



Virtuozzo Hybrid Infrastructure 3.5

Administrator's Command Line Guide

July 21, 2020

Virtuozzo International GmbH

Vordergasse 59

8200 Schaffhausen

Switzerland

Tel: + 41 52 632 0411

Fax: + 41 52 672 2010

<https://virtuozzo.com>

Copyright ©2016-2020 Virtuozzo International GmbH. All rights reserved.

This product is protected by United States and international copyright laws. The product's underlying technology, patents, and trademarks are listed at <https://virtuozzo.com>.

Microsoft, Windows, Windows Server, Windows NT, Windows Vista, and MS-DOS are registered trademarks of Microsoft Corporation.

Apple, Mac, the Mac logo, Mac OS, iPad, iPhone, iPod touch, FaceTime HD camera and iSight are trademarks of Apple Inc., registered in the US and other countries.

Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective owners.

Contents

1. Introduction	1
1.1 Providing Credentials	2
1.2 Managing Tasks	2
2. Managing Storage Cluster	4
2.1 Managing Tokens	4
2.1.1 vinfra node token show	4
2.1.2 vinfra node token create	4
2.1.3 vinfra node token validate	5
2.2 Managing Traffic Types and Networks	5
2.2.1 vinfra cluster traffic-type create	5
2.2.2 vinfra cluster traffic-type list	6
2.2.3 vinfra cluster traffic-type show	7
2.2.4 vinfra cluster traffic-type set	7
2.2.5 vinfra cluster traffic-type delete	8
2.2.6 vinfra cluster network create	8
2.2.7 vinfra cluster network list	9
2.2.8 vinfra cluster network show	9
2.2.9 vinfra cluster network set	10
2.2.10 vinfra cluster network set-bulk	11
2.2.11 vinfra cluster network delete	12
2.3 Managing Storage Nodes	12
2.3.1 vinfra node join	12
2.3.2 vinfra node list	14
2.3.3 vinfra node show	14
2.3.4 vinfra node maintenance precheck	15
2.3.5 vinfra node maintenance start	16

2.3.6	vinfra node maintenance status	17
2.3.7	vinfra node maintenance stop	18
2.3.8	vinfra node release	19
2.3.9	vinfra node forget	20
2.4	Managing Node Network Interfaces	20
2.4.1	vinfra node iface list	20
2.4.2	vinfra node iface show	21
2.4.3	vinfra node iface up	22
2.4.4	vinfra node iface down	23
2.4.5	vinfra node iface set	25
2.4.6	vinfra node iface create-bond	27
2.4.7	vinfra node iface create-vlan	30
2.4.8	vinfra node iface delete	33
2.5	Managing Node Disks	34
2.5.1	vinfra node disk list	34
2.5.2	vinfra node disk show	34
2.5.3	vinfra node disk assign	35
2.5.4	vinfra node disk release	37
2.5.5	vinfra node disk blink on	38
2.5.6	vinfra node disk blink off	38
2.5.7	vinfra node iscsi target add	39
2.5.8	vinfra node iscsi target delete	40
2.6	Creating and Deleting the Storage Cluster	41
2.6.1	vinfra cluster create	41
2.6.2	vinfra cluster delete	42
2.7	Showing Storage Cluster Overview and Details	43
2.7.1	vinfra cluster overview	43
2.7.2	vinfra cluster show	44
3.	Managing Compute Cluster	46
3.1	Creating and Deleting the Compute Cluster	46
3.1.1	vinfra service compute create	46
3.1.2	vinfra service compute delete	48
3.2	Showing Compute Cluster Details and Overview	48
3.2.1	vinfra service compute show	48
3.2.2	vinfra service compute stat	50

3.3	Changing Compute Cluster Parameters	51
3.4	Managing Compute Nodes	52
3.4.1	vinfra service compute node add	52
3.4.2	vinfra service compute node list	53
3.4.3	vinfra service compute node show	53
3.4.4	vinfra service compute node fence	54
3.4.5	vinfra service compute node unfence	54
3.4.6	vinfra service compute node release	55
3.5	Managing Virtual Machines	56
3.5.1	vinfra service compute server create	56
3.5.2	vinfra service compute server list	58
3.5.3	vinfra service compute server show	59
3.5.4	vinfra service compute server stat	60
3.5.5	vinfra service compute server set	60
3.5.6	vinfra service compute server iface attach	62
3.5.7	vinfra service compute server iface list	62
3.5.8	vinfra service compute server iface detach	63
3.5.9	vinfra service compute server volume attach	63
3.5.10	vinfra service compute server volume list	64
3.5.11	vinfra service compute server volume show	65
3.5.12	vinfra service compute server volume detach	65
3.5.13	vinfra service compute server log	66
3.5.14	vinfra service compute server migrate	66
3.5.15	vinfra service compute server resize	67
3.5.16	vinfra service compute server start	67
3.5.17	vinfra service compute server pause	67
3.5.18	vinfra service compute server unpause	68
3.5.19	vinfra service compute server suspend	68
3.5.20	vinfra service compute server resume	69
3.5.21	vinfra service compute server reboot	69
3.5.22	vinfra service compute server reset-state	69
3.5.23	vinfra service compute server stop	70
3.5.24	vinfra service compute server shelve	70
3.5.25	vinfra service compute server unshelve	71
3.5.26	vinfra service compute server evacuate	71

3.5.27	vinfra service compute server delete	71
3.6	Managing Images	72
3.6.1	vinfra service compute image create	72
3.6.2	vinfra service compute image list	73
3.6.3	vinfra service compute image show	74
3.6.4	vinfra service compute image set	74
3.6.5	vinfra service compute image save	76
3.6.6	vinfra service compute image delete	76
3.7	Managing Placements	76
3.7.1	vinfra service compute placement create	76
3.7.2	vinfra service compute placement assign	77
3.7.3	vinfra service compute placement delete assign	78
3.7.4	vinfra service compute placement list	79
3.7.5	vinfra service compute placement show	79
3.7.6	vinfra service compute placement update	80
3.7.7	vinfra service compute placement delete	80
3.8	Managing Flavors	81
3.8.1	vinfra service compute flavor create	81
3.8.2	vinfra service compute flavor list	81
3.8.3	vinfra service compute flavor show	82
3.8.4	vinfra service compute flavor delete	82
3.9	Managing Compute SSH Keys	83
3.9.1	vinfra service compute key create	83
3.9.2	vinfra service compute key list	84
3.9.3	vinfra service compute key show	84
3.9.4	vinfra service compute key delete	85
3.10	Managing Virtual Networks	85
3.10.1	vinfra service compute network create	85
3.10.2	vinfra service compute network list	87
3.10.3	vinfra service compute network show	87
3.10.4	vinfra service compute network set	88
3.10.5	vinfra service compute network delete	89
3.11	Managing Virtual Routers	89
3.11.1	vinfra service compute router create	89
3.11.2	vinfra service compute router list	91

3.11.3	vinfra service compute router show	91
3.11.4	vinfra service compute router set	92
3.11.5	vinfra service compute router iface add	93
3.11.6	vinfra service compute router iface list	94
3.11.7	vinfra service compute router iface remove	94
3.11.8	vinfra service compute router delete	95
3.12	Managing Floating IP Addresses	95
3.12.1	vinfra service compute floatingip create	95
3.12.2	vinfra service compute floatingip list	96
3.12.3	vinfra service compute floatingip show	97
3.12.4	vinfra service compute floatingip set	97
3.12.5	vinfra service compute floatingip delete	98
3.13	Managing Load Balancers	99
3.13.1	vinfra service compute load-balancer create	99
3.13.2	vinfra service compute load-balancer list	100
3.13.3	vinfra service compute load-balancer show	101
3.13.4	vinfra service compute load-balancer stats	101
3.13.5	vinfra service compute load-balancer set	102
3.13.6	vinfra service compute load-balancer pool create	103
3.13.7	vinfra service compute load-balancer pool list	106
3.13.8	vinfra service compute load-balancer pool show	107
3.13.9	vinfra service compute load-balancer pool set	108
3.13.10	vinfra service compute load-balancer pool delete	111
3.13.11	vinfra service compute load-balancer delete	111
3.14	Managing Volumes	111
3.14.1	vinfra service compute volume create	111
3.14.2	vinfra service compute volume list	113
3.14.3	vinfra service compute volume show	113
3.14.4	vinfra service compute volume set	114
3.14.5	vinfra service compute volume extend	116
3.14.6	vinfra service compute volume delete	116
3.15	Managing Volume Snapshots	116
3.15.1	vinfra service compute volume snapshot create	116
3.15.2	vinfra service compute volume snapshot list	117
3.15.3	vinfra service compute volume snapshot show	118

3.15.4	vinfra service compute volume snapshot set	118
3.15.5	vinfra service compute volume snapshot upload-to-image	119
3.15.6	vinfra service compute volume snapshot revert	120
3.15.7	vinfra service compute volume snapshot reset-state	120
3.15.8	vinfra service compute volume snapshot delete	121
3.16	Managing Storage Policies	121
3.16.1	vinfra cluster storage-policy create	122
3.16.2	vinfra cluster storage-policy list	123
3.16.3	vinfra cluster storage-policy show	123
3.16.4	vinfra cluster storage-policy set	124
3.16.5	vinfra cluster storage-policy delete	125
3.17	Managing Kubernetes Clusters	125
3.17.1	vinfra service compute k8saas create	125
3.17.2	vinfra service compute k8saas list	128
3.17.3	vinfra service compute k8saas config	128
3.17.4	vinfra service compute k8saas show	128
3.17.5	vinfra service compute k8saas set	129
3.17.6	vinfra service compute k8saas delete	130
3.18	Managing Compute Quotas	131
3.18.1	vinfra service compute quotas show	131
3.18.2	vinfra service compute quotas update	131
3.19	Managing Updates	132
3.19.1	vinfra software-updates check-for-updates	132
3.19.2	vinfra software-updates eligibility-check	133
3.19.3	vinfra software-updates download	134
3.19.4	vinfra software-updates start	135
3.19.5	vinfra software-updates pause	136
3.19.6	vinfra software-updates resume	136
3.19.7	vinfra software-updates cancel	137
3.19.8	vinfra software-updates status	138
4.	Managing Backup Cluster	139
4.1	Creating, Showing, and Deleting Backup Cluster	139
4.1.1	vinfra service backup cluster create	139
4.1.2	vinfra service backup cluster show	145
4.1.3	vinfra service backup cluster release	145

4.2	Managing Backup Nodes	146
4.2.1	vinfra service backup node add	146
4.2.2	vinfra service backup node list	147
4.2.3	vinfra service backup node release	148
4.3	Updating Backup Cluster Certificates	148
4.4	Changing Storage Parameters	149
4.4.1	vinfra service backup storage-params show	149
4.4.2	vinfra service backup storage-params change	150
4.5	Changing Volume Parameters	153
4.5.1	vinfra service backup volume-params show	153
4.5.2	vinfra service backup volume-params change	154
4.6	Managing Backup Cluster Geo-replication	155
4.6.1	vinfra service backup geo-replication show	156
4.6.2	vinfra service backup geo-replication master setup	156
4.6.3	vinfra service backup geo-replication master download-configs	157
4.6.4	vinfra service backup geo-replication slave setup	158
4.6.5	vinfra service backup geo-replication master establish	158
4.6.6	vinfra service backup geo-replication slave update-certificates	159
4.6.7	vinfra service backup geo-replication master disable	160
4.6.8	vinfra service backup geo-replication slave promote-to-master	160
4.6.9	vinfra service backup geo-replication slave cancel	161
4.6.10	vinfra service backup geo-replication master cancel	162
5.	Managing General Settings	163
5.1	Managing Licenses	163
5.1.1	vinfra cluster license load	163
5.1.2	vinfra cluster license show	164
5.1.3	vinfra cluster license update	164
5.2	Managing Domains	165
5.2.1	vinfra domain create	165
5.2.2	vinfra domain list	166
5.2.3	vinfra domain show	166
5.2.4	vinfra domain set	167
5.2.5	vinfra domain delete	167
5.3	Managing Domain Users	168
5.3.1	vinfra domain user list-available-roles	168

5.3.2	vinfra domain user create	169
5.3.3	vinfra domain user list	170
5.3.4	vinfra domain user show	171
5.3.5	vinfra domain user set	171
5.3.6	vinfra domain user delete	173
5.4	Managing Domain Projects	173
5.4.1	vinfra domain project create	173
5.4.2	vinfra domain project list	174
5.4.3	vinfra domain project show	175
5.4.4	vinfra domain project set	175
5.4.5	vinfra domain project user list	176
5.4.6	vinfra domain project user remove	177
5.4.7	vinfra domain project delete	177
5.5	Managing SSH Keys	178
5.5.1	vinfra cluster sshkey add	178
5.5.2	vinfra cluster sshkey list	179
5.5.3	vinfra cluster sshkey delete	180
5.6	Managing External DNS Servers	180
5.6.1	vinfra cluster settings dns show	180
5.6.2	vinfra cluster settings dns set	181
5.7	Configuring Management Node High Availability	182
5.7.1	vinfra cluster ha create	182
5.7.2	vinfra cluster ha update	183
5.7.3	vinfra cluster ha show	185
5.7.4	vinfra cluster ha delete	185
5.8	Managing Cluster Backups	186
5.8.1	vinfra cluster backup create	186
5.8.2	vinfra cluster backup show	187
5.9	Managing Storage Tier Encryption	187
5.9.1	vinfra cluster settings encryption show	187
5.9.2	vinfra cluster settings encryption set	188
5.10	Managing Automatic Storage Disk Configuration	188
5.10.1	vinfra cluster settings automatic-disk-replacement show	188
5.10.2	vinfra cluster settings automatic-disk-replacement set	189
5.11	Managing Alerts	190

5.11.1	vinfra cluster alert list	190
5.11.2	vinfra cluster alert show	190
5.11.3	vinfra cluster alert delete	191
5.12	Managing Audit Log	192
5.12.1	vinfra cluster auditlog list	192
5.12.2	vinfra cluster auditlog show	193
5.13	Sending Problem Reports	194
6.	Monitoring Storage Cluster	195
6.1	Monitoring General Storage Cluster Parameters	195
6.2	Monitoring Metadata Servers	198
6.3	Monitoring Chunk Servers	199
6.3.1	Understanding Disk Space Usage	202
6.3.1.1	Understanding Allocatable Disk Space	203
6.3.1.2	Viewing Space Occupied by Data Chunks	204
6.3.2	Exploring Chunk States	205
6.4	Monitoring Clients	206
6.5	Monitoring Physical Disks	208
6.6	Monitoring Event Logs	210
6.7	Monitoring Replication Parameters	213
7.	Accessing Storage Clusters via iSCSI	215
7.1	iSCSI Workflow Overview	217
7.2	Configuring CLI Tool	217
7.3	Managing Target Groups	218
7.3.1	Creating Target Groups	218
7.3.2	Starting and Stopping Target Groups	220
7.3.3	Listing Target Groups	220
7.3.4	Printing Details of Target Groups	221
7.3.5	Managing Persistent Reservations of Target Groups	221
7.3.6	Deleting Target Groups	222
7.4	Managing iSCSI Volumes	222
7.4.1	Creating iSCSI Volumes	222
7.4.2	Listing and Printing Details of iSCSI Volumes	222
7.4.3	Attaching iSCSI Volumes to Target Groups	223
7.4.4	Setting the Active/Optimized Path for iSCSI Volumes	223

7.4.5	Viewing iSCSI Volume ALUA Information223
7.4.6	Viewing and Setting iSCSI Volume Parameters224
7.4.7	Increasing iSCSI Volume Size224
7.4.8	Setting iSCSI Volume Limits225
7.4.9	Detaching iSCSI Volumes from Target Groups225
7.4.10	Deleting iSCSI Volumes225
7.5	Managing Nodes in Target Groups225
7.5.1	Adding Nodes to Target Groups226
7.5.2	Listing Nodes in Target Groups226
7.5.3	Enabling and Disabling Nodes in Target Groups227
7.5.4	Deleting Nodes from Target Groups227
7.6	Managing Targets and Portals228
7.6.1	Creating Targets228
7.6.2	Adding and Removing Target Portals228
7.6.3	Deleting Targets229
7.7	Managing CHAP Accounts229
7.7.1	Creating and Listing CHAP Accounts229
7.7.2	Changing CHAP Account Details230
7.7.3	Assigning CHAP Accounts to Target Groups230
7.7.4	Deleting CHAP Accounts230
7.8	Managing LUN Views230
7.8.1	Creating LUN Views231
7.8.2	Listing LUN Views231
7.8.3	Changing LUN View Details231
7.8.4	Deleting LUN Views232
8.	Advanced Tasks233
8.1	Updating Kernel with ReadyKernel233
8.1.1	Installing ReadyKernel Patches Automatically234
8.1.2	Managing ReadyKernel Patches Manually234
8.1.2.1	Downloading, Installing, and Loading ReadyKernel Patches234
8.1.2.2	Loading and Unloading ReadyKernel Patches234
8.1.2.3	Installing and Removing ReadyKernel Patches for Specific Kernels235
8.1.2.4	Downgrading ReadyKernel Patches235
8.1.2.5	Disabling Loading of ReadyKernel Patches on Boot235
8.1.2.6	Managing ReadyKernel Logs236

8.2	Managing Guest Tools236
8.2.1	Installing Guest Tools236
8.2.1.1	Installing Guest Tools in New VMs237
8.2.1.2	Installing Guest Tools in Existing VMs238
8.2.2	Uninstalling Guest Tools240
8.2.2.1	Uninstalling Guest Tools from Linux VMs240
8.2.2.2	Uninstalling Guest Tools from Windows VMs240
8.3	Running Commands in Virtual Machines without Network Connectivity241
8.3.1	Running Commands in Linux Virtual Machines241
8.3.2	Running Commands in Windows Virtual Machines242
8.4	Setting Virtual Machines CPU Model243
8.5	Creating Linux Templates243
8.6	Connecting to OpenStack Command-Line Interface245
8.7	Setting a DNS Name for the Compute API247
8.8	Securing OpenStack API Traffic with SSL248
8.9	Using Metering for Compute Resources249
8.9.1	Enabling Metering Service251
8.9.2	Using Gnocchi Command-line Tool252
8.9.3	Viewing Resource Usage Per Project253
8.10	Configuring Memory Policy for Storage Services254
8.10.1	vinfra memory-policy vstorage-services per-cluster change255
8.10.2	vinfra memory-policy vstorage-services per-cluster show256
8.10.3	vinfra memory-policy vstorage-services per-cluster reset257
8.10.4	vinfra memory-policy vstorage-services per-node change257
8.10.5	vinfra memory-policy vstorage-services per-node show258
8.10.6	vinfra memory-policy vstorage-services per-node reset259
8.11	Migrating Virtual Machines from VMware vCenter260
8.11.1	Deploying the Appliance VM260
8.11.2	Setting Up Authentication in the Appliance VM261
8.11.3	Migrating Virtual Machines to Virtuozzo Hybrid Infrastructure Online262
8.11.4	Migrating Virtual Machines to Virtuozzo Hybrid Infrastructure Offline263
8.12	Changing the Default Load Balancer Flavor264
8.13	Configuring Retention Policy for Prometheus Metrics266
8.14	Viewing Cluster Logs266

CHAPTER 1

Introduction

This guide describes the syntax and parameters of the `vinfra` command-line tool that can be used to manage Virtuozzo Hybrid Infrastructure from console and automate such management tasks.

Note: While the following chapters provide information on specific operations that you can perform with `vinfra`, you can also run `vinfra help` to get a list of all supported commands and their descriptions. For help on a specific command, either run `vinfra help <command>` or `vinfra <command> --help`.

In addition, this guide describes how to use the command line to perform operations unsupported by `vinfra` as of now.

Note that the following operations should not be done from the command line:

- setting custom paths for Virtuozzo Hybrid Infrastructure services, in particular:
 - creating S3 clusters only in `/mnt/vstorage/vols/s3`
 - creating iSCSI targets only in `/mnt/vstorage/vols/iscsi`
- mounting clusters or change cluster mount options
- configuring firewall with `firewall-cmd`
- renaming network connections
- managing MDS/CS
- managing partitions, LVMs, or software RAID
- modifying files in `/mnt/vstorage/vols` and `/mnt/vstorage/webcp/backup` directories

- setting encoding or replication of cluster root

1.1 Providing Credentials

The `vinfra` CLI tool requires the following information:

- IP address or hostname of the management node (set to `backend-api.svc.vstoragedomain` by default).
- Username (`admin` by default).
- Password (created during installation of Virtuozzo Hybrid Infrastructure).

This information can be supplied via the `--vinfra-portal`, `--vinfra-username`, and `--vinfra-password` command-line parameters with each command. Alternatively, you can supply it by setting the environment variables `VINFRA_PORTAL`, `VINFRA_USERNAME`, and `VINFRA_PASSWORD` (e.g., in your `~/.bash_profile`). In this case, you will be able to run the CLI tool without the aforementioned command-line parameters.

As you typically run `vinfra` from the management node as `admin`, the only variable you usually need to set is the password. For example:

```
# export VINFRA_PASSWORD=12345
```

If you installed `vinfra` on a remote machine and/or run it as a different user, you will need to set `VINFRA_PORTAL` and/or `VINFRA_USERNAME` on that machine in addition to `VINFRA_PASSWORD`.

1.2 Managing Tasks

The `vinfra` CLI tool executes some commands immediately, while for other commands (that may take some time to complete) it creates system tasks that are queued. Examples of actions performed via tasks are creating the storage or compute cluster and adding nodes to it.

To keep track of tasks being performed by `vinfra`, use the `vinfra task list` and `vinfra task show` commands. For example:

```
# vinfra task list
+-----+-----+-----+
| task_id           | state  | name                                     |
+-----+-----+-----+
| 8fc27e7a-ba73-471d-9134-e351e1137cf4 | success | backend.tasks.cluster.CreateNewCluster |
| e61377db-9df4-4282-99aa-6a4ae73a7f96 | success | backend.tasks.disks.ApplyDiskRoleTask  |
| a005b748-cb85-40f8-a09d-291a8599bb9c | success | backend.tasks.node.AddNodeInClusterTask |
+-----+-----+-----+
```

```
+-----+-----+-----+
# vinfra task show 8fc27e7a-ba73-471d-9134-e351e1137cf4
+-----+-----+
| Field | Value |
+-----+-----+
| args  | - stor1 |
|       | - 7ffa9540-5a20-41d1-b203-e3f349d62565 |
|       | - null  |
|       | - null  |
| kwargs | {}      |
| name   | backend.tasks.cluster.CreateNewCluster |
| result | cluster_id: 1 |
| state  | success |
| task_id | 8fc27e7a-ba73-471d-9134-e351e1137cf4 |
+-----+-----+-----+
```


CHAPTER 2

Managing Storage Cluster

2.1 Managing Tokens

2.1.1 vinfra node token show

Display the backend token:

```
usage: vinfra node token show
```

Example:

```
# vinfra node token show
+-----+-----+
| Field | Value      |
+-----+-----+
| host  | 10.37.130.101 |
| token | dc56d4d2     |
| ttl   | 86398       |
+-----+-----+
```

This command shows the details of the current token.

