.1.5.0 = INTEGER: 1  
#011iso.3.6.1.2.1.1.5.0 = STRING: "VOSS"
```

13. Backups

13.1. Backups

Backups represent a snapshot of the system, including database, configuration and system applications. Backups can be created manually, scheduled automatically, or created automatically when the system is upgraded.

Note:

- In a multinode environment, database backups are created from the highest priority secondary database node, thereby reducing the system load on the primary database node.
Use the **database config** command to check the secondary node priority.
- The **backup** command on a Generic NBI node (for Billing Data Extract) is not in use.

These backups are encrypted and can be stored on the local file system, or to a remote network location. The encryption key is needed to delete, export, restore and create a backup.

Off-site (non-local) backups are recommended, because this reduces the risk of sabotage or disk failures causing a loss of information.

There is no direct requirement for VMware snapshots. If VMware snapshots are used, also refer to the topic on [VMware Snapshot Maintenance](#).

For examples of backup maintenance commands and output, refer to the topics on Scheduling and Create a Backup.

If the VOSS-4-UC node is not recoverable, due to for example a hardware failure, a new node can be deployed and an existing backup restored to restore the node to service.

13.2. Backup and Restore in a Modular Architecture

Backups that were made *before* topology migration to a modular architecture can not be restored to the topology *after* migration. Manual intervention is required to restore.

A new modular architecture backup must be created immediately after topology migration. This backup will be on the highest priority secondary db node.

Important: Backups for an Application and Web Proxy type node are not required. A backup created on the secondary database node will include all the relevant data, excluding the web certificates.

After a restore is performed by following the documented process: [Restore the Backup](#), the web certificates need to be reinstalled on each of the Application and Web Proxy nodes. For commands, see: [Web Certificate Commands](#). These certificates should already be stored in a safe location after initially installed, therefore do not require inclusion in the backups.

13.3. Backup Destinations

Backups can be made to the local file system or a remote destination. Off-site (non-local) backups are recommended to reduce the risk of the loss of information.

- Display available backup destinations with **backup list**.
- Add a new backup destination with **backup add <location-name> <URI>**.

Local backups are stored on a separate backup volume and the `localbackup` destination is pre-configured.

- Display the list of localbackups with **backup list localbackup**, for example:

```
$ backup list localbackup
localbackup:
  URI: file:///backups
  Backups:
    2016-06-21 13:33
    2016-06-21 13:16
```

If the backup volume is too small, it can be increased - see: [Backup Size Considerations](#). If the `localbackup` destination is removed or renamed, an ISO file upgrade will no longer function. Therefore, it is imperative that this destination is not removed.

Example:

backup add myserverbackup sftp://user:password@server/path

Backups to sftp require ssh key-based authentication to be setup. Refer to [SSH key management](#) for further details. If you select ssh key-based authentication that was set up without a password, you will be prompted to input the password.

```
platform@VOSS-UN-6:~$ backup add remote sftp://dan@182.59.31.201/sftp
No password found. Do you want to use sshkeys? [y/N]: y
```

If a common remote backup point is to be used by all nodes in the cluster, the backup destination needs to be added to each node. This can be automated by using cluster remote execution, for example:

cluster run all backup add myserverbackup sftp://user:password@server/path

The creation of scheduled backups of all nodes is done for failover reasons.

Note: In a multinode configuration, while only the highest priority secondary node backup contains data, in a failover scenario, the new primary node will contain the restored backup data.

13.4. Backup Passphrase

System backups are encrypted. The encryption key is initially set as the platform user's password as set in the installation wizard.

An encryption key is required to delete, export, restore and create a backup. If the backups are on an external system, they can be deleted manually.

It is recommended that this be changed after installation. This can be done by running **backup passphrase**.

The passphrase is subject the same rules as a password. For example, if the minimum password length is set using **user password length**, this also applies to the passphrase. See also Password Strength Rules.

The following example shows the console output:

```
platform@masternode:~$ backup passphrase
Please enter current backup passphrase
Password:
Please enter new backup passphrase
Password:
Please re-enter new backup passphrase
Password:

Backup passphrase successfully changed
```

Keep this password, because restoring the backup to a new system requires this password.

To restore on the new system, run **backup passphrase** and enter the password used to create the backup.

13.5. Backup Size Considerations

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

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

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

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

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

See the Drive Control topic in the Platform Guide.

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

2. Run **voss db_collection_stats all** to show the size of the current database. This command validates the size of the database. This total will be smaller than the suggested backup size.
 - A local backup requires free space of at least twice the size of the database. Preferably add another 30% of this. At any given time it is therefore recommended that available backup space should be approximately twice the database size plus 30% of the database size.

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

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

13.6. Create a Backup

Note:

- On a multi-node system, to reduce the system load on the primary database node, backups are processed by the available secondary unified node with the highest database priority, i.e. the secondary node that is not for example in a recovery state. (To check priorities: **database config**.)

Only if no secondary unified node is available or on a standalone single-node cluster system, is the backup processed on the primary unified node.

- For best performance, it is recommended that remote backups be sent to a SFTP server at the *same data center* as the node processing the backup.

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

Backups can be created using **backup create <location-name>**, for example:

backup create localbackup or **backup create myserverbackup**.

An example of the console output is shown below:

```
platform@myhost:~$ backup create localbackup

... collecting data (step 1/3)
... preparing mongo data backup
... space available: 232 GB
... space required: 93 MB
... creating backup (step 2/3)
... verifying backup (step 3/3)
Backup was successfully created at localbackup::\
058bccead2588a6f11f1dd86678bab68de48691d

WARNING: Backup maintenance of this location is not scheduled
schedule add localbackup -maintain backup clean localbackup keep 5

You have new mail in /var/mail/platform

platform@myhost:~$ backup list localbackup
localbackup:
  URI: file:///backups
  Backups:
    2019-08-13 13:08
```

- Backups contain application data.
- Details of the backup can be seen in the log: **log view platform/backup.log**

Backups can be scheduled to run automatically - refer to the **schedule** command to automated backups.

For example:

- **schedule add mybackups backup create myserverbackup**
- **schedule time mybackups 2 0**
- **schedule enable mybackups**

The creation of scheduled backups of all nodes is therefore done for failover reasons. While only the available secondary unified node with the highest database priority contains data, in a failover scenario, the new secondary unified node with the highest database priority will contain the backup data.

If a common network URI is used as backup destination across the cluster, each node's backup will be uniquely identified by its UID in the remote backup directory.

Note: Off-site backups are recommended. In other words, export a local backup to a remote sftp server. Follow the process as described in the topic called Backup and Import to a New Environment.

13.7. Restore a Backup in a Clustered Environment

In a clustered environment, servers can allow for failures and can keep data intact, because when a server fails, an automatic failover occurs.

If all services are kept running and data remains accessible, a **backup restore** would only be necessary in very specific scenarios.

Restoring a backup in a cluster would only be necessary in the following cases:

- Data Corruption (Bad Data)
- Losing the whole cluster - requiring a redeploy of new servers

13.8. Create Space for a Backup or Restore

If a `No space left on device` message is received during a backup or a restore, carry out the following steps:

1. In VMware, add a disk to the system:
 - a. Click on **VM > Edit Settings...**
 - b. Click **Add...**
 - c. Select **Hard Disk**, then **Create a new virtual disk**.
 - d. Set the size to be the same as the DB disk - 250GB.
 - e. Click **Finish**
2. Log into platform account, and run **drives list**. Make note of the disk under `Unused disks:`.
3. Run `drives reassign <disk from step 2> services:backups`

Once done, all current data would have been moved to new disk and the old one can be removed from VMware. The **restore** command can now be rerun.

13.9. Maintaining Backups

A complete list of backups on a location can be displayed using **backup list <location>**.

Backups can be deleted using the following commands:

- **backup clean <location> keep <N>** will delete older copies so that only N copies are kept. Note: <N> must be larger than 0.
- **backup clean <location> before <yyyy-mm-dd [HH:MM]>** will delete copies older than the specified date.

By default, there is no regular maintenance of backups, and a scheduled job should be created to perform this maintenance, for example:

- **schedule add backuprotate backup clean localbackup keep 5**
- **schedule time backuprotate 3 0**
- **schedule enable backuprotate**

13.10. Backup and Import to a New Environment

The steps below show how to backup and import to a new environment.

1. Export:

- a. On the source system, create a remote backup location, for example location name `sftpbackup`:

backup add sftpbackup <URI>

For example:

```
backup add sftpbackup sftp://sftpusr:sftpw@172.29.41.107/home/sftp
```

If a directory is specified in the <URI>, this will be created during the backup. Backups to sftp require ssh key-based authentication to be set up. Refer to [SSH key management](#) for further details.

Alternatively, enter the password at the prompt, for example:

```
$ backup add sftpbackup sftp://sftpusr:sftpw@172.29.41.107/home/sftp
```

```
No password found. Do you want to use sshkeys? no
What is the host ssh password?
```

```
<type password here>
```

```
Location has been added
```

- b. Create a local backup:

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

backup create localbackup

In a multi-node configuration, the database backup will be created on the secondary node with the highest priority. Use **database config** to check the priority.

- c. List backups to get the date:

backup list

For example:

```
$ backup list
localbackup:
  URI: file:///backups
  Backups:
    1 backups have been created - most recently 2020-03-19 08:21
sftpbackup:
  URI: sftp://sftpusr:*****@172.29.41.107:home/sftp
  Backups:
    No backups created yet
```

- d. Export the local backup to the remote destination created by **backup add <remote_name>**.

- The system ID is appended as a directory to the backup <URI> destination path. This can be checked locally by running **system id**.
- The backup file is called <hostname>_<timestamp>.tar.gz.

Example output:

```
platform@VOSS:~$ backup export localbackup sftpbackup 2020-04-02 11:34
This operation could take a while if the backup is sizeable. Do you wish to_
↪continue? y
Compressing backup files for date 2020-04-02 11:34
Backup files successfully compressed to 202004021134.tar.gz
Backup files successfully compressed to /backups/
↪49940d3feaa39a6a9f36cb5ff533202157c3b77a/VOSS_202004021348.tar.gz
platform@10.120.1.246's password:
platform@10.120.1.246's password:
Export successfully created at platform@10.120.1.246:media/
↪49940d3feaa39a6a9f36cb5ff533202157c3b77a/VOSS_202004021348.tar.gz
```

2. Import:

Note: For large backup files that cannot be imported due to space limitations, see: [Restoring backup files too large for import](#).

- a. From the SFTP server, **scp** the `VOSS_202004021348.tar.gz` file to the new box (for example, `platform@172.29.21.97`). If the file on the SFTP server is in the directory `/backups/49940d3feaa39a6a9f36cb5ff533202157c3b77a`, change to the directory, then:

```
$ ls

VOSS_202004021348.tar.gz

$ scp VOSS_202004021348.tar.gz platform@172.29.21.97:/opt/platform/admin/home/
↪media/
```

- b. Import the file as a local backup, for example:


```
$ backup import localbackup media/VOSS_202004021348.tar.gz
```

- c. Get the file timestamp of the imported backup with **backup list** and restore the backup, for example:

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

```
$ backup restore localbackup 2020-04-02 15:41
```

13.10.1. Restoring backup files too large for import

If a backup file is available and is too large to copy into the local system `/home/media/` directory, the following procedure can be followed. (Available space on `home/media` can be checked using **diag disk**.)

1. On the system where the backup is to be restored, add the remote SFTP site, for example:

```
backup add sftpbackup sftp://sftpusr:sftpw@172.29.41.107/home/sftp
```

2. Obtain the system UUID by running **system id**. It should look like this:

```
49940d3feaa39a6a9f36cb5ff533202157c3b77a
```

3. Given an available backup file, e.g. `VOSS_202004021348.tar.gz`, (with timestamp reference: `2020-04-02 15:41`) this file can be extracted to the *remote* SFTP site configured above as follows:

- a. On the remote SFTP site, locate the sub-directory with the name of the UUID, for example

```
/home/sftp/49940d3feaa39a6a9f36cb5ff533202157c3b77a
```

- b. Extract the backup file in this directory. Ensure that the extracted backup files are listed in the root of this directory (not all in a subdirectory).

4. On the system where the backup is to be restored, run the remote restore, for example:

```
$ backup restore sftpbackup 2020-04-02 15:41
```

13.11. VMware Snapshot Maintenance

It is not recommend to keep more than two snapshots at any one time as this can negatively affect performance. Refer to the VMware Knowledge Base topic on Best practices for virtual machine snapshots in the VMware environment.

Unused or deprecated snapshots on VMware are caused by multiple snapshot creation and deletion. These can be removed to save space. Note however that these snapshots and disk images may not show in the vSphere GUI admin tool (Snapshot Manager).

Follow the steps recommended by VMware to remove deprecated, orphaned, unused and old snapshots via your VMware administrator. Pay special attention to the `.vmtx` configuration file to avoid removing live disks.

Always perform a backup and export of your data. We recommend shutting down your VOSS-4-UC instance(s) during a maintenance window in order to remove all unused images and to create a fresh snapshot.

13.12. Restoring a Backup on a New Environment

13.12.1. Introduction to Restoring Backups on a New Environment

Before restoring a backup on a new environment, take note of the following:

- Every backup made on the VOSS-4-UC platform is encrypted using a passphrase.
- To restore a backup, you need to set the passphrase where the restore will be done.
- The passphrase is initially set on deployment of the environment and uses the platform user's password as the passphrase for backup encryption.
- The passphrase can be set manually using the **backup passphrase** command.
- Note that if a new passphrase has been set on the system, all backups made with the previous passphrase cannot be restored unless the passphrase is set back to the passphrase used to create the backup.
- Currently, two disks can be impacted: the backup drive and the dbroot drive. The backup drive size is initially 50GB and the dbroot size is initially 250GB (60GB on a standalone single-node cluster deployment).
- If the data size you restore is bigger than the size of these drives, you need to reassign these drives to add more space for the restore.

13.12.2. Setting up the Backup Passphrase on a New Environment

To set the backup passphrase to restore on a different environment:

1. Log into the new environment. If this is a cluster deployment, log in on the DB Primary.
2. Run the **backup passphrase** command.
3. Specify the current passphrase. This is normally the password of the platform user set during the deploy of the system.
4. Enter the new passphrase twice.

13.12.3. Adding More Space to Accommodate a Large Restore

1. Right click on the VM in the VMware Client and click **Edit Settings**.
2. On the Hardware tab, click **Add**.
3. Follow the wizard to add a new hard disk to the VM with the correct size.
4. If the restore size exceeds both the backup and dbroot drives size, ensure you add two hard disks to the VM. In a clustered environment, this procedure needs to be performed on all of the DB nodes.

13.12.4. Reassign Current Drives (Backup and DBroot)

1. Once the hard disks are added, reassign the drives using the **drives reassign <disk> <mountpoint name> <size in GB>** command.
2. Use the **drives list** command to list the new drives added through VMware. For example, if the new drive is listed as `sdf`, use the reassign command as follows: **drives reassign sdf services:backups <size in GB>**.
3. Similarly, to reassign the dbroot, use the reassign command as follows: **drives reassign sde mon-godb:dbroot <size in GB>**.

Note: The *<size in GB>* specification means the volume need not be the size of the entire disk. However, a specified size must *not* be more than 90% of the disk size (or more than *<disk size less 10GB>* if the disk size is 100GB or smaller).

13.12.5. Restore the Backup

A complete list of backups on a location can be displayed using **backup list <location>**. To restore on a new system, run the **backup passphrase** command and enter password used to create the backup.

1. Copy the backup to the environment with **scp**. It will be located in the `media/` folder.
2. Once the file is successfully copied, use the **backup import** command to import the backup to a location that was set up, or the default `localbackup`.
3. Once the import is complete, run the **backup list** command as for example:

```
platform@Restore:~$ backup list
localbackup:
  URI: file:///backups
  Backups:
    2019-08-13 13:08
```

4. Run the **backup restore** command as for example:

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

```
platform@Restore:~$ backup restore localbackup 2019-08-13 13:08
Services will be restarted during the restore. Do you wish to continue? y
Check if restore can continue (step 1/4)
Enough space on /opt/platform/apps/mongodb/dbroot: total 93MB / 224GB
Enough space on /backups/appdata: free 934MB / 232GB
Running pre-restore scripts (step 2/4)
Stopping Application while performing database restore
Running backup restore (step 3/4)
System restore starting from \
  file:///backups/058bccead2588a6f11f1dd86678bab68de48691d (1565701713)
Local and Remote metadata are synchronized, no sync needed.
Last full backup date: Tue Aug 13 13:08:33 2019
/backups/appdata/
```

(continues on next page)

(continued from previous page)

```

completed
Running post-restore scripts (step 4/4)
Starting Application after performing database restore
Restarting services

Application processes stopped.

Reconfiguring applications....

If this includes a database restore, it may take some time to sync
Please run 'database config' to check when all nodes are done

Restored successfully
You have new mail in /var/mail/platform

```

13.12.6. Example of a Successful Restore

```

platform@Restore:~$ backup restore localbackup 2015-02-26 00:22
Services will be restarted during the restore. Do you wish to continue? y
Application voss-deviceapi processes stopped.
Stopping Application while performing database restore

--- Restore, ip=172.29.41.240, role=webproxy,application,database, loc=jhb

Application nginx processes stopped.
System restore starting from
  file:///backups/93d19980b574ed743d9b000a7595e42cad6a6d6b (1424910132)
Local and Remote metadata are synchronized, no sync needed.
Last full backup date: Thu Feb 26 00:22:12 2015
Successfully restored to /backups/appdata/restore_temp_1427441507,
  moving to /backups/appdata
Removing temporary files in /backups/appdata/restore_temp_1427441507
local\admin
Dropping database PLATFORM before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/PLATFORM
[object Object]
Repairing database PLATFORM before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/PLATFORM
[object Object]
Dropping database VOSS_FILES before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/VOSS_FILES
[object Object]
Repairing database VOSS_FILES before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/VOSS_FILES
[object Object]
Dropping database VOSS before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/VOSS

```

(continues on next page)

(continued from previous page)

```
[object Object]
Repairing database VOSS before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/VOSS
[object Object]
Trying with oplogReplay
restore successfull
{'172.29.41.240': (200, '\n')}
Starting Application after performing database restore

--- Restore, ip=172.29.41.240, role=webproxy,application,database, loc=jhb

Application services:firewall processes stopped.
Application nginx processes started.
Restarting services

Application processes stopped.

Application processes started.

System settings have changed, please reboot using 'system reboot'

Restored successfully
You have new mail in /var/mail/platform
```

14. System Security

14.1. Security Overview

The VOSS-4-UC platform is not installed with antivirus software or an index of whitelisted applications. The functionality provided by these anti-malware measures is implemented by means of an extensive number of measures to lock down and harden the operating system, platform system and network.

VOSS-4-UC security covers areas such as application and operating system security updates, operating system hardening, file and application encryption, jailed environments for applications, firewalls and user security. This locked down system ensures that platform users cannot install their own packages, binaries, or applications on the system to perform any malicious actions or allow the exploitation of vulnerabilities (such as Meltdown/Spectre).

For details, refer to the topics in the Platform Guide chapters called System Security and Network Security.

Note: Since the VOSS-4-UC application runs in a virtual environment it is important that the underlying VM infrastructure stays up to date to be protected against any vulnerabilities that may compromise the VOSS-4-UC Virtual Machine on the infrastructure.

14.2. Security Patches and Updates

During installation the system will automatically install the application named “security” which is a collection of all the latest security patches available for the various pieces of software in the platform at the time the system was built. Updates to this application are released to customers regularly. The security application provides these updates but does not automatically install them - allowing customers with concerns to verify them on lab machines first for example. Some security updates may also require scheduled downtime to complete and for this reason the final installation of updates is a manually triggered process.

The health command will inform the user if any security updates are currently available but not installed. Users can install security updates at any time by running the command:

security update

Those who would prefer to automate this can create a scheduled command to do so on a regular basis. The security update will install all operating system updates to both the main system and the application jails, but it will not generally contain updates to the core applications themselves - these are shipped separately as new application install versions as they require additional QA to ensure compatibility.

To manage security updates in a *cluster*, two options are available:

1. Run **security update** on *each* node in the cluster.

2. Carry out the update in two steps:

a. From the primary unified node, run:

cluster run notme security update

Wait for security updates to complete on these nodes in the cluster.

b. Then on the primary unified node, run:

security update

Example output:

```
platform@development:~$ security update
You are about to upgrade the system, which may cause services to restart.
Do you wish to continue? y
Application snmp processes stopped.
Installing updates for the main operating system
Starting system security update. This will take a few minutes
Checking packages to start the update process
Updating applications
Application processes stopped.

Application services:firewall processes stopped.
Application services processes started.
Updating /opt/platform/apps/mongodb/chroot
.....
Updating /opt/platform/apps/voss-deviceapi/chroot
.....
Updating /opt/platform/apps/selfservice/chroot
.....
Updating /opt/platform/apps/nginx/chroot
.....
The system is preparing for core security updates.
This is a required step and will require a reboot
Core security updates are now completed, system is configuring updates

Application processes stopped.
Application processes started.
Your system is fully updated and may require a reboot.
Run 'system reboot' or 'cluster run all system reboot' if updates were applied.
platform@development:~$ system reboot
You are about to reboot the system. Do you wish to continue? y
```

14.3. Configuration Encrypted

In order to help protect customer data and service stability the system configuration files are frequently recreated by the platform. This means that even malicious tampering to the platform will generally be undone by a simple restart. The configuration data is stored in the platform's internal files. These files are encrypted using a strong AES encryption layer to make them tamper-proof. They are never decrypted on disk, instead the applications which manage them will decrypt them in memory, read and make modifications as needed and then re-encrypt the data before writing them back to disk.

In this way the risk of tampering or data theft through the configuration system is greatly minimized.

14.4. Backup Encrypted

System backups include copies of the full system configuration as well as the full contents of the database. Thus theft of a backup would effectively constitute theft of all customer data stored on the platform. To mitigate this risk backups are encrypted using a strong 2048-bit in-line GPG encryption.

The encryption key for this is auto-generated by the platform based on a unique machine UUID. While it's possible for support to recover backups from a different machine this process is deliberately hard and only available to official technical support representatives. Backups on shared locations are separated on a per-source-machine basis to prevent conflicts.

14.5. Application Install Files Encrypted

In order to protect the trustworthiness of applications shipped for the platform, all application installers are encrypted files. The strong 2048-bit key needed to decrypt these are shipped with the platform and is different from the per-machine unique keys used for other encryption tasks. This key will only decrypt applications encrypted specifically with the unique key owned by VOSS. The system will refuse to install any application that is not encrypted or encrypted with a different key.

This ensures that only valid, untampered copies of genuine VOSS-released applications can be installed on the system.

14.6. File Integrity

System installation and upgrade binaries, as well as configuration files, are regularly checked for file integrity against a file hash. The types of files and directories to check, is configured.

A scheduled task is configured to initialize and to carry out the regular validation. If audit logging is enabled on a system, this initialization will show in the audit logs as the `EventType FileDetection` and `Audit Details` as `File checksum initialized`.

The Command Line Interface (CLI) diagnostic command **diag filehash** is also available to carry out a manual check for changes to these files of since the previous check. Note that the file check validates all system files and is a time consuming task.

If any files have been changed, removed or added to the configured types and directories, these will be listed in the command output, together with the type of changes.

Also refer to the topics on Diagnostic Tools and Audit Log Format and Details.

14.7. Protected Application Environments (Jails)

VOSS-4-UC runs the service providing applications in secured jail environments. This has significant value for the security and reliability of the system. It prevents applications from cross-interfering which makes the system more stable and reliable. In terms of security it effectively confines all services to dedicated and separate mini file systems with predictable content. In the event that an attacker were to gain access to the system through a vulnerability in a service he would therefore not gain access to the platform but only to the small confined jail in which the service was running. In that environment only the jail itself is vulnerable

and this can be very easily restored if damaged. The underlying system cannot be accessed from the jailed environment.

The VOSS system does not allow direct root access over ssh. If root access is required for debugging purposes, there is a tool called NRS. This tool requires the user to log in as a user with install privileges, who has to run **app install nrs**. The tool generates a key, which can only be deciphered by VOSS. VOSS uses this key to then gain root access in order to proceed with debugging.

14.8. Restricted User Shell

The platform attempts to reduce the risk of unintentional harm to the operation of the software by restricting the actions users can take. This is done using a specially configured setup of the well-known and actively maintained rbash shell.

The shell actively prevents the following:

- Users cannot set environment variables or alter their command path.
- Users cannot change the current directory.
- Users cannot specify a path to a command to run.

The commands users thus are able to run is only what is allowed by the platform setup. The vast majority of these commands use a common execution interface designed to allow only enough privileges to perform the system administration tasks they are created for. The exact list of commands a user can run is determined by his specific privileges and the specific setup of the machine on which he is working (different applications can add their own additional commands). This list is displayed on login and can be redisplayed with the **help** command.

14.9. User Security and Security Policy Management

Upon installation, user passwords are restricted as follows:

- Password length : 8
- Minimum number of days between password change : 1
- Maximum number of days between password change : 60
- Number of days of warning before password expires : 14
- Number number of days between password change: 10

User password and account security settings and policy details can also be configured. Commands are available to manage:

- password length
- automatic account locking after inactivity
- number of days between password change - valid values are from 5 to 20

The following commands are available to show the current length and set the default minimum password length:

- **user password length**
- **user password length <min_length>**

The value of `<min_length>` can be set from 8 to 127 characters. By default, it is 8 characters. For other password rules, refer to Password Strength Rules. The setting also applies to backup passphrases.

By default, any account that is created has the inactive lock set to 35 days.

To set the number of days between user password expiration:

user password expiry <username> [60-365,never]

Valid values for days is from 60 to 365. If `never` is typed in, the password does not expire and when typing **user passwordinfo <username>**, the Maximum number of days between password change value shows as `-1`.

The password re-use frequency default is 10 passwords, which means that the last 10 passwords cannot be re-used, unless this is set, using:

user password history <number-of-passwords>

where `5 <= <number-of-passwords> <= 20`.

To see the current `<number-of-passwords>` re-use frequency:

user password history

The commands below are available to carry out these tasks and to manage users.

- **user passwordinfo <username>**

Show details such as password expiry in days for a user, for example:

```
$ user passwordinfo joebrown
Last password change           : Nov 30, 2015
Password expires               : Feb 28, 2016
Password inactive              : Apr 03, 2016
Account expires                : never
Minimum number of days between password change : 1
Maximum number of days between password change : 60
Number of days of warning before password expires : 14
```

- **user inactivelock <days> <user>**

Set the number of days of inactivity before a user account is locked, for example:

```
$ user inactivelock 35 joebrown
A 35 day inactive logon policy has been set for user: joebrown
```

- **user lock <user>**

Manually lock a user account, for example:

```
$ user lock joebrown
passwd: password expiry information changed.
```

- **user unlock <user>**

Manually unlock a user account, for example:

```
$ user unlock joebrown
passwd: password expiry information changed.
```

To unlock users who exceeded allowed number of failed login attempts:

```
$ system ssh fail_limit reset joebrown
$ user unlock joebrown
passwd: password expiry information changed.
```

- **user password view_lock <user>**

The command output is different in accordance with the event that locked the user account:

Not a manual user lock:

```
$ user password view_lock joebrown
There is no password lock applied for user joebrown.
Please run 'system ssh fail_limit view joebrown' to
ensure the account is not locked because the user has
reached the maximum number of failed attempts .
```

Manual user lock:

```
$ user password view_lock joebrown
The password for user: joebrown has been locked.
Please run 'user unlock joebrown' and
'system ssh fail_limit reset joebrown' to ensure
you unlock and reset lock limits for this user account
```

- **user lastlogon <username>**

Show details of the last logon for:

- a user who has logged in before:

```
$ user lastlogon joebrown
joebrown 172.29.90.74 Thu Dec 3 11:04:54.
```

- a user who has not logged in before:

```
$ user lastlogon joebrown
joebrown logged in***
```

Use the **user help** command to see the general user management options such as user list, add, grant or revoke rights and remove users.

The command **user list** provides rights and security policy details of *all* users, while **user list <username>** provides details for a single user. For example:

```
$ user list
user:
  joebrown:
    rights:
      mail
      app
  janedoe:
    rights: value not set
  billsmith:
    rights: value not set

security_policy:
  user:
    platform:
```

(continues on next page)

(continued from previous page)

```

    auto_inactive_account_lockout: 35
joebrown:
    account_locked: No
    auto_inactive_account_lockout: 35
janedoe:
    auto_inactive_account_lockout: 35
billsmith:
    account_locked: No

```

In addition, a system wide account security setting can be configured and displayed. The setting will then apply to all *new* users and override the default inactive lock setting of 35 days.

The following commands are available:

- **system inactivelock**: show the current system wide inactive lock default:

```

$ system inactivelock
Newly added users will have their inactivity lock set to 35 days.

```

- **system inactivelock <num of days>**: set the system wide inactive lock default for all new user accounts, in other words, for users created *after* the setting of the system wide inactive lock:

```

$ system inactivelock 35
Newly added users will have their inactivity lock set to 35 days.

```

14.10. Creating Additional Users

During installation a user called 'platform' is created which has full access to all allowed commands within the restricted environment. This user (and others with the appropriate rights) can then create additional users who are further restricted to only be able to run certain commands. For example a user could be created who can only run diagnostic and logging commands - able to monitor the health of a system but required to escalate any actions.

Users are created, managed and deleted through the user command. To create a new user use:

user add <username>

The system will create a Unix user with the name specified and set up to use a restricted shell identical to the platform user. Initially this user's password is set to match the username but it must be changed on first login. New users start out with no rights and can only run the very basic system commands provided to all users (such as **ls**).

For SFTP only users, see: [Creating and Managing SFTP Users](#).

14.11. Creating and Managing SFTP Users

Administrators can add and manage users who have SFTP only access. For platform user management, see: [Creating Additional Users](#).

To create a new SFTP only user, use the command:

user sftp add <username>

Add a username and password. See [Password Strength Rules](#).

The system will create a user with the provided name and password provided, with the following restrictions:

- the system can *only* be accessed by SFTP
- user access is restricted to the platform `home/` directory only
- the SFTP user will have a SSH key attached.
- only the administrator can change the SFTP only user password

To attach a SSH key to the SFTP user:

1. Copy the SSH public key for the user onto the system
2. Run **user addkey <username> <keyfile>** to attach the key to the user

See also: [Adding a Key for Automatic User Login](#).

To change the SFTP user password:

user sftp password <username>

To remove the SFTP user:

user del <username>

SFTP users are listed under the `sftp-only-users` group when running the **user list** command.

An SFTP only user can also be added without a password. This option can be used for an SFTP implementation that uses a key only (no password):

user sftp add_nopass

14.12. Granting and revoking user rights

Once a user is added the user needs to be granted access to run commands. The user's command menu will only display those commands to which access have been granted.

To grant access to a command use the 'user grant' command as follows:

user grant <username> <command> [options]

Only one command can be granted at a time, however these can be complex. The more detailed the command, the more fine-grained the privilege becomes. This is best explained by example.

Running the following command:

user grant peter app

Will allow the user peter to execute any command within the 'app' series of commands. However it could be restricted further by instead running a command like:

user grant peter app list

With this version peter will see the **app** command on his menu, but its help will only display 'list' as a sub-command - peter can thus see the list of apps but cannot perform more potentially risky tasks such as installing or restarting applications.

This can be expanded to other subsets by simply running additional grants:

user grant peter app start

Would now allow peter to both see the list of applications or restart applications that failed, however he will not be able to do other app related tasks such as installations. The **grant** command effectively verifies that the start of a command by a user matches one of the privileges granted to that user - so peter will be able to add options to any command he is granted access to.

In order to restrict commands - be sure to determine whether any options should be allowed and if not, only grant access to the specific parameters you wish peter to be able to execute. For example if peter is your database administrator for example you may wish to use:

user grant peter app start mongodb

Instead of giving access to all **app start** commands.

Should you wish to revoke a command privilege from a user you can do this using the following command:

user revoke <username> <full command>

The command being revoked must match exactly one of the commands previously granted to a user. To review the current privileges of a user use:

user list <username>

Which will display the user's entire list of granted commands in full. You can also just run

user list

Without an option to list all users created on your system and their privileges.

14.13. Password Strength Rules

The platform user and any users created are held to strong password rules to help reduce the risk of system penetration. These rules are enforced whenever passwords are changed or set. In order to meet system password strength rules a user's password must:

- By default be at least 8 characters long. This can be modified with the **user password length <min_length>** command. See *User Security and Security Policy Management*.
- Contain at least one capital and one non-capital letter.
- Contain at least one number.
- Contain at least one special character.
- A password change should differ by at least 8 characters from the old password. In other words, if an old password was 8 characters, then *all* new password characters should differ.

14.14. SSH Login Fail Limit

An administrator can view and modify the number of login attempts for a user.

- The default number of failed login attempts for a user is 10 before the account is locked.
- The default duration that an account will be locked, is 15 minutes (900 seconds).
- **system ssh fail_limit set <number>**

Set the number of failed login attempts for all user accounts on this system before account lockout occurs. For example:

```
$ system ssh fail_limit set 3
You are about to set a limit for failed login attempts.
This limit will apply to all user accounts on this system.
Do you wish to continue? Y
```

- **system ssh fail_limit view <username>**

View the current status of a user's failed login attempts. Examples:

```
$ system ssh fail_limit view joebrown
Login      Failures  Latest failure  From
joebrown   0

$ system ssh fail_limit view joebrown
Login      Failures  Latest failure  From
joebrown   1      12/04/15 10:38:00  192.168.0.90
```

- **system ssh fail_limit reset <username>**

Reset the limit back to 0 on a locked out account. This will allow a user to log back in to the system without resetting a password after a lockout occurs. For example:

```
$ system ssh fail_limit view joebrown
Login      Failures  Latest failure  From
joebrown   3      12/04/15 10:38:00  192.168.0.90

$ system ssh fail_limit reset joebrown
You are about to reset the account lockout information for
user: joebrown. This will allow this user to log back in to
the system. Do you wish to continue? y

$ system ssh fail_limit view joebrown
Login      Failures  Latest failure  From
joebrown   0
```

- **system ssh fail_limit unlock_time <seconds>**

Enable the unlock time and set the duration in seconds that an account will be locked for after it has been locked.

To disable the unlock time setting, use the command with the parameter value as zero:

system ssh fail_limit unlock_time 0

The example output below shows the command response for parameter values:

```
$ system ssh fail_limit unlock_time 60
SSH session unlock time has been set.

$ system ssh fail_limit unlock_time 0
SSH session unlock time has been disabled.

$ system ssh fail_limit unlock_time -1
Please enter a valid number for unlock time.
```

- **system ssh fail_limit unlock_time**

Display the status of the unlock time setting.

For example:

```

$ system ssh fail_limit unlock_time
SSH session unlock time is not set.

$ system ssh fail_limit unlock_time 60
SSH session unlock time has been set.

$ system ssh fail_limit unlock_time
SSH session unlock time has been set to 60 seconds.

```

14.15. SSH Session Limit

An administrator can set and modify the number of SSH sessions allowed:

- system-wide (default is 10 if not set)
- for a user (default set to the system-wide setting)

Note: The default number of SSH sessions allowed *per IP source* is limited to 10. This means that if a user SSH session limit is higher than this limit, the user session origin needs to be from a different IP source.

Best practice is to set the system-wide SSH session limit first as this will be the default for any new users created on the system. Also note that the per user SSH session limit cannot be set higher than the system-wide SSH session limit.

To see the current system-wide SSH limit, use:

system ssh_session_limit

To set the system-wide SSH limit:

system ssh_session_limit set <number>

This system wide value will restrict the *per user* limit that can be set.

When a user is added and no session limit is added, the user's number of SSH sessions is set to the default system wide default limit of 10. It is recommended to also set the user's session limit.

To set the SSH session limit for a user:

user credential_policy session_limit <username> <number>

where <number> cannot be larger than the system wide session limit, if it has been set.

The current SSH session limit for users can be seen by using the **user list** command, for example:

```

platform@drp32:~$ user credential_policy session_limit joebrown 5
platform@drp32:~$ user list
user:
  joebrown:
    rights: value not set

security_policy:
  joebrown:
    account_locked: No
    auto_inactive_account_lockout: 35
    ssh_connection_limit: 5

```

(continues on next page)

(continued from previous page)

```
platform@drp32:~$
```

If a user has sessions open while the session limit is set, the limit in affect when new sessions are opened.

14.16. SSH key management

SSH authentication requires maintaining the system SSH keys. This can be done as follows:

- **keys create** creates a local SSH keyset
- **keys add <host>** adds the remote host to the known hosts list allowing outgoing connections
- **keys send <user>@<host>** will send the public key from the local SSH keyset to the remote server, thereby enabling remote SSH authentication.

For example, if you wish to perform a backup to a remote host, first create a local key if necessary with **keys create**. Allow communication with the host using **keys add <host>**. Send the key to the remote host with **keys send <user>@<host>**.

If you select ssh key-based authentication that was set up without a password, you will be prompted to input the password when adding a sftp backup host:

```
platform@VOSS-UN-6:~$ backup add remote sftp://dan@182.59.31.201/sftp
No password found. Do you want to use sshkeys? [y/N]: y
```

The certificates are independent of web servers/proxies.

For more details on SSH key-based authentication, refer to OpenSSH documentation.

14.17. SSH Algorithm Management

SSH algorithms on the VOSS-4-UC platform can be viewed, enabled, disabled or reset to the default list.

The available commands are:

- **system ssh algorithm default** - Reset enabled and disabled algorithms to their default.
- **system ssh algorithm list < all | algorithm-type >** - Display *all* or any of *<cipher/mac/kex/key>* SSH algorithms. The list will also show *enabled* and *disabled* algorithm types.
- **system ssh algorithm disable <algorithm-type> <algorithm-name>** - Disable algorithms of a specific type *<cipher/mac/kex/key>*, by specifying a space separated list of algorithm names.

Note that not all algorithms can be disabled.

For example, to disable two of the `cipher` algorithms, the command would be:

```
system ssh algorithm disable cipher aes128-ctr aes192-ctr
```

- **system ssh algorithm enable <algorithm-type> <algorithm-name>** - Enable algorithms of a specific type *<cipher/mac/kex/key>*, by specifying a space separated list of algorithm names.

Command example to view status of *all* algorithms types:

```
platform@VOSS:~$ system ssh algorithm list all
cipher:
  enabled:
    aes128-ctr
    aes192-ctr
    aes256-ctr
kex:
  enabled:
    diffie-hellman-group1-sha1
    diffie-hellman-group14-sha1
    diffie-hellman-group-exchange-sha1
    diffie-hellman-group-exchange-sha256
    curve25519-sha256@libssh.org
key:
  enabled:
    ssh-ed25519
    ssh-ed25519-cert-v01@openssh.com
    ssh-rsa
    ssh-dss
    ecdsa-sha2-nistp256
    ecdsa-sha2-nistp384
    ecdsa-sha2-nistp521
    ssh-rsa-cert-v01@openssh.com
    ssh-dss-cert-v01@openssh.com
    ecdsa-sha2-nistp256-cert-v01@openssh.com
    ecdsa-sha2-nistp384-cert-v01@openssh.com
    ecdsa-sha2-nistp521-cert-v01@openssh.com
mac:
  enabled:
    hmac-sha1
    hmac-sha2-256
    hmac-sha2-512
    hmac-ripemd160
    hmac-ripemd160@openssh.com
    umac-128@openssh.com
    hmac-sha1-etm@openssh.com
    hmac-sha2-256-etm@openssh.com
    hmac-sha2-512-etm@openssh.com
    hmac-ripemd160-etm@openssh.com
    umac-128-etm@openssh.com
```

14.18. Adding a Key for Automatic User Login

To automate tasks such as backups from remote hosts, it may be necessary to allow for the SSH login on VOSS-4-UC by a user without a password.

This login requires the addition of a public SSH key for the user. The **user addkey** command is available to add a keyfile for a user. The command to run, is of the format:

user addkey <username> <keyfile>

Note that:

- The user who runs this command, should have the <keyfile> available on their local directory. The public keyfile should therefore be copied to VOSS-4-UC.

- The user (<username>) for whom the key is to be added, should exist. If the user does not exist on VOSS-4-UC, a message shows to indicate this.

If the command is successful, the following message is shown:

```
User key added. You should try to ssh now
```

14.19. Prevention of DOS Attacks

The following list shows measures implemented in VOSS-4-UC to protect the system against Denial of Service (DOS) attacks:

- Firewall protection:
 - TCP flood protection against:
 - * the SSH port
 - * web server ports
 - SYN flood protection
- Configurable session limits for the VOSS-4-UC platform SSH access is **Sessions per user** and **Sessions per application**. An administrator can set and modify the number of SSH sessions allowed:
 - system-wide (default is 10 if not set)
 - for a user (default is 10 if not set)

See SSH Session Limit for detailed information.

- The usage of ports, protocols, and services are registered with the DoD PPS Database
- An automated, continuous on-line monitoring of the system is implemented, with:
 - Audit trail creation capability in a format that a log viewing application can immediately alert personnel of any unusual or inappropriate activity with potential Information Assurance (IA) implications.
 - A command line command that a user can automatically disable the system if serious IA violations are detected.
- Applications are monitored and notifications sent when resource conditions reach a predefined threshold indicating there may be attack occurring, for example through SNMP traps and triggers.
- High disk utilization is managed due to error notifications. For log files, disk utilization is managed by:
 - daily log rotation
 - 4 weeks of backlogs
 - the creation of new (empty) log files after rotating old ones
 - log file compression
 - a logging restriction of 20 messages per minute
- A continuous cycle of updating packages during releases is in place with notifications during updates. Commands to carry out a security check or update can be run at any time.

14.20. Memory Dumps and Security

Memory dumps in VOSS-4-UC are restricted to Cisco Administrators. Attackers will therefore not be able to gain access to sensitive data which may appear unencrypted in memory.

14.21. Manage Read-Only Database Users

Remote read-only database access for users can be managed. A username, source IP address and password are required as parameters.

Important: Since the system firewall service is restarted when adding or removing database users, this may affect system operation.

It is therefore strongly recommended that the task be carried out during a maintenance window.

Note:

- When adding a user, a prompt for a password is given. The password rules are:
 - 8 or more characters minimum
 - contains at least 1 uppercase
 - contains at least 1 lower case
 - contains at least 1 digit
 - contains at least 1 punctuation character
- Adding and deleting users require a service restart.

The following commands and parameters are available:

- **database user add <ip> <username>**
- **database user del <ip> <username>**
- **database user list**

Example console output:

```
platform@dev:~$ database user list
No users defined
platform@dev:~$ database user add 192.79.22.52 alex
You are about to restart services. Do you wish to continue? yes
New user password:
Please verify password:
platform@dev:~$ database user list
IP users
192.79.22.52 alex
platform@dev:~$ database user del 192.79.22.52 alex
You are about to restart services. Do you wish to continue? yes
platform@dev:~$ database user list
No users defined
```

15. Network Security

15.1. Network Communications between Nodes within the Cluster

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

All communication between nodes is encrypted.

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

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

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

- Proxy to Proxy Node

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

Communication	Protocol	Port
Cluster Communications	HTTPS	TCP 8443 bi-directional

- Proxy to Unified/Application Node

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

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

- Unified Node to Unified node

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

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

- Application node to Application node

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

Communication	Protocol	Port
Cluster communications	HTTPS	TCP 8443 bi-directional

- Application Node to Database node

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

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

- Database Node to Database node

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

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

15.2. 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 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.

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)
Cisco HCM-F	HTTPS	TCP 8443
MS PowerShell Proxy Node	HTTP, HTTPS	TCP 5985, 5986

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 365/Online	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

15.3. Dynamic Firewall

The most important part of the network security model is the system firewall.

The platform uses a dynamic firewall which does not open a fixed set of ports but adapts to the applications installed, only allowing such traffic as the specific set of running services require.

If an application is stopped, it's ports are automatically closed. This creates a default block list firewall which pinholes only those ports required for the operation of the specific setup in use.

The firewall is one of the very first services the platform brings up and among the very last it shuts down in order maximize the network security.

Where possible, the firewall will also ratelimit connections to services to prevent abuse (see the section: Prevention of DOS attacks for more details).

15.4. Web Certificate Setup Options

The platform installs a self-signed certificate for the web-frontend by default. This provides encryption of the web-traffic but does not provide users with valid authentication that the server is correct or protect against man-in-the-middle attacks.

Two types of certificate setups are supported:

- VOSS-4-UC certificate setup

We strongly advise customers to obtain a trusted CA-signed certificate and install it on the server. A 4096 bit RSA certificate is generated on VOSS-4-UC systems.

Once a signed, trusted certificate is obtained from the CA, copy it to the platform using **scp** and then install the file into the server using:

web cert add <filename>

For details, see: [VOSS-4-UC Setup a Web Certificate](#)

- Own private certificate and generated Subject Alternative Name (SAN) certificate setup

Customers can upload their own private certificate and generated SAN certificates, in other words it is not necessary to run **web cert gen_csr** on the platform CLI. One certificate can therefore be uploaded on all nodes. Note that customers are then responsible for the security of their private keys.

For details, see: [Own Web Certificate Setup](#).

The file to upload should be in a PEM format. PEM certificates typically have extensions like `.pem`, `.crt`, `.cer` and `.key`.

The PEM file must have the correct form of line termination: a single “Line Feed” character. If your PEM file was saved on MS Windows, be sure to remove the `^M` characters from the file, for example in a Linux console with:

```
$ tr -d '\r' < original.pem > fixed.pem
```

In the file, the SAN certificate composition has the private key first and then the certificate and the private key should be *unencrypted* (i.e. the key header text would then not show “BEGIN ENCRYPTED PRIVATE KEY”).

For example:

```
-----BEGIN PRIVATE KEY-----
MIIEvAIBADANBgkqhkiG9w0BAQEFAASCbKYwggSiAgEAAoIBAQDNV1pXvjIiiWuJIABW
[...]
IeJnlBPwDJX6Yo9Q==
-----END PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIIEbTCCAlUCAgPoMA0GCSqGSIb3DQEBCwUAMIGbMQswCQYDVQQGEwJaQTELM
[...]
ulfj0D54fozATLIIdMZSrmImk8CfkDPkmWbIKRce729DTQwHrMG/Oo1ZC2
-----END CERTIFICATE-----
```

Copy the certificate file to the platform `media/` directory using **scp** and then install the file using:

web cert add_san <filename>

For example:


```
platform@host:~$ web cert add_san media/cert.pem
Updating the certificate requires the web server to be restarted.
Do you wish to continue? yes
Restarting nginx
platform@host:~$
```

Note:

- SSO certificate management is carried out on the GUI. Refer to the GUI documentation for details.
- VOSS-4-UC supports wildcards for Common names (CN) in the web browser certificate.
- Only one certificate file can be installed on the platform. For more details on NGINX compatible certificates see the relevant nginx documentation here: [http://nginx.org/en/docs/http/nginx_http_ssl_module.html]
- Please note the importance of ensuring that SSL certificates generated match the assigned network name of the platform.

The list of supported SSL ciphers are as follows. This list may change as ciphers are added or found to be insecure:

- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES256-GCM-SHA384
- DHE-RSA-AES128-GCM-SHA256
- DHE-DSS-AES128-GCM-SHA256
- kEDH+AESGCM
- ECDHE-RSA-AES128-SHA256
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-RSA-AES128-SHA
- ECDHE-ECDSA-AES128-SHA
- ECDHE-RSA-AES256-SHA384
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES256-SHA
- ECDHE-ECDSA-AES256-SHA
- DHE-RSA-AES128-SHA256
- DHE-RSA-AES128-SHA
- DHE-DSS-AES128-SHA256
- DHE-RSA-AES256-SHA256
- DHE-DSS-AES256-SHA
- DHE-RSA-AES256-SHA
- AES128-GCM-SHA256

- AES256-GCM-SHA384
- AES128-SHA256
- AES256-SHA256
- AES128-SHA
- AES256-SHA
- AES
- CAMELLIA

15.5. VOSS-4-UC Setup a Web Certificate

The VOSS-4-UC platform generates a 4096 bit RSA private key file, using the details stored when using the **web cert details edit** command, along with a Certificate Signing Request (.csr) file.

Repeat the steps below for each proxy that requires signed SSL certificates:

1. Check the current certificate details with **web cert details**. Initially, the `User set details` is `Unset`. For example:

```
platform@host:~$ web cert details
Issuer data:
  C: SA
  CN: 11.120.11.100
  L: DeviceAPI
  O: Platform
  ST: WP
Key data:
  C: SA
  CN: 11.120.11.100
  L: DeviceAPI
  O: Platform
  ST: WP
User set details: Unset
```

2. Run **web cert details edit** if needed to edit the details displayed from the server. For example:

```
platform@host:~$ web cert details edit
Country Name (2 letter code): C:IE
State or Province Name (full name): ST:Dublin
Locality Name (eg, city): L:Dublin
Organization Name (eg, company): O:DublinSolutions Ltd.
Organizational Unit Name (eg, section): OU:R&D
Common Name (e.g. server FQDN or IP): CN:dublinsolutions.com
Email Address: platform@dublinsolutions.com
details stored
platform@host:~$
```

Verify the edits by running **web cert details** after editing. For changes, the Issuer details will then not match the User set details.

3. Run **web cert gen_csr** to generate the Certification Signing Request (.csr) file `media/cert_sign_req.csr` for signing.

For example:

```
platform@host:~$ web cert gen_csr
-----BEGIN CERTIFICATE REQUEST-----
M88E8TCCAttrCAQAwgasxCzAJBgNVBAYTA1pBMQswCQYDVQQIDAJXUDERMA8GA1UE
[...]
IIDr1vrepZkFQr+XDah2L5g5v8bI
-----END CERTIFICATE REQUEST-----

=====
Please send the above or the actual file /opt/platform/admin/home/media/cert_sign_
→req.csr to a CA to be signed

platform@host:~$ ls -la media/cert_sign_req.csr
-rw-rw-rw- 1 root platform 1789 Jan 18 11:20 media/cert_sign_req.csr
```

4. Use **scp** on a remote workstation to copy the file off the VOSS-4-UC platform `media/` directory and send it to a Certificate Authority (CA). Request a PEM format file to be returned.

The returned file received from the CA should be a PEM certificate file. PEM certificates typically have extensions like `.pem`, `.crt`, `.cer` and `.key`.

- If you did not receive a combined certificate from the CA, concatenate the reply signed cert and the reply intermediate CA cert into a file.

The signed certificate must be first in the concatenated file.

The PEM must have the correct form of line termination: a single “Line Feed” character. If your PEM file was saved on MS Windows, be sure to remove the `^M` characters from the file, for example in a Linux console with: **\$ tr -d ‘r’ < original.pem > fixed.pem**

- If the received file is a `.p7b` file, it should be converted to a PEM format - refer to the topic: [Convert Web Certificates from P7B to PEM Format](#).
- If the received file is in another format, carry out the required conversion. For example, when a received `.crt` file is opened and is not in the correct format in MS Windows, it may show a message on MS Windows Certificate panel: “Windows does not have enough information to verify the certificate”. Choose the Details tab of the panel, select Copy to File... to open the Export Wizard. Choose Base-64 encoded as export format.

5. Upload the PEM file to the proxy using **sftp** or **scp**. The file will be added to the `media/` directory, for example: `media/cert.pem`.
6. Once the file is uploaded, run **web cert add <filename of uploaded file>**. This command will combine the key and PEM file, and present it to nginx to use for secure (SSL) web communication. For example:

```
platform@host:~$ web cert add media/cert.pem
Updating the certificate requires the web server to be restarted.
Do you wish to continue? yes
Restarting nginx
platform@host:~$
```

15.6. Own Web Certificate Setup

The steps below provide an example of own private certificate and generated Subject Alternative Name (SAN) certificate setup as summarized in [Web Certificate Setup Options](#).

Note: The Subject Alternative Name (`alt_names`) field lets you specify additional host names (sites, IP addresses, common names, etc.) to be protected by a single SAN Certificate.

1. Log into a system that has the **openssl** command set up.
2. Create a bash script file with contents as below:

```
openssl req -new -sha256 -nodes -out cert.csr -newkey rsa:4096 -keyout private.  
↪key -config <(  
cat <<-EOF  
[req]  
default_bits = 2048  
prompt = no  
default_md = sha256  
req_extensions = req_ext  
distinguished_name = dn  
  
[ dn ]  
C=<Country code>  
ST=<County/State>  
L=<City>  
O=<Organization>  
OU=<Org Unit>  
emailAddress=<admin email address>  
CN = <Main DNS Name>  
  
[ req_ext ]  
subjectAltName = @alt_names  
  
[ alt_names ]  
DNS.1 = <Alternate name 1>  
DNS.2 = <Alternate name 2 ... you can add more below, just inc DNS.2>  
EOF  
)
```

3. Edit the sections in < > brackets.
4. Run **bash <scriptfile from above>**
5. Send the file called `cert.csr` to your CA, requesting them to make sure to sign it as a SAN certificate.
6. Take the file that they send back, save it as `signed.crt`
7. Combine the `private.key` file with `signed.crt`:
Run **cat private.key signed.crt > complete.cert**
8. Upload the `complete.cert` file to the VOSS-4-UC system using **sftp** or **scp**. The file will be added to the `media/` directory, for example: `media/complete.cert`
9. On the VOSS-4-UC system, run **web cert add_san media/complete.cert**

15.7. Web Certificate Expiration Notice

If a Web Certificate is due to expire, a notice will display on the status display 30 days before expiration:

```
platform@development:~$ health
```

```
host: AS01, role: webproxy,application,database, LOAD: 3.85
date: 2014-08-28 11:24:22 +00:00, up: 6 days, 3:03
network: 172.29.42.100, ntp: 196.26.5.10
HEALTH: NOT MONITORED
database: 20Gb
application: up
WEB CERT EXPIRES AT: 2014-09-26 11:30:02
```

```
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         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             template - template pack creator
```

Web certificate expiration can also be monitored if scheduled health monitoring is enabled - see:enable_health_monitoring. The health email will then show this message.

If a Web Certificate has expired, the notice on the status displays:

```
WEB CERT EXPIRED AT: 2014-09-26 11:30:02
```

Once the certificate is expired, the system can be used as normal, but the certificate will be expired and for non self-signed certificates (like a Godaddy or Thawte certificates), the data will no longer be properly encrypted.

15.7.1. Renewing Expired Certificates

According to the certificate type in use, refer to the setup steps to manage certificates:

- [Own Web Certificate Setup](#)
- [VOSS-4-UC Setup a Web Certificate](#)

15.8. Convert Web Certificates from P7B to PEM Format

VOSS-4-UC uses web certificates in Privacy Enhanced Mail (PEM) format. PEM certificates typically have extensions like .pem, .crt, .cer and .key.

If a P7B format certificate is received from a Certificate Authority (CA):

1. Copy the files to a workstation with Linux console available (Not the VOSS-4-UC system).
2. Run the following command for each <filename>.p7b, for example:

```
sudo openssl pkcs7 -in <filename>.p7b -inform DER -print_certs -out
<filename>.pem
```

3. Open the PEM file in a text editor. You will see formatting like the example below in the file:

```
subject=/C=GB/ST=West Midlands/L=Coventry/O=Service Coventry/OU=Network Services/
  CN=ccs-cp-v4uc.svcoventry.gov.uk/emailAddress=network@svcoventry.co.uk
issuer=/C=GB/O=Coventry/OU=PKI/CN=BCC Intermediate CA
-----BEGIN CERTIFICATE-----
MIIFxTLmBK2gAwIBAgITXgAAAMBXLQb0/ImKBwALmQAawDANBgkqhkiG9w0BAQsF
ADBOMQswLmYDVLMQEWJHQjETMBEGA1UEQhMKQmlybWluZ2hhbTEMMAoGA1UEQxMD
UETJMRwwGgYDVLMQEXNLM0MgSW50ZXJtZWVpYXRlIENBMB4LmQE2MTAwNTEyNTIx
```

4. Delete all text and blank lines outside the lines:

```
-----BEGIN CERTIFICATE----- and -----END CERTIFICATE-----
```

5. Save the file and make sure that the file is not saved in a DOS format.

15.9. Web Certificate Commands

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

```
web cert add <filename> - Install the certificate from <filename>
                           into the web server
web cert add_san <filename> - Install a SAN certificate from <filename>
                              into the web server
web cert del - Revert to a self-signed certificate
web cert details - Print the certificate details in config
                  system
web cert details edit - Update the certificate details in config
                       system
web cert gen_csr - Create a CSR file in
                  /opt/platform/admin/home/media
web cert gen_selfsigned - Generate a self-signed certificate
web cert print_csr - Create a CSR file in
                    /opt/platform/admin/home/media
web cluster prepnod - Prepares the system so that it can be
                     joined to a cluster as a web proxy
web ssl list - Shows a list of the supported SSL protocols, and
               their current state
web ssl enable <protocol> - Enable an SSL protocol
web ssl disable <protocol> - Disable an SSL protocol
web weight add <server:port> <weight> - Modify the weights of an upstream
                                         service. Higher weights will serve
                                         more requests, while 0 will only be
                                         used if no other servers are available
web weight del <server:port> - Delete the user-defined service weight
                               and use the system default.
web weight list - Display the weights of upstream services
```

15.10. Web TLS Protocol Configuration

Commands are available to list Transport Layer Security (TLS) protocol versions and also to enable or disable TLS versions.

Note:

- The command should be run on all nodes in a cluster.
- When enabling or disabling a TLS protocol version, the web server needs to be restarted. Running the command will show a message and carry out this task.

The following protocols are available in VOSS-4-UC:

- TLSv1.1
- TLSv1.2
- TLSv1.3

Important:

- While TLSv1.1 is still available, you are strongly advised to move to the later versions for security reasons.
- TLSv1.2 is enabled by default upon installation. Upon upgrade, your current protocol is retained.
- TLSv1.2 can only be disabled by enabling TLSv1.3.

- **web ssl list**

Example:

```
$ web ssl list
TLSv1.1: Disabled
TLSv1.3: Disabled
TLSv1.2: Enabled
```

- Enabling or disabling a protocol that is already in that state, will raise an error message.

- **web ssl disable <TLS version>**

- Enabling or disabling a protocol that is already in that state, will raise an error message.

Example:

```
$ web ssl disable TLSv1.1
Disabling the TLSv1.1 protocol requires the web server to be restarted.
Do you wish to continue? yes
TLSv1.1: Disabled
TLSv1.2: Enabled

Restarting nginx for settings to take effect

Application nginx processes stopped.