2.1.2 vinfra node token create

Create the backend token:

```
usage: vinfra node token create [--ttl <ttl>]
```

--ttl <ttl>

Token TTL, in seconds

Example:

```
# vinfra node token create --ttl 86400
+-----+-----+
| Field | Value      |
+-----+-----+
| host  | 10.37.130.101 |
| token | dc56d4d2     |
| ttl   | 86398       |
+-----+-----+
```

This command creates a new token with the time to live (TTL) of 86400 seconds.

2.1.3 vinfra node token validate

Validate the backend token:

```
usage: vinfra node token validate <token>
```

<token>

Token value

Example:

```
# vinfra node token validate dc56d4d2
+-----+-----+
| Field | Value |
+-----+-----+
| status | valid |
+-----+-----+
```

This command validates the token dc56d4d2.

2.2 Managing Traffic Types and Networks

2.2.1 vinfra cluster traffic-type create

Create a new traffic type:

```
usage: vinfra cluster traffic-type create --port <port> <traffic-type-name>
```

--port <port>

Traffic type port

<traffic-type-name>
Traffic type name

Example:

```
# vinfra cluster traffic-type create "MyTrafficType" --port 6900
+-----+-----+
| Field      | Value                |
+-----+-----+
| exclusive  | False                |
| name       | MyTrafficType        |
| port       | 6900                  |
| type       | custom                |
+-----+-----+
```

This command creates a custom traffic type MyTrafficType on port 6900.

2.2.2 vinfra cluster traffic-type list

List available traffic types:

```
usage: vinfra cluster traffic-type list
```

Example:

```
# vinfra cluster traffic-type list
+-----+-----+-----+-----+
| name                | type          | exclusive | port |
+-----+-----+-----+-----+
| Storage             | predefined    | True      |      |
| Internal management | predefined    | True      |      |
| OSTOR private       | predefined    | True      |      |
| S3 public           | predefined    | False     |      |
| iSCSI               | predefined    | False     |      |
| NFS                 | predefined    | False     |      |
| ABGW private        | predefined    | True      |      |
| ABGW public         | predefined    | False     |      |
| Admin panel         | predefined    | False     |      |
| SSH                 | predefined    | False     |      |
| VM public           | predefined    | False     |      |
| VM private          | predefined    | True      |      |
| Compute API         | predefined    | True      |      |
| MyTrafficType       | custom        | False     | 6900 |
+-----+-----+-----+-----+
```

This command lists all traffic types in Virtuozzo Hybrid Infrastructure.

2.2.3 vinfra cluster traffic-type show

Show details of a traffic type:

```
usage: vinfra cluster traffic-type show <traffic-type>
```

<traffic-type>

Traffic type name

Example:

```
# vinfra cluster traffic-type show Storage
+-----+-----+
| Field  | Value  |
+-----+-----+
| exclusive | True  |
| name     | Storage |
| port     |       |
| type     | predefined |
+-----+-----+
```

This command shows the details of the traffic type Storage.

2.2.4 vinfra cluster traffic-type set

Modify traffic type parameters:

```
usage: vinfra cluster traffic-type set [--name <name>] [--port <port>] <traffic-type>
```

--name <name>

A new name for the traffic type

--port <port>

A new port for the traffic type

<traffic-type>

Traffic type name

Example:

```
# vinfra cluster traffic-type set "MyTrafficType" \
--name "MyOtherTrafficType" --port 6901
+-----+-----+
| Field  | Value  |
+-----+-----+
| exclusive | False  |
+-----+-----+
```

```
| name      | MyOtherTrafficType |
| port     | 6901                |
| type     | custom              |
+-----+-----+
```

This command renames the traffic type `MyTrafficType` to `MyOtherTrafficType` and changes its port to 6901.

2.2.5 vinfra cluster traffic-type delete

Delete a traffic type:

```
usage: vinfra cluster traffic-type delete <traffic-type>
```

<traffic-type>

Traffic type name

Example:

```
# vinfra cluster traffic-type delete "MyOtherTrafficType"
Operation successful
```

This command deletes the custom traffic type `MyOtherTrafficType`.

2.2.6 vinfra cluster network create

Create a new network:

```
usage: vinfra cluster network create [--traffic-types <traffic-types>] <network-name>
```

--traffic-types <traffic-types>

A comma-separated list of traffic type IDs or names

<network-name>

Network name

Example:

```
# vinfra cluster network create MyNet --traffic-types ssh
+-----+-----+
| Field | Value |
+-----+-----+
| id    | 03d5eeb3-1833-4626-885d-dd066635f5de |
| name  | MyNet |
| roles | - SSH |
```

```
| type | Custom |
+-----+
```

This command creates a custom network MyNet and assigns the traffic type SSH to it.

2.2.7 vinfra cluster network list

List available networks:

```
usage: vinfra cluster network list
```

Example:

```
# vinfra cluster network list
+-----+-----+-----+
| id | name | roles |
+-----+-----+-----+
| 358bdc39-cd8b-4565-8ebf-e7c12dcd1cf7 | Public | - ABGW public
| | | - iSCSI
| | | - NFS
| | | - S3 public
| | | - SSH
| | | - Admin Panel
| 6095a997-e5f1-493d-a750-41ddf277153b | Private | - ABGW private
| | | - Internal Management
| | | - OSTOR private
| | | - SSH
| | | - Storage
+-----+-----+-----+
```

This command lists all networks in Virtuozzo Hybrid Infrastructure.

2.2.8 vinfra cluster network show

Show details of a network:

```
usage: vinfra cluster network show <network>
```

<network>

Network ID or name

Example:

```
# vinfra cluster network show MyNet
+-----+-----+
| Field | Value |
+-----+-----+
```

```
+-----+-----+
| id     | 03d5eeb3-1833-4626-885d-dd066635f5de |
| name   | MyNet                                   |
| roles  | - SSH                                   |
| type   | Custom                                  |
+-----+-----+
```

This command shows the details of the custom network MyNet.

2.2.9 vinfra cluster network set

Modify network parameters:

```
usage: vinfra cluster network set [--name <network-name>]
                                   [--traffic-types <traffic-types> |
                                   --add-traffic-types <traffic-types> |
                                   --del-traffic-types <traffic-types>]
                                   <network>
```

--name <network-name>

Network name

--traffic-types <traffic-types>

A comma-separated list of traffic type names (overwrites network's current traffic types)

--add-traffic-types <traffic-types>

A comma-separated list of traffic type names (adds the specified traffic types to the network)

--del-traffic-types <traffic-types>

A comma-separated list of traffic type names (removes the specified traffic types from the network)

<network>

Network ID or name

Example:

```
# vinfra cluster network set MyNet --name MyOtherNet --add-traffic-types iscsi,nfs
+-----+-----+
| Field  | Value                                     |
+-----+-----+
| task_id | b29f6f66-37d7-47de-b02e-9f4087ad932b |
+-----+-----+
```

This command creates a task to rename the network MyNet to MyOtherNet and assign to it the traffic types iSCSI and NFS.

Task outcome:

```
# vinfra task show b29f6f66-37d7-47de-b02e-9f4087ad932b
+-----+-----+
| Field | Value |
+-----+-----+
| args  | - 03d5eeb3-1833-4626-885d-dd066635f5de |
| kwargs | name: MyOtherNet |
|        | roles: |
|        | - ssh |
|        | - iscsi |
|        | - nfs |
| name   | backend.presentation.network.roles.tasks.RolesSetChangeTask |
| result | id: 03d5eeb3-1833-4626-885d-dd066635f5de |
|        | name: MyOtherNet |
|        | roles: |
|        | - iSCSI |
|        | - NFS |
|        | - SSH |
|        | type: Custom |
| state  | success |
| task_id | b29f6f66-37d7-47de-b02e-9f4087ad932b |
+-----+-----+
```

2.2.10 vinfra cluster network set-bulk

Modify traffic types of multiple networks:

```
usage: vinfra cluster network set-bulk --network <network>:<traffic-types>
```

```
--network <network>:<traffic-types>
```

Network configuration in the format:

- <network>: network ID or name.
- <traffic-types>: a comma-separated list of traffic type names (this option can be used multiple times).

Example:

```
# vinfra cluster network set-bulk --network MyNet1:snmp --network MyNet2:ssh,snmp
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | c774f55d-c45b-42cd-ac9e-16fc196e9283 |
+-----+-----+
```

This command creates a task to change the traffic type set of the network MyNet1 to SNMP and that of MyNet2 to SSH and SNMP.

Task outcome:

```
# vinfra task show c774f55d-c45b-42cd-ac9e-16fc196e9283
+-----+
| Field | Value |
+-----+
| details |
| name | backend.presentation.network.roles.tasks.RolesSetBulkChangeTask |
| result | - id: adf49487-9deb-4180-bb0c-08a906257981 |
| | name: MyNet1 |
| | roles: |
| | - SNMP |
| | type: Custom |
| | - id: 3f6ff4a3-31bc-440b-a36f-d755c80d5932 |
| | name: MyNet2 |
| | roles: |
| | - SNMP |
| | - SSH |
| | type: Custom |
| state | success |
| task_id | c774f55d-c45b-42cd-ac9e-16fc196e9283 |
+-----+
```

2.2.11 vinfra cluster network delete

Delete a network:

```
usage: vinfra cluster network delete <network>
```

<network>

Network ID or name

Example:

```
# vinfra cluster network delete MyOtherNet
Operation successful
```

This command deletes the network MyOtherNet.

2.3 Managing Storage Nodes

2.3.1 vinfra node join

Join a node to the storage cluster:

```
usage: vinfra node join [--disk <disk>:<role>[:<key=value,...>]] <node>
```

```
--disk <disk>:<role> [:<key=value,...>]
```

Disk configuration in the format:

- <disk>: disk device ID or name
- <role>: disk role (cs, mds, journal, mds-journal, mds-system, cs-system, system)
- comma-separated key=value pairs with keys (optional):
 - tier: disk tier (0, 1, 2 or 3)
 - journal-tier: journal (cache) disk tier (0, 1, 2 or 3)
 - journal-type: journal (cache) disk type (no_cache, inner_cache or external_cache)
 - journal-disk: journal (cache) disk ID or device name
 - journal-size: journal (cache) disk size, in bytes
 - bind-address: bind IP address for the metadata service

E.g., sda:cs:tier=0,journal-type=inner_cache. This option can be used multiple times.

```
<node>
```

Node ID or hostname

Example:

```
# vinfra node join f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 \
--disk sda:mds-system \
--disk sdb:cs \
--disk sdc:cs
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | a2713068-9544-4ea1-8ec8-69a068cf86f2 |
+-----+-----+
```

This command creates a task to add the node with the ID f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 to the storage cluster and assigns roles to disks: mds-system to sda, cs to sdb and sdc.

Task outcome:

```
# vinfra task show a2713068-9544-4ea1-8ec8-69a068cf86f2
+-----+-----+
| Field | Value |
+-----+-----+
```

```

| args      | - f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
|          | - 1 |
| kwargs   | disks: |
|          | - id: 85F32403-94A9-465A-9E6C-C1A2B41294FC |
|          |   role: mds-system |
|          |   service_params: {} |
|          | - id: FE0B5876-E054-489B-B0FD-72429BEFD46A |
|          |   role: cs |
|          |   service_params: {} |
|          | - id: D3BEF4BB-AA3B-4DB6-9376-BC7CDA636700 |
|          |   role: cs |
|          |   service_params: {} |
| name     | backend.tasks.node.AddNodeInClusterTask |
| result   | {} |
| state    | success |
| task_id  | a2713068-9544-4ea1-8ec8-69a068cf86f2 |
+-----+-----+

```

2.3.2 vinfra node list

List storage nodes:

```
usage: vinfra node list
```

Example:

```

# vinfra node list
+-----+-----+-----+-----+-----+-----+
| id          | host          | is_primary | is_online | is_assigned | is_in_ha |
+-----+-----+-----+-----+-----+-----+
| 09bb6b84-70a5-<...> | node001.<...> | True      | True     | True      | False    |
| 187edb11-38c5-<...> | node002.<...> | False     | True     | True      | False    |
| e6255aed-d6e7-<...> | node003.<...> | False     | True     | True      | False    |
+-----+-----+-----+-----+-----+-----+

```

This command lists all nodes registered in Virtuozzo Hybrid Infrastructure (both unassigned and used in the storage cluster).

2.3.3 vinfra node show

Show storage node details:

```
usage: vinfra node show <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra node show 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+
| Field      | Value                                |
+-----+-----+
| cpu_cores  | 2                                    |
| host       | stor-1.example.com.vstagedomain.    |
| id         | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
| ipaddr     | stor-1.example.com.vstagedomain.    |
| is_assigned| False                                |
| is_in_ha   | False                                |
| is_installing| False                                |
| is_online  | True                                  |
| is_primary | True                                  |
| is_virt    | True                                  |
| mem_total  | 8201310208                           |
| roles      | management:                           |
|            |   is_primary: true                    |
| tasks      |                                         |
+-----+-----+
```

This command shows the details of the node with the ID 4f96acf5-3bc8-4094-bcb6-4d1953be7b55.

2.3.4 vinfra node maintenance precheck

Start node maintenance precheck:

```
usage: vinfra node maintenance precheck <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra node maintenance precheck 9dcc9632-911c-4cc5-9a89-5a6fa5db2314
+-----+-----+
| Field    | Value                                |
+-----+-----+
| task_id  | 7c7f0afa-10f4-41b7-9b2e-973f3d392178 |
+-----+-----+
```

This command creates a task to start maintenance precheck for the node with the ID f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4.

Task outcome:

```
# vinfra task show 7c7f0afa-10f4-41b7-9b2e-973f3d392178
+-----+-----+
```

```

| Field | Value |
+-----+-----+
| details |
| name | backend.business.models.maintenance.tasks.MaintenancePrecheckTask |
| result |
| state | success |
| task_id | 7c7f0afa-10f4-41b7-9b2e-973f3d392178 |
+-----+-----+

```

2.3.5 vinfra node maintenance start

Start node maintenance:

```

usage: vinfra node maintenance start [--iscsi-mode <mode>] [--compute-mode <mode>]
                                     [--s3-mode <mode>] [--storage-mode <mode>]
                                     [--alua-mode <mode>] [--nfs-mode <mode>] <node>

```

--iscsi-mode <mode>

Ignore ISCSI evacuation during maintenance (ignore).

--compute-mode <mode>

Ignore compute evacuation during maintenance (ignore).

--s3-mode <mode>

Ignore S3 evacuation during maintenance (ignore).

--storage-mode <mode>

Ignore storage evacuation during maintenance (ignore).

--alua-mode <mode>

Ignore Block Storage target groups during maintenance (ignore).

--nfs-mode <mode>

Ignore NFS evacuation during maintenance (ignore).

<node>

Node ID or hostname

Example:

```

# vinfra node maintenance start 9dcc9632-911c-4cc5-9a89-5a6fa5db2314 \
--iscsi-mode ignore --compute-mode ignore
+-----+-----+
| Field | Value |
+-----+-----+

```

```
| task_id | 3d4c23b6-9f62-412c-ad7c-ec9537a36fa7 |
+-----+-----+
```

This command creates a task to start maintenance for the node with the ID 9dcc9632-911c-4cc5-9a89-5a6fa5db2314 without evacuating its iSCSI and compute services.

Task outcome:

```
# vinfra task show 3d4c23b6-9f62-412c-ad7c-ec9537a36fa7
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.business.models.maintenance.tasks.MaintenanceStartTask |
| result | |
| state | success |
| task_id | 3d4c23b6-9f62-412c-ad7c-ec9537a36fa7 |
+-----+-----+
```

2.3.6 vinfra node maintenance status

Show node maintenance details:

```
usage: vinfra node maintenance status <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra node maintenance status 9dcc9632-911c-4cc5-9a89-5a6fa5db2314
+-----+-----+
| Field | Value |
+-----+-----+
| node_id | 9dcc9632-911c-4cc5-9a89-5a6fa5db2314 |
| params | compute_mode: ignore |
| | iscsi_mode: ignore |
| | nfs_mode: evacuate |
| | s3_mode: evacuate |
| | storage_mode: suspend |
| precheck | |
| resources | compute: |
| | failed: [] |
| | handled: [] |
| | initial: |
| | - id: fa69e3f1-461c-4f4a-b442-3ffb356130c1 |
| | status: ACTIVE |
| | - id: f336f631-cc82-43e6-a582-8dcfcd24a722 |
| | status: ACTIVE |
+-----+-----+
```

```

|           | - id: 48864c3f-d9bf-4c32-abdc-901395c3b5f9 |
|           |   status: ACTIVE                          |
|           | - id: 5fd82e2a-3fef-4171-bfa4-67daa99ae64f |
|           |   status: ACTIVE                          |
|           |   untouched: []                           |
|           | iscsi:                                     |
|           |   failed: []                              |
|           |   handled: []                             |
|           |   initial: []                             |
|           | nfs: null                                  |
| state    | suspended_complete                        |
| task     | flow: complete                            |
|           | id: fe8e5ff5-48c1-4a23-a533-28233aaae4db |
|           | state: success                            |
|           | updated_at: '2019-11-20T13:52:15.849492' |
+-----+

```

This command shows maintenance details for the node with the ID 9dcc9632-911c-4cc5-9a89-5a6fa5db2314.

2.3.7 vinfra node maintenance stop

Return node to operation:

```
usage: vinfra node maintenance stop <node>
```

<node>

Node ID or hostname

Example:

```

# vinfra node maintenance stop 9dcc9632-911c-4cc5-9a89-5a6fa5db2314
+-----+
| Field  | Value                                     |
+-----+
| task_id | 34e0b546-aa2c-466c-93fe-7dff28c543c6 |
+-----+

```

This command creates a task to stop maintenance for the node with the ID 9dcc9632-911c-4cc5-9a89-5a6fa5db2314.

Task outcome:

```

# vinfra task show 34e0b546-aa2c-466c-93fe-7dff28c543c6
+-----+
| Field  | Value                                     |
+-----+
| details |                                           |
| name    | backend.business.models.maintenance.tasks.MaintenanceStopTask |

```

```

| result | |
| state | success |
| task_id | 34e0b546-aa2c-466c-93fe-7dff28c543c6 |
+-----+-----+

```

2.3.8 vinfra node release

Release a node from the storage cluster. Start data migration from the node as well as cluster replication and rebalancing to meet the configured redundancy level:

```
usage: vinfra node release [--force] <node>
```

--force

Release node without data migration

<node>

Node ID or hostname

Example:

```

# vinfra node release f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | c2a653a2-8991-4b3a-8bdf-5c0872aa75b3 |
+-----+-----+

```

This command creates a task to release the node with the ID f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 from the storage cluster with migration of data to maintain the set redundancy mode.

Task outcome:

```

# vinfra task show c2a653a2-8991-4b3a-8bdf-5c0872aa75b3
+-----+-----+
| Field | Value |
+-----+-----+
| args | - f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
| | - false |
| kwargs | {} |
| name | backend.tasks.node.ReleaseNodeTask |
| state | success |
| task_id | c2a653a2-8991-4b3a-8bdf-5c0872aa75b3 |
+-----+-----+

```


2.3.9 vinfra node forget

Remove a node from the storage cluster:

```
usage: vinfra node forget <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra node forget fd1e46de-6e17-4571-bf6b-1ac34ec1c225
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 0eac3b74-e8f5-4974-9efe-a9070187d83c |
+-----+-----+
```

This command creates a task to unregister the node with the ID fd1e46de-6e17-4571-bf6b-1ac34ec1c225 from Virtuozzo Hybrid Infrastructure.

Task outcome:

```
# vinfra task show 0eac3b74-e8f5-4974-9efe-a9070187d83c
+-----+-----+
| Field | Value |
+-----+-----+
| args | - fd1e46de-6e17-4571-bf6b-1ac34ec1c225 |
| kwargs | {} |
| name | backend.tasks.node.DeleteNodeTask |
| state | success |
| task_id | 0eac3b74-e8f5-4974-9efe-a9070187d83c |
+-----+-----+
```

2.4 Managing Node Network Interfaces

2.4.1 vinfra node iface list

List node network interfaces:

```
usage: vinfra node iface list [-a | --node <node>]
```

-a, --all

List all network interfaces on all nodes

--node <node>

Node ID or hostname to list network interfaces on (default: node001.vstoragedomain)

Example:

This command shows network interfaces of the node with the ID 4f96acf5-3bc8-4094-bcb6-4d1953be7b55.

```
# vinfra node iface list --node 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+-----+-----+-----+
| name | node_id | ipv4 | state | network |
+-----+-----+-----+-----+-----+
| eth0 | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 | - 10.94.29.218/16 | up | Public |
| eth1 | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 | - 10.37.130.101/24 | up | Private |
+-----+-----+-----+-----+-----+
```

2.4.2 vinfra node iface show

Show details of a network interface:

```
usage: vinfra node iface show [--node <node>] <iface>
```

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

<iface>

Network interface name

Example:

```
# vinfra node iface show eth0 --node 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+
| Field | Value |
+-----+-----+
| contained_in | |
| dhcp4 | 10.94.29.218 |
| dhcp4_enabled | True |
| dhcp6 | fe80::21c:42ff:fe2a:4fdf |
| dhcp6_enabled | True |
| dns4 | - 127.0.0.1 |
| dns6 | [] |
| duplex | |
| gw4 | 10.94.0.1 |
| gw6 | |
| ignore_auto_dns_v4 | False |
| ignore_auto_dns_v6 | False |
| ignore_auto_routes_v4 | False |
| ignore_auto_routes_v6 | False |
| ipv4 | - 10.94.29.218/16 |
+-----+-----+
```

```

| ipv6           | - fe80::21c:42ff:fe2a:4fdf/64 |
| mac_addr      | 00:1c:42:2a:4f:df             |
| mtu           | 1500                          |
| multicast     | True                           |
| name          | eth0                           |
| node_id       | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
| plugged       | True                           |
| roles_set     | 237e58dd-6c10-49c1-be7f-7ddf7de2efd1 |
| rx_bytes      | 1844502614                     |
| rx_dropped    | 0                               |
| rx_errors     | 0                               |
| rx_overruns   | 0                               |
| rx_packets    | 11543284                       |
| speeds        | current: null                  |
|               | max: null                      |
| state         | up                              |
| tx_bytes      | 28477979                       |
| tx_dropped    | 0                               |
| tx_errors     | 0                               |
| tx_overruns   | 0                               |
| tx_packets    | 107649                         |
| type          | iface                           |
+-----+-----+

```

This command shows the details of the network interface `eth0` located on the node with the ID `4f96acf5-3bc8-4094-bcb6-4d1953be7b55`.