Application services: firewall processes stopped.
Application nginx processes started.
```

- **web ssl enable <TLS version>**

Note:

- When running **web ssl enable TLSv1.3**, it will disable TLSv1.1 and TLSv1.2. Users will not be able to alter web ciphers.
- When running **web ssl enable TLSv1.1** or **web ssl enable TLSv1.2**, it will disable TLSv1.3. Users can change the web ciphers.
- If a user enables TLSv1.1, it will also enable TLSv1.2.

- Enabling or disabling a protocol that is already in that state, will raise an error message.

Example:

```
$ web ssl enable TLSv1.1
Enabling the TLSv1.1 protocol requires the web server to be restarted.
Do you wish to continue? yes
TLSv1.1: Enabled
TLSv1.2: Enabled

Restarting nginx for settings to take effect

Application nginx processes stopped.

Application services: firewall processes stopped.
Application nginx processes started.
```

The table below shows the result of running **web ssl enable** or **web ssl disable** given a specific state (from **web ssl list**).

State			Command	Result		
1.1	1.2	1.3	on/off	1.1	1.2	1.3
off	on	off	1.1 on	on	on	off
off	off	on	1.1 on	on	on	off
off	off	on	1.2 on	off	on	off
off	on	off	1.3 on	off	off	on
on	on	off	1.3 on	off	off	on
on	on	off	1.1 off	off	on	off

15.11. Web TLS Cipher Management

Web TLS ciphers on the VOSS-4-UC platform can be listed and managed. This can be done as follows:

- **web ssl cipher list** will list nginx ciphers grouped by status: `disabled`, `enabled`.
- **web ssl cipher default** will set the default nginx ciphers. This command requires the web server to be restarted.
- **web ssl cipher enable <space separated cipher(s)>** will enable the listed nginx ciphers. This command requires the web server to be restarted.
- **web ssl cipher disable <space separated cipher(s)>** will disable the listed nginx ciphers. This command requires the web server to be restarted.

Note: The enabled ciphers cannot *all* be disabled.

Command examples:

- List:

```
platform@VOSS:~$ web ssl cipher list
enabled:
  ECDHE-RSA-AES256-SHA
  ECDHE-ECDSA-AES256-SHA
  SRP-DSS-AES-256-CBC-SHA
  SRP-RSA-AES-256-CBC-SHA
  SRP-AES-256-CBC-SHA
  DHE-RSA-AES256-SHA
  DHE-DSS-AES256-SHA
  DH-RSA-AES256-SHA
  DH-DSS-AES256-SHA
  DHE-RSA-CAMELLIA256-SHA
  DHE-DSS-CAMELLIA256-SHA
  ...
```

- Disable:

```
platform@VOSS:~$ web ssl cipher disable CAMELLIA256-SHA
Disabling nginx ciphers requires the web server to be restarted.
Do you wish to continue? y

Application services:firewall processes stopped.
Application nginx processes stopped.
Reconfiguring applications...
Application nginx processes started.
disabled:
  CAMELLIA256-SHA
enabled:
  ECDHE-RSA-AES256-GCM-SHA384
  ECDHE-ECDSA-AES256-GCM-SHA384
  ECDHE-RSA-AES256-SHA384
  ...
```

15.12. Network URI specification

All network locations are specified as a URI, for example download locations, backup destinations, notification destinations, and so on.

The following are examples:

```
mailto:user@host
```

```
sftp://user:password@host/directory
```

```
ssh://user:password@host/directory
```

```
snmp://community@host for SNMP v2
```

```
snmp://community:password@host for SNMP v3
```

15.13. Web External Proxies

On a web proxy node only, remote server proxies can be managed.

The commands should be run on the relevant web proxy node.

The commands to add or remove remote server proxies will automatically reconfigure and restart the nginx process, so some downtime will result.

- To *add* remote server proxies and ports, run the command on the relevant node:

web external add <URL> <remote IP:port>

Note the <URL> can include a path, for example:

```
http://172.49.90.53:8888/myurl/set
```

so that when the proxy URL is then entered as `https://<vossIP>/myproxy/`, this will then resolve to `http://172.49.90.53:8888/myurl/set`.

- To *remove* remote server proxies, run the command on the relevant node:

web external del <URL>

- To *list* configured remote proxies on a web proxy node, run the command on the relevant node:

web external list

For example, adding, listing and removing remote servers:

```
platform@VOSS:~$ web external add myproxy http://172.49.90.53:8888
Updating the locations requires the web server to be restarted.
Do you wish to continue? y
You have new mail in /var/mail/platform

platform@VOSS:~$ web external list
external:
  myproxy:
    upstream: http://172.49.90.53:8888

platform@VOSS:~$ web external del myproxy
Updating the locations requires the web server to be restarted.
Do you wish to continue? y
You have new mail in /var/mail/platform

platform@VOSS:~$ web external list
external: value not set
```

16. High Availability and Disaster Recovery (DR)

16.1. High Availability Overview

High Availability (HA) is an approach to IT [system design](#) and configuration that ensures VOSS-4-UC 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 the secondary database node.

This section outlines the configuration steps to deploy a HA enabled VOSS-4-UC Platform on VMware. It presupposes familiarity with the Installation Guide and Platform Guide - for the latter guide, in particular the topics on DR Failover and Recovery.

16.2. Default HA and DR scenario

VOSS-4-UC supports using off-the-shelf VMware tools.

High Availability is implemented using VMware HA clusters, with data accessed via a central storage facility (SAN). VMware monitors the primary server, and should it fail, another instance of the VM is automatically started on a different hardware instance. Since data is shared on the SAN, the new HA instance will have access to the full dataset.

Disaster Recovery is implemented by streaming data updates to a separate DR instance that remains powered on. If the primary server fails, the DR instance can take over operation. The switch-over to DR instance is scripted, but must be invoked manually.

During a HA failover, the HA instance assumes the primary IP address, and no reconfiguration of other UC elements is required. However, in the case of a DR failover, interaction with other UC elements should be considered.

- DNS can be used effectively to provide hostname abstraction of underlying IP addresses. In such a case, a DNS update will allow existing UC elements to seamlessly interact with the new DR instance.
- If DNS is not available, and the UC elements cannot be configured with the IP address of the DR instance, it is necessary for the DR instance to assume the primary IP address. In such a case, the DR and the primary IP addresses can be swapped using the CLI interface. Standard networking practices should be employed to ensure that the IP address is correctly routed, e.g. Stretched layer-2 vLAN, and ensuring that the Primary and DR instances are not operated with the same IP address.

The following failure points should be considered:

- Since the HA instance is started automatically if the primary instance fails, a slight interruption in service is expected, including VMware polling latency in determining that the primary server has failed, and the startup delay of the HA instance. This delay is around 3 minutes

- If data is corrupted on the SAN, the HA instance will start with the same corrupt code and data instances
- Since VMware is checking only for VM liveness, it is not able to check that the primary instance is functionally active.
- Data updates are transported to the DR instance. If data updates cannot be shipped by the primary instance, SNMP traps are generated informing administration of the problem. However, if this is not fixed timeously, it is possible for the DR instance to become out of sync. These delays could result in data loss between the primary and DR instances. Database updates are scheduled every 3 minutes and/or 16MB.
- There are certain manual steps that are required to bring the DR instance online. These steps are documented in the Platform Guide.

16.3. HA and DR scenario with Cisco VMDC geo-redundancy architecture

HA and DR instances can be geo-relocated at will within the capabilities of the underlying network architecture.

For example, it is feasible to extend a VMware HA cluster geographically using high speed data links and layer-2 stretched vLANs.

DR as implemented by the VOSS-4-UC system lends itself to geographical separation with streaming data replication to a second powered-on instance.

Interaction with other UC elements must be considered within the capabilities of the network, using either DNS for seamless transition, or IP reconfiguration either within the UC elements or the VOSS-4-UC system.

16.4. Configuring a HA System Platform on VMware

This is an optional step, however, for production servers it is highly recommended that they are run in a HA deployment configuration. This can be done by the client, but should be checked by a system representative

1. Log into VMware VSphere, then select **File > New >Cluster...**
2. Enter the Name, and select the Turn on VMware HA check box.
3. Make sure that the Enable Host Monitoring check box and Enable: Do not power on VMs that violate availability constraints radio buttons are selected.
4. Select the required default restart priority.
5. Select the VM Monitoring Only option from the VM Monitoring drop-down list, and set the Default Cluster Settings/Monitoring sensitivity to High.
6. Select the Disable EVC radio button, unless you know the exact version of CPU technologies that are enabled on your system.
7. Select the Store the swapfile in the same directory as the virtual machine (recommended) radio button.
8. Ensure the settings are all correct and click the **Finish** button.
9. Drag all of the machines that will be used into the newly created cluster.

10. Once done, they will be listed below the new cluster, with any VM's that were moved into the root of the cluster.
11. Select each of the Machines in the cluster then select the Configuration tab.
12. If Time Configuration is displayed in red, select Properties, then click the **Options** button.
13. Select NTP Settings, and then click the **Add** button.
14. Select the Restart NTP service to apply changes check box, and then click the **OK** button.
15. Select the relevant Cluster, and then select the Summary tab. There should be no configuration issues listed.

The production OVA is deployed as in the hardware specifications of the deployment topologies and installation steps and considerations indicated in the Installation Guide.

16.5. DR Failover and Recovery Overview

16.5.1. DR Failover

The VOSS-4-UC system makes use of database replication facilities during normal operation. During a failover, if 50% or more of the service resources are lost, the system will no longer function without manual intervention.

In this case, the following high level process should be followed.

1. Display the current cluster topology using **cluster status**.
2. Remove the dead nodes using **cluster del <ip>**. Power off the deleted node, or disable its Network Interface Card.
3. Once the cluster topology is adjusted, the cluster must be reprovisioned using **cluster provision**.
4. Afterward, the cluster status can be rechecked with **cluster status**.

See also: *Using the screen command*.

Refer to the appropriate detailed DR scenarios for the complete sequence of DR steps.

16.5.2. Cluster Failure Scenarios

The status of the cluster can be displayed from the command-line on any node using the command:

cluster status

The system can automatically signal email and/or SNMP events in the event that a node is found to be down.

Refer to the diagrams in the Installation Guide section on deployments.

Loss of an Application role The Web Proxy will keep directing traffic to alternate Application role servers. There is no downtime.

Loss of a Web Proxy Communication via the lost Web Proxy will fail, unless some another load balancing infrastructure is in place (DNS, external load balancer, VIP technology). The node can be installed as a HA pair so that the VMware infrastructure will restore the node if it fails. Downtime takes place while updating the DNS entry or returning the Web Proxy to service. For continued service, traffic can be

directed to an alternate Web Proxy or directly to an Application node if available. Traffic can be directed manually (i.e. network elements must be configured to forward traffic to the alternate Web Proxy).

Loss of a Database role If the primary database service is lost, the system will automatically revert to the secondary database. The primary and secondary database nodes can be configured via the Command Line Interface (CLI) using **database weight <ip> <weight>**. For example, the primary can be configured with a weight of 40, and the secondary with a weight of 20. If both the primary and the secondary Database servers are lost, the remaining Database servers will vote to elect a new primary Database server. There is downtime (usually no more than a few seconds) during election and failover, with a possible loss of data in transit (a single transaction). The GUI web-frontend transaction status can be queried to determine if any transactions failed. The downtime for a Primary to Secondary failover is significantly less and the risk of data loss likewise reduced. A full election (with higher downtime and risk) is therefore limited only to cases of severe outages where it is unavoidable.

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 systems have weights that will allow a new primary to be chosen.

Loss of a site Unified and Database nodes have database roles. The status of the roles can be displayed using **cluster status**. If 50% or more of the database roles are down, then there is insufficient availability for the cluster to function as is. Either additional role servers must be added, or the nodes with down roles must be removed from the cluster and the cluster needs to be reprovisioned. If there is insufficient (less than 50% means the system is down) Database role availability, manual intervention is required to reprovision the system – downtime is dependent on the size of the cluster. Refer to the Platform Guide for details on DR Failover. Database role availability can be increased by adding Database roles, providing greater probability of automatic failover. To delete a failed node and replace it with a new one if database primary is for example lost: The node can be deleted using **cluster del <ip>**. Additional nodes can be deployed and added to the cluster with **cluster add <ip>**. The database weights can be adjusted using **database weight <ip> <weight>**. Finally, the cluster can be reprovisioned with **cluster provision** (it is recommended that this step is run in a terminal opened with the **screen** command). This command is the same as **cluster provision fast**. The *fast* parameter is available for backwards compatibility and is the default behavior, which is to run the provisioning on all nodes in parallel. Use the command **cluster provision serial** on systems where the VMware host is under load.

The console output below shows examples of these commands.

The cluster status:

```
platform@cpt-bld2-cluster-01:~$ cluster status

Data Centre: jhb
  application : cpt-bld2-cluster-04 [172.29.21.243]
                cpt-bld2-cluster-03 [172.29.21.242]

  webproxy   : cpt-bld2-cluster-06 [172.29.21.245]
                cpt-bld2-cluster-04 [172.29.21.243]
                cpt-bld2-cluster-03 [172.29.21.242]

  database   : cpt-bld2-cluster-04 [172.29.21.243]
                cpt-bld2-cluster-03 [172.29.21.242]

Data Centre: cpt
  application : cpt-bld2-cluster-02 [172.29.21.241]
```

(continues on next page)

(continued from previous page)

```

cpt-bld2-cluster-01[172.29.21.240] (services down)

webproxy : cpt-bld2-cluster-05[172.29.21.244]
           cpt-bld2-cluster-02[172.29.21.241]
           cpt-bld2-cluster-01[172.29.21.240] (services down)

database : cpt-bld2-cluster-02[172.29.21.241]
           cpt-bld2-cluster-01[172.29.21.240] (services down)

```

Deleting a node:

```

platform@cpt-bld2-cluster-01:~$ cluster del 172.29.21.245
You are about to delete a host from the cluster. Do you wish to continue? y
Cluster successfully deleted node 172.29.21.245

Please run 'cluster provision' to reprovision the services in the cluster

Please note that the remote host may still be part of the database clustering
and should either be shut down or reprovisioned as a single node BEFORE this
cluster is reprovisioned
You have new mail in /var/mail/platform

```

Adding a node:

```

platform@cpt-bld2-cluster-01:~$ cluster add 172.29.21.245

Cluster successfully invited node 172.29.21.245

Please run 'cluster provision' to provision the services in the cluster

```

Database weights: listing and adding

```

platform@cpt-bld2-cluster-01:~$ database weight list
172.29.21.240:
  weight: 5
172.29.21.241:
  weight: 3
172.29.21.243:
  weight: 2
172.29.21.244:
  weight: 1

platform@cpt-bld2-cluster-01:~$ database weight 172.29.21.240 10
172.29.21.240:
  weight: 10
172.29.21.241:
  weight: 3
172.29.21.243:
  weight: 2
172.29.21.244:
  weight: 1

```

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

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

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

- a standalone installation

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

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

- a node installation during multinode installation - see [Notes on Multi-Node Installation](#)

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

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

- or during a failover recovery

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

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

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

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

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

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

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

1. Log in to vSphere to access the ESXi Host.
2. Choose **File > Deploy OVF Template**.
3. Choose Source, browse to the location of the .ova file, and click **Next**.
4. On the Name and Location page, enter a Name for this server.
5. On the Deployment Configuration page, select the appropriate node type. (Refer to the list at the `role` option below.)
6. Choose the resource pool in which to locate the VM.
7. Choose the data store you want to use to deploy the new VM.
8. Choose the disk format to use when deploying the new VM.

In production environments, “thick provisioning” is mandatory.

Thick Provision Eager Zeroed is recommended.

9. On the Network Mapping, choose your network on which this VM will reside.
10. Do not select Power on after deployment.
11. On the Ready to Complete page, click **Finish** to start the deployment.
12. After the VM is created, select the CD ROM configuration and verify the **Connect at Power On** check box is enabled. Also, verify the memory, CPU, and disk settings against the requirements shown in either the Standalone System Hardware Specification or Multinode Cluster Hardware Specification section in the Install Guide.
13. Power on the VM.
14. Configure the options in the installation wizard:

Option	Option name	Description
1	IP	The IP address of the server.
2	netmask	The network mask for the server.
3	gateway	The IP address of the network gateway.
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.
5	NTP	The NTP server is mandatory to ensure that time keeping is accurate and synchronized among nodes in the same cluster.
6	boot password	Enable boot loader configuration password. See the example below.
7	hostname	The hostname, not the fully qualified domain name (FQDN).
8	role	Note: only WebProxy, Application and Database nodes are used for a modular architecture installation. <ul style="list-style-type: none"> • A WebProxy role installs only the front-end web server together with ability to distribute load among multiple middleware nodes. • An Application node is the main transaction processing engine and includes a web server which can operate by itself, or route transactions from a web node. • A Database node provides persistent storage of data. • A Standalone node consists of the Web, Application, and Database roles on one node. • 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.

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

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

Notes on Passwords and Security

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

Password protection can be enabled on the VOSS-4-UC 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  (*****)
```

Notes on Multi-Node Installation

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

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

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

The following Unified nodes are required in the cluster:

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

Note:

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

- *Modular Architecture Topology*

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

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

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

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

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

Also refer to Multinode Installation section in the Install Guide.

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

Finalize the Installation

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

Connect to newly deployed server CLI as the platform user.

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

```

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

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

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

```

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

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

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

16.6. DR Failover and Recovery in a Unified Node Cluster Topology

16.6.1. Election of a New Primary and Failover

In the case where unified nodes fail, the system follows a failover procedure. For details on the failover and DR process, refer to the topics in the Platform Guide.

If the primary database is lost, the failover process involves the election of a new primary database by the remaining database nodes. Each node in a cluster is allocated a number of votes that are used in the failover election of a new primary database - the election of a running node with the highest database weight.

The database weights for a node can be seen as the `priority` value when running the **database config** command. Note that database weight of a node does not necessarily match its number of votes.

```

$ database config
date: 2016-04-25T09:50:34Z
members:
  172.29.21.101:27020:
    priority: 16
    stateStr: PRIMARY

```

(continues on next page)

(continued from previous page)

```
172.29.21.101:27030:
  stateStr: ARBITER
172.29.21.102:27020:
  priority: 8
  stateStr: SECONDARY
172.29.21.102:27030:
  stateStr: ARBITER
172.29.21.103:27020:
  priority: 4
  stateStr: SECONDARY
172.29.21.103:27030:
  stateStr: ARBITER
172.29.21.104:27020:
  priority: 2
  stateStr: SECONDARY
myState: 1
ok: 1
set: DEVICEAPI
```

The maximum number of votes in a cluster should not exceed 7 and arbiter votes are added to nodes to provide a total of 7 votes.

The tables below show the system status and failover for a selection of scenarios for 6 node and 8 node clusters. Also refer to the topics on the specific DR scenarios. The abbreviations used are as follows:

- Pri : Primary site
- DR : DR site
- N : node. Primary node is N1, secondary node is N2.
- w : database weight
- v : vote
- a : arbiter vote

Not all scenarios are listed for 8 node clusters and example weights have been allocated.

- For a 6 node cluster with 4 database nodes and 2 sites, initial votes are as follows:
Primary database node, nodes 2-3: 2 (1 + 1 arbiter) Secondary database nodes 4: 1 (no arbiter)

Pri N1 w:40 v:1 a:1	Pri N2 w:30 v:1 a:1	DR N3 w:20 v:1 a:1	DR N4 w:10 v:1	Votes	System Status under scenario
Up	Up	Up	Up	7	System is functioning normally.
Up	Up	Up	Down	6	Scenario: Loss of a Non-primary Server in the DR Site. System continues functioning normally.
Up	Up	Down	Up	6	Scenario: Loss of a Non-primary Server in the DR Site. System continues functioning normally.
Up	Down	Up	Up	6	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.
Down	Up	Up	Up	5	Scenario: Loss of the Primary Database Server. Some downtime occurs. System automatically fails over to N2.
Down	Down	Up	Up	3	Scenario: Loss of a Primary Site. Manual recovery required
Up	Up	Down	Down	4	System continues functioning normally.
Up	Down	Down	Up	3	Manual recovery required
Up	Down	Up	Down	4	System continues functioning normally.

- For an 8 node cluster with 6 database nodes and 2 sites, initial votes are as follows:

Primary database node: 2 (1 + 1 arbiter voting member) Secondary database nodes total: 5 (no arbiter votes)

The table here shows a representative selection of scenarios.

Pri N1 w:60 v:1 a:1	Pri N2 w:50 v:1	Pri N3 w:40 v:1	Pri N4 w:30 v:1	DR N5 w:20 v:1	DR N6 w:10 v:1	Votes	System Status under scenario
Up	Up	Up	Up	Up	Up	7	System is functioning normally.
Up	Up	Up	Down	Down	Down	4	Scenarios: Loss of a Non-primary Node in the Primary and Secondary Site. System continues functioning normally.
Up	Up	Up	Up	Down	Up	6	Scenario: Loss of a Non-primary Server in the DR Site. System continues functioning normally.
Up	Down	Up	Up	Up	Up	6	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.
Up	Down	Down	Up	Up	Up	6	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.
Down	Up	Up	Up	Up	Up	6	Scenario: Loss of the Primary Database Server. Some downtime occurs. System automatically fails over to N2.
Down	Down	Up	Up	Up	Up	4	Some downtime occurs. System automatically fails over to N3.
Down	Down	Down	Up	Up	Up	3	Manual recovery required
Down	Down	Down	Down	Up	Up	2	Scenario: Loss of a Primary Site. Manual recovery required
Up	Up	Down	Up	Up	Up	6	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.
Up	Up	Down	Down	Up	Up	5	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.
Up	Up	Down	Down	Down	Up	4	Scenarios: Loss of a Non-primary Node in the Primary and Secondary Site. System continues functioning normally.
Up	Up	Down	Down	Down	Down	3	Manual recovery required
Up	Down	Up	Down	Down	Down	3	Manual recovery required

As the representative table above shows, the 8 node status and scenarios are similar for a number of permutations of nodes. For example, the failure of a single node N2, N3 or N4 results in the same scenario:

- Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.

The list below shows individual nodes (N1 to N6) and groups of nodes that will result in the same failover scenario.

Upon recovery, there is typically a delay of 10-20 minutes in the continuance of transaction processing.

- N2, N3, N4
- N5, N6

- N2+N3, N2+N4, N3+N4
- N1+N2+N3, N1+N2+N4, N1+N3+N4
- N1+N5, N1+N6
- N2+N5, N2+N6, N3+N5, N3+N6, N4+N5, N4+N6
- N2+N3+N4
- N2+N3+N5, N2+N3+N6, N2+N4+N5, N2+N4+N6, N3+N4+N5, N3+N4+N6
- N5+N6

A failure in other groupings will require a manual recovery, for example, in such groups as:

- N1+N2+N3, N1+N2+N4, N1+N2+N5, N1+N2+N6, N1+N3+N4, N1+N3+N5, N1+N3+N6, N1+N4+N5, N1+N4+N6, N1+N5+N6
- N2+N3+N4+N5, N2+N3+N4+N6, N3+N4+N5+N6
- N1+N2+N3+N4, N1+N2+N3+N5, N1+N2+N3+N6, N1+N3+N4+N5, N1+N3+N4+N6, N1+N4+N5+N6
- N1+N2+N3+N4+N5, N1+N2+N3+N4+N6

16.6.2. DR Failover and Recovery Scenarios

A number of failover scenarios and recovery steps are shown. In each case, a node topology is assumed: 6 or 8 node clusters in 2 sites - primary and Disaster Recovery (DR). A node failure scenario is indicated and a set of recovery steps are provided.

The following scenarios that are covered:

- Power off of a node
- Loss of a non-primary node in the Primary site
- Loss of a non-primary server in the DR site
- Loss of the Primary Database Server
- Loss of a Primary Site
- Loss of a DR Site

For the scenarios below, the following procedures and definitions apply:

- In the event of a network failure or a temporary network outage affecting a single a node, the node will be inaccessible and the cluster will respond in the same way as if the node had failed. If network connectivity is then restored, no action is required, because the node will again start communicating with the other nodes in the cluster, provided no changes were made to that node during the outage window.
- In a clustered deployment, the datacentre would typically be two different datacentres, for example “Virginia” and “Seattle”. These can be thought of as a primary site and a DR (Disaster Recovery) site in case of a failure in the primary site. These two datacentres can exist on the same physical hardware, so the separation of the cluster is into two sets of three nodes.

When datacentres are defined during installation, the nodes of a cluster may or may not be in the same physical location. The cluster is designed to communicate across all nodes, regardless of their physical location.

- During recovery, the command **cluster provision** must be run every time a node is deleted from or added to a cluster, even if it is a replacement node. It is recommended that this step is run in a terminal opened with the **screen** command. See: [Using the screen command](#).
- During recovery and installation, the command **cluster prenode** must be run on every node.
- During recovery of 8 node clusters, database weights should be deleted and added again.

16.6.3. Scenario: Power Off and On of a Node

The scenario and recovery steps apply to Unified and Proxy nodes.

Node powered off

- Secondary nodes assume primary
- There is no cluster downtime and normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a node is powered off.
- At this point, *all* transactions that are currently in flight at the node are lost and will not recover. The lost transactions have to be rerun.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail. With the node still powered off, replaying the failed transactions is successful

Recovery steps if the node is powered off:

1. Power up the node. The node resyncs. Run the **database config** command to verify the state of the database members. A typical output of the command would be:

```
$ database config
date: 2017-04-25T09:50:34Z
heartbeatIntervalMillis: 2000
members:
  172.29.21.41:27020:
    priority: 60.0
    stateStr: PRIMARY
    storageEngine: WiredTiger
  172.29.21.41:27030:
    priority: 1.0
    stateStr: ARBITER
    storageEngine: WiredTiger
  172.29.21.42:27020:
    priority: 50.0
    stateStr: SECONDARY
    storageEngine: WiredTiger
  172.29.21.43:27020:
    priority: 40.0
    stateStr: SECONDARY
    storageEngine: WiredTiger
  172.29.21.44:27020:
    priority: 30.0
```

(continues on next page)

(continued from previous page)

```

stateStr: SECONDARY
storageEngine: WiredTiger
172.29.21.45:27020:
priority: 20.0
stateStr: SECONDARY
storageEngine: WiredTiger
172.29.21.46:27020:
priority: 10.0
stateStr: SECONDARY
storageEngine: WiredTiger
myState: 1
ok: 1.0
set: DEVICEAPI
term: 38

```

Note that `storageEngine` will show as `WiredTiger` after the database engine upgrade to Wired Tiger when upgrading to VOSS-4-UC 17.4. Otherwise, the value is `MMAPv1`.

In other words, the database should not for example be any of: `STARTUP`, `STARTUP2` or `RECOVERING`. Note however that it is sometimes expected that nodes are recovering or in startup, but then should change to a normal state after a period of time (depending on how far out of sync those members are).

A file system check may take place.

2. If a replacement node is not on standby, rebuild steps such as boot up, adding to cluster, setting database weight and reprovisioning may take 200-300 minutes, depending on hardware specifications.

It is recommended that standby nodes are available to be used for faster recovery.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in `STARTUP2` or `SECONDARY` or `PRIMARY` states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure (it is recommended that this step is run in a terminal opened with the **screen** command). See: [Using the screen command](#).
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
2. Run **database weight del <ip>** from the primary node, <ip> being the IP of the new incoming node.
3. Log into any secondary node (non primary unified node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the new incoming node.

5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.6.4. Scenario: Loss of a Non-primary Node in the Primary Site

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example here is a typical cluster deployment of 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers. In the case where more than one non-primary node is lost on the Primary site, the relevant recovery steps are repeated.

The design is preferably split over 2 physical data centers.

```
Data Centre: jhb
  application : AS01[172.29.42.100]
                AS02[172.29.42.101]

  webproxy :   PS01[172.29.42.102]
                AS01[172.29.42.100]
                AS02[172.29.42.101]

  database :   AS01[172.29.42.100]
                AS02[172.29.42.101]

Data Centre: cpt
  application : AS03[172.29.21.100]
                AS04[172.29.21.101]

  webproxy :   PS02[172.29.21.102]
                AS03[172.29.21.100]
                AS04[172.29.21.101]

  database :   AS03[172.29.21.100]
                AS04[172.29.21.101]
```

Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a non-primary node is experienced. In this 6-node example, AS02 [172.29.42.101] failed while transactions were running.
- Examine the cluster status running **cluster status** to determine the failed state:

```
platform@AS01:~$ cluster status

Data Centre: unknown
application : unknown_172.29.42.101[172.29.42.101] (not responding)

webproxy : unknown_172.29.42.101[172.29.42.101] (not responding)
```

(continues on next page)

(continued from previous page)

```

database : unknown_172.29.42.101[172.29.42.101] (not responding)

Data Centre: jhb
application : AS01[172.29.42.100]

webproxy : PS01[172.29.42.102]
           AS01[172.29.42.100]

database : AS01[172.29.42.100]

Data Centre: cpt
application : AS03[172.29.21.100]
           AS04[172.29.21.101]

webproxy : PS02[172.29.21.102]
           AS03[172.29.21.100]
           AS04[172.29.21.101]

database : AS03[172.29.21.100]
           AS04[172.29.21.101]

```

- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the database server AS02[172.29.42.101] still down, replaying the failed transactions are successful.

Recovery Steps if the server that is lost, is unrecoverable:

1. A new unified node needs to be deployed. Ensure the server name, IP information and data centre name is the same as on the server that was lost.
2. Delete the failed node database weight (**database weight del <ip>**), for example **database weight del 172.29.42.101**
3. Run **cluster del 172.29.42.101**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
4. Run **cluster provision** on the cluster *without* the node to be added and then create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
5. Switch on the newly installed server.
6. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
7. If the node will be a unified or web proxy node, run **cluster prepnode** on it.

8. From the primary unified node, run **cluster add <ip>**, with the IP address of the new unified server to add it to the existing cluster.
9. Add database weights so that the weights distributed throughout the cluster
 - Delete all database weights in the cluster. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**
 - Check weights - either individually for each node, or for the cluster by using the command:
cluster run application database weight list
Make sure all application nodes show correct weights.
10. Run **cluster provision primary <ip of current primary>** to join the new unified node to the cluster communications. It is recommended that this step is run in a terminal opened with the **screen** command.
11. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure. It is recommended that this step is run in a terminal opened with the **screen** command.
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
2. Run **database weight del <ip>** from the primary node, <ip> being the IP of the new incoming node.
3. Log into any secondary node (non primary unified node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.6.5. Scenario: Loss of a Non-primary Server in the DR Site

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a cluster deployment: 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers.

The design is preferably split over 2 physical data centers.

Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a non-primary node is experienced.

In this 6-node example, AS04 [172.29.21.101] failed while transactions were running.

- Examine the cluster status running **cluster status** to determine the failed state:

```
Data Centre: unknown
  application : unknown_172.29.21.101[172.29.21.101] (not responding)

  webproxy : unknown_172.29.21.101[172.29.21.101] (not responding)

  database : unknown_172.29.21.101[172.29.21.101] (not responding)

Data Centre: jhb
  Application : AS01[172.29.42.100]
               AS02[172.29.42.101]

  webproxy :   PS01[172.29.42.102]
               AS01[172.29.42.100]
               AS02[172.29.42.101]

  database :   AS01[172.29.42.100]
               AS02[172.29.42.101]

Data Centre: cpt
  application : AS03[172.29.21.100]

  webproxy :   PS02[172.29.21.102]
               AS03[172.29.21.100]

  database :   AS03[172.29.21.100]
```

- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure

transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the database server `AS04 [172.29.21.101]` still down, replaying the failed transactions are successful.

Recovery Steps if the server that is lost, is unrecoverable:

1. A new unified node needs to be deployed. Ensure the server name, IP information and datacentre name is the same as on the server that was lost.
2. Delete the failed node database weight (**database weight del <ip>**), for example **database weight del 172.29.21.101**
3. Run **cluster del 172.29.21.101**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
4. Run **cluster provision** on the cluster *without* the node to be added and then create the new unified node - see: *Create a New VM Using the Platform-Install OVA* and switch on the newly installed node.
5. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
6. If the node will be a unified or web proxy node, run **cluster prepnod** on it.
7. From the primary unified node, run **cluster add <ip>**, with the IP address of the new unified node to add it to the existing cluster.
8. Add database weights so that the weights distributed throughout the cluster
 - Delete all database weights in the cluster. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**
 - Check weights - either individually for each node, or for the cluster by using the command:
cluster run application database weight list
Make sure all application nodes show correct weights.
9. Run **cluster provision primary <IP of current primary>** to join the new unified node to the cluster communications. It is recommended that this step is run in a terminal opened with the **screen** command.
10. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.
11. If an Active/Passive configuration was enabled prior to failover, this should be reconfigured by logging in on the nodes on the DR site and running the command **voss workers 0**.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
 2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
 3. Run **cluster provision** to mitigate the failure. It is recommended that this step is run in a terminal opened with the **screen** command.
 4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.
-

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
 2. Run **database weight del <ip>** from the primary node, <ip> being the IP of the new incoming node.
 3. Log into any secondary node (non primary unified node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
 4. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the new incoming node.
 5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.
-

16.6.6. Scenario: Loss of the Primary Database Server

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers. If non-primary database servers are also lost on the primary or DR site, then also follow the recovery steps for these nodes.

The design is preferably split over 2 physical data centers.

Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a primary database server is experienced. In this scenario AS01 [172.29.42.100] failed while transactions were running.
- Examine the cluster status running **cluster status** to determine the failed state:


```

Data Centre: unknown
  application : unknown_172.29.42.100[172.29.42.100] (not responding)

  webproxy : unknown_172.29.42.100[172.29.42.100] (not responding)

  database : unknown_172.29.42.100[172.29.42.100] (not responding)

Data Centre: jhb
  application : AS02[172.29.42.101]

  webproxy :   PS01[172.29.42.102]
              AS02[172.29.42.101]

  database :   AS02[172.29.42.101]

Data Centre: cpt
  application : AS03[172.29.21.100]
              AS04[172.29.21.101]

  webproxy :   PS02[172.29.21.102]
              AS03[172.29.21.100]
              AS04[172.29.21.101]

  database :   AS03[172.29.21.100]
              AS04[172.29.21.101]

```

- Some downtime occurs. This can be take up to 15 minutes. To speed up recovery, restart the services: **cluster run all app start**.
- The loss of the Primary database server will cause an election and the node with the highest weighting still running will become primary.
- Check the weights set in the cluster configuration: **database weight list**

```

platform@AS01:~$ database weight list
172.29.21.100:
  weight: 10
172.29.21.101:
  weight: 20
172.29.42.100:
  weight: 50
172.29.42.101:
  weight: 40

```

- The primary node 172.29.42.100 failed and therefore node 172.29.42.101 will become the primary node after election.
- To find the primary database, run **database primary**.

```

platform@AS02:~$ database primary
172.29.42.101

```

- At this point *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure

transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the database server AS01 [172.29.42.100] still down, replaying the failed transactions is successful.

Recovery Steps if the server that is lost, is unrecoverable:

Generally, **cluster provision** must be run every time a node is deleted or added, even if it is a replacement node. It is recommended that this step is run in a terminal opened with the **screen** command.

1. Delete its database weight (**database weight del <ip>**), in other words **database weight del 172.29.42.100**
2. Run **cluster del 172.29.42.100**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
3. Run **cluster provision primary 172.29.42.101** from the current primary node. It is recommended that this step is run in a terminal opened with the **screen** command.

This server should already have the highest weight, and its database weight can be checked with **database weight list**

If all the database weights are deleted and provisioning is run again with **cluster provision**, the CLI message is:

'Please select which of the database should be used as the remaining primary by running "database config", selecting a node to sync from (any node that says primary or secondary and is in a good state, i.e. not in a 'RECOVERING' or 'STARTUP' state) and rerun provisioning with "cluster provision primary <db server ip from command above>"

4. A new unified node needs to be deployed. Ensure the server name, IP information and data centre name is the same as on the server that was lost.
5. Run **cluster provision** on the cluster *without* the node to be added and then create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
6. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
7. Run **cluster prepnode** on *all* servers.
8. Run **cluster add <ip>** from the primary unified node (current), with the IP address of the new unified server to add it to the existing cluster.
9. Check the output of the commands: **cluster list** and **cluster status** from the existing node. If the new node does not show up:
 - a. Run **cluster del <new node>**
 - b. Rerun the add of the node on *another* existing unified node, until the node shows up in **cluster list** and **cluster status**.
 - c. Verify that the node shows up from all existing nodes. The recovery process may be time consuming.
10. Delete all database weights in the cluster. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
11. Re-add all database weights in the cluster. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**, considering the following:

- For the new unified node, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.
- If the lost primary unified node release version is 18.1-V4UC-Patch-Bundle-03b and if it will be the new primary, first set its weight lower than the current primary and re-apply the patch on it:

app install media/18.1-V4UC-Patch-Bundle-03b.script –force

When done, check the database weights - either individually for each node, or for the cluster by using the command:

cluster run application database weight list

Make sure all application nodes show correct weights.

12. Make sure the new node is part of the cluster (run **cluster list**) and run **cluster provision primary 172.29.42.101 from the current primary**. It is recommended that this step is run in a terminal opened with the **screen** command.

During the provision process, the role of primary will then be transferred from the current primary to the node with the highest weight. The role transfer may take a significant amount of time, depending on the database size.

During the process, typing **app status** from the new primary node will still show the database as `not provisioned`:

```
mongodb v11.5.3 (2018-07-01 14:35)
|-arbiter           running
|-database         running (not provisioned)
```

To check the progress of the transfer, the database log can be checked. Type **log follow mongodb/mongod/mongod.log**. When the transfer is complete, an entry will show `sync done` as in the example below:

```
2018-07-09T14:09:48.639986+00:00 un1 mongod.27020[129593]: [initial sync-0]
↪initial sync done; took 5821s.
```

While the primary role transfer is in progress, the system can be used, but bulk database operations should not be carried out, because the sync may fall too far behind to complete.

13. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in `STARTUP2` or `SECONDARY` or `PRIMARY` states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure.

4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
2. Delete all database weights. Run **database weight del <ip>** from the primary node, <ip> being the IP of the nodes, including the new incoming node.
3. Log into any secondary node (non primary unified node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Re-add all database weights. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the nodes, including the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.6.7. Scenario: Loss of a Primary Site

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers.

The design is preferably split over 2 physical data centers.

- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names. In this scenario, a portion of the cluster is in Johannesburg and the other is in Cape Town, South Africa:

```
Data Centre: jhb
  application : AS01 [172.29.42.100]
               AS02 [172.29.42.101]

  webproxy   : AS01 [172.29.42.100]
               AS02 [172.29.42.101]
               PS01 [172.29.42.102]

  database   : AS01 [172.29.42.100]
               AS02 [172.29.42.101]

Data Centre: cpt
  application : AS03 [172.29.21.100]
               AS04 [172.29.21.101]

  webproxy   : PS02 [172.29.21.102]
               AS03 [172.29.21.100]
               AS04 [172.29.21.101]
```

(continues on next page)

(continued from previous page)

```

database : AS03 [172.29.21.100]
          AS04 [172.29.21.101]

```

Primary site failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a Primary site is experienced. In this scenario, AS01 [172.29.42.100], AS02 [172.29.42.101] and PS01 [172.29.42.102] failed while transactions were running.
- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- Examine the cluster status by running **cluster status** to determine the failed state:

```

Data Centre: unknown
  application : unknown_172.29.42.100 [172.29.42.100] (not responding)
               unknown_172.29.42.101 [172.29.42.101] (not responding)

  webproxy :   unknown_172.29.42.100 [172.29.42.100] (not responding)
               unknown_172.29.42.101 [172.29.42.101] (not responding)
               unknown_172.29.42.102 [172.29.42.102] (not responding)

  database :   unknown_172.29.42.100 [172.29.42.100] (not responding)
               unknown_172.29.42.101 [172.29.42.101] (not responding)

Data Centre: jhb
  application :

  webproxy :

  database :

Data Centre: cpt
  application : AS03 [172.29.21.100]
               AS04 [172.29.21.101]

  webproxy :   PS02 [172.29.21.102]
               AS03 [172.29.21.100]
               AS04 [172.29.21.101]

  database :   AS03 [172.29.21.100]
               AS04 [172.29.21.101]

```

- The cluster will be not be operational and manual intervention is needed to recover if a continued flow

of transactions is required with a minimum of downtime.

- If it was possible to recover the lost nodes within a reasonable time frame, the cluster will recover automatically if the nodes that were down were brought back into the cluster array successfully.
- To recover the lost nodes and if they are unrecoverable, carry out the following recovery steps.

Recovery Steps (two options):

Commands should be run on an operational unified node from the DR site. During the recovery of clusters, database weights should be deleted and added again.

1. Delete the failed node database weights from the cluster: **database weight del <ip>**
2. Run **cluster del <ip>** to remove the nodes at the failed primary site. Power off the deleted node, or disable its Network Interface Card.
3. At this point, you have two options:
 - a. Option A: provision half the cluster for a faster uptime of your DR site. Only the DR site will then be operational after the provision. You can also optionally add unified nodes to this cluster.
 - b. Option B: bring the full cluster back up at both the DR site and Primary site. You need to redeploy the Primary site nodes.
4. Option A: provision half the cluster or optionally adding 2 more unified nodes to it.
 - a. If you choose to add 2 more unified nodes to optionally create a cluster with 4 unified nodes, deploy the new nodes as follows.
 - i. Run **cluster provision** on the cluster *without* the node to be added and then create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
 - ii. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
 - iii. Run **cluster prepnode** on all new nodes.
 - iv. From a running unified node, run **cluster add <ip>**, with the IP address of the new unified node to add it to the existing cluster.
 - v. Add the database weights nodes in the cluster at the DR site.
 - Delete all database weights in the cluster of the DR site. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster of the DR site. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**, considering the following:
For the new unified node, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.
 - b. Run **cluster provision primary <ip>** (current primary IP) It is recommended that this step is run in a terminal opened with the **screen** command.
 - c. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

- d. Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
 - **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh unifieds)
 - **cluster run application cluster list** (make sure all application nodes show 3 or 5 nodes)
 - **cluster run application database weight list** (make sure all application nodes show correct weights)
5. Option B: bring the full cluster back up at both the DR site and Primary site. You need to redeploy the Primary site nodes.
 - a. Deploy 3 nodes: 2 as unified nodes and 1 as a proxy node. For an 8-node topology, deploy the number of Primary site unified nodes and the web proxy node that were lost.
 - i. Run **cluster provision** on the cluster *without* the node to be added and then create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
 - ii. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
 - iii. Run **cluster prepnode** on all new nodes.
 - iv. Run **cluster add <ip>** from the current primary unified node, with the IP address of the new unified node to add it to the existing cluster.
 - v. Ensure the database weights are added back:
 - Delete all database weights in the cluster. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**, considering the following:
For a new unified node, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.
 - vi. Run **cluster provision primary <ip>** (current primary IP), It is recommended that this step is run in a terminal opened with the **screen** command.
After provisioning, the node with the largest database weight will be the primary server.
 - vii. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.
 - b. Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
 - **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh unifieds)

- **cluster run application cluster list** (make sure all application nodes show 6 nodes - or 8 nodes for an 8-node topology).
 - **cluster run application database weight list** (make sure all application nodes show correct weights)
- c. Run **cluster provision primary <ip>**, where <ip> is *the current primary in the DR site*. It is recommended that this step is run in a terminal opened with the **screen** command. The six node (or eight node) cluster then pulls the data from this <ip> into the new primary server at the Primary site.

After provisioning, the database configuration can then be checked with **database config** to verify the primary node in the Primary site.

16.6.8. Scenario: Loss of a DR Site

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example here is a cluster deployment of 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers.

The design is preferably split over 2 physical data centers.

- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names. In this scenario, a portion of the cluster is in Johannesburg and the other is in Cape Town, South Africa:

```
Data Centre: jhb
  application : AS02 [172.29.42.101]

  webproxy   : PS01 [172.29.42.102]
              AS02 [172.29.42.101]

  database   : AS02 [172.29.42.101]

Data Centre: cpt
  application : AS03 [172.29.21.100]
              AS04 [172.29.21.101]

  webproxy   : PS02 [172.29.21.102]
              AS03 [172.29.21.100]
              AS04 [172.29.21.101]

  database   : AS03 [172.29.21.100]
              AS04 [172.29.21.101]
```

DR site failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a DR site is experienced. In this scenario, AS03 [172.29.21.100], AS04 [172.29.21.101] and PS02 [172.29.21.100] failed while transactions were running.

- At this point, *all* transactions that are currently in flight are lost and will not recover. The lost transactions have to be rerun.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the DR site still down, replaying the failed transactions is successful
- Examine the cluster status by running **cluster status** to determine the failed state:

```
Data Centre: unknown
  application : unknown_172.29.21.100[172.29.21.100] (not responding)
                unknown_172.29.21.101[172.29.21.101] (not responding)

  webproxy   : unknown_172.29.21.100[172.29.21.100] (not responding)
                unknown_172.29.21.101[172.29.21.101] (not responding)
                unknown_172.29.21.102[172.29.21.102] (not responding)

  database   : unknown_172.29.21.100[172.29.21.100] (not responding)
                unknown_172.29.21.101[172.29.21.101] (not responding)

Data Centre: jhb
  application : AS01[172.29.42.100]
                AS02[172.29.42.101]

  webproxy   : PS01[172.29.42.102]
                AS01[172.29.42.100]
                AS02[172.29.42.101]

  database   : AS01[172.29.42.100]
                AS02[172.29.42.101]

Data Centre: cpt
  application :

  webproxy   :

  database   :
```

- The cluster will be operational, but only on the Primary Site.
- You need to recover the lost nodes and if they are unrecoverable. Follow the recovery steps below.

Recovery Steps

1. Remove the database weights of the failed nodes from the cluster: **database weight del <ip>**
2. Run **cluster del <ip>** to remove the failed nodes from the existing half of the cluster. Power off the deleted node, or disable its Network Interface Card.
3. Run **cluster provision primary <ip>** before a new server is added. It is recommended that this step is run in a terminal opened with the **screen** command.

4. Redeploy the failed DR site nodes if the nodes are unrecoverable. Deploy 3 nodes: 2 as unified nodes and 1 as a proxy node. This applies to the DR site of a 6 node deployment or 8 node deployment.
5. Run **cluster provision primary** on the cluster *without* the node to be added and then create the new unified node - see: [Create a New VM Using the Platform-Install OVA](#).
6. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added unified node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
7. If a node will be a unified or web proxy node, run **cluster prepnode** on it.
8. From the primary unified node, after the redeployment, run **cluster add <ip>** with the IP address of the new unified node to add it to the existing cluster. Run **cluster list** to make sure the nodes added in cluster.
9. Add the database weights nodes in the cluster.
 - Delete all database weights in the cluster. On a selected unified node, *for each unified node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each unified node*, for each unified node IP, run **database weight add <IP> <weight>**
10. Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
 - **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh unifieds)
 - **cluster run application cluster list** (make sure all application nodes show 6 nodes - or 8 nodes for an 8-node topology)
 - **cluster run application database weight list** (make sure all application nodes show correct weights)
11. Run **cluster provision primary <ip>** to ensure that a primary is selected for the provisioning stage. It is recommended that this step is run in a terminal opened with the **screen** command.
After provisioning, the database configuration can then be checked with the command **database config**.
12. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.
13. If an Active/Passive configuration was enabled prior to failover, this should be reconfigured by logging in on the nodes on the DR site and running the command **voss workers 0**.

16.6.9. DR Failover and Recovery in a 2 Node Cluster

Important: A 2 node cluster will *not* fail over automatically.

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.

Scenario: Loss of Primary Node

- The administrator deployed the 2-node cluster.

```
$ cluster status

Data Centre: jhb
  application : AS01[172.29.42.100]
                AS02[172.29.42.101]

  webproxy   : AS01[172.29.42.100]
                AS02[172.29.42.101]

  database   : AS01[172.29.42.100]
                AS02[172.29.42.101]
```

Example database weights:

```
$ database weight list
172.29.42.100:
  weight: 20
172.29.42.101:
  weight: 10
```

- **Node Failure:** in the case where the primary node is lost on the Primary site:

```
$ cluster status

Data Centre: unknown
  application : unknown_172.29.248.100[172.29.248.100] (not responding)

  webproxy   : unknown_172.29.248.100[172.29.248.100] (not responding)

  database   : unknown_172.29.248.100[172.29.248.100] (not responding)

Data Centre: jhb
  application : AS02[172.29.248.101]

  webproxy   : AS02[172.29.248.101]

  database   : AS02[172.29.248.101]
```

Recovery Steps

The primary node server is lost.

- A. It is decided to fail over to the secondary node:

1. *On the secondary node*, remove the lost server from the cluster:

```
cluster del 172.29.248.100
```

2. *On the secondary node*, run **cluster provision** (it is recommended that this step is run in a terminal opened with the **screen** command). See: [Using the screen command](#).

On the secondary node, check:

```
$ cluster status

Data Centre: jhb
              application : AS02[172.29.248.101]

              webproxy   : AS02[172.29.248.101]

              database   : AS02[172.29.248.101]
```

- B. It is decided to recover the primary node:

1. *On the secondary node*, remove the lost server from the cluster:

```
cluster del 172.29.248.100
```

2. *On the secondary node*, run **cluster provision** (it is recommended that this step is run in a terminal opened with the **screen** command).

On the secondary node, check:

```
$ cluster status

Data Centre: jhb
              application : AS02[172.29.248.101]

              webproxy   : AS02[172.29.248.101]

              database   : AS02[172.29.248.101]
```

3. Switch on the newly installed server.

On the secondary node, add the server. Run **cluster add 172.29.42.100**.

On either node, check:

```
$ cluster status

Data Centre: jhb
              application : AS01[172.29.42.100]
                          AS02[172.29.42.101]

              webproxy   : AS01[172.29.42.100]
                          AS02[172.29.42.101]

              database   : AS01[172.29.42.100]
                          AS02[172.29.42.101]
```

4. Configure the primary database. *On the newly installed server*, run **cluster provision primary 172.29.42.100** (it is recommended that this step is run in a terminal opened with the **screen** command).

Check database configuration on both nodes, for example:

```

$ database config
  date:
    $date: 1549450382862
  heartbeatIntervalMillis: 2000
  members:
    172.29.42.100:27020:
      priority: 20.0
      stateStr: PRIMARY
      storageEngine: WiredTiger
    172.29.42.100:27030:
      priority: 1.0
      stateStr: ARBITER
      storageEngine: Unknown
    172.29.42.101:27020:
      priority: 10.0
      stateStr: SECONDARY
      storageEngine: WiredTiger
  myState: 1
  ok: 1.0
  set: DEVICEAPI
  term: 8

```

16.6.10. Scenario: Loss of Full Cluster

Background

- The administrator deployed a single-node cluster or deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 6 nodes, where 4 nodes are database servers and 2 nodes are proxy servers.

However, this scenario also applies to a cluster deployment of 8 nodes: 6 database servers and 2 proxy servers.

The design is preferably split over 2 physical data centers.

- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names.

Full cluster failure

- In this scenario, *all* nodes failed while transactions were running.
- At this point, *all* transactions that were in flight are lost and will not recover.
- The lost transactions have to be rerun.
- The cluster will not be operational and manual intervention is needed to recover.
- To recover the cluster, carry out the Recovery Steps.

Recovery Steps

Important:

- Prerequisite: a system backup exported to a remote backup location. The backup file on the remote location would typically have a format *<timestamp>.tar.gz*. This recovery procedure will *only* succeed if you have a valid recent backup to restore.
- For details, considerations and specific commands at each step below, refer to the “Standalone (single-node cluster) Installation” or “Multinode Installation” topic in the *Installation Guide*.

1. Ensure all traces of the previous nodes have been removed from the VMware environment.
2. Deploy fresh nodes as per the original topology.
 - Check topologies and hardware requirements in the *Installation Guide*.
 - For new node deployment, see: [Create a New VM Using the Platform-Install OVA](#).
 - For the steps below, follow either the “Standalone (single-node cluster) Installation” or “Multinode Installation” topic in the *Installation Guide*:
3. Add each non-primary node to the cluster by running **cluster prepnode**.
4. From the primary node, add each node to the cluster using the **cluster add <IP address of node>** command.
5. For multi-node clusters:
 - a. On the primary node, set the database weights for each database node using the **database weight add <IP address of node> <weight>** command.
 - b. Restore a backup made from the highest weighted secondary database node in the original cluster.

Follow the Import steps here: [Backup and Import to a New Environment](#).

Note: It is not necessary to run **cluster provision** again on the primary node. This action is included in the backup restore process.
6. Ensure all services are up and running:

Run **cluster run all app status** to check if all the services are up and running after the restore completes.

Note: For multi-node clusters:

Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure.
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: For multi-node clusters:

If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
2. Delete all database weights. Run **database weight del <ip>** from the primary node, <ip> being the IP of the nodes, including the new incoming node.
3. Log into any secondary node (non primary unified node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Re-add all database weights. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the nodes, including the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.7. DR Failover and Recovery in a Modular Cluster Topology

16.7.1. Modular Cluster: Election of a New Primary and Failover

In the case where nodes fail, the system follows a failover procedure. For details on the failover and DR process, refer to the topics in the Platform Guide.

If the primary database is lost, the failover process involves the election of a new primary database by the remaining database nodes. Each node in a cluster is allocated a number of votes that are used in the failover election of a new primary database - the election of a running node with the highest database weight.

The database weights for a node can be seen as the `priority` value when running the **database config** command.

Note: The database weight of a node does not necessarily match its number of votes.

For voting on a modular system, arbiters are added to the database nodes (N1 and N2) on the primary site, but the secondary site database node (N3) arbiter will be unused.

```
members:
  192.168.100.4:27020:
    priority: 30.0
    stateStr: SECONDARY
    storageEngine: WiredTiger
  192.168.100.4:27030:
    priority: 30.0
    stateStr: ARBITER
    storageEngine: WiredTiger
  192.168.100.6:27020:
    priority: 40.0
    stateStr: PRIMARY
```

(continues on next page)

(continued from previous page)

```

storageEngine: WiredTiger
192.168.100.6:27030:
  priority: 40.0
  stateStr: ARBITER
  storageEngine: WiredTiger
192.168.100.8:27020:
  priority: 10.0
  stateStr: SECONDARY
  storageEngine: WiredTiger

```

The maximum number of votes in a cluster should not exceed 5 and arbiter votes are added to nodes to provide a total of 5 votes.

The tables below show the system status and failover for a selection of scenarios for a 6 node cluster. Also refer to the topics on the specific DR scenarios. The abbreviations used are as follows:

- Pri : Primary site
- DR : DR site
- N : node. Primary node is N1, secondary node is N2.
- w : database weight
- v : vote
- a : arbiter vote
- For example, for a 6 node cluster with 3 database nodes and 2 sites, initial votes per node are as follows:

Primary database nodes (N1, N2): 2 (each 1 + 1 arbiter) Secondary database node (N3): 1 (no arbiter vote)

Pri N1 w:40 v:1 a:1	Pri N2 w:30 v:1 a:1	DR N3 w:10 v:1	Votes	System Status under scenario
Up	Up	Up	5	System is functioning normally.
Up	Up	Down	4	Scenario: Loss of a Non-primary Server in the DR Site. System continues functioning normally.
Up	Up	Up	5	Scenario: Loss of an Application Server. System continues functioning normally. Some transactions may hang.
Up	Down	Up	3	Scenario: Loss of a Non-primary Node in the Primary Site. System continues functioning normally.

Pri N1 w:40 v:1 a:1	Pri N2 w:30 v:1 a:1	DR N3 w:10 v:1	Votes	System Status under scenario
Down	Up	Up	3	Scenario: Loss of the Primary Database Server. Some downtime occurs. System automatically fails over to N2.
Down	Down	Up	1	Scenario: Loss of a Primary Site. Manual recovery required.
Up	Down	Down	2	Scenario: Loss of all secondary nodes. Manual recovery required.
Down	Up	Down	2	Scenario: Loss of all Primary and DR nodes. Manual recovery required.

16.7.2. DR Failover and Recovery Scenarios in a Modular Cluster

A number of failover scenarios and recovery steps are shown. In each case, a node topology is assumed: 8 node clusters in 2 sites - primary and Disaster Recovery (DR).

- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

A node failure scenario is indicated and a set of recovery steps are provided.

The following scenarios that are covered:

- Power Off and On of a Node in a Modular Cluster
- Loss of an app node: Modular Cluster
- Loss of the Primary Database Server in a Modular Cluster
- Loss of a non-primary database: Modular Cluster
- Loss of a Primary Site in a Modular Cluster
- Loss of a DR Site in a Modular Cluster
- Loss of Full Cluster in a Modular Cluster

Background

For the scenarios below, the following procedures and definitions apply:

- In the event of a network failure or a temporary network outage affecting a single a node, the node will be inaccessible and the cluster will respond in the same way as if the node had failed. If network connectivity is then restored, no action is required, because the node will again start communicating with the other nodes in the cluster, provided no changes were made to that node during the outage window.
- In a clustered deployment, the data center would typically be two different data centers, for example “Virginia” and “Seattle”. These can be thought of as a primary site and a DR (Disaster Recovery) site in case of a failure in the primary site. These two data centers can exist on the same physical hardware, so the separation of the cluster is into two sets of three nodes.