2.4.3 vinfra node iface up

Bring a network interface up:

```
usage: vinfra node iface up [--node <node>] <iface>
```

`--node <node>`

Node ID or hostname (default: `node001.vstoragedomain`)

`<iface>`

Network interface name

Example:

```

# vinfra node iface up eth2 --node 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+
| Field          | Value                          |
+-----+-----+
| contained_in   |                                 |
| dhcp4          | 10.37.130.138                  |
| dhcp4_enabled  | True                           |

```

```

| dhcp6                | fe80::21c:42ff:fe8:5b90 |
| dhcp6_enabled       | True                    |
| dns4                 | - 127.0.0.1            |
| dns6                 | []                      |
| duplex               |                         |
| gw4                  | 10.94.0.1              |
| gw6                  |                         |
| ignore_auto_dns_v4  | False                   |
| ignore_auto_dns_v6  | False                   |
| ignore_auto_routes_v4 | False                   |
| ignore_auto_routes_v6 | False                   |
| ipv4                 | - 10.37.130.138/24     |
| ipv6                 | - fe80::21c:42ff:fe8:5b90/64 |
| mac_addr             | 00:1c:42:f8:5b:90     |
| mtu                  | 1500                   |
| multicast            | True                    |
| name                 | eth2                    |
| node_id              | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
| plugged              | True                    |
| roles_set            |                         |
| rx_bytes             | 97632                  |
| rx_dropped           | 0                       |
| rx_errors            | 0                       |
| rx_overruns          | 0                       |
| rx_packets           | 1258                   |
| speeds               | current: null          |
|                      | max: null               |
| state                | up                      |
| tx_bytes             | 1116                   |
| tx_dropped           | 0                       |
| tx_errors            | 0                       |
| tx_overruns          | 0                       |
| tx_packets           | 8                       |
| type                 | iface                   |
+-----+-----+

```

This command brings up the network interface eth2 located on the node with the ID 4f96acf5-3bc8-4094-bcb6-4d1953be7b55.

2.4.4 vinfra node iface down

Bring a network interface down:

```
usage: vinfra node iface down [--node <node>] <iface>
```

--node <node>

Node ID or hostname

<iface>

Network interface name

Example:

```
# vinfra node iface down eth2 --node 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+
| Field          | Value          |
+-----+-----+
| contained_in   |                |
| dhcp4          |                |
| dhcp4_enabled  | True           |
| dhcp6          |                |
| dhcp6_enabled  | True           |
| dns4           | - 127.0.0.1   |
| dns6           | []             |
| duplex         |                |
| gw4            | 10.94.0.1     |
| gw6            |                |
| ignore_auto_dns_v4 | False         |
| ignore_auto_dns_v6 | False         |
| ignore_auto_routes_v4 | False         |
| ignore_auto_routes_v6 | False         |
| ipv4           | []             |
| ipv6           | []             |
| mac_addr       | 00:1c:42:f8:5b:90 |
| mtu            | 1500           |
| multicast      | True           |
| name           | eth2           |
| node_id        | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
| plugged        | False          |
| roles_set      |                |
| rx_bytes       | 97984          |
| rx_dropped     | 0              |
| rx_errors      | 0              |
| rx_overruns    | 0              |
| rx_packets     | 1264           |
| speeds         | current: null  |
|                | max: null      |
| state          | down           |
| tx_bytes       | 1116           |
| tx_dropped     | 0              |
| tx_errors      | 0              |
| tx_overruns    | 0              |
| tx_packets     | 8              |
| type           | iface          |
+-----+-----+
```

This command brings down the network interface eth2 located on the node with the ID 4f96acf5-3bc8-4094-bcb6-4d1953be7b55.

2.4.5 vinfra node iface set

Modify network interface parameters (overwrite the omitted options with the default values for the interface):

```
usage: vinfra node iface set [--ipv4 <ipv4>] [--ipv6 <ipv6>] [--gw4 <gw4>] [--gw6 <gw6>]
                             [--mtu <mtu>] [--dhcp4 | --no-dhcp4] [--dhcp6 | --no-dhcp6]
                             [--auto-routes-v4 | --ignore-auto-routes-v4]
                             [--auto-routes-v6 | --ignore-auto-routes-v6]
                             [--network <network> | --no-network]
                             [--connected-mode | --datagram-mode] [--node <node>] <iface>
```

`--ipv4 <ipv4>`

A comma-separated list of IPv4 addresses

`--ipv6 <ipv6>`

A comma-separated list of IPv6 addresses

`--gw4 <gw4>`

Gateway IPv4 address

`--gw6 <gw6>`

Gateway IPv6 address

`--mtu <mtu>`

MTU interface value

`--dhcp4`

Enable DHCPv4

`--no-dhcp4`

Disable DHCPv4

`--dhcp6`

Enable DHCPv6

`--no-dhcp6`

Disable DHCPv6

`--auto-routes-v4`

Enable automatic IPv4 routes

`--ignore-auto-routes-v4`

Ignore automatic IPv4 routes

```
--auto-routes-v6
    Enable automatic IPv6 routes

--ignore-auto-routes-v6
    Ignore automatic IPv6 routes

--network <network>
    Network ID or name

--no-network
    Remove a network from the interface

--connected-mode
    Enable connected mode (InfiniBand interfaces only)

--datagram-mode
    Enable datagram mode (InfiniBand interfaces only)

--node <node>
    Node ID or hostname (default: node001.vstoragedomain)

<iface>
    Network interface name
```

Example:

```
# vinfra node iface set eth2 --network Private \
--node 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 8a378098-6760-4fe9-ac20-1f18a8ed9d2e |
+-----+-----+
```

This command creates a task to assign the network interface eth2 located on the node with the ID 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 to the network Private.

Task outcome:

```
# vinfra task show 8a378098-6760-4fe9-ac20-1f18a8ed9d2e
+-----+-----+
| Field | Value |
+-----+-----+
| args  | - 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
|       | - eth2 |
| kwargs | roles_set: 6095a997-e5f1-493d-a750-41ddf277153b |
| name  | backend.presentation.network.tasks.NetworkInterfaceChangeTask |
+-----+-----+
```

```

| result | contained_in: null
|         | dhcp4: null
|         | dhcp4_enabled: false
|         | dhcp6: null
|         | dhcp6_enabled: false
|         | duplex: null
|         | gw4: null
|         | gw6: null
|         | ignore_auto_routes_v4: true
|         | ignore_auto_routes_v6: true
|         | ipv4:
|         | - 10.37.130.103/24
|         | ipv6:
|         | - fe80::21c:42ff:fe75:7c4d/64
|         | mac_addr: 00:1c:42:75:7c:4d
|         | mtu: 1500
|         | multicast: true
|         | name: eth2
|         | node_id: 4f96acf5-3bc8-4094-bcb6-4d1953be7b55
|         | plugged: true
|         | roles_set: 6095a997-e5f1-493d-a750-41ddf277153b
|         | rx_bytes: 38156
|         | rx_dropped: 0
|         | rx_errors: 0
|         | rx_overruns: 0
|         | rx_packets: 225
|         | speeds:
|         |   current: null
|         |   max: null
|         | state: up
|         | tx_bytes: 13087
|         | tx_dropped: 0
|         | tx_errors: 0
|         | tx_overruns: 0
|         | tx_packets: 145
|         | type: iface
| state   | success
| task_id | 8a378098-6760-4fe9-ac20-1f18a8ed9d2e
+-----+

```

2.4.6 vinfra node iface create-bond

Create a network bonding:

```

usage: vinfra node iface create-bond [--ipv4 <ipv4>] [--ipv6 <ipv6>] [--gw4 <gw4>]
                                     [--gw6 <gw6>] [--mtu <mtu>] [--dhcp4 | --no-dhcp4]
                                     [--dhcp6 | --no-dhcp6] [--network <network>]
                                     [--auto-routes-v4 | --ignore-auto-routes-v4]
                                     [--auto-routes-v6 | --ignore-auto-routes-v6]

```



```
[--bonding-opts <bonding_opts>] [--node <node>]  
--bond-type <bond-type> --ifaces <ifaces>
```

--ipv4 <ipv4>

A comma-separated list of IPv4 addresses

--ipv6 <ipv6>

A comma-separated list of IPv6 addresses

--gw4 <gw4>

Gateway IPv4 address

--gw6 <gw6>

Gateway IPv6 address

--mtu <mtu>

MTU interface value

--dhcp4

Enable DHCPv4

--no-dhcp4

Disable DHCPv4

--dhcp6

Enable DHCPv6

--no-dhcp6

Disable DHCPv6

--auto-routes-v4

Enable automatic IPv4 routes

--ignore-auto-routes-v4

Ignore automatic IPv4 routes

--auto-routes-v6

Enable automatic IPv6 routes

--ignore-auto-routes-v6

Ignore automatic IPv6 routes

--network <network>

Network ID or name

```
--bonding-opts <bonding_opts>
```

Additional bonding options

```
--bond-type <bond-type>
```

Bond type (balance-rr, active-backup, balance-xor, broadcast, 802.3ad, balance-tlb, balance-alb)

```
--node <node>
```

Node ID or hostname (default: node001.vstoragedomain)

```
--ifaces <ifaces>
```

A comma-separated list of network interface names, e.g., iface1,iface2,...,iface<N>

Example:

```
# vinfra node iface create-bond --ifaces eth2,eth3 --bond-type balance-xor \
--dhcp4 --node fd1e46de-6e17-4571-bf6b-1ac34ec1c225
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | becf96ad-9e39-4bec-b82c-4e1219a196de |
+-----+-----+
```

This command creates a task to bond network interfaces eth2 and eth3 into bond0 of the type balance-xor on the node with the ID fd1e46de-6e17-4571-bf6b-1ac34ec1c225.

Task outcome:

```
# vinfra task show becf96ad-9e39-4bec-b82c-4e1219a196de
+-----+-----+
| Field  | Value                                |
+-----+-----+
| args    | - fd1e46de-6e17-4571-bf6b-1ac34ec1c225 |
| kwargs  | bond_type: balance-xor                |
|         | ifaces:                                |
|         | - eth2                                 |
|         | - eth3                                 |
|         | registration_token: 3102ed1a          |
| name    | backend.presentation.network.tasks.NetworkInterfaceCreateBondingTask |
| result  | bond_type: balance-xor                |
|         | dhcp4: 10.37.130.117                  |
|         | dhcp4_enabled: true                   |
|         | dhcp6: fe80::21c:42ff:fe81:27d0       |
|         | dhcp6_enabled: true                   |
|         | duplex: null                           |
|         | gw4: 10.94.0.1                         |
|         | gw6: null                              |
|         | ignore_auto_routes_v4: false          |
|         | ignore_auto_routes_v6: false          |
|         | ipv4:                                  |
+-----+-----+
```

```

|         | - 10.37.130.117/24
|         | ipv6:
|         | - fe80::21c:42ff:fe81:27d0/64
|         | mac_addr: 00:1c:42:81:27:d0
|         | mtu: 1500
|         | multicast: true
|         | name: bond0
|         | node_id: fd1e46de-6e17-4571-bf6b-1ac34ec1c225
|         | plugged: true
|         | roles_set: ''
|         | rx_bytes: 3048
|         | rx_dropped: 0
|         | rx_errors: 0
|         | rx_overruns: 0
|         | rx_packets: 22
|         | speeds:
|         |   current: null
|         |   max: null
|         | state: up
|         | tx_bytes: 1782
|         | tx_dropped: 0
|         | tx_errors: 0
|         | tx_overruns: 0
|         | tx_packets: 13
|         | type: bonding
| state   | success
| task_id | becf96ad-9e39-4bec-b82c-4e1219a196de
+-----+

```

2.4.7 vinfra node iface create-vlan

Create a VLAN:

```

usage: vinfra node iface create-vlan [--ipv4 <ipv4>] [--ipv6 <ipv6>] [--gw4 <gw4>]
                                     [--gw6 <gw6>] [--mtu <mtu>] [--dhcp4 | --no-dhcp4]
                                     [--dhcp6 | --no-dhcp6] [--network <network>]
                                     [--auto-routes-v4 | --ignore-auto-routes-v4]
                                     [--auto-routes-v6 | --ignore-auto-routes-v6]
                                     [--node <node>] --iface <iface> --tag <tag>

```

--ipv4 <ipv4>

A comma-separated list of IPv4 addresses

--ipv6 <ipv6>

A comma-separated list of IPv6 addresses

--gw4 <gw4>

Gateway IPv4 address

```
--gw6 <gw6>  
    Gateway IPv6 address  
  
--mtu <mtu>  
    MTU interface value  
  
--dhcp4  
    Enable DHCPv4  
  
--no-dhcp4  
    Disable DHCPv4  
  
--dhcp6  
    Enable DHCPv6  
  
--no-dhcp6  
    Disable DHCPv6  
  
--auto-routes-v4  
    Enable automatic IPv4 routes  
  
--ignore-auto-routes-v4  
    Ignore automatic IPv4 routes  
  
--auto-routes-v6  
    Enable automatic IPv6 routes  
  
--ignore-auto-routes-v6  
    Ignore automatic IPv6 routes  
  
--network <network>  
    Network ID or name  
  
--node <node>  
    Node ID or hostname (default: node001.vstoragedomain)  
  
--iface <iface>  
    Interface name  
  
--tag <tag>  
    VLAN tag number
```

Example:

```
# vinfra node iface create-vlan --iface eth2 --tag 100 --dhcp4 \
--node fd1e46de-6e17-4571-bf6b-1ac34ec1c225
+-----+
| Field | Value |
+-----+
| task_id | 0b978acd-367b-47ad-8572-4f4e6ffb8877 |
+-----+
```

This command creates a task to create a VLAN with the tag 100 on the network interface eth2 on the node with the ID fd1e46de-6e17-4571-bf6b-1ac34ec1c225.

Task outcome:

```
# vinfra task show 0b978acd-367b-47ad-8572-4f4e6ffb8877
+-----+
| Field | Value |
+-----+
| args | - fd1e46de-6e17-4571-bf6b-1ac34ec1c225 | |
| kwargs | iface: eth2 |
| | tag: 100 |
| name | backend.presentation.network.tasks.NetworkInterfaceCreateVlanTask |
| result | built_on: eth2 |
| | dhcp4: null |
| | dhcp4_enabled: false |
| | dhcp6: null |
| | dhcp6_enabled: false |
| | duplex: null |
| | gw4: null |
| | gw6: null |
| | ignore_auto_routes_v4: true |
| | ignore_auto_routes_v6: true |
| | ipv4: [] |
| | ipv6: |
| | - fe80::21c:42ff:fe81:27d0/64 |
| | mac_addr: 00:1c:42:81:27:d0 |
| | mtu: 1500 |
| | multicast: true |
| | name: eth2.100 |
| | node_id: fd1e46de-6e17-4571-bf6b-1ac34ec1c225 |
| | plugged: true |
| | roles_set: '' |
| | rx_bytes: 0 |
| | rx_dropped: 0 |
| | rx_errors: 0 |
| | rx_overruns: 0 |
| | rx_packets: 0 |
| | speeds: |
| | | current: null |
| | | max: null |
| | state: up |
| | tag: 100 |
+-----+
```

```

|         | tx_bytes: 738 |
|         | tx_dropped: 0 |
|         | tx_errors: 0 |
|         | tx_overruns: 0 |
|         | tx_packets: 7 |
|         | type: vlan |
| state   | success |
| task_id | 0b978acd-367b-47ad-8572-4f4e6ffb8877 |
+-----+-----+

```

2.4.8 vinfra node iface delete

Delete a network interface:

```
usage: vinfra node iface delete [--node <node>] <iface>
```

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

<iface>

Network interface name

Example:

```

# vinfra node iface delete --node fd1e46de-6e17-4571-bf6b-1ac34ec1c225 eth2.100
+-----+-----+
| Field  | Value |
+-----+-----+
| task_id | 16503616-6c1c-48f9-999a-9d87b617d9ee |
+-----+-----+

```

This command creates a task to delete a VLAN interface eth1.100 from the node with the ID fd1e46de-6e17-4571-bf6b-1ac34ec1c225.

Task outcome:

```

# vinfra task show 16503616-6c1c-48f9-999a-9d87b617d9ee
+-----+-----+
| Field  | Value |
+-----+-----+
| details | |
| name    | backend.presentation.network.tasks.NetworkInterfaceRemoveTask |
| result  | |
| state   | success |
| task_id | 16503616-6c1c-48f9-999a-9d87b617d9ee |
+-----+-----+

```

2.5 Managing Node Disks

2.5.1 vinfra node disk list

List node disks:

```
usage: vinfra node disk list [-a | --node <node>]
```

`-a, --all`

List disks on all nodes

`--node <node>`

Node ID or hostname to list disks on (default: `node001.vstoragedomain`)

Example:

```
# vinfra node disk list --node 94d58604-6f30-4339-8578-adb7903b7277 \
-c id -c node_id -c device -c used -c size -c role
+-----+-----+-----+-----+-----+-----+
| id           | node_id      | device | used  | size  | role      |
+-----+-----+-----+-----+-----+-----+
| E0B7CE6F-<...> | 94d58604-<...> | sda    | 5.5GiB | 239.1GiB | mds-system |
| EAC7DF5D-<...> | 94d58604-<...> | sdb    | 2.1GiB | 1007.8GiB | cs         |
| 49D792CA-<...> | 94d58604-<...> | sdc    | 2.1GiB | 1007.8GiB | cs         |
+-----+-----+-----+-----+-----+-----+
```

This command lists disks on the node with the ID `94d58604-6f30-4339-8578-adb7903b7277`. (The output is abridged to fit on page.)

2.5.2 vinfra node disk show

Show details of a disk:

```
usage: vinfra node disk show [--node <node>] <disk>
```

`--node <node>`

Node ID or hostname

`<disk>`

Disk ID or device name (default: `node001.vstoragedomain`)

Example:

```
# vinfra node disk show EAC7DF5D-9E60-4444-85F7-5CA5738399CC \
--node 94d58604-6f30-4339-8578-adb7903b7277
```

Field	Value
being_released	False
device	sdb
disk_status	ok
encryption	
id	EAC7DF5D-9E60-4444-85F7-5CA5738399CC
is_blink_available	False
is_blinking	False
latency	
lun_id	
model	Vz_HARDDISK2
mountpoint	/vstorage/33aac2d5
node_id	94d58604-6f30-4339-8578-adb7903b7277
role	cs
rpm	
serial_number	45589b5823ce4c188b55
service_id	1026
service_params	journal_type: inner_cache tier: 0
service_status	ok
slot	
smart_status	not_supported
space	full_size: 1099511627776 size: 1082101518336 used: 2246164480
tasks	
temperature	0.0
transport	
type	hdd

This command shows the details of the disk with the ID EAC7DF5D-9E60-4444-85F7-5CA5738399CC attached to the node with the ID 94d58604-6f30-4339-8578-adb7903b7277.

2.5.3 vinfra node disk assign

Add multiple disks to the storage cluster:

```
usage: vinfra node disk assign --disk <disk>:<role>[:<key=value,...>]
      [--node <node>]
```

```
--disk <disk>:<role>[:<key=value,...>]
```

Disk configuration in the format:

- <disk>: disk device ID or name

- <role>: disk role (cs, mds, journal, mds-journal, mds-system, cs-system, system)
- comma-separated key=value pairs with keys (optional):
 - tier: disk tier (0, 1, 2 or 3)
 - journal-tier: journal (cache) disk tier (0, 1, 2 or 3)
 - journal-type: journal (cache) disk type (no_cache, inner_cache OR external_cache)
 - journal-disk: journal (cache) disk ID or device name
 - journal-size: journal (cache) disk size, in bytes
 - bind-address: bind IP address for the metadata service

E.g., sda:cs:tier=0,journal-type=inner_cache. This option can be used multiple times.

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

Example:

```
# vinfra node disk assign --disk sdc:cs --node f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 080337ba-0508-44a0-9363-eddcd9df9f0d |
+-----+-----+
```

This command creates a task to assign the role cs to the disk sdc on the node with the ID f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4

Task outcome:

```
# vinfra task show 080337ba-0508-44a0-9363-eddcd9df9f0d
+-----+-----+
| Field | Value |
+-----+-----+
| args | [] |
| kwargs | cluster_id: 1 |
| | disks: |
| | - id: D3BEF4BB-AA3B-4DB6-9376-BC7CDA636700 |
| | role: cs |
| | service_params: {} |
| | logger: |
| | __classname: backend.logger.tracer.TracingLogger |
| | __dict: |
| | prefix: POST /api/v2/1/nodes/f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4/disks/ |
| | token: '3215629651314950' |
+-----+-----+
```

```

|         | node_id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
| name    | backend.tasks.disks.BulkAssignDiskTask         |
| result  | {}                                              |
| state   | success                                         |
| task_id | 080337ba-0508-44a0-9363-eddc9df9f0d          |
+-----+-----+

```

2.5.4 vinfra node disk release

Release a disk from the storage cluster. Start data migration from the node as well as cluster replication and rebalancing to meet the configured redundancy level:

```
usage: vinfra node disk release [--force] [--node <node>] <disk>
```

--force

Release without data migration

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

<disk>

Disk ID or device name

Example:

```

# vinfra node disk release sdc --node f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4
+-----+-----+
| Field  | Value |
+-----+-----+
| task_id | 587a936d-3953-481c-a2cd-b1223b890bec |
+-----+-----+

```

This command creates a task to release the role `cs` from the disk `sd` on the node with the ID `f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4`.

Task outcome:

```

# vinfra task show 587a936d-3953-481c-a2cd-b1223b890bec
+-----+-----+
| Field  | Value |
+-----+-----+
| args   | []    |
| kwargs | cluster_id: 1
|        | disk_id: 43EF3400-EA95-43DE-B624-3D7ED0F9DDDD
|        | force: false
|        | logger:
+-----+-----+

```

```

|         | __classname: backend.logger.tracer.TracingLogger |
|         | __dict: |
|         |   prefix: POST /api/v2/1/nodes/f59dabdb- |
|         | bd1c-4944-8af2-26b8fe9ff8d4/disks/43EF3400-EA95-43DE-B624-3D7ED0F9DDDD/release/ |
|         |   token: '3217122839314940' |
|         |   node_id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
| name    | backend.tasks.disks.ReleaseDiskTask |
| state   | success |
| task_id | 587a936d-3953-481c-a2cd-b1223b890bec |
+-----+-----+

```

2.5.5 vinfra node disk blink on

Start blinking the specified disk bay to identify disk for maintenance purposes:

```
usage: vinfra node disk blink on [--node <node>] <disk>
```

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

<disk>

Disk ID or device name

Example:

```
# vinfra node disk blink on sda --node f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4
```

This command starts blinking the disk `sda` on the node with the ID `f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4`.

2.5.6 vinfra node disk blink off

Stop blinking the specified disk bay:

```
usage: vinfra node disk blink off [--node <node>] <disk>
```

--node <node>

Node ID or hostname (default: node001.vstoragedomain)

<disk>

Disk ID or device name

Example:

```
# vinfra node disk blink off sda --node f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4
```

This command stops blinking the disk sda on the node with the ID f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4.