When data centers are defined during installation, the nodes of a cluster may or may not be in the same physical location. The cluster is designed to communicate across all nodes, regardless of their physical location.

- During recovery, the command **cluster provision** must be run every time a node is deleted from or added to a cluster, even if it is a replacement node. It is recommended that this step is run in a terminal opened with the **screen** command. See: [Using the screen command](#).
- During recovery and installation, the command **cluster prenode** must be run on every node.
- During recovery of 8 node clusters, database weights should be deleted and added again.

16.7.3. Scenario: Power Off and On of a Node in a Modular Cluster

The scenario and recovery steps apply to database, application and Proxy nodes.

Node powered off

- Secondary database node assumes primary
- There is no cluster downtime and normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a node is powered off.
- At this point, *all* transactions that are currently in flight at the node are lost and will not recover. The lost transactions have to be rerun.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on an application node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail. With the node still powered off, replaying the failed transactions is successful

Recovery steps if the node is powered off:

1. Power up the node. The node resyncs.

For a database node, run the **database config** command to verify the state of the database members. A typical output of the command would be:

```
$ database config
  date: 2017-04-25T09:50:34Z
  heartbeatIntervalMillis: 2000
  members:
    172.29.21.41:27020:
      priority: 60.0
      stateStr: PRIMARY
      storageEngine: WiredTiger
    172.29.21.41:27030:
      priority: 1.0
      stateStr: ARBITER
      storageEngine: WiredTiger
    172.29.21.42:27020:
      priority: 50.0
      stateStr: SECONDARY
      storageEngine: WiredTiger
```

(continues on next page)

(continued from previous page)

```

172.29.21.43:27020:
  priority: 40.0
  stateStr: SECONDARY
  storageEngine: WiredTiger
172.29.21.44:27020:
  priority: 30.0
  stateStr: SECONDARY
  storageEngine: WiredTiger
172.29.21.45:27020:
  priority: 20.0
  stateStr: SECONDARY
  storageEngine: WiredTiger
172.29.21.46:27020:
  priority: 10.0
  stateStr: SECONDARY
  storageEngine: WiredTiger
myState: 1
ok: 1.0
set: DEVICEAPI
term: 38

```

Note that `storageEngine` will show as `WiredTiger` after the database engine upgrade to Wired Tiger when upgrading to VOSS-4-UC 17.4. Otherwise, the value is `MMAPv1`.

In other words, the database should not for example be any of: `STARTUP`, `STARTUP2` or `RECOVERING`. Note however that it is sometimes expected that nodes are recovering or in startup, but then should change to a normal state after a period of time (depending on how far out of sync those members are).

A file system check may take place.

2. If a replacement node is not on standby, rebuild steps such as boot up, adding to cluster, setting database weight and reprovisioning may take 200-300 minutes, depending on hardware specifications.

It is recommended that standby nodes are available to be used for faster recovery.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in `STARTUP2` or `SECONDARY` or `PRIMARY` states with correct arbiter placement.
 2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
 3. Run **cluster provision** to mitigate the failure (it is recommended that this step is run in a terminal opened with the **screen** command). See: [Using the screen command](#).
 4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.
-

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary database node, <ip> being the IP of the new incoming node.
2. For database nodes, run **database weight del <ip>** from the primary database node, <ip> being the IP of the new incoming node.

3. Log into primary database node and run **cluster add <ip>** ,<ip> being the IP of the new incoming node.
4. For database nodes, run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.7.4. Scenario: Loss of an app node: Modular Cluster

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

The design is preferably split over 2 physical data centers.

Application Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a app node is experienced. In this 8-node example, AS02 [172.29.42.101] failed while transactions were running.
- Examine the cluster status running **cluster status** to determine the failed state:

```
Data Centre: unknown
      application : unknown_172.29.42.101[172.29.42.101] (not responding)

Data Centre: jhb
      application : AS01[172.29.42.100]
      webproxy   : PS01[172.29.42.102]
      database   : DB01[172.29.42.103]
                  DB02[172.29.42.104]

Data Centre: cpt
      application : AS03[172.29.21.100]
      webproxy   : PS02[172.29.21.102]
      database   : DB03[172.29.21.101]
```

- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on an application node in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

Recovery steps

If the server that is lost, is unrecoverable:

1. A new app node needs to be deployed. Ensure the server name, IP information and data centre name is the same as on the server that was lost.
2. Run **cluster del <IP of lost app node>**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
3. Run **cluster provision** on the cluster *without* the node to be added and then create the *new app node* at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#).
4. Switch on the newly installed server.
5. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added app node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
6. Run **cluster prepnode** on the new app node.
7. From the primary database node, run **cluster add <ip>**, with the IP address of the new app node to add it to the existing cluster.
8. From the primary database node, run **cluster provision** to join the new app node to the cluster communications. It is recommended that this step is run in a terminal opened with the **screen** command.
9. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

Important: Re-apply any patches and services (for example, Phone Based Registration) to this node that were added after the initial Delta Bundle upgrade.

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

10. On the new app node, check the number of queues using **voss queues** and if the number is *less than* 2, set the queues to 2 with **voss queues 2**.

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

11. If the app node was replaced on the DR site and an Active/Passive configuration was enabled prior to failover, this should be reconfigured by logging in on the nodes on the DR site and running the command **voss workers 0**.

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.

2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure. It is recommended that this step is run in a terminal opened with the **screen** command.
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary database node, <ip> being the IP of the new incoming node.
2. Log into any other node (not the new node) and run **cluster add <ip>**, <ip> being the IP of the new incoming node.

16.7.5. Scenario: Loss of the Primary Database Server in a Modular Cluster

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

The design is preferably split over 2 physical data centers.

Database Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a primary database server is experienced. In this scenario DB01 [172.29.42.103] failed while transactions were running.
- Examine the cluster status running **cluster status** to determine the failed state:

```
Data Centre: unknown
    database : unknown_172.29.42.103[172.29.42.103] (not responding)

Data Centre: jhb
    application : AS01[172.29.42.100]
                AS02[172.29.42.101]

    webproxy   : PS01[172.29.42.102]

    database   : DB02[172.29.42.104]

Data Centre: cpt
    application : AS03[172.29.21.100]

    webproxy   : PS02[172.29.21.102]

    database   : DB03[172.29.21.101]
```

- Some downtime occurs. This can be take up to 15 minutes. To speed up recovery, restart the services: **cluster run all app start**.
- The loss of the primary database server will cause an election and the database node with the highest weighting still running will become primary.
- Check the weights set in the cluster configuration: **database weight list**

```
[ new example for modular needed ]
platform@AS01:~$ database weight list
 172.29.21.101:
   weight: 20
 172.29.42.103:
   weight: 50
 172.29.42.104:
   weight: 40
```

- The primary database node `172.29.42.103` failed and therefore node `172.29.42.104` will become the primary database node after election.
- To find the primary database, run **database primary**.

```
platform@AS02:~$ database primary
172.29.42.104
```

- At this point *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the database server `DB01 [172.29.42.103]` still down, replaying the failed transactions is successful.

Recovery Steps

If the server that is lost, is unrecoverable:

Generally, **cluster provision** must be run every time a node is deleted or added, even if it is a replacement node. It is recommended that this step is run in a terminal opened with the **screen** command.

1. Delete its database weight (**database weight del <ip>**), in other words **database weight del 172.29.42.103**
2. Run **cluster del 172.29.42.103**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
3. *Only* run **cluster provision primary 172.29.42.104** from the current primary database node if **database config** shows *no* primary. Else only run **cluster provision** from the current primary node. It is recommended that this step is run in a terminal opened with the **screen** command.

This server should already have the highest weight, and its database weight can be checked with **database weight list**

If all the database weights are deleted and provisioning is run again with **cluster provision**, the CLI message is:

‘Please select which of the database should be used as the remaining primary by running “database config”, selecting a node to sync from (any node that says primary or secondary and is in a good state, i.e. not in a ‘RECOVERING’ or ‘STARTUP’ state) and rerun provisioning with “cluster provision primary <db server ip from command above>”

4. A new database node needs to be deployed. Ensure the server name, IP information and data centre name is the same as on the server that was lost.
5. Run **cluster provision** on the cluster *without* the node to be added.

Create the *new database node* at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#).

6. Run **cluster prepnode** on *all* servers.
7. Run **cluster add <ip>** from an existing node, with the IP address of the new database server to add it to the existing cluster.
8. Check the output of the commands: **cluster list** and **cluster status** from the existing node. If the new node does not show up:
 - a. Run **cluster del <new node>**
 - b. Rerun the add of the node on *another* node, until the node shows up in **cluster list** and **cluster status**.
 - c. Verify that the node shows up from all existing nodes. The recovery process may be time consuming.
9. Delete all database weights in the cluster. On a selected database node, *for each database node IP*, run **database weight del <IP>**.
10. Re-add all database weights in the cluster. *On each database node*, for each database node IP, run **database weight add <IP> <weight>**, considering the following:
 - *For the new database node*, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.

When done, check the database weights - either individually for each node, or for the cluster by using the command:

cluster run application database weight list

Make sure all database nodes show correct weights.

11. Make sure the new node is part of the cluster (run **cluster list**) and run **cluster provision primary 172.29.42.104** *from the current primary* (where 172.29.42.104 is an example). It is recommended that this step is run in a terminal opened with the **screen** command.

During the provision process, the role of primary will then be transferred from the current primary to the node with the highest weight. The role transfer may take a significant amount of time, depending on the database size.

During the process, typing **app status** from the new primary node will still show the database as `not provisioned`:

```
mongodb v21.1.1 (2021-05-09 13:36)
|-arbiter           running
|-database         running (not provisioned)
```

To check the progress of the transfer, the database log can be checked. Type **log follow mongod/mongod/mongod.log**. When the transfer is complete, an entry will show `sync done` as in the example below:


```
2021-05-10T14:09:48.639986+00:00 un1 mongod.27020[129593]: [initial sync-0]
↪initial sync done; took 5821s.
```

While the primary role transfer is in progress, the system can be used, but bulk database operations should not be carried out, because the sync may fall too far behind to complete.

12. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure.
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
2. Delete all database weights. Run **database weight del <ip>** from the primary database node, <ip> being the IP of the nodes, including the new incoming node.
3. Log into a non primary database node and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Re-add all database weights. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the nodes, including the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.7.6. Scenario: Loss of a non-primary database: Modular Cluster

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

The design is preferably split over 2 physical data centers.

Database Node Failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a non-primary node is experienced.

In this example, DB02 [172.29.42.104] failed while transactions were running.

- Examine the cluster status running **cluster status** to determine the failed state:

```
Data Centre: unknown
    database : unknown_172.29.42.104[172.29.42.104] (not responding)

Data Centre: jhb
    application : AS01[172.29.42.100]
                  AS02[172.29.42.101]

    webproxy   : PS01[172.29.42.102]

    database   : DB01[172.29.42.103]

Data Centre: cpt
    application : AS03[172.29.21.100]

    webproxy   : PS02[172.29.21.102]

    database   : DB03[172.29.21.101]
```

- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on an application node in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the database server DB02 [172.29.42.104] still down, replaying the failed transactions are successful.

Recovery steps

If the server that is lost, is unrecoverable:

1. A new database node needs to be deployed. Ensure the server name, IP information and data center name is the same as on the server that was lost.
2. Delete the failed node database weight (**database weight del <ip>**), for example **database weight del 172.29.42.104**
3. Run **cluster del 172.29.42.104**, because this server no longer exists. Power off the deleted node, or disable its Network Interface Card.
4. Run **cluster provision** on the cluster *without* the node to be added.

Then create the *new database node* at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#) and switch on the newly installed node.

5. From the primary database node, run **cluster add <ip>**, with the IP address of the new database node to add it to the existing cluster.
6. Add database weights so that the weights distributed throughout the cluster
 - Delete all database weights in the cluster. On the primary database node, *for each database node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each database node*, for each database node IP, run **database weight add <IP> <weight>**
 - Check weights - either individually for each node, or for the cluster by using the command:
cluster run application database weight list
7. Run **cluster provision** to join the new database node to the cluster communications. It is recommended that this step is run in a terminal opened with the **screen** command.
8. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new database node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

Note: Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:

1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
3. Run **cluster provision** to mitigate the failure. It is recommended that this step is run in a terminal opened with the **screen** command.
4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.

Note: If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:

1. Run **cluster del <ip>** from the primary database node, <ip> being the IP of the new incoming node.
2. Run **database weight del <ip>** from the primary database node, <ip> being the IP of the new incoming node.
3. Log into any non primary database node and run **cluster add <ip>**, <ip> being the IP of the new incoming node.
4. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the new incoming node.
5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.

16.7.7. Scenario: Loss of a Primary Site in a Modular Cluster

- The administrator deployed the cluster into a primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

The design is preferably split over 2 physical data centers.

- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names. In this scenario, a portion of the cluster is in Johannesburg and the other is in Cape Town, South Africa.

Primary site failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a primary site is experienced. In this scenario, the following nodes failed while transactions were running:

- AS01[172.29.42.100]
- AS02[172.29.42.101]
- PS01[172.29.42.102]
- DB01[172.29.42.103]
- DB02[172.29.42.104]

- At this point, *all* transactions that are currently in flight are lost and will not recover.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- Examine the cluster status by running **cluster status** to determine the failed state:

```
Data Centre: unknown
  application : unknown_172.29.42.100 [172.29.42.100] (not responding)
               unknown_172.29.42.101 [172.29.42.101] (not responding)

  webproxy   : unknown_172.29.42.102 [172.29.42.102] (not responding)

  database   : unknown_172.29.42.103 [172.29.42.104] (not responding)
               unknown_172.29.42.103 [172.29.42.103] (not responding)

Data Centre: jhb
  application :
```

(continues on next page)

(continued from previous page)

```

webproxy :
database :
Data Centre: cpt
application : AS03[172.29.21.100]
webproxy : PS02[172.29.21.102]
database : DB03[172.29.21.101]

```

- The cluster will not be operational and manual intervention is needed to recover if a continued flow of transactions is required with a minimum of downtime.
- If it was possible to recover the lost nodes within a reasonable time frame, the cluster will recover automatically if the nodes that were down were brought back into the cluster array successfully.
- To recover the lost nodes and if they are unrecoverable, carry out the following recovery steps.

Recovery Steps (two options):

Commands should be run on an operational unified node from the DR site. During the recovery of clusters, database weights should be deleted and added again.

1. Delete the failed node database weights from the cluster: **database weight del <ip>**
2. Run **cluster del <ip>** to remove the nodes at the failed primary site. Power off the deleted node, or disable its Network Interface Card.
3. At this point, you have two options:
 - a. Option A: provision half the cluster for a faster uptime of your DR site. Only the DR site will then be operational after the provision. You can also optionally add nodes to this cluster.
 - b. Option B: bring the full cluster back up at both the DR site and primary site. You need to redeploy the primary site nodes.
4. Option A: provision half the cluster or optionally adding 2 more nodes to it.
 - a. If you choose to add 2 more nodes to optionally create a cluster with 2 application and 2 database nodes, deploy the new nodes as follows.
 - i. Run **cluster provision** on the cluster *without* the node to be added and then create the *new application and database nodes* at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#).
 - ii. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added application node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
 - iii. Run **cluster prepnod** on all new nodes.
 - iv. From a running database node, run **cluster add <ip>**, with the IP address of each new node to add it to the existing cluster.
 - v. Add the database weights nodes in the cluster at the DR site.
 - Delete all database weights in the cluster of the DR site. On a selected database node, for each database node IP, run **database weight del <IP>**.

- Re-add all database weights in the cluster of the DR site. *On each database node, for each database node IP, run **database weight add <IP> <weight>**, considering the following:*

For the new database node, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.

- Run **database config** to determine if you have a primary database. If not, run **cluster provision primary <ip>** (current primary IP) It is recommended that this step is run in a terminal opened with the **screen** command. If you do have a primary database, only run **cluster provision**.
- If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

- Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
 - **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh database nodes)
 - **cluster run application cluster list** (make sure all nodes show)
 - **cluster run application database weight list** (make sure all database nodes show correct weights)
- Option B: bring the full cluster back up at both the DR site and primary site. You need to redeploy the primary site nodes.
 - Deploy 5 nodes: 2 database nodes, 2 application nodes and 1 proxy node.
 - Run **cluster provision** on the cluster *without* the node to be added and then create the new application, proxy and database nodes at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#).
 - An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added application node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
 - Run **cluster prepnode** on all new nodes.
 - Run **cluster add <ip>** from the current primary database node, with the IP address of each new node to add it to the existing cluster.
 - Ensure the database weights are added back:
 - Delete all database weights in the cluster. On a selected database node, *for each database node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each database node, for each database node IP, run **database weight add <IP> <weight>**, considering the following:*
For a new database node, add a database weight lower than that of the weight of the current primary if this will be a secondary, or higher if this will be the new primary.

- vi. Since the primary database node is newly added, run **cluster provision primary <ip>** (current primary IP), It is recommended that this step is run in a terminal opened with the **screen** command.

After provisioning, the node with the largest database weight will be the primary server.

- vii. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.

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

See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.

- b. Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
- **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh database nodes)
 - **cluster run application cluster list** (make sure all nodes show)
 - **cluster run application database weight list** (make sure all database nodes show correct weights)
- c. Run **cluster provision primary <ip>**, where <ip> is *the current primary database in the DR site*. It is recommended that this step is run in a terminal opened with the **screen** command. The six node (or eight node) cluster then pulls the data from this <ip> into the new primary database server at the primary site.

After provisioning, the database configuration can then be checked with **database config** to verify the primary database node in the primary site.

- d. On the new app nodes, check the number of queues using **voss queues** and if the number is *less than 2*, set the queues to 2 with **voss queues 2**.

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

16.7.8. Scenario: Loss of a DR Site in a Modular Cluster

- The administrator deployed the cluster into a Primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.

The design is preferably split over 2 physical data centers.

- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names. In this scenario, a portion of the cluster is in `jhb` and the other is in `cpt`.

DR site failure

- Normal operations continue where the cluster is processing requests and transactions are committed successfully up to the point where a loss of a DR site is experienced. In this scenario, the following nodes failed while transactions were running:

- AS03[172.29.21.100]
- PS02[172.29.21.101]
- DB03[172.29.21.102]

- At this point, *all* transactions that are currently in flight are lost and will not recover. The lost transactions have to be rerun.
- The lost transactions have to be replayed or rerun.

Bulk load transactions cannot be replayed and have to be rerun. Before resubmitting a failed Bulk load job, carry out the following command on the primary node CLI in order to manually clear each failure transaction that still has a Processing status *after a service restart*. Use the command:

voss finalize_transaction <Trans ID>

The failed transaction status then changes from Processing to Fail.

- With the DR site still down, replaying the failed transactions is successful
- Examine the cluster status by running **cluster status** to determine the failed state:

```
Data Centre: unknown
  application : unknown_172.29.21.100[172.29.21.100] (not responding)

  webproxy   :   unknown_172.29.21.101[172.29.21.101] (not responding)

  database   :   unknown_172.29.21.102[172.29.21.102] (not responding)

Data Centre: jhb
  application : AS01[172.29.42.100]
                AS02[172.29.42.101]

  webproxy   :   PS01[172.29.42.102]

  database   :   DB01[172.29.42.103]
                DB02[172.29.42.104]

Data Centre: cpt
  application :

  webproxy   :

  database   :
```

- The cluster will be operational, but only on the primary site.
- You need to recover the lost nodes and if they are unrecoverable. Follow the recovery steps below.

Recovery Steps

1. Remove the database weights of the failed database nodes from the cluster: **database weight del <ip>**

2. Run **cluster del <ip>** to remove the failed nodes from the existing half of the cluster. Power off the deleted node, or disable its Network Interface Card.
3. Run **cluster provision** from the primary database node before a new server is added. It is recommended that this step is run in a terminal opened with the **screen** command.
4. Redeploy the failed DR site nodes if the nodes are unrecoverable. Deploy 3 nodes: 1 application, 1 database and 1 proxy node.
5. Run **cluster provision** from the primary database node on the cluster *without* the node to be added and then create the new node at the *required data center* - see: [Create a New VM Using the Platform-Install OVA](#).
6. An extra functions file (`extra_functions.py`) that is installed on the existing cluster needs to be re-installed *on each added application node*. Request the `Macro_Update_<version>.template` file from VOSS Level 2 support and run the command **app template Macro_Update_<version>.template**.
7. Run **cluster prepnode** on all nodes.
8. From the primary database node, after the redeployment, run **cluster add <ip>** with the IP address of the new nodes to add them to the existing cluster. Run **cluster list** to make sure the nodes added in cluster.
9. Add the database weights nodes in the cluster.
 - Delete all database weights in the cluster. On a selected database node, *for each database node IP*, run **database weight del <IP>**.
 - Re-add all database weights in the cluster. *On each database node*, for each database node IP, run **database weight add <IP> <weight>**
10. Check all services, nodes and weights - either individually for each node, or for the cluster by using the commands:
 - **cluster run all app status** (make sure no services are stopped/broken - the message 'suspended waiting for mongo' is normal on the fresh database nodes)
 - **cluster run application cluster list** (make sure all nodes show)
 - **cluster run application database weight list** (make sure all database nodes show correct weights)
11. Run **cluster provision** on the primary database node to ensure that a primary is selected for the provisioning stage. It is recommended that this step is run in a terminal opened with the **screen** command.
 After provisioning, the database configuration can then be checked with the command **database config**.
12. If an OVA file was not available for your current release and you used the most recent release OVA for which there is an upgrade path to your release to create the new unified node, *re-apply* the Delta Bundle upgrade to the cluster.
 Note that the new node version mismatch in the cluster can be ignored, since this upgrade step aligns the versions.
 See the "Upgrade" step in the "Upgrade a Multinode Environment with the Delta Bundle" topic of the Upgrade Guide with Delta Bundle.
13. If an Active/Passive configuration was enabled prior to failover, this should be reconfigured by logging in on the *application* nodes on the DR site and running the command **voss workers 0**.

14. On the new app nodes, check the number of queues using **voss queues** and if the number is *less than 2*, set the queues to 2 with **voss queues 2**.

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

16.7.9. Scenario: Loss of Full Cluster in a Modular Cluster

Background

- The administrator deployed the cluster into a primary and DR site.
- The cluster is deployed following the Installation Guide.
- The example is a typical cluster deployment: 8 nodes, where 3 nodes are database servers, 3 nodes are application nodes and 2 nodes are proxy servers.
The design is preferably split over 2 physical data centers.
- The cluster might also be in two geographically dispersed areas. The cluster has to be installed in two different site names or data center names.

Full cluster failure

- In this scenario, *all* nodes failed while transactions were running.
- At this point, *all* transactions that were in flight are lost and will not recover.
- The lost transactions have to be rerun.
- The cluster will not be operational and manual intervention is needed to recover.
- To recover the cluster, carry out the Recovery Steps.

Recovery Steps

Important:

- Prerequisite: a system backup exported to a remote backup location. The backup file on the remote location would typically have a format `<timestamp>.tar.gz`. This recovery procedure will *only* succeed if you have a valid recent backup to restore.
 - For details, considerations and specific commands at each step below, refer to the “Modular Cluster Multinode Installation” topic in the *Installation Guide*.
-

1. Ensure all traces of the previous nodes have been removed from the VMware environment.
2. Deploy fresh nodes as per the original topology.
 - Check topologies and hardware requirements in the *Installation Guide*.
 - For new *node type* deployment at the *required data center*, see: [Create a New VM Using the Platform-Install OVA](#).
 - For the steps below, follow the “Modular Cluster Installation” topics in the *Installation Guide*:
3. Add each node to the cluster by running **cluster prenode**.

4. From the primary database node, add each node to the cluster using the **cluster add <IP address of node>** command.
5. On the primary database node, set the database weights for each database node using the **database weight add <IP address of node> <weight>** command.
6. Restore a backup made from the highest weighted secondary database node in the original cluster.
Follow the Import steps here: [Backup and Import to a New Environment](#).
Note: It is not necessary to run **cluster provision** again on the primary node. This action is included in the backup restore process.
7. On the new app nodes, check the number of queues using **voss queues** and if the number is *less than 2*, set the queues to 2 with **voss queues 2**.

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

8. Ensure all services are up and running:
Run **cluster run all app status** to check if all the services are up and running after the restore completes.

Note:

- Upon cluster provision failure at any of the proxy nodes during provisioning, the following steps illustrate the cluster provisioning:
 1. Run **database config** and check if nodes are either in STARTUP2 or SECONDARY or PRIMARY states with correct arbiter placement.
 2. Login to web proxy on both primary and secondary site and add a web weight using **web weight add <ip>:443 1** for all those nodes that you want to provide a web weight of 1 on the respective proxies.
 3. Run **cluster provision** to mitigate the failure.
 4. Run **cluster run all app status** to check if all the services are up and running after cluster provisioning completes.
 - If the existing nodes in the cluster do not see the new incoming cluster after **cluster add**, try the following steps:
 1. Run **cluster del <ip>** from the primary node, <ip> being the IP of the new incoming node.
 2. Delete all database weights. Run **database weight del <ip>** from the primary node, <ip> being the IP of the nodes, including the new incoming node.
 3. Log into any secondary node (non primary unified node) and run **cluster add <ip>** ,<ip> being the IP of the new incoming node.
 4. Re-add all database weights. Run **database weight add <ip> <weight>** from the same session, <ip> being the IP of the nodes, including the new incoming node.
 5. Use **cluster run database cluster list** to check if all nodes see the new incoming nodes inside the cluster.
-

17. Troubleshooting

17.1. Platform User Password Recovery Procedure

The steps below describe how to reset the VOSS-4-UC platform user password if you forget the password and you are not able to access the CLI via platform user.

1. Log in to VMWare and choose the VOSS-4-UC Virtual Machine (VM).
2. Right-click the VM and choose **Edit Settings**.
3. Disconnect the network adapter by un-checking the **Connected** check box. This ensures that transactions are not lost.
4. Click the **VM Options** tab, **During the next boot, force entry into the BIOS setup screen** check-box is checked.
5. Click the **OK** button to apply the settings.
6. Open the VOSS-4-UC display (**Launch Virtual Machine Console**).
7. Under the **Power** menu option, click the reboot button (**Restart Guest**).
8. In the VM console, press F10 and **YES** to exit the BIOS (Do not make any changes in the BIOS). The next step needs to be performed quickly before the VOSS-4-UC system boots.
9. While the cursor highlights the first GRUB console entry (Ubuntu), press e.
10. Navigate to the second to the last line which starts with `linux` and ends with `fsck.repair=yes`.
11. Navigate to the end of the line after `=yes`, add a space and add `init=/bin/bash`.
12. Press Ctrl-x in order to boot the system.
13. When the system has booted, on the console at the `root@(none) :/#` prompt, enter commands as follows:
mount -o remount,rw /
passwd platform to type in and confirm a new password. Check for the success message.
mount -o rw /var/log to allow counters to be reset.
/sbin/pam_tally2 --reset --user platform to reset the failed password attempt counter.
sync to force a file system sync.
exit to exit the console.
14. Reconnect the VM network adapter under the **Edit Settings** option.
15. Power off the VM and power on the VM again.

16. When the system boots, choose the default highlighted GRUB entry (not recovery mode).
17. Allow the disk checks to complete if they do run.
18. You can now log in as platform user with the password set above.

17.2. 'No Space Left on Device' Error

You receive the following error message while backing up or restoring VOSS-4-UC on a virtual machine: 'No Space Left on Device.' You can create a new virtual disk on the node with the primary database and then reassign the VOSS-4-UC data to the new disk. The new disk has enough space for you to perform the backup or restore operation.

Important: If you wish to revert to a *smaller* disk size or back to your original disk size after following the steps below, contact VOSS support.

1. In VMware, add a disk on the node that contains the primary database:
 - a. From the VM menu, click **Edit Settings**.
 - b. Click **Add**. The Add Hardware Wizard opens.
 - c. Select Hard Disk and then click Next.
 - d. Select Create a new virtual disk and then click Next.
 - e. Set the capacity to be the same as the database disk: 250 GB.
 - f. Accept the default file name and location, or click **Browse** to select a different location.
 - g. Click **Finish**.

Your guest operating system recognizes the new virtual disk as a new, blank hard disk.

Note: If the steps above do not succeed on your version of VMWare, first turn off the virtual machine that contains the primary database, carry out the steps and then turn on the virtual machine.

2. Log in to the platform account on the virtual machine and run the **drives list** command.
 3. In the command output, note the following information, which you will use in the next step:
 - The name of the new disk in the 'Unused disks' section
 - The identifier of the current disk, 'services:backups,' in the 'Used disks and mountpoints' section
 4. Run the following command: **drives reassign <new disk name> services:backups**
- All current data is moved to the new disk. You can continue with your backup or restore operation.