2.5.7 vinfra node iscsi target add

Add an iSCSI target as a disk to a node:

```
usage: vinfra node iscsi target add [--auth-username <auth-username>]
                                     [--auth-password <auth-password>]
                                     --portal <portal> --node <node> <target-name>
```

```
--auth-username <auth-username>
```

User name

```
--auth-password <auth-password>
```

User password

```
--portal <portal>
```

Portal IP address in the format IP:port (this option can be specified multiple times)

```
--node <node>
```

Node ID or hostname

```
<target-name>
```

Target name

Example:

```
# vinfra node iscsi target add iqn.2014-06.com.vstorage:target1 \
--portal 172.16.24.244:3260 --node f1931be7-0a01-4977-bfef-51a392adcd94
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | c42bfbe5-7292-41c2-91cb-446795535ab9 |
+-----+-----+
```

This command creates a task to connect a remote iSCSI target iqn.2014-06.com.vstorage:target1 with the IP address 172.16.24.244 and port 3260 to the node with the ID f1931be7-0a01-4977-bfef-51a392adcd94.

Task outcome:

```
# vinfra task show c42bfbe5-7292-41c2-91cb-446795535ab9
+-----+-----+
| Field | Value |
+-----+-----+
```

```

| args      | - f1931be7-0a01-4977-bfef-51a392adcd94 |
| kwargs   | portals:                                |
|          | - address: 172.16.24.244                |
|          | port: 3260                              |
|          | target_name: iqn.2014-06.com.vstorage:target1 |
| name     | backend.presentation.nodes.iscsi_initiators.tasks.ConnectTask |
| result   | connected: true                          |
|          | portals:                                |
|          | - address: 172.16.24.244                |
|          | port: 3260                              |
|          | state: connected                        |
|          | target_name: iqn.2014-06.com.vstorage:target1 |
| state    | success                                  |
| task_id  | c42bfbe5-7292-41c2-91cb-446795535ab9   |
+-----+-----+

```

2.5.8 vinfra node iscsi target delete

Delete an iSCSI target from a node:

```
usage: vinfra node iscsi target delete --node <node> <target-name>
```

--node <node>

Node ID or hostname

<target-name>

Target name

Example:

```

# vinfra node iscsi target delete iqn.2014-06.com.vstorage:target1 \
--node f1931be7-0a01-4977-bfef-51a392adcd94
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | c8dc74ee-86d6-4b89-8b6f-153ff1e78cb7 |
+-----+-----+

```

This command creates a task to disconnect a remote iSCSI target `iqn.2014-06.com.vstorage:target1` from the node with the ID `f1931be7-0a01-4977-bfef-51a392adcd94`.

Task outcome:

```

# vinfra task show c8dc74ee-86d6-4b89-8b6f-153ff1e78cb7
+-----+-----+
| Field  | Value                                |
+-----+-----+
| args   | - f1931be7-0a01-4977-bfef-51a392adcd94 |

```

```
| kwargs | target_name: iqn.2014-06.com.vstorage:target1 |
| name   | backend.presentation.nodes.iscsi_initiators.tasks.DisconnectTask |
| state  | success |
| task_id | c8dc74ee-86d6-4b89-8b6f-153ff1e78cb7 |
+-----+
```

2.6 Creating and Deleting the Storage Cluster

2.6.1 vinfra cluster create

Create a storage cluster:

```
usage: vinfra cluster create [--disk <disk>:<role>[:<key=value,...>]]
                             [--tier-encryption {0,1,2,3}] --node <node> <cluster-name>
```

`--disk <disk>:<role> [:<key=value,...>]`

Disk configuration in the format:

- `<disk>`: disk device ID or name
- `<role>`: disk role (cs, mds, journal, mds-journal, mds-system, cs-system, system)
- comma-separated key=value pairs with keys (optional):
 - `tier`: disk tier (0, 1, 2 or 3)
 - `journal-tier`: journal (cache) disk tier (0, 1, 2 or 3)
 - `journal-type`: journal (cache) disk type (no_cache, inner_cache or external_cache)
 - `journal-disk`: journal (cache) disk ID or device name
 - `journal-size`: journal (cache) disk size, in bytes
 - `bind-address`: bind IP address for the metadata service

E.g., `sda:cs:tier=0,journal-type=inner_cache`. This option can be used multiple times.

`--tier-encryption {0,1,2,3}`

Enable encryption for storage cluster tiers. Encryption is disabled by default. This option can be used multiple times.

`--node <node>`

Node ID or hostname

<cluster-name>

Storage cluster name

Example:

```
# vinfra cluster create stor1 --node 94d58604-6f30-4339-8578-adb7903b7277
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | d9ca8e1d-8ac8-4459-898b-2d803efd7bc6 |
+-----+-----+
```

This command creates a task to create the storage cluster `stor1` on the node with the ID `94d58604-6f30-4339-8578-adb7903b7277`. As disk roles are not explicitly specified, they are assigned automatically: `mds-system` to the system disk, and `cs` to all other disks.

Task outcome:

```
# vinfra task show d9ca8e1d-8ac8-4459-898b-2d803efd7bc6
+-----+-----+
| Field | Value |
+-----+-----+
| args  | - stor1
|       | - 94d58604-6f30-4339-8578-adb7903b7277
|       | - null
|       | - null
| kwargs | {}
| name   | backend.tasks.cluster.CreateNewCluster
| result | cluster_id: 1
| state  | success
| task_id | d9ca8e1d-8ac8-4459-898b-2d803efd7bc6
+-----+-----+
```

2.6.2 vinfra cluster delete

Delete the storage cluster:

```
usage: vinfra cluster delete
```

Example:

```
# vinfra cluster delete
Operation waiting (timeout=600s) [Elapsed Time: 0:01:09] ... |
Operation successful
```

This command releases all nodes from the storage cluster.

2.7 Showing Storage Cluster Overview and Details

2.7.1 vinfra cluster overview

Show storage cluster overview:

```
usage: vinfra cluster overview
```

Example:

```
# vinfra cluster overview
+-----+-----+
| Field          | Value          |
+-----+-----+
| chunks         | blocked: 0     |
|                | degraded: 0    |
|                | deleting: 0    |
|                | healthy: 2     |
|                | offline: 0     |
|                | overcommitted: 0 |
|                | pending: 0     |
|                | replicating: 0 |
|                | standby: 0     |
|                | total: 2       |
|                | unique: 2      |
|                | urgent: 0      |
|                | void: 0        |
| fs_stat        | chunk_maps: 2  |
|                | chunk_nodes: 2 |
|                | file_maps: 2   |
|                | files: 9       |
|                | inodes: 9      |
|                | used_size: 11335680 |
| id             | 1              |
| license        | capacity: 1099511627776 |
|                | expiration_ts: null |
|                | keynumber: null |
|                | status: 0      |
|                | used_size: 11335680 |
| logic_space    | free: 1099500292096 |
|                | total: 1099511627776 |
|                | used: 11335680  |
| name           | stor1          |
| repl           | eta: null      |
|                | reads: 0       |
|                | writes: 0      |
```



```

| resistance      | to_lose: 0      |
|                 | total: 1       |
| space_per_service | abgw: null     |
|                 | compute: null  |
|                 | iscsi: null    |
|                 | nfs: null      |
|                 | s3: null       |
| status         | healthy        |
| tiers          | - id: 0        |
|                 | phys_space:    |
|                 |   free: 2164191700992 |
|                 |   total: 2164203036672 |
|                 |   used: 11335680 |
+-----+-----+

```

This command shows an overview of the cluster.

2.7.2 vinfra cluster show

Show cluster details:

```
usage: vinfra cluster show
```

Example:

```

# vinfra cluster show
+-----+-----+
| Field | Value |
+-----+-----+
| id    | 1     |
| name  | stor1 |
| nodes | - host: stor-4.example.com.vstagedomain |
|       | id: 4b83a87d-9adf-472c-91f0-782c47b2d5f1 |
|       | is_installing: false |
|       | is_releasing: false |
|       | - host: stor-3.example.com.vstagedomain |
|       | id: 7d7d37b8-4c06-4f1a-b3a6-4b54257d70ce |
|       | is_installing: false |
|       | is_releasing: false |
|       | - host: stor-5.example.com.vstagedomain |
|       | id: fd1e46de-6e17-4571-bf6b-1ac34ec1c225 |
|       | is_installing: false |
|       | is_releasing: false |
|       | - host: stor-1.example.com.vstagedomain |
|       | id: 94d58604-6f30-4339-8578-adb7903b7277 |
|       | is_installing: false |
|       | is_releasing: false |
|       | - host: stor-2.example.com.vstagedomain |
|       | id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |

```

```
|      | is_installing: false |
|      | is_releasing: false  |
+-----+-----+-----+
```

This command shows cluster details.

CHAPTER 3

Managing Compute Cluster

3.1 Creating and Deleting the Compute Cluster

3.1.1 `vinfra service compute create`

Create a compute cluster:

```
usage: vinfra service compute create [--public-network <network>]
                                     [--subnet cidr=CIDR[,key=value,...]]
                                     [--cpu-model <cpu-model>] [--force]
                                     [--enable-k8saas] [--enable-lbaas]
                                     [--enable-metering] --nodes <nodes>
```

`--public-network <network>`

A physical network to connect the public virtual network to. It must include the 'VM public' traffic type.

`--subnet cidr=CIDR[,key=value,...]`

Subnet for IP address management in the public virtual network (the `--public-network` option is required):

- `cidr`: subnet range in CIDR notation;
- comma-separated `key=value` pairs with keys (optional):
 - `gateway`: gateway IP address.
 - `dhcp`: enable/disable the virtual DHCP server.
 - `allocation-pool1`: allocation pool of IP addresses from CIDR in the format `ip1-ip2`, where `ip1` and `ip2` are starting and ending IP addresses. Specify the key multiple times to create multiple IP pools.

- `dns-server`: DNS server IP address, specify multiple times to set multiple DNS servers.

Example: `--subnet cidr=192.168.5.0/24,dhcp=enable`.

`--cpu-model <cpu-model>`

CPU model for virtual machines. View the list of available CPU models using `compute cluster show`.

`--force`

Skip checks for minimal hardware requirements.

`--enable-k8saas`

Enable Kubernetes-as-a-Service services.

`--enable-lbaas`

Enable Load-Balancing-as-a-Service services.

`--enable-metering`

Enable metering services.

`--nodes <nodes>`

A comma-separated list of node IDs or hostnames.

Example:

```
# vinfra service compute create --nodes 7ffa9540-5a20-41d1-b203-e3f349d62565,\
02ff64ae-5800-4090-b958-18b1fe8f5060,6e8afc28-7f71-4848-bdbe-7c5de64c5013,\
37c70bfb-c289-4794-8be4-b7a40c2b6d95,827a1f4e-56e5-404f-9113-88748c18f0c2 \
--public-network Public --subnet cidr=10.94.0.0/16,dhcp=enable,\
gateway=10.94.0.1,allocation-pool=10.94.129.64-10.94.129.79,\
dns-server=10.30.0.27,dns-server=10.30.0.28
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | be517afa-fae0-457e-819c-f4d6399f3ae2 |
+-----+-----+
```

This command creates a task to create the compute cluster from five nodes specified by ID. It also specifies the public network for VMs, the gateway, the allocation pool of IP addresses to assign to VMs, and the DNS servers to use.

Task outcome:

```
# vinfra task show be517afa-fae0-457e-819c-f4d6399f3ae2
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
+-----+-----+
```

```

| name      | backend.presentation.compute.tasks.DeployComputeClusterTask |
| progress  | 100                                                            |
| result    |                                                                |
| state     | success                                                        |
| task_id   | be517afa-fae0-457e-819c-f4d6399f3ae2                        |
+-----+-----+

```

3.1.2 vinfra service compute delete

Delete all nodes from the compute cluster:

```
usage: vinfra service compute delete
```

Example:

```

# vinfra service compute delete
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | 063e8a15-fcfe-4629-865f-b5e5fa44b38f |
+-----+-----+

```

This command creates a task to release nodes from the compute cluster.

Task outcome:

```

# vinfra task show 063e8a15-fcfe-4629-865f-b5e5fa44b38f
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                        |
| name    | backend.presentation.compute.tasks.DestroyComputeClusterTask |
| result  |                                        |
| state   | success                              |
| task_id | 063e8a15-fcfe-4629-865f-b5e5fa44b38f |
+-----+-----+

```

3.2 Showing Compute Cluster Details and Overview

3.2.1 vinfra service compute show

Display compute cluster details:

```
usage: vinfra service compute show
```

Example:

```
# vinfra service compute show
+-----+-----+
| Field      | Value |
+-----+-----+
| capabilities | cpu_models:
|              | - Nehalem
|              | - Nehalem-IBRS
|              | - SandyBridge
|              | - SandyBridge-IBRS
|              | - IvyBridge
|              | - IvyBridge-IBRS
|              | - Haswell
|              | - Haswell-IBRS
|              | - Haswell-noTSX
|              | - Haswell-noTSX-IBRS
|              | - Broadwell
|              | - Broadwell-IBRS
|              | - Broadwell-noTSX
|              | - Broadwell-noTSX-IBRS
|              | - Skylake-Client
|              | - Skylake-Client-IBRS
|              | - Skylake-Server
|              | - Skylake-Server-IBRS
|              | - HostPassthrough
|              | os_distributions:
|              | - id: linux
|              |   os_type: linux
|              |   title: Generic Linux
|              | - id: centos8
|              |   os_type: linux
|              |   title: CentOS 8
|              | - id: centos7
|              |   os_type: linux
|              |   title: CentOS 7
|              | - id: centos6
|              |   os_type: linux
|              |   title: CentOS 6
|              | - id: rhel8
|              |   os_type: linux
|              |   title: Red Hat Enterprise Linux 8
|              | - id: rhel7
|              |   os_type: linux
|              |   title: Red Hat Enterprise Linux 7
|              | - id: ubuntu20.04
|              |   os_type: linux
|              |   title: Ubuntu 20.04
|              | - id: ubuntu18.04
|              |   os_type: linux
```

```

|         | title: Ubuntu 18.04 |
|         | - id: ubuntu16.04  |
|         |   os_type: linux   |
|         |   title: Ubuntu 16.04 |
|         | - id: debian9      |
|         |   os_type: linux   |
|         |   title: Debian 9   |
|         | - id: windows      |
|         |   os_type: windows |
|         |   title: Generic Windows |
|         | - id: win2k19      |
|         |   os_type: windows |
|         |   title: Windows Server 2019 |
|         | - id: win2k16      |
|         |   os_type: windows |
|         |   title: Windows Server 2016 |
|         | - id: win2k12r2    |
|         |   os_type: windows |
|         |   title: Windows Server 2012 R2 |
|         | - id: win2k12      |
|         |   os_type: windows |
|         |   title: Windows Server 2012 |
|         | - id: win2k8r2     |
|         |   os_type: windows |
|         |   title: Windows Server 2008 R2 |
|         | - id: win2k8       |
|         |   os_type: windows |
|         |   title: Windows Server 2008 |
|         | - id: win10        |
|         |   os_type: windows |
|         |   title: Windows 10 |
|         | - id: win8.1       |
|         |   os_type: windows |
|         |   title: Windows 8.1 |
|         | - id: win7         |
|         |   os_type: windows |
|         |   title: Windows 7 |
| options | cpu_model: null    |
| status  | active            |
+-----+

```

This command shows the status and capabilities of the compute cluster.

3.2.2 vinfra service compute stat

Display compute cluster statistics:

```
usage: vinfra service compute stat
```

Example:

```
# vinfra service compute stat
+-----+-----+
| Field | Value |
+-----+-----+
| compute | block_capacity: 0
|         | block_usage: 0
|         | cpu_usage: 0.0
|         | mem_total: 0
|         | mem_usage: 0
|         | vcpus: 0
| datetime | 2018-09-11T15:50:18.758258
| physical | block_capacity: 1099511627776
|         | block_free: 1099498911464
|         | cpu_cores: 10
|         | mem_total: 41006247936
| reserved | cpus: 5
|         | memory: 17721982976
| servers  | count: 0
|         | error: 0
|         | in_progress: 0
|         | running: 0
|         | stopped: 0
|         | top:
|         |   disk: []
|         |   memory: []
|         |   vcpus: []
+-----+-----+
```

This command shows the overview of the compute cluster.

3.3 Changing Compute Cluster Parameters

Change compute cluster parameters:

```
usage: vinfra service compute set [--cpu-model <cpu-model>] [--enable-k8saas]
                                     [--enable-lbaas] [--enable-metering]
```

`--cpu-model <cpu-model>`

Set the default CPU model for virtual machines. View the list of available CPU models using `compute cluster show`.

`--enable-k8saas`

Enable Kubernetes-as-a-Service services.

`--enable-lbaas`

Enable Load-Balancing-as-a-Service services.

`--enable-metering`

Enable metering services.

Example:

```
# vnfra service compute set --cpu-model Haswell
```

This command sets the default CPU model for VMs to Haswell.

3.4 Managing Compute Nodes

3.4.1 vnfra service compute node add

Add a node to the compute cluster:

```
usage: vnfra service compute node add [--compute] [--controller] [--force] <node>
```

`--compute`

Compute node role

`--controller`

Compute controller node role

`--force`

Skip checks for minimal hardware requirements

`<node>`

Node ID or hostname

Example:

```
# vnfra service compute node add 827a1f4e-56e5-404f-9113-88748c18f0c2 --compute
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 4c58e63c-31b6-406a-8070-9197445ec794 |
+-----+-----+
```

This command creates a task to add the node with the ID 827a1f4e-56e5-404f-9113-88748c18f0c2 to the compute cluster with the compute role.

Task outcome:

```
# vinfra task show 4c58e63c-31b6-406a-8070-9197445ec794
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.presentation.compute.tasks.AddComputeNodesTask |
| result |
| state | success |
| task_id | 4c58e63c-31b6-406a-8070-9197445ec794 |
+-----+-----+
```

3.4.2 vinfra service compute node list

List compute nodes:

```
usage: vinfra service compute node list
```

Example:

```
# vinfra service compute node list
+-----+-----+-----+-----+
| id | hypervisor_hostname | state | vms |
+-----+-----+-----+-----+
| 7ffa9540-5a20-41d1-b203-e3f349d62565 | stor-1.example.com.vstoragedomain | up | 1 |
| 6e8afc28-7f71-4848-bdbe-7c5de64c5013 | stor-3.example.com.vstoragedomain | up | 1 |
| 02ff64ae-5800-4090-b958-18b1fe8f5060 | stor-2.example.com.vstoragedomain | up | 1 |
| 827a1f4e-56e5-404f-9113-88748c18f0c2 | stor-5.example.com.vstoragedomain | up | 0 |
| 37c70bfb-c289-4794-8be4-b7a40c2b6d95 | stor-4.example.com.vstoragedomain | up | 1 |
+-----+-----+-----+-----+
```

This command lists nodes in the compute cluster.

3.4.3 vinfra service compute node show

Display compute node details:

```
usage: vinfra service compute node show <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra service compute node show 7ffa9540-5a20-41d1-b203-e3f349d62565
+-----+-----+
| Field | Value |
+-----+-----+
```

```

+-----+
| host      | stor-1.example.com.vstoragedomain |
| host_ip   | 10.37.130.101                     |
| hypervisor| id: 86f1ca2c-71c7-47a0-9c7f-bb9dd705e67e |
|           | state: up                          |
|           | status: enabled                    |
|           | vms: 0                             |
| id        | 7ffa9540-5a20-41d1-b203-e3f349d62565 |
| orig_hostname | stor-1.example.com                |
| roles     | - controller                       |
|           | - compute                          |
| state     | healthy                             |
+-----+

```

This command shows the details of the compute node with the ID 7ffa9540-5a20-41d1-b203-e3f349d62565.

3.4.4 vinfra service compute node fence

Fence a compute node:

```
usage: vinfra service compute node fence <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra service compute node fence e6255aed-d6e7-41b2-ba90-86164c1cd9a6
Operation successful
```

This command fences the node with the ID e6255aed-d6e7-41b2-ba90-86164c1cd9a6.

3.4.5 vinfra service compute node unfence

Unfence a compute node:

```
usage: vinfra service compute node unfence <node>
```

<node>

Node ID or hostname

Example:

```
# vinfra service compute node unfence e6255aed-d6e7-41b2-ba90-86164c1cd9a6
Operation successful
```

This command unfences the node with the ID e6255aed-d6e7-41b2-ba90-86164c1cd9a6.

3.4.6 vinfra service compute node release

Release a node from the compute cluster:

```
usage: vinfra service compute node release [--compute] [--controller] <node>
```

--compute

Compute node role

--controller

Compute controller node role

<node>

Node ID or hostname

Example:

```
# vinfra service compute node release 827a1f4e-56e5-404f-9113-88748c18f0c2
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 3b39738c-80a6-40a6-a50d-c3c8118ed212 |
+-----+-----+
```

This command creates a task to release the node with the ID 827a1f4e-56e5-404f-9113-88748c18f0c2 from the compute cluster.

Task outcome:

```
# vinfra task show 3b39738c-80a6-40a6-a50d-c3c8118ed212
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.compute.tasks.DeleteComputeNodesTask |
| result | |
| state | success |
| task_id | 3b39738c-80a6-40a6-a50d-c3c8118ed212 |
+-----+-----+
```

3.5 Managing Virtual Machines

3.5.1 `vinfra service compute server create`

Create a new compute server:

```
usage: vinfra service compute server create [--description <description>]
                                           [--metadata <metadata>]
                                           [--user-data <user-data>]
                                           [--key-name <key-name>]
                                           [--config-drive] [--count <count>]
                                           [--ha-enabled {true,false}]
                                           [--placements <placements>]
                                           --network <id=id[,key=value,...]>
                                           --volume <source=source[,key=value,...]>
                                           --flavor <flavor> <server-name>
```

`--description <description>`

Server description

`--metadata <metadata>`

Server metadata

`--user-data <user-data>`

User data file

`--key-name <key-name>`

Key pair to inject

`--config-drive`

Use an ephemeral drive

`--count <count>`

If count is specified and greater than 1, the name argument is treated as a naming pattern.

`--ha-enabled {true,false}`

Enable or disable HA for the compute server.

`--placements <placements>`

Names or IDs of placements to add the compute server to.

`--network <id=id[,key=value,...]>`

Create a compute server with a specified network. Specify this option multiple times to create multiple networks.

- `id`: attach network interface to a specified network (ID or name)
- comma-separated `key=value` pairs with keys (optional):
 - `mac`: MAC address for network interface
 - `fixed-ip`: fixed IP address for network interface
 - `spoofing-protection`: enable or disable spoofing protection for network interface (on or off)

`--volume <source=source[,key=value,...]>`

Create a compute server with a specified volume. Specify this option multiple times to create multiple volumes.