17.3. Loss of the whole cluster and redeploying new servers

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

The high level redeploy and backup restore steps are as follows:

- Redeploy the cluster.
- Store the backup you want to restore in a different location.
- Recreate the remote backups on the primary node using **backup create <loc-name> <URI>**.
- Copy the saved backup under the new UID folder on the remote backup server.
- Do a **backup list**

For example:

```
pxetest:
  URI: sftp://sftpusr:*****@172.29.42.249/AS03
  Backups:
    1 backups have been created - most recently 2014-08-21 10:24
```

A **backup restore** can now be run on the primary.

The example console output below shows the steps and process:

Identifying the database primary:

```
platform@AS01:~$ database primary
172.29.42.100
```

Listing the backups:

```
platform@AS01:~$ backup list
localbackup:
  URI: file:///backups
  Backups:
    2 backups have been created - most recently 2014-08-21 17:59
pxetest:
  URI: sftp://sftpusr:*****@172.29.42.249/AS01
  Backups:
    2 backups have been created - most recently 2014-08-21 12:54

You have new mail in /var/mail/platform
```

Restoring the backup:

```
platform@AS01:~$ backup restore pxetest 2014-08-21 12:54
Services will be restarted during the restore. Do you wish to continue? y
Application <name>-deviceapi processes stopped.
Stopping Application while performing database restore

----- AS02, ip=172.29.42.101, role=webproxy,application,database, loc=cpt

Stopping nginx:proxy

----- AS01, ip=172.29.42.100, role=webproxy,application,database, loc=cpt

Application nginx processes stopped.

----- AS02, ip=172.29.42.101, role=webproxy,application,database, loc=cpt
```

(continues on next page)

(continued from previous page)

```
Application nginx processes stopped.

----- AS04, ip=172.29.21.191, role=webproxy,application,database, loc=jhb

Application nginx processes stopped.

----- AS03, ip=172.29.21.190, role=webproxy,application,database, loc=jhb

Application nginx processes stopped.
System restore starting from
  sftp://sftpusr:sftpusr@172.29.42.249/AS01/bale37deff1309edcc2595bf46c6bfc2a99ca164
  (1408625665)
Local and Remote metadata are synchronized, no sync needed.
Last full backup date: Thu Aug 21 12:54:25 2014
Successfully restored to
  /backups/appdata/restore_temp_1408699183, moving to /backups/appdata
Removing temporary files in /backups/appdata/restore_temp_1408699183
local
Dropping database <name>_FILES before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>_FILES
[object Object]
Repairing database <name>_FILES before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>_FILES
[object Object]
Dropping database PLATFORM before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/PLATFORM
[object Object]
Repairing database PLATFORM before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/PLATFORM
[object Object]
Dropping database <name> before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>
[object Object]
Repairing database <name> before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>
[object Object]
Dropping database <name>_LOCKING before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>_LOCKING
[object Object]
Repairing database <name>_LOCKING before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/<name>_LOCKING
[object Object]
```

(continues on next page)

(continued from previous page)

```

Dropping database admin before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/admin
[object Object]
Repairing database admin before restoring
MongoDB shell version: 2.6.1
connecting to: 127.0.0.1:27020/admin
[object Object]
Trying with oplogReplay
Trying without oplogReplay
restore successfull
Restarting services

Application processes stopped.

Application processes started.

System settings have changed, please reboot using 'system reboot'

```

17.4. Memory (RAM) Increase for Large End User Capacity

Performance on a VOSS-4-UC installation with less than 16GB RAM may degrade when a large end user capacity is reached.

Memory on unified nodes of a cluster with a database role should be 16GB.

For each site with a multinode cluster, for each non-primary database server:

1. Shut down the non-primary database server.
2. In VMware, increase the memory for each non-primary database server node at the site as follows:
 - a. In the Inventory panel, select the virtual machine.
 - b. In the Hardware list on the Summary tab, select **Memory > Edit**.
 - c. On the Resources tab, **Memory > Resource Allocation** and ensure the minimum allocation is 16GB.
 - d. Click **OK**.
4. Restart the database server.
5. Repeat the process described above for the primary database server.

17.5. Error Messages

The tables below provide:

- an error code range reference
- message details of the error codes

To inspect application log messages from the command line, set the debug level on and view the app log. Refer to the Platform Guide for more details.


```
voss set_debug 1
log view voss-deviceapi/app.log
```

The message strings are shown in their template format: references to specific properties are shown as placeholders that are represented by {} .

Note: For AuthError codes, the following rules apply:

- For API version 11.5.3 and below, only the **AuthError_11_5_3** table messages apply.
- For API greater than 11.5.3, **AuthError** table messages override the corresponding **AuthError_11_5_3** table messages, while the unchanged **AuthError_11_5_3** table messages still apply.

RuleError	Message	HTTP Code
15000	Invalid hierarchy for this operation. Please select new hierarchy.	449
15001	Multiple devices found at this Hierarchy level. Please select device.	449
15002	Multiple network device lists (NDL) found at this Hierarchy. Please select a NDL.	449
15003	Network device list reference (NDLR) not found at this Hierarchy.	449
15004	Network device list (NDL) with pkid [{}] not found in available list. Please check NDL rule at the Hierarchy	400
15005	No network device lists (NDL) found at this Hierarchy.	449
15999	Error, (UNHANDLED_ERROR)	400

TransactionError	Message	HTTP Code
23000	Unable to determine Transaction ID.	400
23001	Transaction must be registered with valid user details.	400
23002	Transaction not found.	404
23003	Transaction must be viewed with valid user details.	400
23004	{ } (MAX_INSTANCES_EXCEEDED)	400
23005	Invalid Transaction State: { }	400
23006	Transaction canceled.	400
23007	Transaction must be registered with the hierarchy in which it is executing.	400
23008	Transaction must be registered with model_type if pkid is provided.	400
23010	The current filter caused a long running request. Please add more filter fields, use Case Sensitive or change the criteria types to one of { }.	400
23011	Invalid choices field [{}].	400
23012	The [{}] condition on field [{}], is not allowed.	400
23013	Invalid start and end date range provided in filter.	400
23014	Invalid start and end ID range provided in filter.	400
23015	Invalid ID value in filter	400
23999	Error, { } (UNHANDLED_ERROR)	400

ListUtilError	Message	HTTP Code
20000	Invalid query dictionary, expected 1 key!	400
20999	Error, (UNHANDLED_ERROR)	400

AllError	Message	HTTP Code
999999	All Error	400

ForeignKeyError	Message	HTTP Code
24000	Could not resolve foreign key to {model_type} with "{attr_name}": {attr_value}.	400
24999	Error, {} (UNHANDLED_ERROR)	400

ChoicesError	Message	HTTP Code
26000	Instance context for choices not valid, instance: {instance}	400
26999	Error, {} (UNHANDLED_ERROR)	400

CnfError	Message	HTTP Code
40000	Device change notifications are not supported for device {}.	400
40001	Device change notification data for device {} has been lost. Tracking data has been repaired and collector process will continue. Some changes may have been lost, please run a full sync on the device.	400
40002	Device change notification tracking data for device {} has become corrupted. Tracking data has been repaired and collector process will continue. Some changes may have been lost, please run a full sync on the device.	400
40003	Device change notification tracking DB write for device {} failed. The collector process will continue to attempt DB writes. Please investigate the database write failure. {}	400
40004	Device change notification data DB write for device {} failed. The collector process will continue to attempt DB writes. Please investigate the database write failure. {}	400
40005	Unable to repair device change notification tracking data for device {}. {}	400
40006	Too many unprocessed changes recorded for device {}. No new changes will be recorded until at least {} changes are processed. Please configure and run the necessary data syncs.	400
40008	Could not update pending changes data for device {}. {}.	400
40010	Unable to clear device change notifications for device {}. {}.	400

PackageError	Message	HTTP Code
17000	Unable to load package. Package ({}) depends on ({}) but it does not exist.	400
17001	Unable to load package. Package ({}) requires ({} {}) but {} is currently loaded.	400
17999	Unable to load package. {}	400

CascadeDeleteError	Message	HTTP Code
13000	Hierarchy path or pkid required	400
13001	Could not delete {} out of {} resources.	400
13002	Could not move the following resources that failed to delete: {}.	400
13999	Error, (UNHANDLED_ERROR)	400

WebExError	Message	HTTP Code
31000	[{}] Site Name or Site ID must at least be specified	400

CertificateError	Message	HTTP Code
25001	Certificate request cannot be exported while "Generate Certificate Signing Request" is not set.	400
25002	Certificate can only be imported when "Generate Certificate Signing Request" is set.	400
25003	Certificate upload failed.	400
25004	Uploaded file is not a certificate in .pem format.	400
25005	The SSL certificate expired.	400
25999	Error, {} (UNHANDLED_ERROR)	400

FileUploadError	Message	HTTP Code
39000	Can not determine supported file extensions.	400
39001	'{}' does not have a valid file extension.	400
39002	File is too large. Maximum permitted file size is {} bytes.	400

BulkLoadError	Message	HTTP Code
10000	File Upload Error for File Name : ({})	400

Continued on next page

Table 1 -- continued from previous page

BulkLoadError	Message	HTTP Code
10001	File Encoding Error : ({})	400
10002	Only valid Excel xlsx files are accepted	400
10003	General Error; ({})	400
10004	{success} out of {total} items loaded successfully.	400
10005	Resource data was not found in worksheet '{worksheet}'.	400
10006	Both parallel and serial are not allowed in '{worksheet}'.	400
10007	Differing parallel_transaction_limit values are not allowed in '{worksheet}'.	400
10008	Invalid value of '{limit}' for parallel_transaction_limit header in '{worksheet}', should be left blank or a number between 1 and 100(inclusive).	400
10010	Data does not conform to schema; ({})	400
10011	Hierarchy not specified for row with data; ({})	400
10012	'{user}' is not permitted access to resources at '{hierarchy}'.	403
10020	Hierarchy '{hierarchy}' was not found.	400
10021	Action '{action}' not allowed.	400
10022	Action '{action}' not allowed for model '{model}'.	400
10030	User '{username}' is not allowed to {operation} {model_type}.	403
10040	Fields do not exist in {model}: {fields}.	400
10041	No search fields specified in row.	400
10042	More than one resource found. Search fields '{search}'.	400
10043	Resource not found. Search fields '{search}'.	400
10044	Malformed search fields: {fields}.	400
10045	Malformed fields{message}: {fields}.	400
10046	Can not find meta actions for specified resource instance.	400
10047	Malformed entity header '{header}' in cell '{cell}' worksheet '{sheet}'.	400
10050	Can not enforce data type '{data_type}' on '{data}'. Row data: {row_data}	400
10051	An internal error occurred while processing workbook '{filename}'{note}	400
10052	The specified meta_prefix '{meta_prefix}' in sheet '{sheet_name}' is invalid.	400
10053	The specified meta_prefix '{meta_prefix}' in sheet '{sheet_name}' was not found in base headers.	400
10054	The following base headers '{headers}' in '{sheet_name}' are prefixed, but meta_prefix is not specified.	400
10061	No match for device '{device}'.	400
10062	XLSX File Error: ({})	400

CnfWarning	Message	HTTP Code
45000	Unprocessed changes at 75%% of limit for device {}. Please configure and run the necessary data syncs.	400

DataSyncError	Message	HTTP Code
29000	Could not find user executing data sync operation.	500
29001	User [{}] does not have {} {} permissions.	403
29002	Could not establish a test connection to the device. Verify that your device connection details are correct.	400
29003	Aborting operation. Reason: {}	400
29004	{} (CRITICAL_SUBTRANSACTION_ERROR)	400
29005	Auth Error while testing connection to device	400
29999	Error, {} (UNHANDLED_ERROR)	500

WorkflowError	Message	HTTP Code
7000	Workflow not found	400
7001	Maximum workflow recursion depth exceeded	400
7002	Invalid workflow script identifier {}	400
7003	Specified workflow script name {} not found	400
7004	Error looking up workflow script names against API	400
7005	Invalid workflow action	400
7006	{} (FAILED)	400
7007	Advanced Find Options invalid - Resource not found with options {}	400
7008	{} (CONDITION_CONSTRAINT)	400
7009	Advanced Find Options invalid - More than one resource found with options {}	400
7010	Network Device List {} does not contain an entry for type {}	400
7011	Workflow operation Sync not supported for type {}	400
7012	No target device found for Workflow Sync operation	400
7999	Unexpected error occurred.	400

ExpectError	Message	HTTP Code
35000	The expect binary is not present in the path on the server	500
35001	There was an error executing the expect script : {}	500

ResourceError	Message	HTTP Code
4000	Error, Cannot delete Hierarchy until all resources under it are removed	400
4001	Error, Duplicate Resource Found. {}	400
4002	Resource Not Found {}	404
4003	Failed to save {}. {}	400
4004	Failed to save {}. {}	400
4005	Model Type cannot be None when adding a new Resource	400
4006	Resource Parent {} not found	400
4007	Resource Meta structure corrupt for {}	400
4008	Cannot create a Resource without a Parent Hierarchy	400
4009	Failed to save {}. {}	400
4010	Cannot find Resource relation {}	400
4011	Cannot find target device for model type {} in current hierarchy context	400
4012	Cannot find summary attr [{}] in schema root	400
4013	Cannot perform operation, model {} already has one or more instances	400
4014	Cannot perform operation, resource is part of domain model {}	400
4015	Resource Meta structure corrupt. {}	400
4016	Badly-formed schema; "properties" missing for data type "object"	400
4017	Cannot perform operation, model {} is already referenced by one or more resources: {}	400
4018	Failed to execute {}. {}	400
4019	One or more errors occurred during import	400
4020	Transaction resource failed with errors {}	400
4021	Resources are not of the same type	400
4022	Model type for Resources not found	400
4023	Cannot move Hierarchy Node {} to {}	400
4024	Resource move failed with error {}	400
4025	Invalid business key {}, expected {}	400
4026	Cascade delete failed with error {}	400
4027	Invalid business key for import. Did not expect path, found {}.	400
4028	Resource move failed, Device at source hierarchy [{}] is different from the target hierarchy [{}]	400
4029	Resource [{}] cannot be accessed by user [{}]	403
4030	Cannot perform operation. Hierarchy Node Type [{}] is reserved.	400
4031	Search index is not up to date. Please notify your administrator before proceeding	400
4032	Attempting to create hierarchy node '{}' is not permitted.	403
4033	Could not update reference cache, from: {}, reference: {}, error: {}	403
4034	Resource move failed, hierarchy [{}] of type [{}] does not contain an NDLR	400

Continued on next page

Table 2 -- continued from previous page

ResourceError	Message	HTTP Code
4999	Unhandled Resource Error	400

MacroError	Message	HTTP Code
6000	Template must be a dictionary - got {}	400
6001	No hierarchy supplied	400
6002	Invalid macro specified: {}	400
6003	Macro lookup of {} failed at hierarchy {}	400
6004	Macro lookup of {} returned multiple values {} at hierarchy {}	400
6005	Macro lookup of {} failed when fetching from {} at hierarchy {}	400
6006	Macro lookup failed for field {} in context {}	400
6007	Macro lookup failed for field {} in context {}, type str or int expected not type dict {}	400
6008	Macro function {} not found	400
6009	Macro function arguments error - {}	400
6010	Macro function error - {}	400
6011	Unexpected business key format - {}	400
6012	Conditional Logic error occurred - {}	400
6013	Custom Macro function {} not found	400
6014	Custom Macro function {} not secure or contains invalid strings	400
6015	Could not parse the WhereClause Error: {} WhereClause: {} Please check quotation	400
6016	Lookup field {} not supported/permited.	400
6017	Filter field: {} not in fields: {}.	400
6018	Incorrect hierarchy direction, {}. Allowed: {}.	400
6019	Error in macro function '{}' - {}	400
6999	Error, (UNHANDLED_ERROR)	400

InternalError	Message	HTTP Code
1000	Cannot import Python model name {}	404
1001	Python Type error	400
1002	{} must be an integer	400
1003	Improperly configured settings, {}	400

GraphLookupError	Message	HTTP Code
37000	Cannot perform operation, Resource with pkid [{}] cannot be accessed.	403

AuthError	Message	HTTP Code
27000	{ } (INCORRECT_PASSWORD_ERROR)	401
27001	{ } (PASSWORD_VERIFICATION_ERROR)	401
27009	Please enter a valid username and password.	401
27013	External (SSO or LDAP) authentication is required.	401
27014	Please enter valid answers to security questions.	401

ModelError	Message	HTTP Code
5000	[{}] Child model exists; ({})	400
5001	[{}] Model already exists; ({})	400
5002	One or more data sync errors occurred; ({})	400
5003	[{}] The helper cannot instantiate a model it does not recognize; ({})	400
5004	[{}] The specified resource could not be found; ({})	404
5005	[{}] A single model instance was expected but more than one was found; ({})	404
5006	[{}] Attempt to modify a read-only model failed; ({})	400
5007	[{}] Attempt to modify a read-only model field failed; ({})	400
5008	[{}] Data does not conform to schema; {}	400
5009	[{}] Validation failed; {}	400
5010	[{}] Error manipulating schema; ({})	400
5011	[{}] Error generating schema; ({})	400
5012	[{}] Invalid foreign key to {} for business keys {}	400
5013	[{}] Badly-formed schema; ({})	400
5014	[{}] Error deriving field value; {}	400
5015	Singleton constraint violated: Only one instance of [{}] is allowed per {}.	400
5016	The existing device in [{}] model cannot be modified, it is referenced by other resources.	400
5017	[{}] Invalid foreign key to {} for value {}	400
5018	[{}] Operation not supported for model instance; ({})	405
5019	[{}] Operation not supported; ({})	405
5020	Unable to determine workflow for operation "{}"	400
5021	Workflow "{}" not found	400
5022	Workflow operation "{}" clashes with an existing model attribute/method	400
5023	Unable to execute {} workflow. {}	400
5024	Unable to compile data for provisioning workflow for {}, error {}	400
5025	[{}] Connection timeout error after ({}) seconds	400
5026	[{}] Connection error; ({})	400

Continued on next page

Table 3 -- continued from previous page

ModelError	Message	HTTP Code
5027	API retry error; {}	503
5028	Authentication error; {}	400
5029	Attempt to add a contradicting rule; {}	400
5030	Phones of this type must be added as gateway endpoints	400
5031	Unable to add NDLR to hierarchy node containing device models belonging to devices not referenced by NDLR	400
5032	Unable to query API with available data {}	400
5033	Retries exhausted; {}	400
5050	Password cannot be reused.	400
5051	New password must have {} characters different from old password.	400
5052	User cannot change their password more than once within {} day(s). Please contact your administrator.	400
5053	Password does not meet minimum length required.	400
5054	Password {}.	400
5200	Invalid connection parameters for {}. Username and Password must specified for BASIC authentication method.	400
5201	Invalid connection parameters for {}. Token must specified for OAUTH authentication method.	400
5202	{} {} Unable to render model template {}. TEMPLATE: {} CONTEXT: {}	400
5203	{} {} Unable to parse API response. RESPONSE: {}	400
5204	Invalid connection parameters for {}. Hierarchy must be specified.	400
5205	Invalid paging parameters: page_size {} page_offset {}	400
5206	Paging required: page_size {} page_offset {}	400
5207	External response exceeded memory limit {} {} {}	400
5208	Template output exceeded memory limit {} {}	400
5209	Bad override for {}	400
5210	Session expired. The session cache has been cleared and the next request will go through successfully.	400
5211	Unable to authenticate using session based auth. {}	400
5212	Cannot add device {}	400
5215	Disallowed input {}	400
5998	{0} {1}	400
5999	Error, {}. (UNHANDLED_ERROR)	400

ApiError	Message	HTTP Code
3000	Hierarchy context may not be None, please select Hierarchy	400
3001	Error, Incorrect request format	400

Continued on next page

Table 4 -- continued from previous page

ApiError	Message	HTTP Code
3002	Error, Unhandled method for URL	400
3003	Invalid import file specified. {}	400
3004	Invalid export URL specified. {}	400
3005	Error, Invalid list view sort key [{}]. Valid options are {}	400
3006	Error, Invalid list direction [{}]. Valid options are {}	400
3007	Error, No schema available during list view	400
3008	Provisioning Workflow error [{}]	400
3009	Nothing to export	400
3010	List delete failed, error [{}]	400
3011	List size not allowed, requested [{}], maximum [{}]	400
3012	List sort by hierarchy path not allowed	400
3013	Function not implemented	400
3014	Attribute field name required	400
3015	Hierarchy path [{}] not found.	400
3016	Model type list [{}] not found at or above the current hierarchy.	400
3017	Bulk update failed, error [{}].	400
3018	Bulk operation {} failed, error [{}].	400
3019	Schemas of data being imported have cyclic foreign keys {}.	400
3020	Imported {} out of {} items successfully.	400
3021	{} is a required GET parameter.	400
3022	Invalid Range HTTP header: {}	400
3023	{} is an invalid GET parameter.	400
3024	Resource pkid(s) must be specified	400
3025	Request was throttled.	429
3026	Invalid UTC date format given: {0}, requires: {1} or {2}	400
3027	The current filter caused a long running request. Please add more filter fields, use Case Sensitive or change the criteria types to one of {}.	400
3028	Model Instance Filter [{}] not found at or above the current hierarchy.	400
3029	Purge failed, error [{}]	400
3030	Model Type List of [{}] type not valid for [{}] sync.	400
3031	Model Instance Filter of [{}] type not valid for [{}] sync.	400
3032	{} GET parameter has an invalid value.	400
3999	Unhandled API Error	400

AuthError_11_5_3	Message	HTTP Code
27000	{ } (INCORRECT_PASSWORD_ERROR)	403
27001	{ } (PASSWORD_VERIFICATION_ERROR)	403
27002	{ } (USER_NOT_FOUND_ERROR)	404
27003	{ } (LOGIN_NOT_ALLOWED_ERROR)	403
27004	Account locked. Please contact your administrator.	403
27005	Too many failed login attempts for this user account. Try again later.	403
27006	Too many failed login attempts from this computer. Try again later.	403
27007	Your Web browser doesn't appear to have cookies enabled. Cookies are required for logging in.	400
27008	User is not allowed to log in.	403
27009	Please enter a valid username and password.	403
27010	This account is inactive.	403
27011	User account password must be changed before any API requests are authorized.	403
27012	{ } (ACCOUNT_DISABLED)	403
27013	External (SSO or LDAP) authentication is required.	403
27014	Please enter valid answers to security questions.	403
27015	Password reset is not available for user.	403
27016	Security questions and answers not set up.	403
27017	User can not log in to this interface.	403
27018	User is disabled due to inactivity	403
27019	User is not allowed to login. Please contact your administrator.	403
27020	Login is currently disabled due to a temporary overload. Please try again later.	503
27021	User is not allowed to log in. Maximum user login sessions has been reached.	403

DatabaseError	Message	HTTP Code
2000	Cannot setup Mongo DB collection {}	400
2001	Find failed with spec={}, fields={}, skip={}, limit={}, sort_by={}, err={}	400
2002	Find one failed with spec={}, fields={}, err={}	400
2003	Get archive history failed with spec={}, fields={}, skip={}, limit={}, err={}	400
2004	Remove failed with spec={}, err={}	400
2005	Find and modify failed with spec={}, modify={}, err={}	400
2006	Save failed with spec={}, modify={}, err={}	400
2007	Count failed for {}	400
2008	Find failed with spec={}, fields={}, err={}	400
2009	Duplicate error with spec={}, modify={}, err={}	400
2010	Found more than one record with spec={}	400
2100	Error, Cannot connect to RESOURCE database collection	400
2101	Error, Cannot connect to DATA database collection	400
2102	Error, Cannot connect to ARCHIVE database collection	400
2103	Aggregate failed with group_by={}, match={}, aggregations={}, sort={}, err={}	400
2104	Bulk insert failed, err={}	400
2106	Bulk write failed, err={}	400
2999	Unhandled Database Error	400

AuthenticationProxyError	Message	HTTP Code
32000	Cannot decode target user from authentication proxy. Error: {}	400
32001	Insufficient target user details specified by authentication proxy. Target user details must be contained in a JSON-formatted object with an email attribute.	400
32002	User [{}] is not a valid authentication proxy.	400
32003	Proxy user must be at a hierarchy above that of the target user.	400
32004	Error, {} (UNHANDLED_ERROR)	500

LibSchemaError	Message	HTTP Code
9000	Unhandled schema property error: [{}]	400
9001	Unhandled schema and data processing error: [{}]	400
9002	Data type incorrect, property: {}, not of type: {}	400
9999	Error, (UNHANDLED_ERROR)	400

RbacError	Message	HTTP Code
16000	Permission denied: {}.	400
16001	User not found.	400
16002	Role not specified; User [{}]	400
16003	Access profile not specified; User [{}], Role [{}]	400
16004	Role not found; User [{}], Role [{}]	400
16005	Access profile not found; User [{}], Role [{}], Access Profile [{}]	400
16006	User [{}username] is not allowed to {operation} attribute(s) of {model_type} resource [{}pkid]. Attribute(s) in breach: {breach_attrs}. This operation must be performed by the user's administrator.	403
16007	User [{}username] is not allowed to {operation} {model_type} resource [{}pkid]. This operation must be performed by the user's administrator.	403
16008	Invalid authorization token detected.	403
16009	Role not found; Hierarchy [{}], Role [{}]	400
16010	Access profile [{}] not found for Role [{}] in or above Hierarchy [{}]	400
16011	Access profile of role [{}] is not a subset of the request user's.	400
16012	SelfService Access Profile [{}] for Role [{}] at Hierarchy [{}] must not be created outside 'sys' hierarchy.	400
16999	Error, (UNHANDLED_ERROR)	400

SsoSettingsError	Message	HTTP Code
30000	Invalid certificate file found.	400
30001	Invalid key file found.	400
30002	Validity must not be negative or larger than {} hours ({} years).	400

ApiVersionError	Message	HTTP Code
38000	Invalid API header version specified: {}.	400
38001	No API version mapping defined.	400
38002	API header version: {} and API parameter version: {} mismatch	400

ExportError	Message	HTTP Code
36000	The export format is not specified in request.	400
36001	The specified export format is not supported.	415
36002	The worksheet was not initialized and can not be exported.	500

DataImportError	Message	HTTP Code
11000	Multiple json files {} found in zip archive root; only 1 expected	400
11001	Import file validation failed with: {}	400
11999	Error, (UNHANDLED_ERROR)	400

InterfaceError	Message	HTTP Code
50000	Invalid interface value [{}] for header 'X_INTERFACE'	403
50001	No access profile associated with Interface [{}]	403

BulkLoadMacroError	Message	HTTP Code
60000	Data type must be {}	400
60001	Invalid bulk load macro format {}. Supported format: {}	400

MigrationError	Message	HTTP Code
21000	Post condition failed. {}	400
21999	Error, {} (UNHANDLED_ERROR)	400

CryptoError	Message	HTTP Code
19000	Cryptography validation failed; {}.	400
19999	Error, (UNHANDLED_ERROR)	400

Saml2SsoError	Message	HTTP Code
14000	Could not find SSO settings; Hierarchy: {}.	400
14001	Found multiple SSO settings, only one expected; Hierarchy: {}.	400
14002	Could not find SSO Identity Provider; Hierarchy: {}, IDP uri: {}.	400
14003	Could not resolve SSO Identity Provider; Hierarchy: {}, IDP uri: {}.	400
14004	System generated certificate expected but not specified in data/SsoSettings.	400
14005	System generated certificate has an invalid private key.	400
14006	System generated certificate has an invalid certificate.	400
14007	Unknown principal: {}.	400
14008	Unsupported binding: {}.	400
14009	Verification error: {}.	400
14010	SubjectConfirmation is used but there is no NotOnOrAfter attribute	400
14012	NotBefore and NotOnOrAfter should be present when using either in Condition	400
14013	OneTimeUse element should be present when neither NotBefore nor NotOnOrAfter attributes in Condition	400
14014	Only one OneTimeUse element should be present in Condition	400
14015	Unencrypted assertions are not allowed	400
14016	The session cannot be used yet	400
14999	Error: {}. (UNHANDLED_ERROR)	400

ScriptError	Message	HTTP Code
8000	Script not found	400
8002	Syntax error on line {}	400
8003	Could not connect to {}	400
8004	Authentication failed {}	400
8999	Error, (UNHANDLED_ERROR)	400

HierarchyBasedAccessError	Message	HTTP Code
22000	Invalid traversal argument: '{}'; Traversal must be one of {}.	400
22001	{model_type} with {attr_name} "{attr_value}" is only permitted at the following hierarchy type(s): {hierarchy_types}.	403
22999	Error, {} (UNHANDLED_ERROR)	400

TestConnectionErrorMessage		HTTP Code
12000	Please specify the model type of the device connection parameters	400
12999	Error, (UNHANDLED_ERROR)	400

SysError	Message	HTTP Code
0	Error, Mongo service not started	400
1	Error, Server too busy	400
2	Error, Celery service not started	400

PlatformError	Message	HTTP Code
28000	Could not execute platform command; Exit code: {}	500
28999	Error, {} (UNHANDLED_ERROR)	500

InternalApiUserErrorMessage		HTTP Code
18000	Authorization user [{}] not found.	400
18999	Error, (UNHANDLED_ERROR)	400

SystemMonitoringErrorMessage	Message	HTTP Code
70000	Aggregate {} is not supported by {}	400

RisApiError	Message	HTTP Code
80000	RIS API data collection failed for {}	400

ThemeError	Message	HTTP Code
90000	Theme name {} is reserved for system use. Please choose another name. RIS API data collection failed for {}	400

ClientShapeError	Message	HTTP Code
100000	File operation error: '{}'	400

ClientShapeWarningMessage	HTTP Code
105000	System has not yet been linked to an account in ClientShape 400
105001	System is already linked to an account in ClientShape 400
105002	System has already been registered in ClientShape 400

18. Appendices

18.1. MIBs

18.1.1. MIB List

Important: The VOSS-4-UC system uses standard MIBs that are usually deployed as part of a Network Management System (NMS). No VOSS-4-UC specific MIBs are added. The standard MIBs can for example be inspected from on-line resources, such as <http://www.mibdepot.com>.

The default net-SNMP packages that ship with VOSS-4-UC include:

- ACCOUNTING-CONTROL-MIB
- ADSL-LINE-EXT-MIB
- ADSL-LINE-MIB
- ADSL-TC-MIB
- ADSL2-LINE-MIB
- ADSL2-LINE-TC-MIB
- AGENTX-MIB
- AGGREGATE-MIB
- ALARM-MIB
- APM-MIB
- APPC-MIB
- APPLETALK-MIB
- APPLICATION-MIB
- APPN-DLUR-MIB
- APPN-MIB
- APPN-TRAP-MIB
- APS-MIB
- ARC-MIB
- ATM-ACCOUNTING-INFORMATION-MIB

- ATM-MIB
- ATM-TC-MIB
- ATM2-MIB
- BGP4-MIB
- BRIDGE-MIB
- CAPWAP-BASE-MIB
- CAPWAP-DOT11-MIB
- CHARACTER-MIB
- CIRCUIT-IF-MIB
- CLNS-MIB
- COPS-CLIENT-MIB
- DECNET-PHIV-MIB
- DIAL-CONTROL-MIB
- DIFFSERV-CONFIG-MIB
- DIFFSERV-DSCP-TC
- DIFFSERV-MIB
- DIRECTORY-SERVER-MIB
- DISMAN-EVENT-MIB
- DISMAN-EXPRESSION-MIB
- DISMAN-NSLOOKUP-MIB
- DISMAN-PING-MIB
- DISMAN-SCHEDULE-MIB
- DISMAN-SCRIPT-MIB
- DISMAN-TRACEROUTE-MIB
- DLSW-MIB
- DNS-RESOLVER-MIB
- DNS-SERVER-MIB
- DOCS-BPI-MIB
- DOCS-CABLE-DEVICE-MIB
- DOCS-IETF-BPI2-MIB
- DOCS-IETF-CABLE-DEVICE-NOTIFICATION-MIB
- DOCS-IETF-QOS-MIB
- DOCS-IETF-SUBMGT-MIB
- DOCS-IF-MIB
- DOT12-IF-MIB
- DOT12-RPTR-MIB

- DOT3-EPON-MIB
- DOT3-OAM-MIB
- DPI20-MIB
- DS0-MIB
- DS0BUNDLE-MIB
- DS1-MIB
- DS3-MIB
- DSA-MIB
- DSMON-MIB
- DVB-RCS-MIB
- EBN-MIB
- EFM-CU-MIB
- ENTITY-MIB
- ENTITY-SENSOR-MIB
- ENTITY-STATE-MIB
- ENTITY-STATE-TC-MIB
- ETHER-CHIPSET-MIB
- EtherLike-MIB
- FC-MGMT-MIB
- FCIP-MGMT-MIB
- FDDI-SMT73-MIB
- FIBRE-CHANNEL-FE-MIB
- FLOW-METER-MIB
- FORCES-MIB
- FR-ATM-PVC-SERVICE-IWF-MIB
- FR-MFR-MIB
- FRAME-RELAY-DTE-MIB
- FRNETSERV-MIB
- FRSLD-MIB
- Finisher-MIB
- GMPLS-LABEL-STD-MIB
- GMPLS-LSR-STD-MIB
- GMPLS-TC-STD-MIB
- GMPLS-TE-STD-MIB
- GSMP-MIB
- HC-ALARM-MIB

- HC-PerfHist-TC-MIB
- HC-RMON-MIB
- HCNUM-TC
- HDLSL2-SHDSL-LINE-MIB
- HOST-RESOURCES-MIB
- HOST-RESOURCES-TYPES
- HPR-IP-MIB
- HPR-MIB
- IBM-6611-APPN-MIB
- IF-CAP-STACK-MIB
- IF-INVERTED-STACK-MIB
- IF-MIB
- IFCP-MGMT-MIB
- IGMP-STD-MIB
- INET-ADDRESS-MIB
- INTEGRATED-SERVICES-GUARANTEED-MIB
- INTEGRATED-SERVICES-MIB
- INTERFACETOPN-MIB
- IP-FORWARD-MIB
- IP-MIB
- IPATM-IPMC-MIB
- IPIX-MIB
- IPMCAST-MIB
- IPMROUTE-STD-MIB
- IPOA-MIB
- IPS-AUTH-MIB
- IPSEC-SPD-MIB
- IPV6-FLOW-LABEL-MIB
- IPV6-ICMP-MIB
- IPV6-MIB
- IPV6-MLD-MIB
- IPV6-TC
- IPV6-TCP-MIB
- IPV6-UDP-MIB
- ISCSI-MIB
- ISDN-MIB

- ISIS-MIB
- ISNS-MIB
- ITU-ALARM-MIB
- ITU-ALARM-TC-MIB
- Job-Monitoring-MIB
- L2TP-MIB
- LANGTAG-TC-MIB
- LM-SENSORS-MIB
- LMP-MIB
- MALLOC-MIB
- MAU-MIB
- MGMD-STD-MIB
- MIDCOM-MIB
- MIOX25-MIB
- MIP-MIB
- MOBILEIPV6-MIB
- MPLS-FTN-STD-MIB
- MPLS-L3VPN-STD-MIB
- MPLS-LC-ATM-STD-MIB
- MPLS-LC-FR-STD-MIB
- MPLS-LDP-ATM-STD-MIB
- MPLS-LDP-FRAME-RELAY-STD-MIB
- MPLS-LDP-GENERIC-STD-MIB
- MPLS-LDP-STD-MIB
- MPLS-LSR-STD-MIB
- MPLS-TC-STD-MIB
- MPLS-TE-STD-MIB
- MSDP-MIB
- MTA-MIB
- Modem-MIB
- NAT-MIB
- NEMO-MIB
- NET-SNMP-AGENT-MIB
- NET-SNMP-EXAMPLES-MIB
- NET-SNMP-EXTEND-MIB
- NET-SNMP-MIB

- NET-SNMP-MONITOR-MIB
- NET-SNMP-PASS-MIB
- NET-SNMP-SYSTEM-MIB
- NET-SNMP-TC
- NET-SNMP-VACM-MIB
- NETWORK-SERVICES-MIB
- NHRP-MIB
- NOTIFICATION-LOG-MIB
- OPT-IF-MIB
- OSPF-MIB
- OSPF-TRAP-MIB
- OSPFV3-MIB
- P-BRIDGE-MIB
- PARALLEL-MIB
- PIM-BSR-MIB
- PIM-MIB
- PIM-STD-MIB
- PINT-MIB
- PKTC-IETF-EVENT-MIB
- PKTC-IETF-MTA-MIB
- PKTC-IETF-SIG-MIB
- POLICY-BASED-MANAGEMENT-MIB
- POWER-ETHERNET-MIB
- PPP-BRIDGE-NCP-MIB
- PPP-IP-NCP-MIB
- PPP-LCP-MIB
- PPP-SEC-MIB
- PTOPO-MIB
- PW-ATM-MIB
- PW-ENET-STD-MIB
- PW-MPLS-STD-MIB
- PW-STD-MIB
- PW-TC-STD-MIB
- PW-TDM-MIB
- PerfHist-TC-MIB
- Printer-MIB

- Q-BRIDGE-MIB
- RADIUS-ACC-CLIENT-MIB
- RADIUS-ACC-SERVER-MIB
- RADIUS-AUTH-CLIENT-MIB
- RADIUS-AUTH-SERVER-MIB
- RADIUS-DYNAUTH-CLIENT-MIB
- RADIUS-DYNAUTH-SERVER-MIB
- RAQMON-MIB
- RAQMON-RDS-MIB
- RDBMS-MIB
- RFC1155-SMI
- RFC1213-MIB
- RFC1381-MIB
- RFC1382-MIB
- RFC1414-MIB
- RIPv2-MIB
- RMON-MIB
- RMON2-MIB
- ROHC-MIB
- ROHC-RTP-MIB
- ROHC-UNCOMPRESSED-MIB
- RS-232-MIB
- RSERPOOL-MIB
- RSTP-MIB
- RSVP-MIB
- RTP-MIB
- SCSI-MIB
- SCTP-MIB
- SFLOW-MIB
- SIP-COMMON-MIB
- SIP-MIB
- SIP-SERVER-MIB
- SIP-TC-MIB
- SIP-UA-MIB
- SLAPM-MIB
- SMON-MIB

- SMUX-MIB
- SNA-NAU-MIB
- SNA-SDLC-MIB
- SNMP-COMMUNITY-MIB
- SNMP-FRAMEWORK-MIB
- SNMP-IEEE802-TM-MIB
- SNMP-MPD-MIB
- SNMP-NOTIFICATION-MIB
- SNMP-PROXY-MIB
- SNMP-REPEATER-MIB
- SNMP-SSH-TM-MIB
- SNMP-TARGET-MIB
- SNMP-TSM-MIB
- SNMP-USER-BASED-SM-MIB
- SNMP-USM-AES-MIB
- SNMP-USM-DH-OBJECTS-MIB
- SNMP-VIEW-BASED-ACM-MIB
- SNMPv2-CONF
- SNMPv2-M2M-MIB
- SNMPv2-MIB
- SNMPv2-PARTY-MIB
- SNMPv2-PDU
- SNMPv2-SMI
- SNMPv2-TC
- SNMPv2-TM
- SNMPv2-USEC-MIB
- SONET-MIB
- SOURCE-ROUTING-MIB
- SSPM-MIB
- SYSAPPL-MIB
- SYSLOG-MSG-MIB
- SYSLOG-TC-MIB
- T11-FC-FABRIC-ADDR-MGR-MIB
- T11-FC-FABRIC-CONFIG-SERVER-MIB
- T11-FC-FABRIC-LOCK-MIB
- T11-FC-FSPF-MIB

- T11-FC-NAME-SERVER-MIB
- T11-FC-ROUTE-MIB
- T11-FC-RSCN-MIB
- T11-FC-SP-AUTHENTICATION-MIB
- T11-FC-SP-POLICY-MIB
- T11-FC-SP-SA-MIB
- T11-FC-SP-TC-MIB
- T11-FC-SP-ZONING-MIB
- T11-FC-VIRTUAL-FABRIC-MIB
- T11-FC-ZONE-SERVER-MIB
- T11-TC-MIB
- TCP-ESTATS-MIB
- TCP-MIB
- TCPIPX-MIB
- TE-LINK-STD-MIB
- TE-MIB
- TIME-AGGREGATE-MIB
- TN3270E-MIB
- TN3270E-RT-MIB
- TOKEN-RING-RMON-MIB
- TOKENRING-MIB
- TOKENRING-STATION-SR-MIB
- TPM-MIB
- TRANSPORT-ADDRESS-MIB
- TRIP-MIB
- TRIP-TC-MIB
- TUNNEL-MIB
- UCD-DEMO-MIB.inc
- UCD-DEMO-MIB
- UCD-DISKIO-MIB.inc
- UCD-DISKIO-MIB
- UCD-DLMOD-MIB.inc
- UCD-DLMOD-MIB
- UCD-IPFILTER-MIB.inc
- UCD-IPFILTER-MIB
- UCD-IPFWACC-MIB.inc

- UCD-IPFWACC-MIB
- UCD-SNMP-MIB-OLD
- UCD-SNMP-MIB.inc
- UCD-SNMP-MIB
- UDP-MIB
- UDPLITE-MIB
- UPS-MIB
- URI-TC-MIB
- VDSL-LINE-EXT-MCM-MIB
- VDSL-LINE-EXT-SCM-MIB
- VDSL-LINE-MIB
- VDSL2-LINE-MIB
- VDSL2-LINE-TC-MIB
- VPN-TC-STD-MIB
- VRRP-MIB
- WWW-MIB
- IANA-ADDRESS-FAMILY-NUMBERS-MIB
- IANA-CHARSET-MIB
- IANA-FINISHER-MIB
- IANA-GMPLS-TC-MIB
- IANA-IPPM-METRICS-REGISTRY-MIB
- IANA-ITU-ALARM-TC-MIB
- IANA-LANGUAGE-MIB
- IANA-MALLOC-MIB
- IANA-MAU-MIB
- IANA-PRINTER-MIB
- IANA-PWE3-MIB
- IANA-RTPROTO-MIB
- IANATn3270eTC-MIB
- IANAifType-MIB
- IPFIX-SELECTOR-MIB

For further information on how to add a MIB, see:

http://www.net-snmp.org/wiki/index.php/TUT:Using_and_loading_MIBS

18.2. Data Export Types

18.2.1. Analogue line MGCP Data Export

Filename: <YYYY-MM-DD_HHMM>_analogue_line_mgcp.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
usernames	List of users assigned to the analog port	Array of strings	v1
gateway	name of the gateway that the port is on	string	v1
port_number	gateway port for this configuration	string	v1
port_type	the type of port for this gateway (typically FXS for analog)	string	v1
description	description of the gateway	string	v1
cucm_dn	Internal Number assigned to the device profile (as configured in the PBX)	string	v1
E164	External Number (E164 number) assigned to the device profile	string	v1

Example

(* marked fields are new in version 2)

```
[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-NB",
    "customer_name": "InGen",
    "division_name": "",
    "location_name": "StandardSite1",
    "usernames": [
      "SSUser33"
    ],
    "description": "",
    "hierarchy": "sys.hcs.CS-P.CS-NB.InGen.StandardSite1",
    "port_number": 4,
    "port_type": "Cisco MGCP FXS Port",
    "E164": "\\+441425204033",
    "cucm_dn": "81214033",
  }
]
```

(continues on next page)

(continued from previous page)

```

    "gateway": "SKIGW9981220001"
  }
]

```

18.2.2. Analogue Line SCCP Data Export

Filename: <YYYY-MM-DD_HHMM>_analogue_line_sccp.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
usernames	List of users assigned to the analog port	Array of strings	v1
gateway	name of the gateway that the port is on	string	v1
port_number	gateway port for this configuration	string	v1
port_type	the type of port for this gateway (typically FXS for analog)	string	v1
description	description of the gateway	string	v1
cucm_dn	Internal Number assigned to the device profile (as configured in the PBX)	string	v1
E164	External Number (E164 number) assigned to the device profile	string	v1

Example

(* marked fields are new in version 2)

```

[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-NB",
    "customer_name": "InGen",
    "division_name": "",
    "location_name": "StandardSite1",
    "usernames": [
      "SSUser33"
    ],
    "description": "",
    "hierarchy": "sys.hcs.CS-P.CS-NB.InGen.StandardSite1",
    "port_number": 0,
    "port_type": "Analog Phone",

```

(continues on next page)

(continued from previous page)

```

    "E164": "",
    "cucm_dn": "81214050",
    "gateway": "SKIGW9981212041"
  }
]

```

18.2.3. Call Pickup Group Data Export

(New report in version 2)

Filename: <YYYY-MM-DD_HHMM>_call_pickup_group.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	Name of the Customer	string	v2
division_name	Intermediate Node (e.g Division or other node)	string	v2
location_name	Name of the Site	string	v2
hierarchy	The full hierarchy path for the item being exported	string	v2
pickup_group_name	The name of the Call Pickup Group	string	v2
pickup_group_number	The DN for the Call Pickup Group	string	v2
pickup_group_partition	The route partition for the Call Pickup Group DN	string	v2
member	Array of member lines	array	v2
member.cucm_dn	Description of the directory number and partition	string	v2
member.partition	Route partition associated with the member directory number	string	v2

Example

```

[
  {
    "provider_name": "CS-P",
    "reseller_name": "CS-NB",
    "customer_name": "CustomerName",
    "division_name": "",
    "location_name": "AAA-Boston",
    "hierarchy": "sys.hcs.CS-P.CS-NB.CustomerName.AAA-Boston",
    "pickup_group_name": "Support",
    "pickup_group_number": "80000",
    "pickup_group_partition": "Cul-AllowVm-PT",
    "member": [
      {
        "cucm_dn": "50409",
        "partition": "Cul-AllowVm-PT"
      }
    ]
  }
]

```

(continues on next page)

(continued from previous page)

```

    ]
  }
]

```

18.2.4. Contact Center Enterprise Data Export

Filename: <YYYY-MM-DD_HHMM>_contact_center_enterprise.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2.2
reseller_name	Name of the Reseller	string	v2.2
customer_name	name of the customer	string	v2.2
division_name	Intermediate Node (e.g Division or other node)	string	v2.2
location_name	Site Name	string	v2.2
hierarchy	The full hierarchy path for the item being exported	string	v2.2
Name	Contact Center Username	string	v2.2
PeripheralNumber	Skill group peripheral number	integer	v2.2
Supervisor	User type	boolean	v2.2

Example

```

[
  {
    "division_name": "",
    "Supervisor": false,
    "Name": "standalone_ccdm_user_2",
    "hierarchy": "sys.hcs.Provider_01.Reseller_01.Customer_01.Site_01",
    "reseller_name": "Reseller_01",
    "location_name": "Site_01",
    "provider_name": "Provider_01",
    "PeripheralNumber": 2,
    "customer_name": "Customer_01"
  }
]

```

18.2.5. Contact Center Express Data Export

Filename: <YYYY-MM-DD_HHMM>_contact_center_express.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2.2
reseller_name	Name of the Reseller	string	v2.2
customer_name	name of the customer	string	v2.2
division_name	Intermediate Node (e.g Division or other node)	string	v2.2
location_name	Site Name	string	v2.2
hierarchy	The full hierarchy path for the item being exported	string	v2.2
username	Contact Center Express username	string	v2.2
userID	CUCM user ID	string	v2.2
teamName	Contact Center Express team name	string	v2.2
type	Contact Center Express user type	string	v2.2
autoAvailable	Availability status of the user	boolean	v2.2

Example

```
[
  {
    "division_name": "",
    "location_name": "Site_01",
    "firstName": "user_46",
    "extension": 2,
    "hierarchy": "sys.hcs.Provider_01.Reseller_01.Customer_01.Site_01",
    "lastName": "Latame",
    "userID": "user_46",
    "teamName": "Default",
    "reseller_name": "Reseller_01",
    "provider_name": "Provider_01",
    "customer_name": "Customer_01",
    "type": "Agent",
    "autoAvailable": false
  }
]
```

18.2.6. Customer Data Export

(New report in version 2)

Filename: <YYYY-MM-DD_HHMM>_customer.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	Name of the Customer	string	v2
hierarchy	The full hierarchy path for the item being exported	string	v2
account_id	The customer's account identifier	string	v2
external_id	An externally defined identifier for the customer	string	v2

Example

```
[
  {
    "provider_name": "CS-P",
    "reseller_name": "CS-NB",
    "customer_name": "Customer1",
    "hierarchy": "sys.hcs.CS-P.CS-NB.Customer1",
    "account_id": "ABCXYZ",
    "external_id": ""
  }
]
```

18.2.7. Extension Mobility Data Export

Filename: <YYYY-MM-DD_HHMM>_extension_mobility.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
username	the username of the owner of device profile	string	v1
device_type	Model the extension mobility profile is setup as	string	v1
lines	Array of objects containing line information	array	v1
lines.cucm_dn	Internal Number assigned to the device profile (as configured in the PBX)	string	v1
lines.E164	External Number (E164 number) assigned to the device profile	string	v1
lines.line_order	Line index.	integer	v2
device_profile_name	Name of the extension mobility profile	string	v2

Example:

(* marked fields are new in version 2)

```
[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-NB",
    "customer_name": "AAAGlobal",
    "division_name": "",
    "location_name": "AAA-Boston",
    "username": "ba_user4",
    "hierarchy": "sys.hcs.CS-P.CS-NB.AAAGlobal.AAA-Boston",
    * "device_profile_name": "FirstnameLastname-UDP",
    "lines": [
      {
        * "line_order": 1,
        "cucm_dn": "50409",
        "E164": "\\+18575550409"
      }
    ],
    "device_type": "Cisco 9971"
  }
]
```

18.2.8. FMC Data Export

(New report in version 2)

This report includes users who have the FMC feature configured. The report includes the destination configured and an indication of whether the service is currently enabled or disabled (based on v2 FMC with CIM-based FMC). Any users without the FMC feature configured will not appear in the file. This report is only populated if the FMC adaptation is installed on the system - the file will be blank on systems without any users configured or if the adaptation is not installed.

Filename: <YYYY-MM-DD_HHMM>_fmc.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	Name of the Customer	string	v2
division_name	Intermediate Node (e.g Division or other node)	string	v2
location_name	Name of the Site	string	v2
hierarchy	The full hierarchy path for the item being exported	string	v2
username	The userid of the remote destination profile	string	v2
destination_number	The mobile number associated with CIM device	string	v2
fmc_enabled	An indication of whether fixed mobile convergence is enabled for the destination	boolean	v2

Example

```
[
  {
    "provider_name": "CS-P",
    "reseller_name": "CS-NB",
    "customer_name": "AAAGlobal",
    "division_name": "",
    "location_name": "AAA-Boston",
    "hierarchy": "sys.hcs.CS-P.CS-NB.AAAGlobal.AAA-Boston"
    "username": "ba_user4",
    "destination": "08212345678",
    "fmc_enabled": true
  }
]
```

18.2.9. Hunt Group Data Export

Filename: <YYYY-MM-DD_HHMM>_hunt_group.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
location_name	Site Name	string	v1
hunt_group_name	Name assigned to the hunt group	string	v1
pilot_number	the internal number assigned as the pilot for the hunt group (as configured in the PBX)	string	v1
E164	the external number (Full E164 format) assigned as the pilot for the hunt group (as configured in the PBX)	string	v1
lines	Array of objects containing line information	array	v1
lines.cucm_dn	Internal Number assigned to the device profile (as configured in the PBX)	string	v1
partition	The route partition to which the Hunt Pilot number belongs	string	v2
line_group_name	Name of the line group	string	v2

Example

(* marked fields are new in version 2)

```
[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-AB",
    "customer_name": "AB_Group",
    "division_name": "AB-C",
```

(continues on next page)

(continued from previous page)

```

    "location_name": "CL1-AB-C-St_Nazaire",
    "hierarchy":
"sys.hcs.CS-P.CS-AB.AB_Group.AB-C.CL1-AB-C-St_Nazaire",
    "lines": [
      {
        "cucm_dn": "8134808",
        * "line_group_name": "BackOffice"
      }
    ],
    "hunt_group_name": "HL-825",
    "E164": "\\+33228544825",
    "pilot_number": "8134825",
    * "partition": "Cul-AllowVm-PT"
  }
]

```

18.2.10. Line Data Export

(New report in version 2)

Filename: <YYYY-MM-DD_HHMM>_line.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	Name of the Customer	string	v2
division_name	Intermediate Node (e.g Division or other node)	string	v2
location_name	Name of the Site	string	v2
hierarchy	The full hierarchy path for the item being exported	string	v2
cucm_dn	Internal Number of this line	string	v2
partition	The route partition to which the number belongs	string	v2
description	Description of the directory number and partition	string	v2
calling_search_space	This is mapped to the shareLineAppearanceCss- Name of the line	string	v2

Example

```

[
  {
    "provider_name": "CS-P",
    "reseller_name": "CS-NB",
    "customer_name": "CustomerName",
    "division_name": "",
    "location_name": "AAA-Boston",
    "hierarchy": "sys.hcs.CS-P.CS-NB.CustomerName.AAA-Boston",

```

(continues on next page)

(continued from previous page)

```

    "cucm_dn": "50409",
    "partition": "Cu1-AllowVm-PT",
    "description": "Front Desk",
    "calling_search_space": "Cu1-ANumAnaly-CSS"
  }
]

```

18.2.11. Phones Data Export

Filename: <YYYY-MM-DD_HHMM>_phones.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
usernames	list of usernames associated to the phones via Unified CM user, associated devices	array	v1
device_name	the name of the device (includes mac address if hardphone, softclients no mac)	string	v1
description	Text field attached to the device	string	v3
device_type	the model of the phone	string	v1
lines	Array of objects containing line information	array	v1
lines.cucm_dn	Internal Number assigned to the device profile (as configured in the PBX)	string	v1
lines.E164	External Number (E164 number) assigned to the device profile	string	v1
lines.line_order	Line index.	integer	v2
device_css	Calling search space of the phone	string	v2
add_on_modules	Array of phone addon modules, incl. name, model, position	array	v4

Example:

(* marked fields are new in version 2 and version 3)

(** marked fields are new in version 4)

```

[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-NB",

```

(continues on next page)

(continued from previous page)

```

    "customer_name": "AAAGlobal",
    "division_name": "",
    "location_name": "AAA-Boston",
    "hierarchy": "sys.hcs.CS-P.CS-NB.AAAGlobal.AAA-Boston",
    "usernames": [
      "ba_user4"
    ],
    "lines": [
      *
      {
        "line_order": 1,
        "cucm_dn": "50409",
        "E164": "\\+18575550409"
      }
    ],
    **
    "add_on_modules": [
      {
        "position": "1",
        "name": "Cisco Unified IP Color Key Expansion Module",
        "model": "CP-CKEM-C"
      }
    ],
    "device_name": "SEP99887777788",

    "description": "Meeting Room Phone",
    "device_type": "Cisco 9971",
    "device_css": "Cu1Sil-USADP-Emer-CSS"
  }
]

```

18.2.12. Site Data Export

Filename: <YYYY-MM-DD_HHMM>_site.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VER.
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	Name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
customer_address1	Address string 1 for the customer	string	v1
customer_address2	Address string 2 for the customer	string	v1
customer_address3	Address string 3 for the customer	string	v1
location_address1	Address string 1 for the site	string	v1
location_address2	Address string 2 for the site	string	v1
location_address3	Address string 3 for the site	string	v1
emergency_number	External emergency callback number assigned to the site	string	v1
ndl	The NDL name that the site uses	string	v1
inter_site_prefix	Digit dialled to prefix intersite calls (if the dial plan is setup that way)	string	v1
external_access_prefix	Digit dialled to make external calls (if the dial plan is setup that way)	string	v1
site_code	Dial Plan site code assigned to the site (if the dial plan is setup that way)	string	v1
published_number	External published callback number assigned to the site	string	v1
country_code	Country code identifying the site	string	v1
voice_bandwidth	voice bandwidth allocation for the site	string	v1
video_bandwidth	video bandwidth allocation for the site	string	v1
external_id	An externally defined ID for the site	string	v2
extended_name	An expanded name for the site	string	v2

Example:

(* marked fields are new in version 2)

```
[
  {
    * "provider_name": "CS-P",
    * "reseller_name": "CS-NB",
    "customer_name": "Varidion",
    "division_name": "",
    "location_name": "Varidion-Reading",
    "hierarchy": "sys.hcs.CS-P.CS-NB.Varidion.Varidion-Reading",
    "customer_address1": "Varidion New York (Head Office)",
    "customer_address2": "L23, 33 Central Square",
    "customer_address3": "Dallas, TX, USA",
    "ndl": "GS-R3-VDN-CL1-NDL",
```

(continues on next page)

(continued from previous page)

```
"inter_site_prefix": "",
"site_code": "",
"video_bandwidth": "",
"emergency_number": "",
"voice_bandwidth": "",
"country_code": "44",
"external_access_prefix": "",
"location_address1": "Varidion Reading",
"location_address3": "Reading, Berkshire",
"location_address2": "Atlantic House, Imperial Way",
"published_number": "",
* "external_id": "ABCXYZ",
* "extended_name": "UK IT"
}
]
```