- `source`: source type (volume, image, snapshot, or blank)
- comma-separated `key=value` pairs with keys (optional):
 - `id`: resource ID or name for the specified source type (required for source types volume, image, and snapshot)
 - `size`: block device size, in gigabytes (required for source types image and blank)
 - `boot-index`: block device boot index (required for multiple volumes with source type volume)
 - `bus`: block device controller type (scsi)
 - `type`: block device type (disk or cdrom)
 - `rm`: remove block device on compute server termination (yes or no)
 - `storage-policy`: block device storage policy

`--flavor <flavor>`

Flavor ID or name

`<server-name>`

A new name for the compute server

Example:

```
# vinfra service compute server create myvm \
--network id=private,fixed-ip=192.168.128.100 \
--volume source=image,id=cirros,size=1 --flavor tiny
+-----+-----+
| Field      | Value                |
+-----+-----+
| config_drive |
```

```

| created      | 2019-05-29T11:24:04Z |
| description  |                       |
| flavor      | disk: 0               |
|              | ephemeral: 0         |
|              | extra_specs: {}      |
|              | original_name: tiny  |
|              | ram: 512              |
|              | swap: 0               |
|              | vcpus: 1              |
| ha_enabled   | True                  |
| host        |                       |
| id          | 8cd29296-8bee-4efb-828d-0e522d816c6e |
| key_name    |                       |
| metadata    | {}                    |
| name        | myvm                  |
| networks    | []                    |
| power_state | NOSTATE               |
| project_id  | b4267de6fd0c442da99542cd20f5932c |
| status      | BUILD                 |
| task_state  | scheduling             |
| updated     | 2019-05-29T11:24:21Z |
| user_data   |                       |
| volumes     | []                    |
+-----+-----+

```

This command creates a virtual machine `myvm` based on the default Cirros image and the flavor `tiny`, connects it to the network `private` with the fixed IP address `192.168.128.100`, and enables HA for it.

3.5.2 vinfra service compute server list

List compute servers:

```
usage: vinfra service compute server list
```

Example:

```

# vinfra service compute server list
+-----+-----+-----+-----+
| id          | name | status | host          |
+-----+-----+-----+-----+
| 8cd29296-8bee-4efb-828d-0e522d816c6e | myvm | ACTIVE | node001.vstoragedomain |
+-----+-----+-----+-----+

```

This command lists all virtual machines in the compute cluster.

3.5.3 vinfra service compute server show

Display compute server details:

```
usage: vinfra service compute server show <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server show myvm
+-----+-----+
| Field          | Value                                |
+-----+-----+
| config_drive   |                                       |
| created        | 2019-05-29T11:24:04Z                |
| description    |                                       |
| flavor         | disk: 0                              |
|                | ephemeral: 0                        |
|                | extra_specs: {}                     |
|                | original_name: tiny                 |
|                | ram: 512                             |
|                | swap: 0                              |
|                | vcpus: 1                            |
| ha_enabled     | True                                 |
| host           | node001.vstoragedomain              |
| id             | 8cd29296-8bee-4efb-828d-0e522d816c6e |
| key_name       |                                       |
| metadata       | {}                                   |
| name          | myvm                                 |
| networks       | - id: 79b3da71-c6a2-49e8-97f8-9431a065bed7 |
|                | ipam_enabled: true                  |
|                | ips:                                |
|                | - 192.168.128.100                   |
|                | mac_addr: fa:16:3e:d8:42:f6         |
|                | name: private                       |
|                | spoofing_protection: true           |
| orig_hostname  | node001                              |
| power_state    | RUNNING                              |
| project_id     | b4267de6fd0c442da99542cd20f5932c   |
| status         | ACTIVE                               |
| task_state     |                                       |
| updated        | 2019-05-29T11:24:21Z                |
| user_data      |                                       |
| volumes        | - delete_on_termination: false      |
|                | id: edd3df0a-95f5-4892-9053-2793a3976f94 |
+-----+-----+
```

This command shows the details of the virtual machine `myvm`.

3.5.4 vinfra service compute server stat

Display compute server statistics:

```
usage: vinfra service compute server stat <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server stat myvm
+-----+-----+
| Field   | Value                                     |
+-----+-----+
| datetime| 2019-05-29T11:39:46.429000+00:00        |
| metrics | block_capacity: 1073741824              |
|         | block_usage: 268435456                  |
|         | cpu_usage: 0                            |
|         | mem_usage: 149876736                    |
+-----+-----+
```

This command shows the statistics for the virtual machine `myvm`.

3.5.5 vinfra service compute server set

Modify compute server parameters:

```
usage: vinfra service compute server set [--name <name>] [--description <description>]
      [--ha-enabled <ha_enabled>]
      [--no-placements | --placement placement]
      <server>
```

`--name <name>`

A new name for the compute server

`--description <description>`

A new description for the compute server

`--ha-enabled {true,false}`

Enable or disable HA for the compute server.

`--no-placements`

Clean up placements from the compute server.

--placement placement

Placement name or ID to add the compute server to. Specify this option multiple times to add the compute server to multiple placements.

<server>

Compute server ID or name

Example:

```
# vinfra service compute server set myvm --description "My new VM" --ha-enabled false
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| config_drive   |                                           |
| created        | 2019-05-29T11:24:04Z                     |
| description     | My new VM                                |
| flavor         | disk: 0                                  |
|                | ephemeral: 0                             |
|                | extra_specs: {}                          |
|                | original_name: tiny                      |
|                | ram: 512                                  |
|                | swap: 0                                   |
|                | vcpus: 1                                  |
| ha_enabled     | False                                    |
| host           | node001.vstoragedomain                   |
| id             | 8cd29296-8bee-4efb-828d-0e522d816c6e     |
| key_name       |                                           |
| metadata       | {}                                        |
| name           | myvm                                      |
| networks       | - id: 79b3da71-c6a2-49e8-97f8-9431a065bed7 |
|                | ipam_enabled: true                       |
|                | ips:                                     |
|                | - 192.168.128.100                         |
|                | mac_addr: fa:16:3e:d8:42:f6              |
|                | name: private                            |
|                | spoofing_protection: true                |
| orig_hostname  | node001                                   |
| power_state    | RUNNING                                  |
| project_id     | b4267de6fd0c442da99542cd20f5932c       |
| status         | ACTIVE                                   |
| task_state     |                                           |
| updated        | 2019-05-29T11:24:21Z                     |
| user_data      |                                           |
| volumes        | - delete_on_termination: false           |
|                | id: edd3df0a-95f5-4892-9053-2793a3976f94 |
+-----+-----+
```

This command adds a description to the virtual machine `myvm` and disables HA for it.

3.5.6 vinfra service compute server iface attach

Attach a network to a compute server:

```
usage: vinfra service compute server iface attach [--mac <mac>] [--ip <ip-address>]
        [--spoofing-protection {on,off}]
        --server <server> --network <network>
```

--mac <mac>

MAC address

--ip <ip-address>

IP address

--spoofing-protection {on|off}

Enable spoofing protection for the network interface

--server <server>

Compute server ID or name

--network <network>

Network ID or name

Example:

```
# vinfra service compute server iface attach --network myprivnet --server myvm
+-----+-----+
| Field          | Value                                |
+-----+-----+
| fixed_ip       | 192.168.129.8                        |
| id             | 690ed3f2-2301-40e2-879a-126db2ecb57b |
| mac_address    | fa:16:3e:54:59:08                    |
| network_id     | 0710372e-2bdf-4dfe-b413-eb763da37e68 |
| spoofing<...> | False                                 |
+-----+-----+
```

This command attaches the private network `myprivnet` to the virtual machine `myvm`.

3.5.7 vinfra service compute server iface list

List compute server networks:

```
usage: vinfra service compute server iface list --server <server>
```

```
--server <server>
    Compute server ID or name
```

Example:

```
# vinfra service compute server iface list --server myvm
+-----+-----+-----+-----+
| id          | network_id  | mac_address  | fixed_ip    |
+-----+-----+-----+-----+
| 690ed3f2-<...> | 0710372e-<...> | fa:16:3e:54:59:08 | 192.168.129.8 |
| a5b13bf3-<...> | 1bf2c9da-<...> | fa:16:3e:b9:33:bb | 192.168.128.100 |
+-----+-----+-----+-----+
```

This command lists the virtual networks that the virtual machine `myvm` is attached to. It also shows VM's IP address in each network.

3.5.8 vinfra service compute server iface detach

Detach a network interface from a compute server:

```
usage: vinfra service compute server iface detach --server <server> <interface>
```

```
--server <server>
    Compute server ID or name

<interface>
    Network interface ID
```

Example:

```
# vinfra service compute server iface detach 471e37fd-13ae-4b8f-b70c-90ac02cc4386 \
--server 6c80b07f-da46-4a8a-89a4-eeeb8faceb27
Operation successful
```

This command detaches the network interface with the ID `471e37fd-13ae-4b8f-b70c-90ac02cc4386` from the VM with the ID `6c80b07f-da46-4a8a-89a4-eeeb8faceb27`.

3.5.9 vinfra service compute server volume attach

Attach a volume to a compute server:

```
usage: vinfra service compute server volume attach --server <server> <volume>
```

```
--server <server>
    Compute server ID or name
```

```
<volume>
    Volume ID or name
```

Example:

```
# vinfra service compute server volume attach e4cb5363-1fb2-41f5-b24b-18f98a388cba \
--server 871fef54-519b-4111-b18d-d2039e2410a8
+-----+-----+
| Field | Value |
+-----+-----+
| device | /dev/vdb |
| id     | e4cb5363-1fb2-41f5-b24b-18f98a388cba |
+-----+-----+
```

This command attaches the available volume with the ID e4cb5363-1fb2-41f5-b24b-18f98a388cba to the VM with the ID 871fef54-519b-4111-b18d-d2039e2410a8.

3.5.10 vinfra service compute server volume list

List compute server volumes:

```
usage: vinfra service compute server volume list --server <server>
```

```
--server <server>
    Compute server ID or name
```

Example:

```
# vinfra service compute server volume list --server myvm
+-----+-----+
| id           | device |
+-----+-----+
| e4cb5363-1fb2-41f5-b24b-18f98a388cba | /dev/vdb |
| b325cc6e-8de1-4b6c-9807-5a497e3da7e3 | /dev/vda |
+-----+-----+
```

This command lists the volumes attached to the virtual machine myvm.

3.5.11 vinfra service compute server volume show

Show details of a compute server volume:

```
usage: vinfra service compute server volume show --server <server> <volume>
```

```
--server <server>
```

Compute server ID or name

```
<volume>
```

Volume ID or name

Example:

```
# vinfra service compute server volume show --server myvm \
e4cb5363-1fb2-41f5-b24b-18f98a388cba
+-----+-----+
| Field | Value |
+-----+-----+
| device | /dev/vdb |
| id     | e4cb5363-1fb2-41f5-b24b-18f98a388cba |
+-----+-----+
```

This command shows the details for the volume with the ID e4cb5363-1fb2-41f5-b24b-18f98a388cba attached to the virtual machine myvm.

3.5.12 vinfra service compute server volume detach

Detach a volume from a compute server:

```
usage: vinfra service compute server volume detach --server <server> <volume>
```

```
--server <server>
```

Compute server ID or name

```
<volume>
```

Volume ID or name

Example:

```
# vinfra service compute server volume detach e4cb5363-1fb2-41f5-b24b-18f98a388cba \
--server 871fef54-519b-4111-b18d-d2039e2410a8
Operation successful
```

This command detaches the volume with the ID e4cb5363-1fb2-41f5-b24b-18f98a388cba from the VM with the ID 871fef54-519b-4111-b18d-d2039e2410a8.

3.5.13 vinfra service compute server log

Display compute server log:

```
usage: vinfra service compute server log <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server log myvm > myvm.log
```

This command prints the log of the virtual machine `myvm` to the file `myvm.log`.

3.5.14 vinfra service compute server migrate

Migrate a compute server to another host:

```
usage: vinfra service compute server migrate [--cold] [--node <node>] <server>
```

--cold

Perform cold migration. If not set, the migration type is determined automatically.

--node <node>

Destination node ID or hostname

<server>

Compute server ID or name

Example:

```
# vinfra service compute server migrate 6c80b07f-da46-4a8a-89a4-eeeb8faceb27 \
--node e6255aed-d6e7-41b2-ba90-86164c1cd9a6
Operation successful
```

This command starts migration of the VM with the ID 6c80b07f-da46-4a8a-89a4-eeeb8faceb27 to the compute node with the ID e6255aed-d6e7-41b2-ba90-86164c1cd9a6.

3.5.15 vinfra service compute server resize

Resize a compute server:

```
usage: vinfra service compute server resize --flavor <flavor> <server>
```

`--flavor <flavor>`

Apply flavor with ID or name

`<server>`

Compute server ID or name

Example:

```
# vinfra service compute server resize myvm --flavor small
Operation successful
```

This command changes the flavor of the virtual machine `myvm` to `small`.

3.5.16 vinfra service compute server start

Start a compute server:

```
usage: vinfra service compute server start <server>
```

`<server>`

Compute server ID or name

Example:

```
# vinfra service compute server start myvm
Operation successful
```

This command starts the virtual machine `myvm`.

3.5.17 vinfra service compute server pause

Pause a compute server:

```
usage: vinfra service compute server pause <server>
```

`<server>`

Compute server ID or name

Example:

```
# vinfra service compute server pause myvm
```

This command pauses the running virtual machine `myvm`.

3.5.18 vinfra service compute server unpause

Unpause a compute server:

```
usage: vinfra service compute server unpause <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server unpause myvm
```

This command unpauses the paused virtual machine `myvm`.

3.5.19 vinfra service compute server suspend

Suspend a compute server:

```
usage: vinfra service compute server suspend <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server suspend myvm  
Operation successful
```

This command suspends the running virtual machine `myvm`.

3.5.20 vinfra service compute server resume

Resume a compute server:

```
usage: vinfra service compute server resume <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server resume myvm
Operation successful
```

This command resumes the suspended virtual machine `myvm`.

3.5.21 vinfra service compute server reboot

Reboot a compute server:

```
usage: vinfra service compute server reboot [--hard] <server>
```

--hard

Perform hard reboot

<server>

Compute server ID or name

Example:

```
# vinfra service compute server reboot myvm
Operation successful
```

This command reboots the virtual machine `myvm`.

3.5.22 vinfra service compute server reset-state

Reset compute server state:

```
usage: vinfra service compute server reset-state [--state-error] <server>
```

--state-error

Reset server to 'ERROR' state

<server>

Compute server ID or name

Example:

```
# vinfra service compute server reset-state myvm
Operation successful
```

This command resets the transitional state of the virtual machine `myvm` to the previous one.

3.5.23 vinfra service compute server stop

Shut down a compute server:

```
usage: vinfra service compute server stop [--hard] <server>
```

--hard

Power off a compute server

<server>

Compute server ID or name

Example:

```
# vinfra service compute server stop myvm
Operation successful
```

This command stops the virtual machine `myvm`.

3.5.24 vinfra service compute server shelve

Shelve a compute server:

```
usage: vinfra service compute server shelve <server>
```

<server>

Compute server ID or name.

Example:

```
# vinfra service compute server shelve myvm
```

This command unbinds the virtual machine `myvm` from the node it is hosted on and releases its reserved resources such as CPU and RAM.

3.5.25 vinfra service compute server unshelve

Unshelve a compute server:

```
usage: vinfra service compute server unshelve <server>
```

<server>

Compute server ID or name.

Example:

```
# vinfra service compute server unshelve myvm
```

This command spawns the virtual machine `myvm` on a node with enough resources to host it.

3.5.26 vinfra service compute server evacuate

Evacuate a stopped compute server from a failed host:

```
usage: vinfra service compute server evacuate <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server evacuate myvm  
Operation successful
```

This command evacuates the stopped VM `myvm` from its node to another, healthy compute node.

3.5.27 vinfra service compute server delete

Delete a compute server:

```
usage: vinfra service compute server delete <server>
```

<server>

Compute server ID or name

Example:

```
# vinfra service compute server delete myvm
Operation successful
```

This command deletes the virtual machine `myvm`.

3.6 Managing Images

3.6.1 `vinfra service compute image create`

Create a new compute image:

```
usage: vinfra service compute image create [--min-disk <size-gb>] [--min-ram <size-mb>]
      [--os-distro <os-distro>] [--protected]
      [--disk-format <disk_format>]
      [--container-format <format>]
      --file <file> <image-name>
```

`--min-disk <size-gb>`

Minimum disk size required to boot from image, in gigabytes

`--min-ram <size-mb>`

Minimum RAM size required to boot from image, in megabytes

`--os-distro <os-distro>`

OS distribution. To list available distributions, run `service compute cluster show`.

`--protected`

Protect image from deletion

`--disk-format <disk_format>`

Disk format `aki`, `ami`, `ari`, `detect`, `iso`, `ploop`, `qcow2`, `raw`, `vdi`, `vhd`, `vhdx`, `vmdk` (default: `detect`)

`--container-format <format>`

Container format: `aki`, `ami`, `ari`, `bare`, `docker`, `ovf`, `ova` (default: `bare`)

`--file <file>`

Create image from a local file

`<image-name>`

Image name

Example:

```
# vinfra service compute image create mycirrosimg \
--file /distr/cirros-0.4.0-x86_64-disk.img
Uploading image to server [elapsed time: 0:00:04]... |
+-----+-----+
| Field   | Value                                     |
+-----+-----+
| task_id | 03874663-d03f-4891-a10b-64837e7faf43 |
+-----+-----+
```

This command creates a task to create a Cirros image from the local file and upload it to Virtuozzo Hybrid Infrastructure.

Task outcome:

```
# vinfra task show 03874663-d03f-4891-a10b-64837e7faf43
+-----+-----+
| Field   | Value                                     |
+-----+-----+
| details | |
| name    | backend.presentation.compute.images.tasks.ImportComputeImageTask |
| result  | id: 179f45ef-c5d6-4270-b0c0-085b542544c5 |
| state   | success |
| task_id | 03874663-d03f-4891-a10b-64837e7faf43 |
+-----+-----+
```

3.6.2 vinfra service compute image list

List compute images:

```
usage: vinfra service compute image list
```

Example:

This command lists images available to the compute cluster.

```
# vinfra service compute image list
+-----+-----+-----+-----+-----+
| id                | name      | size  | status | disk_format |
+-----+-----+-----+-----+-----+
| 179f45ef-c5d6-4270-b0c0-085b542544c5 | mycirrosimg | 12716032 | active | qcow2 |
| 4741274f-5cca-4205-8f66-a2e89fb346cc | cirros      | 12716032 | active | qcow2 |
+-----+-----+-----+-----+-----+
```

3.6.3 vinfra service compute image show

Display compute image details:

```
usage: vinfra service compute image show <image>
```

<image>

Image ID or name

Example:

```
# vinfra service compute image show 4741274f-5cca-4205-8f66-a2e89fb346cc
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| checksum       | 443b7623e27ecf03dc9e01ee93f67afe       |
| container_format | bare                                     |
| created_at     | 2018-09-11T13:29:10Z                    |
| disk_format    | qcow2                                    |
| file           | /api/v2/compute/images/4741274f-5cca-4205-8f66-a2e89fb346cc/file/ |
| id             | 4741274f-5cca-4205-8f66-a2e89fb346cc   |
| min_disk       | 1                                        |
| min_ram        | 0                                        |
| name           | cirros                                   |
| os_distro      | linux                                    |
| os_type        | linux                                    |
| project_id     | 72a5db3a033c403a86756021e601ef34       |
| protected      | False                                    |
| size           | 12716032                                 |
| status         | active                                   |
| tags           | []                                       |
| updated_at     | 2018-09-11T13:29:13Z                    |
| virtual_size   |                                           |
| visibility     | public                                   |
+-----+-----+
```

This command shows the details of the default Cirros image.

3.6.4 vinfra service compute image set

Modify compute image parameters:

```
usage: vinfra service compute image set [--min-disk <size-gb>] [--min-ram <size-mb>]
      [--os-distro <os-distro>] [--protected]
      [--name <name>] <image>
```

--min-disk <size-gb>

Minimum disk size required to boot from image, in gigabytes

`--min-ram <size-mb>`

Minimum RAM size required to boot from image, in megabytes

`--os-distro <os-distro>`

OS distribution. To list available distributions, run `service compute cluster show`.

`--protected`

Protect image from deletion

`--name <name>`

Image name

`<image>`

Image ID or name

Example:

```
# vinfra service compute image set 4741274f-5cca-4205-8f66-a2e89fb346cc --protected --min-ram 1
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| checksum       | 443b7623e27ecf03dc9e01ee93f67afe       |
| container_format | bare                                     |
| created_at     | 2018-09-11T13:29:10Z                   |
| disk_format    | qcow2                                    |
| file           | /api/v2/compute/images/4741274f-5cca-4205-8f66-a2e89fb346cc/file/ |
| id             | 4741274f-5cca-4205-8f66-a2e89fb346cc   |
| min_disk       | 1                                        |
| min_ram        | 1                                        |
| name           | cirros                                   |
| os_distro      | linux                                    |
| os_type        | linux                                    |
| project_id     | 72a5db3a033c403a86756021e601ef34     |
| protected      | True                                     |
| size           | 12716032                                 |
| status         | active                                   |
| tags           | []                                       |
| updated_at     | 2018-09-12T09:26:29Z                   |
| virtual_size   |                                           |
| visibility     | public                                   |
+-----+-----+
```

This command protects the default Cirros image and sets the minimum RAM size for it to 1 GB.

3.6.5 vinfra service compute image save

Download a compute image:

```
usage: vinfra service compute image save [--file <filename>] <image>
```

`--file <filename>`

File to save the image to (default: stdout)

`<image>`

Image ID or name

Example:

```
# vinfra service compute image save 4741274f-5cca-4205-8f66-a2e89fb346cc --file cirros.qcow2
Operation successful
```

This command downloads the default Cirros image to the local disk as `cirros.qcow2`.

3.6.6 vinfra service compute image delete

Delete a compute image:

```
usage: vinfra service compute image delete <image>
```

`<image>`

Image ID or name

Example:

```
# vinfra service compute image delete 179f45ef-c5d6-4270-b0c0-085b542544c5
Operation successful
```

This command deletes the image with the ID `179f45ef-c5d6-4270-b0c0-085b542544c5`.

3.7 Managing Placements

3.7.1 vinfra service compute placement create

Create a new compute placement:

```
usage: vinfra service compute placement create [--description <description>]
                                             [--nodes <nodes>]
                                             [--images <images>]
                                             <placement-name>
```

--description <description>

Placement description

--nodes <nodes>

A comma-separated list of compute node IDs to assign to a compute placement

--images <images>

A comma-separated list of image IDs to assign to a compute placement

<placement-name>

Placement name

Example:

```
# vinfra service compute placement create placement1
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| description |                                           |
| id          | e4230b75-a858-404c-be3b-4b3f2dedb057 |
| images      | 0                                         |
| name        | placement1                               |
| nodes       | 0                                         |
| servers     | 0                                         |
+-----+-----+
```

This command creates a placement called `placement1`.

3.7.2 vinfra service compute placement assign

Assign nodes or images to a placement:

```
usage: vinfra service compute placement assign (--images <images> | --nodes <nodes>)
                                             <placement>
```

--images <images>

A comma-separated list of image IDs to assign to a compute placement

--nodes <nodes>

A comma-separated list of compute node IDs to assign to a compute placement

<placement>

Placement ID or name

If the node has virtual machines on it, they will not automatically inherit the changes to node's placement configuration. That is, if you add a node with VMs to a placement, the VMs will not be assigned the same placement. If you delete such a node from a placement, the VMs will still keep that placement. You will have to edit virtual machines' placement configuration accordingly using the `vinfra service compute server set` command. Make sure that the node and its VMs have the same placement configuration.