18.2.13. Subscriber Data Export

Filename: <YYYY-MM-DD_HHMM>_subscriber.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2
reseller_name	Name of the Reseller	string	v2
customer_name	name of the customer	string	v1
division_name	Intermediate Node (e.g Division or other node)	string	v1
location_name	Site Name	string	v1
hierarchy	The full hierarchy path for the item being exported	string	v1
username	username of the user	string	v1
first_name	First name of the user	string	v1
middle_name	Middle name of the user	string	v3
last_name	Last name of the user	string	v1
email	email address of the user	string	v1
entitlement_profile	the profile assigned to the user that defines the features they are enabled to have configured	string	v1
role	The role assigned to the user - defines privileges in the portal	string	v1
credential_policy	The security profile assigned to the user - defined credential and other security rules for portal access	string	v1
snr	Does the user have the SNR service configured	boolean	v1
voicemail	Does the user have a voicemail box configured	boolean	v1
title	Subscriber's title	string	v2
department	Subscriber's department	string	v2
telephone_number	Subscriber's telephone number as configured in the CUCM user record	string	v2
pager_number	Subscriber's pager number	string	v3

Example

(* marked fields are new in version 2, ** marked fields are new in version 3)

```
[
  {
    *   "provider_name": "CS-P",
    *   "reseller_name": "CS-NB",
        "customer_name": "AAAGlobal",
        "division_name": "",
        "location_name": "AAA-Boston",
        "hierarchy": "sys.hcs.CS-P.CS-NB.AAAGlobal.AAA-Boston",
        "username": "ba_user4",
        "first_name": "Dean",
    **  "middle_name": "John",
        "last_name": "Daniels",
        "voicemail": false,
        "entitlement_profile": "AAAGlobal-Foundation-EP",
        "snr": false,
        "credential_policy": "HcsCredentialPolicy",
        "role": "AAA-BostonSelfService",
```

(continues on next page)

(continued from previous page)

```

        "email": "email@theinternet.com",
    *   "title": "Dr.",
    *   "department": "R&D",
    *   "telephone_number": "0215252020"
    **  "pager_number": "5551234545"
    }
]

```

Hybrid System Subscriber Data Export

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
username	username of the user	string	v4
service_type	Hybrid type	string	v4
lines	array of lines, CoS	array	v4

Example

```

[
  {
    "username": "user_1",
    "hierarchy": "sys.9F73F4303A93.Provider1Hierarchy.Reseller1Hierarchy.
↔Customer1Hierarchy",
    "lines": [
      {
        "e164": "\\+441184025574",
        "class_of_service": "International-24Hrs-Enhanced",
        "extension": "8445574"
      },
      {
        "e164": "\\+441184025576",
        "class_of_service": "International-24Hrs-Enhanced",
        "extension": "8445576"
      }
    ],
    "service_type": "Cisco-MS",
    "reseller_name": "Reseller1Hierarchy",
    "provider_name": "Provider1Hierarchy",
    "customer_name": "Customer1Hierarchy"
  }
]

```

MS Office 365 Subscriber Data Export

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
username	username of the user	string	v4
first_name	First name of the user	string	v4
last_name	Last name of the user	string	v4
o365_username	user name on O365	string	v4
display_name	displayed name of user	string	v4
is_licensed	User licensing status	string	v4
licenses	List of licenses	array	v4

Example

```
[
  {
    "username": "user_1",
    "first_name": "FirstName_1",
    "last_name": "LastName_1",
    "display_name": "DisplayName",
    "hierarchy": "sys.CF24424CF914.Provider1Hierarchy.Reseller1Hierarchy.
↪Customer1Hierarchy",
    "is_licensed": "Y",
    "o365_username": "user_1@dummy-emailaccount.com",
    "licenses": [
      {
        "account_sku_id": "42"
      }
    ],
    "reseller_name": "Reseller1Hierarchy",
    "provider_name": "Provider1Hierarchy",
    "customer_name": "Customer1Hierarchy"
  }
]
```

MS Teams Subscriber Data Export

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
ms_teams_username	UserPrincipalName	string	v4
alias	user alias	string	v4
username	username of the user	string	v4
first_name	First name of user	string	v4
last_name	Last name of user	string	v4
department	User department	string	v4
enterprise_voice	EnterpriseVoiceEnabled	string	v4
line	OnPremLineURI	string	v4
hosted_voicemail	HostedVoiceMail	boolean	v4
voice_routing_policy	OnlineVoiceRoutingPolicy	string	v4
tenant_dialplan	TenantDialPlan	string	v4
voicemail_policy	HostedVoicemailPolicy	string	v4

Example

```
[
  {
    "username": "user_1",
    "first_name": "FirstName_1",
    "last_name": "LastName_1",
    "enterprise_voice": false,
    "hierarchy": "sys.4E0EF94F90B9.Provider1Hierarchy.Reseller1Hierarchy.
↔Customer1Hierarchy",
    "ms_teams_username": "user_1@dummy-emailaccount.com",
    "voicemail_policy": "BusinessVoice",
    "voice_routing_policy": "Global",
    "alias": "user_1",
    "tenant_dialplan": "",
    "department": "R&D",
    "reseller_name": "Reseller1Hierarchy",
    "provider_name": "Provider1Hierarchy",
    "hosted_voicemail": false,
    "line": "18694400002",
    "customer_name": "Customer1Hierarchy"
  }
]
```

Pexip Subscriber Data Export

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
name	First name of user	string	v4
description	Last name of user	string	v4
owner_email	User email address	string	v4
type	service type	string	v4

Example

```
[
  {
    "name": "user_1",
    "hierarchy": "sys.0ECD98831FCF.Provider1Hierarchy.Reseller1Hierarchy.
↔Customer1Hierarchy",
    "description": "Description_PexIp_1",
    "owner_email": "user_1@dummy-emailaccount.com",
    "reseller_name": "Reseller1Hierarchy",
    "provider_name": "Provider1Hierarchy",
    "type": "conference",
    "customer_name": "Customer1Hierarchy"
  }
]
```

VOSS Phone Servers Subscriber Data Export

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
mac	MAC address of phone server	string	v4
phone_vendor	Vendor name	string	v4
phone_model	Model Name	string	v4
lines	Lines names and CoS	array	v4

Example

```
[
  {
    "phone_model": "Cisco 6921",
    "hierarchy": "sys.A49D7D3D9752.Provider1Hierarchy.Reseller1Hierarchy.
↔Customer1Hierarchy",
    "lines": [
      {
        "class_of_service": "International-24Hrs-Enhanced",
        "name": "8445574"
      },
      {
        "class_of_service": "International-24Hrs-Enhanced",
        "name": "8445576"
      }
    ],
    "mac": "2C:54:91:88:C9:01",
    "phone_vendor": "Cisco",
    "reseller_name": "Reseller1Hierarchy",
    "provider_name": "Provider1Hierarchy",
    "customer_name": "Customer1Hierarchy"
  }
]
```

18.2.14. Webex Teams Data Export

Filename: <YYYY-MM-DD_HHMM>_webex_teams.json.gz

Layout:

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
provider_name	Name of the Provider	string	v2.2
reseller_name	Name of the Reseller	string	v2.2
customer_name	name of the customer	string	v2.2
division_name	Intermediate Node (e.g. Division or other node)	string	v2.2
location_name	Site Name	string	v2.2
hierarchy	The full hierarchy path for the item being exported	string	v2.2
firstName	First name of user	string	v2.2
lastName	Last name of user	string	v2.2
email	User email address	string	v2.2
line	User line	string	v2.2
messaging	Webex Teams Messaging	boolean	v2.2

Note: Services reported on are dynamically included. The reference material and JSON snippet here are examples.

ELE-MENT	DESCRIPTION	DATA TYPE	VER-SION
hybrid_call_services	Users' incoming calls will ring their work phones and the Cisco Webex Teams app. Users can call their colleagues from either their phones or the app, too. Aware must be enabled before the user can be enabled for Connect.	boolean	v2.2
connect_aware	Users can share content from the Cisco Webex Teams app during a call from their work phones and view their call history in the app.	boolean	v2.2

ELEMENT	DESCRIPTION	DATA TYPE	VER-SION
hybrid_calendar_services:	Google Calendar	boolean	v2.2
google	Microsoft Exchange/Office 365	boolean	v2.2
microsoft_exchange			v2.2

ELEMENT	DESCRIPTION	DATA TYPE	VERSION
meeting	Named User Licence. Each Named User license allows 1 user to be entitled as a meeting host. Named users can hold unlimited meetings.	N/A	N/A
we-bex_enterprise_200		boolean	v2.2
we-bex_support_center		boolean	v2.2
we-bex_meeting_center		boolean	v2.2
webex_cmr		boolean	v2.2
we-bex_event_center		boolean	v2.2
we-bex_training_center		boolean	v2.2
meeting		boolean	v2.2

Example

```
[
  {
    "division_name": "",
    "status": "",
    "location_name": "Site_03",
    "firstName": "",
    "hierarchy": "sys.hcs.Provider_01.Reseller_01.Customer_02.Site_03",
    "lastName": "",
    "provider_name": "Provider_01",
    "services": {
      "hybrid_call_services": {
        "connect": false,
        "aware": false
      },
      "message": {
        "messaging": false
      },
      "meeting": {
        "webex_enterprise_200": false,
        "webex_support_center": false,
        "webex_meeting_center": false,
        "webex_cmr": false,
        "webex_event_center": false,
        "webex_training_center": false,
        "meeting": false
      },
      "hybrid_calendar_services": {
        "google": false,
        "microsoft_exchange": false
      },
      "hybrid_message_services": {
        "message": false
      }
    }
  }
]
```

(continues on next page)

(continued from previous page)

```

    },
    "reseller_name": "Reseller_01",
    "line": "",
    "email": "spark_user_36@emailaccount.com",
    "customer_name": "Customer_02"
  },
]

```

18.3. Command Examples

18.3.1. diag Command Examples

For details, see: *Diagnostic Tools*.

diag disk

```

platform@VOSS:~$ diag disk
Filesystem                Size  Used Avail Use% Mounted on
udev                      3.9G   0  3.9G  0% /dev
tmpfs                     1.6G   80M  1.5G  5% /run
/dev/sda1                 18G   7.9G  8.9G  47% /
tmpfs                     3.9G   3.8M  3.9G  1% /dev/shm
tmpfs                     5.0M   0  5.0M  0% /run/lock
tmpfs                     3.9G   0  3.9G  0% /sys/fs/cgroup
none                      3.9G   8.0K  3.9G  1% /tmp
none                      3.9G   0  3.9G  0% /run/shm
/dev/sdb1                 9.9G  846M  8.6G  9% /var/log
/dev/sdb2                 40G   15G   23G  39% /opt/platform
cgmfs                     100K   0  100K  0% /run/cgmanager/fs
/dev/sdc1                 50G   9.2G   38G  20% /backups
[...]

```

diag monitor

```

platform@VOSS:~$ diag monitor
2021-02-15 12:59:32.325225
 4810.5 Mb used, 2252.8 Mb free, 291 processes
Private + Shared = RAM used Files Sockets CPU usage Program
-0.3 + 0.8 = 0.5 4 0 0 launch.true (20)
-0.0 + 0.5 = 0.5 10 3 0 qmgr
-0.0 + 0.6 = 0.6 10 3 0 pickup
0.0 + 0.2 = 0.2 3 10 0 dbus-daemon
0.0 + 0.1 = 0.1 7 7 0 systemd-udev
0.1 + 0.2 = 0.3 4 0 0 rbash
0.1 + 0.2 = 0.3 3 21 0 ntpd
0.7 + 0.7 = 1.4 4 3 0 (sd-pam)
[...]

```


diag monitor list

```
platform@VOSS:~$ diag monitor list
The following monitor findings were recorded:
-----
launch.true:
  max num_threads = 21 at 2019-10-23 07:10:03.2

systemd:
  avg num_sockets = 31
  max num_sockets = 33 at 2020-10-02 17:30:03.1
  avg num_files = 22
  max num_files = 27 at 2020-10-02 17:30:03.2

metric-collecto:
  max num_files = 27 at 2021-02-04 11:00:03.8

dockerd-ce:
  max mem_private = 94 at 2021-01-27 17:10:03.2
  avg num_sockets = 58
  max num_sockets = 65 at 2019-10-23 08:30:04.8
  avg num_files = 67
  max num_files = 78 at 2019-10-23 07:20:03.5
  max mem_total = 94 at 2021-01-27 17:10:03.2

scripts.py:
  max mem_private = 55 at 2019-10-23 06:20:03.4
  max mem_total = 58 at 2019-10-23 06:20:03.4

nginx:
[...]
```

diag nicstat

```
platform@VOSS:~$ diag nicstat
  Time      Int      rKB/s    wKB/s    rPk/s    wPk/s    rAvs    wAvs %Util    Sat
13:01:28 veth64e996f  1.23    6.86    4.29    2.98    292.7  2357.4  0.00  0.00
13:01:28   ens32    0.06    0.60    0.22    0.38    300.5  1621.6  0.00  0.00
13:01:28 veth7b06047  0.00    0.00    0.00    0.00    0.00   72.52  0.00  0.00
13:01:28 veth143e619  0.45    4.08    0.74    0.68    621.0  6172.7  0.00  0.00
13:01:28    lo    0.01    0.01    0.24    0.24    61.94   61.94  0.00  0.00
13:01:28 veth8b143f0 29.02    8.53   28.12   50.87   1056.9  171.8  0.02  0.00
13:01:28 veth6774425  0.04    0.00    0.00    0.00  15564.9  98.22  0.00  0.00
13:01:28 veth7f61555  0.00    0.00    0.00    0.00    0.00   72.52  0.00  0.00
13:01:28 veth1694cf5  0.00    0.00    0.00    0.00    0.00   72.11  0.00  0.00
13:01:28 veth655de40  0.03    0.05    0.28    0.15   123.8   324.7  0.00  0.00
13:01:28 veth76b16c9  0.04    0.09    0.05    0.06   855.2  1523.7  0.00  0.00
13:01:28 docker0  36.91   37.69   78.75   78.69   480.0   490.4  0.00  0.00
13:01:28 veth6f85a4b  0.35    3.94    1.63    1.17   218.6  3460.7  0.00  0.00
13:01:28 veth41ddd78  6.82   14.12   43.64   22.78   160.0   634.7  0.01  0.00
```

diag free

```
platform@VOSS:~$ diag free
          total      used      free      shared  buff/cache  available
Mem:      8174732    5024144    127472     59840    3023116    2714864
Swap:     2096124     270168    1825956
```

diag iostat

```
platform@VOSS:~$ diag iostat
Linux 4.4.0-151-generic (VOSS)          02/15/2021      _x86_64_      (4 CPU)

avg-cpu:  %user   %nice %system %iowait  %steal   %idle
           1.86    0.00   1.08   0.09   0.00   96.97

Device:            tps    kB_read/s    kB_wrtn/s    kB_read    kB_wrtn
sdb                 2.31         37.31         6.40    47828330    8197356
sdc                 0.90        114.40         0.01   146633661    14176
sdd                 2.10         0.42        18.24     533207    23383024
sda                 1.40         26.37         6.98    33802387    8949616
dm-0                2.53         0.41        18.24     530915    23383024
```

diag top

```
platform@VOSS:~$ diag top
top - 12:48:53 up 14 days, 20:03,  1 user,  load average: 0.53, 0.48, 0.30
Tasks: 297 total,  2 running, 295 sleeping,  0 stopped,  0 zombie
%Cpu(s):  1.9 us,  1.0 sy,  0.0 ni, 97.0 id,  0.1 wa,  0.0 hi,  0.1 si,  0.0 st
KiB Mem : 8174732 total, 149012 free, 5032732 used, 2992988 buff/cache
KiB Swap: 2096124 total, 1825764 free, 270360 used. 2706380 avail Mem

   PID USER      PR  NI  VIRT  RES  SHR S  %CPU  %MEM    TIME+  COMMAND
3165010 root        20   0   40656 3868 3180 R   11.1   0.0   0:00.04 top
   3783 snmp        20   0   59296 5968 2960 S    5.6   0.1  27:33.32 snmpd
   4915 root        20   0   15944 3184 1352 S    5.6   0.0 134:38.96 docker-pt+
   8213 www-data   20   0 649084 332456 9660 S    5.6   4.1 445:08.94 python
     1 root        20   0   38036 4952 3108 S    0.0   0.1   1:07.13 systemd
     2 root        20   0      0      0      0 S    0.0   0.0   0:00.41 kthreadd
```

diag config platform

```
platform@VOSS:~$ diag config platform
platform:
  backup:
    default: localbackup
    locations:
      localbackup:
        uri: file:///backups
        passphrase: sftp://nobody:*****@nowhere/
  credstore: *****
  disa:
    inactive_lock_days: 35
    max_password_age: 60
    max_simultaneous_logins: 10
```

(continues on next page)

(continued from previous page)

```

|           |-{in:impstats}
|           |-{in:imuxsock}
|           `-{rs:main Q:Reg}
|-snmpd
|-snmptrapd
|-sshd---sshd---sshd---rssh---diag---sh---sudo---ui-real.py---execute---
↪pstree
|-stunnel4
|-systemd---(sd-pam)
|-systemd-journal
|-systemd-logind
|-systemd-udev
`-vmttoolsd---{gmain}

```

diag stats

```

platform@VOSS:~$ diag stats
CONTAINER ID   NAME                                CPU %   MEM USAGE / LIMIT   MEM %   NET I/O
↪   BLOCK I/O      PIDS
selenium       selenium                            0.18%   58.8MiB / 7.796GiB   0.74%   32.1kB / 0B
↪   78.9MB / 0B     30
selfservice_   selfservice_node                    0.00%   101.9MiB / 7.796GiB  1.28%   32.1kB / 0B
↪   32.6MB / 0B     31
voss-portal    voss-portal                          0.00%   14.81MiB / 7.796GiB  0.19%   349kB / 58.
↪9MB  31MB / 0B       5
voss-wsgi      voss-wsgi                            0.90%   570.2MiB / 7.796GiB  7.14%   5.17GB /
↪456MB 186MB / 0B     10
voss-monitor   voss-monitoring                     1.29%   189MiB / 7.796GiB   2.37%   62.5MB / 43.
↪8MB  34.8MB / 0B     23
voss-risapi_   voss-risapi_collector                0.10%   106MiB / 7.796GiB   1.33%   5.35GB /
↪589MB 4.71MB / 0B     3
voss-queue     voss-queue                           3.57%   315.3MiB / 7.796GiB  3.95%   18.5GB / 8.
↪96GB  5.89MB / 0B     4
voss-cnf_col   voss-cnf_collector                   0.15%   105.1MiB / 7.796GiB  1.32%   9.02GB / 1.
↪61GB  31MB / 0B       3
mongoarbiter   mongoarbiter                         1.45%   31.75MiB / 7.796GiB  0.40%   32.3kB / 0B
↪   151MB / 520kB  27
mongodb        mongodb                              3.22%   2.98GiB / 7.796GiB   38.23%  11.2GB / 38.
↪1GB  1.2GB / 37.4GB  109
nginx          nginx                                 0.03%   21.03MiB / 7.796GiB  0.26%   118MB / 56.
↪1MB  28.4MB / 0B     6

```

diag tasks

```

platform@VOSS:~$ diag tasks
top - 13:04:13 up 14 days, 20:18,  1 user,  load average: 0.07, 0.30, 0.27
Tasks: 294 total,  1 running, 293 sleeping,  0 stopped,  0 zombie
%Cpu(s):  1.9 us,  1.0 sy,  0.0 ni, 97.0 id,  0.1 wa,  0.0 hi,  0.1 si,  0.0 st
KiB Mem : 8174732 total, 390904 free, 5046048 used, 2737780 buff/cache
KiB Swap: 2096124 total, 1828992 free, 267132 used. 2655756 avail Mem

   PID USER      PR  NI   VIRT   RES   SHR  S  %CPU  %MEM     TIME+ COMMAND
3168836 root        20   0  40680   3916  3200  R   22.2   0.0   0:00.07 /usr/bin/top -b -
↪n 1 -c -w 230

```

(continues on next page)

(continued from previous page)

```

6680 www-data 20 0 578060 8168 5536 S 5.6 0.1 136:09.89 /opt/voss-
↳deviceapi/go/bin/sysmon
8213 www-data 20 0 649084 332120 9332 S 5.6 4.1 445:41.70 /usr/bin/python /
↳opt/voss-deviceapi/bin/python /opt/voss-deviceapi/src/deviceapi/background/queue.py
↳30
1 root 20 0 38036 4896 3108 S 0.0 0.1 1:07.29 /sbin/init
2 root 20 0 0 0 0 S 0.0 0.0 0:00.41 [kthreadd]
3 root 20 0 0 0 0 S 0.0 0.0 0:35.52 [ksoftirqd/0]
7 root 20 0 0 0 0 S 0.0 0.0 6:54.66 [rcu_sched]
8 root 20 0 0 0 0 S 0.0 0.0 0:00.00 [rcu_bh]
9 root rt 0 0 0 0 S 0.0 0.0 0:08.10 [migration/0]
10 root rt 0 0 0 0 S 0.0 0.0 0:04.14 [watchdog/0]
11 root rt 0 0 0 0 S 0.0 0.0 0:04.03 [watchdog/1]
12 root rt 0 0 0 0 S 0.0 0.0 0:07.98 [migration/1]
13 root 20 0 0 0 0 S 0.0 0.0 0:34.08 [ksoftirqd/1]
15 root 0 -20 0 0 0 S 0.0 0.0 0:00.00 [kworker/1:0H]
[...]
```

diag top

```

platform@VOSS:~$ diag top
top - 13:04:43 up 14 days, 20:19, 1 user, load average: 0.04, 0.27, 0.26
Tasks: 295 total, 1 running, 294 sleeping, 0 stopped, 0 zombie
%Cpu(s): 1.9 us, 1.0 sy, 0.0 ni, 97.0 id, 0.1 wa, 0.0 hi, 0.1 si, 0.0 st
KiB Mem : 8174732 total, 379256 free, 5053288 used, 2742188 buff/cache
KiB Swap: 2096124 total, 1835344 free, 260780 used. 2644160 avail Mem

  PID USER      PR  NI   VIRT   RES   SHR  S  %CPU  %MEM     TIME+ COMMAND
3168863 root      20   0  40656   3860  3176  R  23.5   0.0   0:00.06 top
1 root     20   0  38036   4896  3108  S   0.0   0.1   1:07.29 systemd
2 root     20   0     0     0     0  S   0.0   0.0   0:00.41 kthreadd
3 root     20   0     0     0     0  S   0.0   0.0   0:35.52 ksoftirq+
7 root     20   0     0     0     0  S   0.0   0.0   6:54.69 rcu_sched
8 root     20   0     0     0     0  S   0.0   0.0   0:00.00 rcu_bh
9 root     rt    0     0     0     0  S   0.0   0.0   0:08.10 migratio+
10 root    rt    0     0     0     0  S   0.0   0.0   0:04.14 watchdog+
11 root    rt    0     0     0     0  S   0.0   0.0   0:04.03 watchdog+
12 root    rt    0     0     0     0  S   0.0   0.0   0:07.98 migratio+
13 root    20   0     0     0     0  S   0.0   0.0   0:34.08 ksoftirq+
[...]
```

diag vmstat

```

platform@VOSS:~$ diag vmstat
procs -----memory----- --swap-- -----io----- -system-- -----cpu-----
 r b swpd free buff cache si so bi bo in cs us sy id wa st
1 0 260780 378512 273412 2468844 0 0 45 8 4 6 2 1 97 0 0
```

Index

A

app

- app install, 45, 55
- app install nrs, 162
- app list, 45, 55
- app multi_install, 46
- app remove-role, 34
- app start, 44
- app status, 43, 44, 92
- app stop, 44
- app template, 7

B

backup

- backup add, 47, 149, 153
- backup clean, 153
- backup create, 150, 151, 153, 247
- backup default, 149
- backup import, 157
- backup list, 149, 157, 247
- backup passphrase, 150, 156, 157
- backup restore, 157

C

cluster, 49

- cluster add, 32, 220
- cluster check, 62
- cluster del, 32, 40, 203, 205, 208, 210, 220, 228, 230, 232, 235
- cluster job, 48
- cluster list, 32, 210, 232
- cluster prepnode, 5, 205, 208, 210, 218, 230, 241
- cluster primary, 38
- cluster primary role, 34
- cluster primary role application, 38
- cluster primary role database, 38
- cluster provision, 7, 32, 39, 203, 205, 208, 210, 214, 218, 220, 223, 228, 230, 232, 235, 238, 241, 244
- cluster run, 5, 42, 44, 48, 149, 160, 203, 205, 208, 210, 214, 218, 223, 228, 230, 232, 235, 238, 241, 244
- cluster status, 32, 191, 220

- cluster upgrade, 7, 47, 55

D

database

- database config, 198, 203, 205, 208, 210, 214, 218, 220, 223, 225, 228, 232, 235, 238, 241, 244
- database primary, 210, 232
- database user, 174
- database weight, 191, 203, 205, 208, 210, 214, 218, 220, 223, 228, 232, 235, 238, 241, 244

diag, 298

- diag config app snmp, 301
- diag config platform, 300
- diag disk, 55, 91, 92, 298
- diag filehash, 91, 162
- diag free, 91, 92, 299
- diag health, 91, 92
- diag iostat, 300
- diag mem, 91
- diag monitor, 298
- diag monitor list, 299
- diag nicstat, 299
- diag ping, 91
- diag proc, 301
- diag resolve, 91
- diag stats, 302
- diag tasks, 302
- diag test_connection, 91
- diag top, 91, 92, 300, 303
- diag unittests, 91
- diag vmstat, 303

drives

- drives add, 56, 150
- drives alignswap, 56
- drives checkswap, 56
- drives create_volume, 56
- drives list, 56, 152, 247
- drives offset, 56
- drives reassign, 56, 157, 247
- drives remove_logical, 34, 56
- drives remove_volume, 34, 56

H

health, 5, 96, 146, 182

help, 5

K

keys

keys add, 171

keys create, 69, 171

keys send, 69, 171

L

log

log audit, 67, 77

log audit;log audit ruleset, 74

log cert, 89

log collect, 69

log event, 67, 77

log follow, 68, 77, 193

log list, 67, 69

log purge, 67

log send, 69

log send output, 69

log sendnewer, 69

log ssl, 89

log stream, 77, 87

log view, 68, 77, 92

M

mail

mail del, 90

mail list, 90

mail read, 90

N

network

network <interface-name>, 40

network container range, 42

network dns, 41, 135

network domain, 41, 136

network interfaces, 40

network name, 40

network ntp, 40

network routes, 40

network search, 41

notify

notify add, 96, 99

notify emailrelay, 65, 69, 96

notify test, 96

S

schedule

schedule add, 65, 66, 117, 145, 151, 153

schedule enable, 145, 151

schedule time, 65, 66, 117, 145, 151

screen, 7

security

security update, 61, 160, 193

selfservice

selfservice get_translation_template,
50

selfservice import_translation, 50

snmp

snmp contact, 99

snmp list, 99

snmp load, 96, 99

snmp location, 99

snmp name, 99

snmp query, 99

system, 5

system banner, 59

system boot password, 53

system checksum, 59

system download, 55

system fips, 54

system history, 66

system id, 153

system inactivelock, 163

system password, 5, 52

system reboot, 52, 54

system shutdown, 52

system ssh, 168, 170

system ssh algorithm, 171

system ssh_session_limit, 52

U

user

user add, 52, 72, 166

user addkey, 166, 172

user credential_policy, 170

user del, 72, 166

user grant, 72, 167

user inactivelock, 163

user lastlogon, 163

user list, 72, 163, 166, 167, 170

user lock, 163

user password, 150, 163, 168

user password expiry, 163

user password history, 163

user passwordinfo, 163

user revoke, 72, 167

user sftp add, 166

user sftp add_nopass, 166

user sftp password, 166

user unlock, 163

V

voss, 8

- voss clear_device_pending_changes, 20
- voss cleardown, 9, 193
- voss db_collection_cap, 10
- voss db_collection_cap_check, 10
- voss db_collection_stats, 9, 150
- voss db_index_stats, 10
- voss export, 28, 29
- voss finalize_transaction, 19, 203, 205, 208, 210, 214, 218, 223, 228, 230, 232, 235, 238, 241, 244
- voss get_extra_functions_version, 9, 19
- voss migrate_summary_attributes, 9, 19, 48
- voss post-upgrade-migrations, 20
- voss queues, 17, 230, 238, 241, 244
- voss report api, 26
- voss report transaction, 22
- voss reset_device_concurrency, 19
- voss session-limits, 15, 138
- voss set_debug, 20, 250
- voss throttle-rates, 13, 138
- voss transaction archive, 12, 21, 116, 117
- voss transaction count, 11, 21
- voss transaction delete, 12, 21, 116, 117
- voss transaction export, 12, 21
- voss unlock_sysadmin_account, 20, 72
- voss update_device_schemas, 20
- voss upgrade_db, 193
- voss worker, 18
- voss workers, 17, 58, 208, 218, 230, 235, 241
- voss export
 - voss export group, 7
 - voss export type, 7
- voss subscriber_data_export, 7

W

- web
 - web cert, 178, 180, 184
 - web external, 188
 - web service, 51, 193
 - web ssl, 184, 186
 - web weight, 203, 205, 208, 210, 228, 230, 232, 235
- web external
 - web external add; web external del; web external list, 188
- web ssl
 - web ssl cipher, 186