Example:

```
# vinfra service compute placement assign \
--images b23e23e8-7338-4a09-a827-3c9c509cf35c placement1
Operation successful.
```

This command assigns the image with the ID `b23e23e8-7338-4a09-a827-3c9c509cf35c` to placement `placement1`.

3.7.3 vinfra service compute placement delete assign

Remove images and nodes from a compute placement:

```
usage: vinfra service compute placement delete assign
      (--image-id <images> | --node-id <nodes>)
      <placement>
```

--image-id <images>

An image ID to remove from a compute placement

--node-id <nodes>

A compute node ID to remove from a compute placement

<placement>

Placement ID or name

If the node has virtual machines on it, they will not automatically inherit the changes to node's placement configuration. That is, if you add a node with VMs to a placement, the VMs will not be assigned the same placement. If you delete such a node from a placement, the VMs will still keep that placement. You will have to edit virtual machines' placement configuration accordingly using the `vinfra service compute server set` command. Make sure that the node and its VMs have the same placement configuration.

Example:

```
# vinfra service compute placement delete assign \
--image b23e23e8-7338-4a09-a827-3c9c509cf35c placement1
Operation successful.
```

This command removes the image with the ID b23e23e8-7338-4a09-a827-3c9c509cf35c from placement placement1.

3.7.4 vinfra service compute placement list

List compute placements:

```
usage: vinfra service compute placement list
```

Example:

```
# vinfra service compute placement list
+-----+-----+-----+-----+-----+-----+
| id          | name          | description | nodes | images | servers |
+-----+-----+-----+-----+-----+-----+
| 2d152d33-<...> | placement1    |             | 0     | 1     | 0     |
| 11857e11-<...> | placement2    |             | 1     | 0     | 0     |
+-----+-----+-----+-----+-----+-----+
```

This command lists available compute placements.

3.7.5 vinfra service compute placement show

Display compute placement details:

```
usage: vinfra service compute placement show <placement>
```

<placement>

Placement ID or name

Example:

```
# vinfra service compute placement show placement1
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| description |                                           |
| id          | e4230b75-a858-404c-be3b-4b3f2dedb057 |
| images      | 0                                         |
| name        | placement1                               |
| nodes       | 0                                         |
```

```
| servers      | 0 |
+-----+-----+
```

This command shows the details of the placement placement1.

3.7.6 vinfra service compute placement update

Update a compute placement:

```
usage: vinfra service compute placement update [--name <placement-name>]
       [--description <placement-description>]
       <placement>
```

`--name <placement-name>`

A new name for the placement

`--description <placement-description>`

A new description for the placement

`<placement>`

Placement ID or name

Example:

```
# vinfra service compute placement update --name placement1-UPD placement1
Operation successful
```

This command renames the placement placement1 to placement1-UPD.

3.7.7 vinfra service compute placement delete

Delete a compute placement:

```
usage: vinfra service compute placement delete <placement>
```

`<placement>`

Placement ID or name

Example:

```
# vinfra service compute placement delete placement1
Operation successful
```

This command deletes the placement placement1.

3.8 Managing Flavors

3.8.1 vinfra service compute flavor create

Create a new compute flavor:

```
usage: vinfra service compute flavor create [--swap <size-mb>] --vcpus <vcpus>
      --ram <size-mb> <flavor-name>
```

--swap <size-mb>

Swap space size, in megabytes

--vcpus <vcpus>

Number of virtual CPUs

--ram <size-mb>

Memory size, in megabytes

<flavor-name>

Flavor name

Example:

```
# vinfra service compute flavor create myflavor --vcpus 1 --ram 3072
+-----+-----+
| Field | Value |
+-----+-----+
| id    | 561a48ea-0c1c-4152-8b7d-e4b4af276c2d |
| name  | myflavor |
| ram   | 3072 |
| swap  | 0 |
| vcpus | 1 |
+-----+-----+
```

This command creates a flavor `myflavor` with 1 vCPU and 3 GB RAM.

3.8.2 vinfra service compute flavor list

List compute flavors:

```
usage: vinfra service compute flavor list
```

Example:

```
# vinfra service compute flavor list
+-----+-----+-----+-----+
| id           | name   | ram  | swap | vcpus |
+-----+-----+-----+-----+
| 100          | tiny   | 512  | 0    | 1     |
| 101          | small  | 2048 | 0    | 1     |
| 102          | medium | 4096 | 0    | 2     |
| 103          | large  | 8192 | 0    | 4     |
| 104          | xlarge | 16384| 0    | 8     |
| 561a48ea-0c1c-4152-8b7d-e4b4af276c2d | myflavor | 3072 | 0    | 1     |
+-----+-----+-----+-----+
```

This command lists all flavors.

3.8.3 vinfra service compute flavor show

Display compute flavor details:

```
usage: vinfra service compute flavor show <flavor>
```

<flavor>

Flavor ID or name

Example:

```
# vinfra service compute flavor show myflavor
+-----+-----+
| Field | Value |
+-----+-----+
| id    | 561a48ea-0c1c-4152-8b7d-e4b4af276c2d |
| name  | myflavor |
| ram   | 3072 |
| swap  | 0 |
| vcpus | 1 |
+-----+-----+
```

This command shows the details of the flavor myflavor.

3.8.4 vinfra service compute flavor delete

Delete a compute flavor:

```
usage: vinfra service compute flavor delete <flavor>
```

<flavor>

Flavor ID or name

Example:

```
# vinfra service compute flavor delete myflavor
Operation successful
```

This command deletes the flavor `myflavor`.

3.9 Managing Compute SSH Keys

3.9.1 vinfra service compute key create

Create a new compute SSH key:

```
usage: vinfra service compute key create --public-key <public-key>
      [--description <description>] <ssh-key>
```

`--public-key <public-key>`

Filename for a public key to upload

`--description <description>`

SSH key description

`<ssh-key>`

SSH key name

Example:

```
# vinfra service compute key create publickey --public-key /root/.ssh/id_rsa.pub \
--description 'public key'
+-----+
| Field      | Value                                |
+-----+
| created_at | 2019-04-25T13:41:14.241736+00:00    |
| description | public key                           |
| name       | publickey                            |
+-----+
```

This command creates a public SSH key `publickey`.

3.9.2 vinfra service compute key list

List compute SSH keys:

```
usage: vinfra service compute key list
```

Example:

```
# vinfra service compute key list
+-----+-----+-----+
| name      | description | created_at |
+-----+-----+-----+
| testkey   | test key   | 2019-04-24T13:41:05.209837+00:00 |
| publickey | public key | 2019-04-25T13:41:14.241736+00:00 |
+-----+-----+-----+
```

This command lists all SSH keys.

3.9.3 vinfra service compute key show

Display compute SSH key details:

```
usage: vinfra service compute key show <ssh-key>
```

<ssh-key>

SSH key name

Example:

```
# vinfra service compute key show publickey
+-----+-----+
| Field          | Value |
+-----+-----+
| created_at     | 2019-04-25T13:41:14.241736+00:00 |
| description    | public key |
| name           | publickey |
| public_key_fingerprint | 1a:fb:de:d8:1e:0a:84:30:fc:ff:e4:fd:89:e7:96:a9 |
+-----+-----+
```

This command shows the details of the SSH key publickey.

3.9.4 vinfra service compute key delete

Delete a compute SSH key:

```
usage: vinfra service compute key delete <ssh-key>
```

<ssh-key>

SSH key name

Example:

```
# vinfra service compute key delete publickey
Operation successful
```

This command deletes the SSH key publickey.

3.10 Managing Virtual Networks

3.10.1 vinfra service compute network create

Create a compute network:

```
usage: vinfra service compute network create [--dhcp | --no-dhcp]
                                             [--dns-nameserver <dns-nameserver>]
                                             [--allocation-pool <allocation-pool>]
                                             [--gateway <gateway> | --no-gateway]
                                             [--ip-version <ip-version>]
                                             [--physical-network <physical-network>]
                                             [--cidr <cidr>] <network-name>
```

--dhcp

Enable DHCP.

--no-dhcp

Disable DHCP.

--dns-nameserver <dns-nameserver>

DNS server IP address. This option can be used multiple times.

--allocation-pool <allocation-pool>

Allocation pool to create inside the network in the format: ip_addr_start-ip_addr_end. This option can be used multiple times.

`--gateway <gateway>`
 Gateway IP address

`--no-gateway`
 Do not configure a gateway for this network.

`--ip-version <ip-version>`
 Network IP version

`--physical-network <physical-network>`
 A physical network to link to a public network

`--cidr <cidr>`
 Subnet range in CIDR notation

`<network-name>`
 Network name

Example:

```
# vnfra service compute network create myprivnet --type vxlan \
--cidr 192.128.128.0/24 --gateway 192.128.128.1
```

Field	Value
id	3848fb5d-bc98-4320-acd0-cde2df7c5bdd
name	myprivnet
physical_network	
project_id	72a5db3a033c403a86756021e601ef34
subnet	allocation_pools: - end: 192.128.128.254 start: 192.128.128.2 cidr: 192.128.128.0/24
	dns_nameservers: []
	enable_dhcp: true
	gateway_ip: 192.128.128.1
	ip_version: 4
type	vxlan

This command creates a private network `myprivnet` with the specific CIDR and gateway.

3.10.2 vinfra service compute network list

List compute networks:

```
usage: vinfra service compute network list
```

Example:

```
# vinfra service compute network list -c id -c name -c cidr -c type -c allocation_pools
+-----+-----+-----+-----+-----+
| id          | name      | type  | cidr          | allocation_pools |
+-----+-----+-----+-----+-----+
| 1bf2c9da-<...> | private  | vxlan | 192.168.128.0/24 | - end: 192.168.128.254 |
|               |          |       |                 | start: 192.168.128.2 |
| 3848fb5d-<...> | myprivnet | vxlan | 192.128.128.0/24 | - end: 192.128.128.254 |
|               |          |       |                 | start: 192.128.128.2 |
| 417606ac-<...> | public   | flat  | 10.94.0.0/16    | - end: 10.94.129.79   |
|               |          |       |                 | start: 10.94.129.64   |
+-----+-----+-----+-----+-----+
```

This command lists networks used in the compute cluster. (The output is abridged to fit on page.)

3.10.3 vinfra service compute network show

Display compute network details:

```
usage: vinfra service compute network show <network>
```

<network>

Network ID or name

Example:

```
# vinfra service compute network show 417606ac-1dbe-426a-844d-e047831ddce9
+-----+-----+
| Field          | Value |
+-----+-----+
| allocation_pools |       |
| cidr           |       |
| dns_nameservers |       |
| enable_dhcp     |       |
| gateway_ip      |       |
| id             | 417606ac-1dbe-426a-844d-e047831ddce9 |
| ip_version      |       |
| name           | public |
| physical_network | Public |
| project_id      | 72a5db3a033c403a86756021e601ef34 |
+-----+-----+
```

```
| type          | flat          |
+-----+-----+
```

This command shows the details of the network with the ID 417606ac-1dbe-426a-844d-e047831ddce9.

3.10.4 vinfra service compute network set

Modify compute network parameters:

```
usage: vinfra service compute network set [--dhcp | --no-dhcp]
                                           [--dns-nameserver <dns-nameserver>]
                                           [--allocation-pool <allocation-pool>]
                                           [--gateway <gateway> | --no-gateway]
                                           [--name <name>] <network>
```

--dhcp

Enable DHCP.

--no-dhcp

Disable DHCP.

--dns-nameserver <dns-nameserver>

DNS server IP address. This option can be used multiple times.

--allocation-pool <allocation-pool>

Allocation pool to create inside the network in the format: ip_addr_start-ip_addr_end. This option can be used multiple times.

--gateway <gateway>

Gateway IP address

--no-gateway

Do not configure a gateway for this network.

--name <name>

A new name for the network

<network>

Network ID or name

Example:

```
# vinfra service compute network set myprivnet --no-dhcp
+-----+-----+
```

Field	Value
id	3848fb5d-bc98-4320-acd0-cde2df7c5bdd
name	myprivnet
physical_network	
project_id	72a5db3a033c403a86756021e601ef34
subnet	allocation_pools:
	- end: 192.128.128.254
	start: 192.128.128.2
	cidr: 192.128.128.0/24
	dns_nameservers: []
	enable_dhcp: false
	gateway_ip: 192.128.128.1
	ip_version: 4
type	vxlan

This command disables DHCP for the private network `myprivnet`.

3.10.5 vinfra service compute network delete

Delete a compute network:

```
usage: vinfra service compute network delete <network>
```

<network>

Network ID or name

Example:

```
# vinfra service compute network delete myprivnet
Operation successful
```

This command deletes the private network `myprivnet`.

3.11 Managing Virtual Routers

3.11.1 vinfra service compute router create

Create a virtual router:

```
usage: vinfra service compute router create [--external-gateway <network>]
      [--enable-snat | --disable-snat]
      [--fixed-ip <fixid-ip>]
```

```
[--internal-interface <network=network,
ip-addr=ip-addr>|<network>] <router-name>
```

```
--external-gateway <network>
```

Specify a public network to be used as the router's external gateway (name or ID)

```
--enable-snat
```

Enable source NAT on the external gateway

```
--disable-snat
```

Disable source NAT on the external gateway

```
--fixed-ip <fixid-ip>
```

Desired IP on the external gateway

```
--internal-interface <network=network,ip-addr=ip-addr>|<network>
```

Specify an internal interface. This option can be used multiple times.

- `network`: name of a private virtual network.
- `ip-addr`: an unused IP address from the selected private network to assign to the interface; specify if the default gateway of the selected private network is in use.

```
<router-name>
```

Virtual router name

Example:

```
# vinfra service compute router create myrouter --external-gateway public \
--internal-interface private --enable-snat
```

Field	Value
external_gateway_info	enable_snat: true ip_addresses: - 10.94.129.76 network_id: 720e45bc-4225-49de-9346-26513d8d1262
id	b9d8b000-5d06-4768-9f65-2715250cda53
name	myrouter
project_id	894696133031439f8aaa7e4868dcbd4d
routes	[]
status	ACTIVE

This command creates a router `myrouter` between the public network `public` and the private network `private` with enabled SNAT on the external gateway.

3.11.2 vinfra service compute router list

List virtual routers:

```
usage: vinfra service compute router list
```

Example:

```
# vinfra service compute router list -c id -c external_gateway_info -c name -c status
+-----+-----+-----+-----+
| id          | external_gateway_info | name   | status |
+-----+-----+-----+-----+
| b9d8b000-5d06-<...> | enable_snat: true     | myrouter | ACTIVE |
|                | ip_addresses:         |         |         |
|                | - 10.94.129.76       |         |         |
|                | network_id: 720e45bc-4225-<...> |         |         |
+-----+-----+-----+-----+
```

This command lists virtual routers used in the compute cluster. (The output is abridged to fit on page.)

3.11.3 vinfra service compute router show

Display information about a virtual router:

```
usage: vinfra service compute router show <router>
```

<router>

Virtual router name

Example:

```
# vinfra service compute router show myrouter
+-----+-----+
| Field          | Value |
+-----+-----+
| external_gateway_info | enable_snat: true |
|                  | ip_addresses:    |
|                  | - 10.94.129.76  |
|                  | network_id: 720e45bc-4225-49de-9346-26513d8d1262 |
| id             | b9d8b000-5d06-4768-9f65-2715250cda53 |
| name           | myrouter |
| project_id     | 894696133031439f8aaa7e4868dcbd4d |
| routes         | [] |
| status         | ACTIVE |
+-----+-----+
```

This command shows the details of the virtual router `myrouter`.

3.11.4 vinfra service compute router set

Modify virtual router parameters:

```
usage: vinfra service compute router set [--name <name>] [--external-gateway <network> |
--no-external-gateway] [--fixed-ip <fixed-ip>]
[--enable-snat | --disable-snat]
[--route <destination=destination,nextthop=nextthop> |
--no-route] <router>
```

--name <name>

Virtual router name

--external-gateway <network>

Specify a public network to be used as the router's external gateway (name or ID)

--no-external-gateway

Remove the external gateway from the router

--enable-snat

Enable source NAT on the external gateway

--disable-snat

Disable source NAT on the external gateway

--fixed-ip <fixed-ip>

Desired IP on the external gateway

--route <destination=destination,nextthop=nextthop>

A static route for the router. This option can be used multiple times.

- destination: destination subnet range in CIDR notation.
- nextthop: next hop IP address from one of the networks that the router is connected to.

--no-route

Clear routes associated with the router

<router>

Virtual router name or ID

Example:

```
# vinfra service compute router set myrouter --disable-snat --external-gateway public
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Field	Value
external_gateway_info	enable_snat: false
	ip_addresses:
	- 10.94.129.76
	network_id: 720e45bc-4225-49de-9346-26513d8d1262
id	b9d8b000-5d06-4768-9f65-2715250cda53
name	myrouter
project_id	894696133031439f8aaa7e4868dcbd4d
routes	[]
status	ACTIVE

This command disables SNAT on the external gateway of the virtual router `myrouter`.

3.11.5 vinfra service compute router iface add

Add an interface to a virtual router:

```
usage: vinfra service compute router iface add [--ip-address <ip-address>]
      --interface <network> router
```

`--ip-address <ip-address>`

IP address

`--interface <network>`

Network name or ID

`router`

Virtual router name or ID

Example:

```
# vinfra service compute router iface add myrouter --interface private2 \
--ip-address 192.168.30.3
```

network_id	is_external	ip_addresses	status
720e45bc-4225-49de-9346-26513d8d1262	True	- 10.94.129.76	ACTIVE
e6f146ce-a6d0-48b2-9e4f-64a128ce97ae	False	- 192.168.128.1	ACTIVE
86803e07-a6d7-4809-9566-1cbe4a89adfd	False	- 192.168.30.3	DOWN

This command adds an interface from the virtual network `private2` to the virtual router `myrouter` with the IP address `192.168.30.3`.

3.11.6 vinfra service compute router iface list

List router interfaces:

```
usage: vinfra service compute router iface list router
```

router

Virtual router name or ID

Example:

```
# vinfra service compute router iface list myrouter
+-----+-----+-----+-----+
| network_id | is_external | ip_addresses | status |
+-----+-----+-----+-----+
| 720e45bc-4225-49de-9346-26513d8d1262 (public) | True | - 10.94.129.76 | ACTIVE |
| e6f146ce-a6d0-48b2-9e4f-64a128ce97ae (private) | False | - 192.168.128.1 | ACTIVE |
| 86803e07-a6d7-4809-9566-1cbe4a89adfd (private2) | False | - 192.168.30.3 | ACTIVE |
+-----+-----+-----+-----+
```

This command lists interfaces of the virtual router myrouter.

3.11.7 vinfra service compute router iface remove

Remove an interface from a virtual router:

```
usage: vinfra service compute router iface remove --interface <network> router
```

--interface <network>

Network name or ID

router

Virtual router name or ID

Example:

```
# vinfra service compute router iface remove myrouter --interface private2
+-----+-----+-----+-----+
| network_id | is_external | ip_addresses | status |
+-----+-----+-----+-----+
| 720e45bc-4225-49de-9346-26513d8d1262 | True | - 10.94.129.76 | ACTIVE |
| e6f146ce-a6d0-48b2-9e4f-64a128ce97ae | False | - 192.168.128.1 | ACTIVE |
+-----+-----+-----+-----+
```

This command removes the interface from the virtual network private2 from the virtual router myrouter.

3.11.8 vinfra service compute router delete

Delete a virtual router:

```
usage: vinfra service compute router delete <router>
```

<router>

Virtual router ID or name

Example:

```
# vinfra service compute router delete myrouter
Operation successful
```

This command deletes the virtual router `myrouter`.

3.12 Managing Floating IP Addresses

3.12.1 vinfra service compute floatingip create

Create a floating IP address:

```
usage: vinfra service compute floatingip create [--floating-ip-address <floating-ip-address>]
                                               [--port-id <port-id>]
                                               [--fixed-ip-address <fixed-ip-address>]
                                               [--description description] --network <network>
```

`--floating-ip-address <floating-ip-address>`

Floating IP address

`--port-id <port-id>`

ID of the port to be associated with the floating IP address. To learn the port ID of the selected server, use the command *vinfra service compute server iface list* (page 62).

`--fixed-ip-address <fixed-ip-address>`

Port IP address (required only if the port has multiple IP addresses)

`--description description`

Description of the floating IP address

`--network <network>`

ID or name of the network from which to allocate the floating IP

Example:

```
# vinfra service compute floatingip create 720e45bc-4225-49de-9346-26513d8d1262 \
--port-id 418c8c9e-aaa5-42f2-8da7-24bfead6f28b --fixed-ip-address 192.168.128.5
+-----+
| Field          | Value                               |
+-----+
| attached_to    | a172cb6a-1c7b-4157-9e86-035f3077646f |
| description    |                                       |
| fixed_ip_address | 192.168.128.5                       |
| floating_ip_address | 10.94.129.72                         |
| floating_network_id | 720e45bc-4225-49de-9346-26513d8d1262 |
| id             | a709f884-c43f-4a9a-a243-a340d7682ef8 |
| port_id        | 418c8c9e-aaa5-42f2-8da7-24bfead6f28b |
| project_id     | 894696133031439f8aaa7e4868dcbd4d     |
| router_id      | f7f86029-a553-4d61-b7ec-6f581d9c5f5f |
| status         | DOWN                                  |
+-----+
```

This command creates a floating IP address from the public network with the ID 720e45bc-4225-49de-9346-26513d8d1262 and assigns it to a server on port with the ID 418c8c9e-aaa5-42f2-8da7-24bfead6f28b and the private IP address 192.168.128.5.

3.12.2 vinfra service compute floatingip list

List floating IP addresses:

```
usage: vinfra service compute floatingip list
```

Example:

```
# vinfra service compute floatingip list -c id -c fixed_ip_address -c port_id \
-c floating_ip_address -c floating_network_id
+-----+
| id          | fixed_ip_address | port_id          | floating_ip_address | floating_network_id |
+-----+
| a709f884-...> | 192.168.128.5   | 418c8c9e-...> | 10.94.129.72       | 720e45bc-...>     |
+-----+
```

This command lists floating IP addresses used in the compute cluster. (The output is abridged to fit on page.)

3.12.3 vinfra service compute floatingip show

Display information about a floating IP address:

```
usage: vinfra service compute floatingip show <floatingip>
```

<floatingip>

ID of the floating IP address

Example:

```
# vinfra service compute floatingip show a709f884-c43f-4a9a-a243-a340d7682ef8
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| attached_to    | a172cb6a-1c7b-4157-9e86-035f3077646f |
| description    |                                         |
| fixed_ip_address | 192.168.128.5                           |
| floating_ip_address | 10.94.129.72                             |
| floating_network_id | 720e45bc-4225-49de-9346-26513d8d1262 |
| id             | a709f884-c43f-4a9a-a243-a340d7682ef8 |
| port_id        | 418c8c9e-aaa5-42f2-8da7-24bfead6f28b |
| project_id     | 894696133031439f8aaa7e4868dcbd4d       |
| router_id      | f7f86029-a553-4d61-b7ec-6f581d9c5f5f |
| status         | ACTIVE                                   |
+-----+-----+
```

This command shows the details of the floating IP address with the ID a709f884-c43f-4a9a-a243-a340d7682ef8.

3.12.4 vinfra service compute floatingip set

Modify parameters of a floating IP address:

```
usage: vinfra service compute floatingip set [--port-id <port-id>]
                                             [--fixed-ip-address <fixed-ip-address>]
                                             [--description <description>] <floatingip>
```

--port-id <port-id>

ID of the port to be associated with the floating IP address

--fixed-ip-address <fixed-ip-address>

Port IP address (required only if the port has multiple IP addresses)

--description <description>

Description of the floating IP address

<floatingip>

ID of the floating IP address

Example:

```
# vinfra service compute floatingip set a709f884-c43f-4a9a-a243-a340d7682ef8 \
--description "Floating IP for myvm"
+-----+
| Field          | Value                                     |
+-----+
| attached_to    | a172cb6a-1c7b-4157-9e86-035f3077646f    |
| description    | Floating IP for myvm                    |
| fixed_ip_address | 192.168.128.5                           |
| floating_ip_address | 10.94.129.72                             |
| floating_network_id | 720e45bc-4225-49de-9346-26513d8d1262    |
| id             | a709f884-c43f-4a9a-a243-a340d7682ef8    |
| port_id        | 418c8c9e-aaa5-42f2-8da7-24bfead6f28b    |
| project_id     | 894696133031439f8aaa7e4868dcbd4d        |
| router_id      | f7f86029-a553-4d61-b7ec-6f581d9c5f5f    |
| status         | ACTIVE                                   |
+-----+
```

This command adds a description for the floating IP address with the ID a709f884-c43f-4a9a-a243-a340d7682ef8.

3.12.5 vinfra service compute floatingip delete

Delete a floating IP address:

```
usage: vinfra service compute floatingip delete <floatingip>
```

<floatingip>

ID of the floating IP address

Example:

```
# vinfra service compute floatingip delete a709f884-c43f-4a9a-a243-a340d7682ef8
Operation successful
```

This command deletes the floating IP address with the ID a709f884-c43f-4a9a-a243-a340d7682ef8.

3.13 Managing Load Balancers

3.13.1 vinfra service compute load-balancer create

Create a load balancer:

```
usage: vinfra service compute load-balancer create [--description <description>]
                                                [--enable | --disable]
                                                [--address <address>]
                                                [--floating-ip <floating-ip>]
                                                [--pools-config <pools>]
                                                <name> <network>
```

--description <description>

Load balancer description

--enable

Enable the load balancer.

--disable

Disable the load balancer.

--address <address>

The IP address the load balancer will try to allocate in the network.

--floating-ip <floating-ip>

The floating IP address that will be used to connect to the load balancer from public networks.

--pools-config <pools>

Pool configuration file

Below is an example of a pool configuration file in the YAML format:

```
- backend_protocol: HTTPS
  backend_protocol_port: 443
  healthmonitor: {delay: 5, max_retries: 3, max_retries_down: 3, timeout: 5, type: PING,
                  url_path: /}
  lb_algorithm: ROUND_ROBIN
  members:
  - {address: 192.168.30.49, compute_server_id: 5fd82e2a-3fef-4171-bfa4-67daa99ae64f}
  - {address: 192.168.30.15, compute_server_id: 4f7262d0-2031-4772-9984-2f1066ac166b}
  name: pool1
  protocol: HTTPS
  protocol_port: 443
  sticky_session: False
```


<name>

Load balancer name

<network>

The ID or name of network the load balancer will operate in.

Example:

```
# vinfra service compute load-balancer create mylbaas private1 --floating-ip 10.94.129.70
+-----+
| Field      | Value                                |
+-----+
| address    | 192.168.30.230                      |
| amphorae   |                                       |
| created_at | 2019-11-18T12:59:08.243413          |
| description |                                       |
| enabled    | True                                  |
| floating_ip | 10.94.129.70                        |
| ha_enabled  |                                       |
| id         | 941bf637-2d55-40f0-92c0-e65d6567b468 |
| members_count | 0                                    |
| name       | mylbaas                              |
| network_id | 2b821d00-e428-4a76-b1ae-d181c9f5ae7f |
| pools     | []                                    |
| port_id    | 2d8ab88a-847c-4396-857e-11eaa80e1b24 |
| project_id | e4e059c67dee4736851df14d4519a5a5    |
| status     | CREATING                             |
| updated_at |                                       |
+-----+
```

This command creates a load balancer `mylbaas` without balancing pools that will operate in the network `private` with the floating IP address `10.94.129.70`.

3.13.2 vinfra service compute load-balancer list

List load balancers:

```
usage: vinfra service compute load-balancer list
```

Example:

```
# vinfra service compute load-balancer list
+-----+
| id              | name  |
+-----+
| 941bf637-2d55-40f0-92c0-e65d6567b468 | mylbaas |
+-----+
```

This command lists load balancers in the compute cluster.

3.13.3 vinfra service compute load-balancer show

Display load balancer details:

```
usage: vinfra service compute load-balancer show <load-balancer>
```

<load-balancer>

Load balancer ID or name

Example:

```
# vinfra service compute load-balancer show mylbaas
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| address        | 192.168.30.230                           |
| amphorae       | - active: true                           |
|                | compute_id: b0c4793f-e1b1-4251-91c2-94e34787f537 |
|                | created_at: '2019-11-18T12:59:12.742446'    |
|                | id: b7b23106-a87b-412d-9ce6-7c69b5594342     |
|                | image_id: 6d1ba6f9-cf86-4ea4-a32d-f138868a9742 |
|                | role: STANDALONE                           |
|                | status: ALLOCATED                           |
|                | updated_at: '2019-11-18T13:01:07.601184'    |
| created_at     | 2019-11-18T12:59:08.243413                |
| description    |                                             |
| enabled        | True                                       |
| floating_ip    | 10.94.129.70                              |
| ha_enabled     | False                                      |
| id             | 941bf637-2d55-40f0-92c0-e65d6567b468     |
| members_count  | 0                                          |
| name           | mylbaas                                    |
| network_id     | 2b821d00-e428-4a76-b1ae-d181c9f5ae7f     |
| pools         | []                                         |
| port_id       | 2d8ab88a-847c-4396-857e-11eaa80e1b24     |
| project_id     | e4e059c67dee4736851df14d4519a5a5        |
| status        | ACTIVE                                    |
| updated_at     | 2019-11-18T13:01:10.983144                |
+-----+-----+
```

This command shows the details of the load balancer `mylbaas`.

3.13.4 vinfra service compute load-balancer stats

Show statistics for a load balancer:

```
usage: vinfra service compute load-balancer stats <load-balancer>
```

<load-balancer>

Load balancer ID or name

Example:

```
# vinfra service compute load-balancer stats mylbaas
+-----+-----+-----+
| Field | Value |
+-----+-----+-----+
| stats | active_connections: 0 |
|       | bytes_in: 0 |
|       | bytes_out: 0 |
|       | listeners: null |
|       | loadbalancer_id: 17cfa86f-c374-4ca3-8cd6-f638a5234fe7 |
|       | request_errors: 0 |
|       | total_connections: 0 |
+-----+-----+-----+
```

This command shows statistics for the load balancer `mylbaas`.

3.13.5 vinfra service compute load-balancer set

Modify a load balancer:

```
usage: vinfra service compute load-balancer set [--description <description>]
        [--enable | --disable] [--name <name>]
        <load-balancer>
```

`--description <description>`

Load balancer description

`--enable`

Enable the load balancer.

`--disable`

Disable the load balancer.

`--name <name>`

Load balancer name

<load-balancer>

Load balancer ID or name

Example:

```
# vinfra service compute load-balancer set mylbaas --disable \
--description "Disabled load balancer"
+-----+-----+
| Field      | Value                                |
+-----+-----+
| address    | 192.168.30.230                       |
| amphorae   |                                        |
| created_at | 2019-11-18T12:59:08.243413           |
| description | Disabled load balancer               |
| enabled    | False                                 |
| floating_ip |                                        |
| ha_enabled |                                        |
| id         | 941bf637-2d55-40f0-92c0-e65d6567b468 |
| members_count | 0                                     |
| name       | mylbaas                               |
| network_id | 2b821d00-e428-4a76-b1ae-d181c9f5ae7f |
| pools      | []                                     |
| port_id    | 2d8ab88a-847c-4396-857e-11eaa80e1b24 |
| project_id | e4e059c67dee4736851df14d4519a5a5    |
| status     | DISABLED                             |
| updated_at | 2019-11-18T13:09:09.151442          |
+-----+-----+
```

This command disables the load balancer `mylbaas` and adds a description to it.

3.13.6 vinfra service compute load-balancer pool create

Create a load balancer pool:

```
usage: vinfra service compute load-balancer pool create --name <name>
--protocol {HTTP,HTTPS}
--port <port>
--algorithm <algorithm>
--backend-protocol {HTTP,HTTPS}
--backend-port <backend_port>
[--certificate-file <cert_file>]
[--connection-limit <limit>]
[--description <description>]
[--healthmonitor type=<type>,
url_path=<url>[,key=value,...]]
[--member address=<ip>
[,key=value,...]]
[--privatekey-file <key>]
[--enable-sticky-session |
--disable-sticky-session]
[--enable | --disable]
<load-balancer>
```

`--name <name>`
Pool name

`--protocol {HTTP,HTTPS}`
The protocol for incoming connections

`--port <port>`
The port for incoming connections

`--algorithm <algorithm>`
Load balancing algorithm (LEAST_CONNECTIONS, ROUND_ROBIN, or SOURCE_IP)

`--backend-protocol {HTTP,HTTPS}`
The protocol for destination connections

`--backend-port <backend_port>`
The port for destination connections

`--certificate-file <cert_file>`
An x.509 certificate file in the PEM format. Required for TLS-terminated HTTPS->HTTP load balancers.

`--connection-limit <limit>`
The maximum number of connections permitted for this pool. The default value is -1 (infinite connections).

`--description <description>`
Pool description

`--healthmonitor type=<type>,url_path=<url>[,key=value,...]`
Health monitor parameters:

- `type`: the health monitor type (HTTP, HTTPS, PING, or TCP)
- `url_path`: the URL path to the health monitor
- comma-separated `key=value` pairs with keys (optional):
 - `delay`: the time, in seconds, between sending probes to members.
 - `enabled`: declares whether the health monitor is enabled or not (true or false).
 - `max_retries`: the number of successful checks required to change member status to 'HEALTHY'. Ranges from 1 to 10.

- `max_retries_down`: the number of unsuccessful checks required to change member status to 'UNHEALTHY'. Ranges from 1 to 10.
- `timeout`: the maximum time, in seconds, that a monitor waits to connect before it times out. This value must be less than the `delay` value.

`--member address=<ip>[,key=value,...]`

Member parameters:

- `address`: an IPv4 address of the compute server
- comma-separated `key=value` pairs with keys (optional):
 - `enabled`: declares whether the member is enabled or not (true or false).
 - `weight`: determines the share of connections that the member services compared to the other pool members. For example, a weight of 10 means that the member handles five times as many connections than a member with a weight of 2. A weight of 0 means that the member does not receive new connections but continues to service existing ones. Ranges from 0 to 256. The default value is 1. This option can be used multiple times.

`--privatekey-file <key>`

A private TLS key file in the PEM format. Required for TLS-terminated HTTPS->HTTP load balancers.

`--enable-sticky-session`

Enable session persistence.

`--disable-sticky-session`

Disable session persistence.

`--enable`

Enable the pool.

`--disable`

Disable the pool.

`<load-balancer>`

Load balancer ID or name

Example:

```
# vinfra service compute load-balancer pool create mylbaas --name mypool --protocol HTTP \
--port 80 --backend-protocol HTTP --backend-port 80 --algorithm LEAST_CONNECTIONS \
--member address=192.168.31.153, compute_server_id=d51c10a7-6187-4a5a-a838-de5fc78a688a \
```

```

--member address=192.168.31.22,compute_server_id=54603109-8963-49f2-8c49-332537c57e90 \
--enable-sticky-session
+-----+-----+
| Field          | Value                               |
+-----+-----+
| backend_protocol | HTTP                                |
| backend_protocol_port | 80                                  |
| certificate      |                                     |
| connection_limit | -1                                  |
| created_at      | 2019-11-18T13:11:27.982129         |
| description     |                                     |
| enabled         | True                                 |
| healthmonitor   |                                     |
| id              | fa40e282-b29a-465a-afaa-2c702d2bde17 |
| lb_algorithm    | LEAST_CONNECTIONS                  |
| listener_id     | 66cc714e-af7f-40eb-9db8-67b8b6b6d23c |
| loadbalancer_id | 941bf637-2d55-40f0-92c0-e65d6567b468 |
| members         | []                                   |
| name           | mypool                              |
| private_key     |                                     |
| project_id      | e4e059c67dee4736851df14d4519a5a5   |
| protocol        | HTTP                                |
| protocol_port   | 80                                  |
| status          | CREATING                            |
| sticky_session  | True                                 |
| updated_at     |                                     |
+-----+-----+

```

This command adds a balancing pool `mypool` to the load balancer `mylbaas` with the following parameters:

- “HTTP on port 80 -> HTTP on port 80” forwarding rule
- the `LEAST_CONNECTIONS` balancing algorithm
- two members in the pool
- enabled sticky session

3.13.7 vinfra service compute load-balancer pool list

List load balancer pools:

```
usage: vinfra service compute load-balancer pool list
```

Example:

```

# vinfra service compute load-balancer pool list
+-----+-----+-----+-----+
| id              | name          | loadbalancer_id |
+-----+-----+-----+-----+

```

```
+-----+-----+-----+
| fa40e282-b29a-465a-afaa-2c702d2bde17 | mypool | 941bf637-2d55-40f0-92c0-e65d6567b468 |
+-----+-----+-----+
```

This command lists load balancer pools in the compute cluster.

3.13.8 vinfra service compute load-balancer pool show

Display load balancer pool details:

```
usage: vinfra service compute load-balancer pool show <pool>
```

<pool>

Load balancer pool ID or name

Example:

```
# vinfra service compute load-balancer pool show mypool
+-----+-----+-----+
| Field                | Value                                     |
+-----+-----+-----+
| backend_protocol     | HTTP                                     |
| backend_protocol_port | 80                                       |
| certificate           |                                           |
| connection_limit     | -1                                       |
| created_at           | 2019-11-18T13:11:27.982129              |
| description          |                                           |
| enabled              | True                                     |
| healthmonitor        |                                           |
| id                   | fa40e282-b29a-465a-afaa-2c702d2bde17   |
| lb_algorithm         | LEAST_CONNECTIONS                       |
| listener_id          | 66cc714e-af7f-40eb-9db8-67b8b6b6d23c   |
| loadbalancer_id     | 941bf637-2d55-40f0-92c0-e65d6567b468   |
| members              | - address: 192.168.31.153               |
|                      |   compute_server_id: d51c10a7-6187-4a5a-a838-de5fc78a688a |
|                      |   created_at: '2019-11-18T13:11:59.681101' |
|                      |   enabled: true                           |
|                      |   id: 3fd5dcc5-6e2c-4e22-8d0a-8e94e20a122f |
|                      |   name: ''                                |
|                      |   pool_id: null                           |
|                      |   status: HEALTHY                         |
|                      |   updated_at: '2019-11-18T13:12:01.467306' |
|                      |   weight: 1                               |
|                      | - address: 192.168.31.22                  |
|                      |   compute_server_id: 54603109-8963-49f2-8c49-332537c57e90 |
|                      |   created_at: '2019-11-18T13:12:10.176853' |
|                      |   enabled: true                           |
|                      |   id: ccb645b3-63c7-44f8-b861-b197c85506d4 |
|                      |   name: ''                                |
|                      |
```



```

|           | pool_id: null |
|           | status: HEALTHY |
|           | updated_at: '2019-11-18T13:12:12.281578' |
|           | weight: 1 |
| name      | mypool |
| private_key | |
| project_id | e4e059c67dee4736851df14d4519a5a5 |
| protocol   | HTTP |
| protocol_port | 80 |
| status     | ACTIVE |
| sticky_session | True |
| updated_at | 2019-11-18T13:12:12.305509 |
+-----+-----+

```

This command shows the details of the load balancer pool `mypool`.

3.13.9 `vinfra service compute load-balancer pool set`

Modify a load balancer pool:

```

usage: vinfra service compute load-balancer pool set --name <name> --protocol {HTTP,HTTPS}
                                                    --port <port> --algorithm <algorithm>
                                                    --backend-protocol {HTTP,HTTPS}
                                                    --backend-port <backend_port>
                                                    [--certificate-file <cert_file>]
                                                    [--connection-limit <limit>]
                                                    [--description <description>]
                                                    [--healthmonitor type=<type>,
url_path=<url>[,key=value,...]]
                                                    [--member address=<ip>,
compute_server_id=<id>
[,key=value,...]]
                                                    [--privatekey-file <key>]
                                                    [--enable-sticky-session |
--disable-sticky-session]
                                                    [--enable | --disable] <pool>

```

`--name <name>`

Pool name

`--protocol {HTTP,HTTPS}`

The protocol for incoming connections

`--port <port>`

The port for incoming connections

`--algorithm <algorithm>`

Load balancing algorithm (LEAST_CONNECTIONS, ROUND_ROBIN, or SOURCE_IP)

`--backend-protocol {HTTP,HTTPS}`

The protocol for destination connections

`--backend-port <backend_port>`

The port for destination connections

`--certificate-file <cert_file>`

An x.509 certificate file in the PEM format. Required for TLS-terminated HTTPS->HTTP load balancers.

`--connection-limit <limit>`

The maximum number of connections permitted for this pool. The default value is -1 (infinite connections).

`--description <description>`

Pool description

`--healthmonitor type=<type>,url_path=<url>[,key=value,...]`

Health monitor parameters:

- `type`: the health monitor type (HTTP, HTTPS, PING, or TCP)
- `url_path`: the URL path to the health monitor
- comma-separated `key=value` pairs with keys (optional):
 - `delay`: the time, in seconds, between sending probes to members.
 - `enabled`: declares whether the health monitor is enabled or not (true or false).
 - `max_retries`: the number of successful checks required to change member status to 'HEALTHY'. Ranges from 1 to 10.
 - `max_retries_down`: the number of unsuccessful checks required to change member status to 'UNHEALTHY'. Ranges from 1 to 10.
 - `timeout`: the maximum time, in seconds, that a monitor waits to connect before it times out. This value must be less than the `delay` value.

`--member address=<ip>,compute_server_id=<id>[,key=value,...]`

Member parameters:

- `address`: an IPv4 address of the compute server
- `compute_server_id`: the compute server UUID
- comma-separated `key=value` pairs with keys (optional):

- `enabled`: declares whether the member is enabled or not. Can be 'true' or 'false'.
- `weight`: determines the share of connections that the member services compared to the other pool members. For example, a weight of 10 means that the member handles five times as many connections than a member with a weight of 2. A weight of 0 means that the member does not receive new connections but continues to service existing ones. Ranges from 0 to 256. The default value is 1. This option can be used multiple times.

`--privatekey-file <key>`

A private TLS key file in the PEM format. Required for TLS-terminated HTTPS->HTTP load balancers.

`--enable-sticky-session`

Enable session persistence.

`--disable-sticky-session`

Disable session persistence.

`--enable`

Enable the pool.

`--disable`

Disable the pool.

`<pool>`

Load balancer pool ID or name

Example:

```
# vinfra service compute load-balancer pool set mypool --algorithm ROUND_ROBIN \
--member address=192.168.31.153,compute_server_id=d51c10a7-6187-4a5a-a838-de5fc78a688a \
--member address=192.168.31.22,compute_server_id=54603109-8963-49f2-8c49-332537c57e90 \
--member address=192.168.31.51,compute_server_id=5eb49a67-f0ff-44ec-b4e6-105e16e2db07 \
--disable-sticky-session
Operation accepted.
```

This command changes the parameters for the balancing pool `mypool` as follows:

- sets the balancing algorithm to `ROUND_ROBIN`
- adds the third member to the pool
- disables sticky session

3.13.10 vinfra service compute load-balancer pool delete

Delete a load balancer pool:

```
usage: vinfra service compute load-balancer pool delete <pool>
```

<pool>

Load balancer pool ID or name

Example:

```
# vinfra service compute load-balancer pool delete mypool
Operation successful.
```

This command removes the load balancer pool `mypool`.

3.13.11 vinfra service compute load-balancer delete

Delete a load balancer:

```
usage: vinfra service compute load-balancer delete <load-balancer>
```

<load-balancer>

Load balancer ID or name

Example:

```
# vinfra service compute load-balancer delete mylbaas
Operation accepted.
```

This command deletes the load balancer `mylbaas`.

3.14 Managing Volumes

3.14.1 vinfra service compute volume create

Create a new compute volume:

```
usage: vinfra service compute volume create [--description <description>]
                                           [--network-install <network_install>]
                                           [--image <image>] [--snapshot <snapshot>]
```

```
--storage-policy <storage_policy>
--size <size-gb> <volume-name>
```

```
--description <description>
```

Volume description

```
--network-install <network_install>
```

Perform network installation (true or false).

```
--image <image>
```

Source compute image ID or name

```
--snapshot <snapshot>
```

Source compute volume snapshot ID or name

```
--storage-policy <storage_policy>
```

Storage policy ID or name

```
--size <size-gb>
```

Volume size, in gigabytes

```
<volume-name>
```

Volume name

Example:

```
# vinfra service compute volume create myvolume --storage-policy default --size 8
+-----+-----+
| Field | Value |
+-----+-----+
| attachments | [] |
| availability_zone | nova |
| bootable | False |
| consistencygroup_id | |
| created_at | 2018-09-12T12:30:12.665916 |
| description | |
| encrypted | False |
| id | c9c0e9e7-ce7a-4566-99d5-d7e40f2987ab |
| imageRef | |
| migration_status | |
| multiattach | False |
| name | myvolume |
| network_install | False |
| os-vol-host-attr:host | |
| os-vol-mig-status-attr:migstat | |
| os-vol-mig-status-attr:name_id | |
| project_id | 72a5db3a033c403a86756021e601ef34 |
| replication_status | |
```

```

| size | 8 |
| snapshot_id | |
| source_volid | |
| status | creating |
| storage_policy_name | default |
| updated_at | |
| user_id | 98bf389983c24c07af9677b931783143 |
| volume_image_metadata | |
+-----+-----+

```

This command creates a volume `myvolume` sized 8 GB and chooses the default storage policy for it.

3.14.2 vinfra service compute volume list

List compute volumes:

```
usage: vinfra service compute volume list
```

Example:

```

# vinfra service compute volume list -c id -c name -c size -c status
+-----+-----+-----+-----+
| id | name | size | status |
+-----+-----+-----+-----+
| c9c0e9e7-ce7a-4566-99d5-d7e40f2987ab | myvolume | 8 | available |
+-----+-----+-----+-----+

```

This command lists volumes available to the compute cluster. (The output is abridged to fit on page.)

3.14.3 vinfra service compute volume show

Display compute volume details:

```
usage: vinfra service compute volume show <volume>
```

<volume>

Volume ID or name

Example:

```

# vinfra service compute volume show myvolume
+-----+-----+
| Field | Value |
+-----+-----+
| attachments | [] |
| availability_zone | nova |
+-----+-----+

```

```

| bootable | False |
| consistencygroup_id | |
| created_at | 2018-09-12T12:30:12.665916 |
| description | |
| encrypted | False |
| id | c9c0e9e7-ce7a-4566-99d5-d7e40f2987ab |
| imageRef | |
| migration_status | |
| multiattach | False |
| name | myvolume |
| network_install | False |
| os-vol-host-attr:host | stor-1.example.com.vstoragedomain@vstorage#vstorage |
| os-vol-mig-status-attr:migstat | |
| os-vol-mig-status-attr:name_id | |
| project_id | 72a5db3a033c403a86756021e601ef34 |
| replication_status | |
| size | 8 |
| snapshot_id | |
| source_volid | |
| status | available |
| storage_policy_name | default |
| updated_at | 2018-09-12T12:30:33.167654 |
| user_id | 98bf389983c24c07af9677b931783143 |
| volume_image_metadata | |
+-----+-----+

```

This command shows the details for the volume `myvolume`.

3.14.4 vinfra service compute volume set

Modify volume parameters:

```

usage: vinfra service compute volume set [--description <description>]
                                         [--network-install <network_install>]
                                         [--storage-policy <storage_policy>]
                                         [--bootable <bootable>]
                                         [--name <name>] <volume>

```

`--description <description>`

Volume description

`--network-install <network_install>`

Perform network install (true or false)

`--storage-policy <storage_policy>`

Storage policy ID or name

--bootable <bootable>

Make bootable (true or false)

--name <name>

A new name for the volume

<volume>

Volume ID or name

Example:

```
# vinfra service compute volume set myvolume --storage-policy mystorpolicy
+-----+-----+
| Field | Value |
+-----+-----+
| attachments | [] |
| availability_zone | nova |
| bootable | False |
| consistencygroup_id | |
| created_at | 2018-09-12T12:30:12.665916 |
| description | |
| encrypted | False |
| id | c9c0e9e7-ce7a-4566-99d5-d7e40f2987ab |
| imageRef | |
| migration_status | |
| multiattach | False |
| name | myvolume |
| network_install | False |
| os-vol-host-attr:host | stor-1.example.com.vstagedomain@vstorage#vstorage |
| os-vol-mig-status-attr:migstat | |
| os-vol-mig-status-attr:name_id | |
| project_id | 72a5db3a033c403a86756021e601ef34 |
| replication_status | |
| size | 8 |
| snapshot_id | |
| source_volid | |
| status | available |
| storage_policy_name | mystorpolicy |
| updated_at | 2018-09-12T12:55:29.298717 |
| user_id | 98bf389983c24c07af9677b931783143 |
| volume_image_metadata | |
+-----+-----+
```

This command changes the storage policy of the volume myvolume to mystorpolicy.

3.14.5 vinfra service compute volume extend

Extend a compute volume:

```
usage: vinfra service compute volume extend --size <size_gb> <volume>
```

<volume>

Volume ID or name

Example:

```
# vinfra service compute volume extend myvolume --size 16
Operation successful
```

This command extends the volume `myvolume` to 16 GB.

3.14.6 vinfra service compute volume delete

Delete a compute volume:

```
usage: vinfra service compute volume delete <volume>
```

<volume>

Volume ID or name

Example:

```
# vinfra service compute volume delete myvolume2
Operation successful
```

This command deletes the volume `myvolume2`.

3.15 Managing Volume Snapshots

3.15.1 vinfra service compute volume snapshot create

Create a snapshot of a volume:

```
usage: vinfra service compute volume snapshot create [--description <description>]
           --volume <volume>
           <volume-snapshot-name>
```

```
--description <description>
    Volume snapshot description

--volume <volume>
    Volume ID or name

<volume-snapshot-name>
    Volume snapshot name
```

Example:

```
# vinfra service compute volume snapshot create mysnapshot --volume myvolume
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| created_at | 2019-04-30T13:12:54.297629+00:00        |
| description |                                           |
| id         | 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71    |
| metadata   | {}                                        |
| name       | mysnapshot                              |
| project_id | fd0ae61496d04ef6bb637bc3167b7eaf       |
| size       | 8                                        |
| status     | creating                                 |
| volume_id  | 92dc3bd7-713d-42bf-83cd-4de40c24fed9    |
+-----+-----+
```

This command initiates creation of a snapshot `mysnapshot` of the volume `myvolume`.

3.15.2 vinfra service compute volume snapshot list

List volume snapshots:

```
usage: vinfra service compute volume snapshot list
```

Example:

```
# vinfra service compute volume snapshot list -c id -c name -c size -c status
+-----+-----+-----+-----+
| id           | name           | status   |         |
+-----+-----+-----+-----+
| 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71 | mysnapshot    | available |         |
+-----+-----+-----+-----+
```

This command lists volume snapshots available to the compute cluster. (The output is abridged to fit on page.)

3.15.3 vinfra service compute volume snapshot show

Display details of a volume snapshot:

```
usage: vinfra service compute volume snapshot show <volume-snapshot>
```

<volume-snapshot>

Volume snapshot ID or name

Example:

```
# vinfra service compute volume snapshot show mysnapshot
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| created_at | 2019-04-30T13:12:54.297629+00:00        |
| description |                                           |
| id         | 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71    |
| metadata   | {}                                        |
| name       | mysnapshot                              |
| project_id | fd0ae61496d04ef6bb637bc3167b7eaf       |
| size       | 8                                        |
| status     | available                                |
| volume_id  | 92dc3bd7-713d-42bf-83cd-4de40c24fed9    |
+-----+-----+
```

This command shows the details for the volume snapshot `mysnapshot`.

3.15.4 vinfra service compute volume snapshot set

Modify volume snapshot parameters:

```
usage: vinfra service compute volume snapshot set [--description <description>]
                                                [--name <name>] <volume-snapshot>
```

--description <description>

Volume snapshot description

--name <name>

A new name for the volume snapshot

<volume-snapshot>

Volume snapshot ID or name

Example:

```
# vinfra service compute volume snapshot set mysnapshot --name mynewsnapshot
+-----+-----+
| Field      | Value                                |
+-----+-----+
| created_at | 2019-04-30T13:12:54.297629+00:00    |
| description |                                       |
| id         | 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71 |
| metadata   | {}                                    |
| name       | mynewsnapshot                        |
| project_id | fd0ae61496d04ef6bb637bc3167b7eaf    |
| size       | 8                                     |
| status     | available                             |
| volume_id  | 92dc3bd7-713d-42bf-83cd-4de40c24fed9 |
+-----+-----+
```

This command changes the name of the volume snapshot `mysnapshot` to `mynewsnapshot`.

3.15.5 vinfra service compute volume snapshot upload-to-image

Create a compute image from a compute volume snapshot:

```
usage: vinfra service compute volume snapshot upload-to-image [--name <name>]
                                     <volume-snapshot>
```

`--name <name>`

Image name

`<volume-snapshot>`

Volume snapshot ID or name

Example:

```
# vinfra service compute volume snapshot upload-to-image --name myvm-image \
mynewsnapshot
+-----+-----+
| Field      | Value                                |
+-----+-----+
| checksum   |                                       |
| container_format | bare                                |
| created_at |                                       |
| disk_format | qcow2                               |
| id         | 6a7a78c1-7168-4387-9b55-23fd477fdaa0 |
| min_disk   |                                       |
| min_ram    |                                       |
| name       | myvm-image                          |
| os_distro  | linux                                |
| os_type    | linux                                |
| project_id |                                       |
+-----+-----+
```

```

| protected      | False      |
| public        | False      |
| size          | 1          |
| status        | uploading  |
| tags          |            |
| updated_at    | 2019-06-07T12:30:43.462707 |
| virtual_size  |            |
+-----+-----+

```

This command creates the compute image `myvm-image` from the volume snapshot `mynewsnapshot`.

3.15.6 vinfra service compute volume snapshot revert

Revert a volume to the specified snapshot:

```
usage: vinfra service compute volume snapshot revert <volume-snapshot>
```

<volume-snapshot>

Volume snapshot ID or name

Example:

```

# vinfra service compute volume snapshot revert mynewsnapshot
+-----+-----+
| Field      | Value      |
+-----+-----+
| created_at | 2019-04-30T13:12:54.297629+00:00 |
| description |           |
| id         | 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71 |
| metadata   | {}         |
| name       | mynewsnapshot |
| project_id | fd0ae61496d04ef6bb637bc3167b7eaf |
| size       | 8          |
| status     | available  |
| volume_id  | 92dc3bd7-713d-42bf-83cd-4de40c24fed9 |
+-----+-----+

```

This command reverts the volume to its snapshot `mynewsnapshot`.

3.15.7 vinfra service compute volume snapshot reset-state

Reset a volume snapshot stuck in the “Error” state or one of transitional states to the “Available” state:

```
usage: vinfra service compute volume snapshot reset-state <volume-snapshot>
```

<volume-snapshot>

Volume snapshot ID or name

Example:

```
# vinfra service compute volume snapshot reset-state mynewsnapshot
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| created_at | 2019-04-30T13:12:54.297629+00:00       |
| description |                                           |
| id         | 3fdfe5d6-8bd2-4bf5-8599-a9cef50e5b71   |
| metadata   | {}                                       |
| name       | mynewsnapshot                          |
| project_id | fd0ae61496d04ef6bb637bc3167b7eaf      |
| size       | 8                                        |
| status     | available                               |
| volume_id  | 92dc3bd7-713d-42bf-83cd-4de40c24fed9   |
+-----+-----+
```

This command resets the state of the volume snapshot mynewsnapshot.

3.15.8 vinfra service compute volume snapshot delete

Delete a volume snapshot:

```
usage: vinfra service compute volume snapshot delete <volume-snapshot>
```

<volume-snapshot>

Volume snapshot ID or name

Example:

```
# vinfra service compute volume snapshot delete mynewsnapshot
Operation successful
```

This command deletes the volume snapshot mynewsnapshot.

3.16 Managing Storage Policies

You can manage storage policies only after creating the compute cluster.

3.16.1 vinfra cluster storage-policy create

Create a new storage policy:

```
usage: vinfra cluster storage-policy create --tier {0,1,2,3}
      (--replicas <norm>[:<min>] |
      --encoding <M>+<N>) --failure-domain
      {disk,host,rack,row,room} <name>
```

--tier {0,1,2,3}

Storage tier

--replicas <norm>[:<min>]

Storage replication mapping in the format:

- norm: the number of replicas to maintain
- min: the minimum required number of replicas (optional)

--encoding <M>+<N>

Storage erasure encoding mapping in the format:

- M: the number of data blocks
- N: the number of parity blocks

--failure-domain {disk,host,rack,row,room}

Storage failure domain

<name>

Storage policy name

Example:

```
# vinfra cluster storage-policy create mystorpolicy --tier 3 \
--encoding 3+2 --failure-domain host
```

```
+-----+
| Field          | Value                               |
+-----+
| failure_domain | host                                |
| id             | 2199e71e-ce8a-4ba9-81cd-75502f0344ca |
| name          | mystorpolicy                        |
| redundancy    | encoding=3+2                        |
| tier           | 3                                    |
+-----+
+-----+
+-----+
```

This command creates a storage policy `mystorpolicy` with the tier set to 3, redundancy scheme to erasure coding 3+2, and failure domain set to host.

3.16.2 vinfra cluster storage-policy list

List existing storage policies:

```
usage: vinfra cluster storage-policy list
```

Example:

```
# vinfra cluster storage-policy list
+-----+-----+-----+-----+-----+
| id          | name          | tier | redundancy | failure_domain |
+-----+-----+-----+-----+-----+
| 2199e71e-... | mystorpolicy | 3   | encoding=3+2 | host           |
| 4274d6fd-... | default       | 0   | replicas=3   | host           |
+-----+-----+-----+-----+-----+
```

This command lists storage policies available to the compute cluster.

3.16.3 vinfra cluster storage-policy show

Show details of a storage policy:

```
usage: vinfra cluster storage-policy show <storage-policy>
```

<storage-policy>

Storage policy ID or name

Example:

```
# vinfra cluster storage-policy show mystorpolicy
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| failure_domain | host                                     |
| id             | 2199e71e-ce8a-4ba9-81cd-75502f0344ca    |
| name           | mystorpolicy                             |
| redundancy     | encoding=3+2                             |
| tier            | 3                                         |
+-----+-----+
```

This command shows the details of the storage policy `mystorpolicy`.

3.16.4 vinfra cluster storage-policy set

Modify storage policy parameters:

```
usage: vinfra cluster storage-policy set [--name <name>] [--tier {0,1,2,3}]
                                         [--replicas <norm>[:<min>] |
                                         --encoding <M>+<N>] [--failure-domain
                                         {disk,host,rack,row,room}] <storage-policy>
```

`--name <name>`

A new name for the storage policy

`--tier {0,1,2,3}`

Storage tier

`--replicas <norm>[:<min>]`

Storage replication mapping in the format:

- `norm`: the number of replicas to maintain
- `min`: the minimum required number of replicas (optional)

`--encoding <M>+<N>`

Storage erasure encoding mapping in the format:

- `M`: the number of data blocks
- `N`: the number of parity blocks

`--failure-domain {disk,host,rack,row,room}`

Storage failure domain

`<storage-policy>`

Storage policy ID or name

Example:

```
# vinfra cluster storage-policy set mystorpolicy --encoding 5+2
+-----+-----+
| Field          | Value                               |
+-----+-----+
| failure_domain | host                                 |
| id             | 2199e71e-ce8a-4ba9-81cd-75502f0344ca |
| name          | mystorpolicy                        |
| redundancy     | encoding=5+2                        |
| tier           | 3                                    |
+-----+-----+
```

This command changes the redundancy type for the storage policy `mystorpolicy` from erasure coding 3+2 to 5+2.

3.16.5 `vinfra cluster storage-policy delete`

The default policy cannot be deleted.

Remove an existing storage policy:

```
usage: vinfra cluster storage-policy delete <storage-policy>
```

<storage-policy>

Storage policy ID or name

Example:

```
# vinfra cluster storage-policy delete mystorpolicy
Operation successful
```

This command deletes the storage policy `mystorpolicy`.

3.17 Managing Kubernetes Clusters

3.17.1 `vinfra service compute k8saas create`

Create a new Kubernetes cluster:

```
usage: vinfra service compute k8saas create [--master-node-count <count>]
      [--node-count <count>]
      [--volume-storage-policy <policy>]
      [--kubernetes-version <version>]
      --master-flavor <flavor> --flavor
      <flavor> [--volume-size <size>]
      --external-network <network>
      [--network <network>] --key-name
      <key-name>
      [--use-floating-ip <use-floating-ip>]
      <name>
```

<name>

Kubernetes cluster name

`--master-node-count <count>`
The amount of master nodes in the Kubernetes cluster

`--node-count <count>`
The amount of worker nodes in the Kubernetes cluster

`--volume-storage-policy <policy>`
The ID or name of the storage policy for the volume where containers will reside.

`--kubernetes-version <version>`
Kubernetes version

`--master-flavor <flavor>`
The flavor to use for Kubernetes master nodes.

`--flavor <flavor>`
The flavor to use for Kubernetes worker nodes.

`--volume-size <size>`
The size of the storage volume on each Kubernetes node

`--external-network <network>`
The ID or name of a public network that will provide Internet access to Kubernetes nodes.

`--network <network>`
The ID or name of a private network that will provide networking between Kubernetes nodes.

`--key-name <key-name>`
The key pair to use for accessing the Kubernetes nodes.

`--use-floating-ip <use-floating-ip>`
Assign floating IP addresses to master and worker nodes ('true' or 'false').

The prerequisites for creating a Kubernetes cluster are:

- The Kubernetes-as-a-service component. It can be deployed along with the compute cluster or later (see [Creating the Compute Cluster](#) or [Managing Add-On Services](#)).
- A private network that will interconnect Kubernetes nodes. It needs to have a gateway and a DNS server specified.
- A public network with Internet access (for reaching the [etcd discovery service](#)) and the traffic types **Compute API** and **VM public** assigned to it.

- An SSH key that will be installed on both the master and worker nodes.
- Enough resources for all of the Kubernetes nodes, taking their flavors into account.

Example:

```
# vinfra service compute k8saas create --master-node-count 1 --node-count 3 \
--volume-storage-policy default --kubernetes-version v1.15.6 --master-flavor medium \
--flavor small --volume-size 10 --external-network public1 --network private1 \
--key-name key1 --use-floating-ip true k8s1 --vinfra-username user1 \
--vinfra-password password --vinfra-domain domain1 --vinfra-project project1
```

Field	Value
boot_volume_size	10
boot_volume_storage_policy	default
containers_volume_size	10
containers_volume_storage_policy	default
create_timeout	60
external_network_id	7006065f-9067-4aed-b888-d89baa7004b8
id	c0754d99-6066-4675-8062-e62602939cf3
key_name	key1
master_flavor	medium
master_node_count	1
name	k8s1
network_id	d037623b-0db7-40c2-b38a-9ac34fbd1cc5
project_id	c734b9832e9540bd8f79bc2272c167e6
status	CREATING
user_id	c2cba773dc824125b07720744d0e49e2
worker_pools	- flavor: small node_count: 3

This command, run as user1 from domain1 > project1, starts creation of the Kubernetes cluster k8s1 with these parameters:

- Kubernetes version 1.15.6
- 1 master node based on the medium flavor and 3 worker nodes based on the small flavor
- 10 GB storage volumes covered by the default storage policy
- private network private1 that will connect to the Internet via the public network public1
- floating IP addresses for each node picked from the specified public network
- public SSH key key1

3.17.2 vinfra service compute k8saas list

List Kubernetes clusters:

```
usage: vinfra service compute k8saas list
```

Example:

```
# vinfra service compute k8saas list
+-----+-----+-----+
| id                | name | status |
+-----+-----+-----+
| f3e71ee8-8583-4b6a-abce-0132818f5108 | k8s1 | ACTIVE |
+-----+-----+-----+
```

This command displays the list of Kubernetes clusters.

3.17.3 vinfra service compute k8saas config

Print Kubernetes cluster configuration (must be run as the user who created the Kubernetes cluster):

```
usage: vinfra service compute k8saas config <cluster>
```

cluster

Cluster ID or name

Example:

```
# vinfra service compute k8saas config k8s1 --vinfra-domain domain1 \
--vinfra-project project1 --vinfra-username user1 --vinfra-password password \
> kubeconfig
```

This command prints the configuration of the Kubernetes cluster k8s1 to the file kubeconfig.

3.17.4 vinfra service compute k8saas show

Display Kubernetes cluster details:

```
usage: vinfra service compute k8saas show <cluster>
```

cluster

Cluster ID or name

Example:

```
# vinfra service compute k8saas show k8s1
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| boot_volume_size | 10                                       |
| boot_volume_storage_policy | default                                 |
| containers_volume_size | 10                                       |
| containers_volume_storage_policy | default                                 |
| create_timeout   | 60                                       |
| external_network_id | 7006065f-9067-4aed-b888-d89baa7004b8 |
| id               | c0754d99-6066-4675-8062-e62602939cf3 |
| key_name         | key1                                     |
| master_flavor    | medium                                  |
| master_node_count | 1                                       |
| name             | k8s1                                     |
| network_id       | d037623b-0db7-40c2-b38a-9ac34fbd1cc5 |
| project_id       | c734b9832e9540bd8f79bc2272c167e6     |
| stack_id         | 3ef9ec9d-fde4-4358-bdb6-91205cd8ca52 |
| status           | ACTIVE                                  |
| user_id          | c2cba773dc824125b07720744d0e49e2     |
| version          | v1.15.6                                 |
| worker_pools     | - flavor: small                         |
|                  |   node_count: 3                         |
+-----+-----+
```

This command displays the details of the Kubernetes cluster k8s1.

3.17.5 vinfra service compute k8saas set

Modify Kubernetes cluster parameters (must be run as the user who created the Kubernetes cluster):

```
usage: vinfra service compute k8saas set [--node-count <count>]
                                           [--volume-storage-policy <policy>]
                                           <cluster>
```

cluster

Cluster ID or name

--node-count <count>

The amount of worker nodes in the Kubernetes cluster

--volume-storage-policy <policy>

The ID or name of the storage policy for the volume where containers will reside.

Example:

```
# vinfra service compute k8saas set --node-count 5 k8s1 \
--vinfra-domain domain1 --vinfra-project project1 \
--vinfra-username user1 --vinfra-password password
```

Field	Value
boot_volume_size	10
boot_volume_storage_policy	default
containers_volume_size	10
containers_volume_storage_policy	default
create_timeout	60
external_network_id	7006065f-9067-4aed-b888-d89baa7004b8
id	c0754d99-6066-4675-8062-e62602939cf3
key_name	key1
master_flavor	medium
master_node_count	1
name	k8s1
network_id	d037623b-0db7-40c2-b38a-9ac34fbd1cc5
project_id	c734b9832e9540bd8f79bc2272c167e6
stack_id	3ef9ec9d-fde4-4358-bdb6-91205cd8ca52
status	ACTIVE
user_id	c2cba773dc824125b07720744d0e49e2
version	v1.15.6
worker_pools	- flavor: small node_count: 3

This command starts configuring the Kubernetes cluster k8s1 to bring the number of worker nodes to 5.

3.17.6 vinfra service compute k8saas delete

Delete a Kubernetes cluster:

```
usage: vinfra service compute k8saas delete <cluster>
```

cluster

Cluster ID or name

Example:

```
# vinfra service compute k8saas delete k8s1
Operation accepted.
```

This command deletes the Kubernetes cluster k8s1.

3.18 Managing Compute Quotas

3.18.1 vinfra service compute quotas show

List compute quotas:

```
usage: vinfra service compute quotas show [--usage] <project_id>
```

`--usage`

Include quota usage.

`<project_id>`

Project ID

Example:

```
# vinfra service compute quotas show 6ef6f48f01b640ccb8ff53117b830fa3 --usage
+-----+-----+
| Field                               | Value |
+-----+-----+
| compute.cores.limit                 | 20    |
| compute.cores.used                   | 2     |
| compute.ram.limit                   | 40960 |
| compute.ram.used                     | 4096  |
| k8saas.cluster.limit                | 10    |
| k8saas.cluster.used                  | 0     |
| lbaas.loadbalancer.limit             | 10    |
| lbaas.loadbalancer.used              | 0     |
| network.floatingip.limit            | 10    |
| network.floatingip.used              | 0     |
| storage.gigabytes.default.limit     | 1024  |
| storage.gigabytes.default.used      | 66    |
+-----+-----+
```

This command shows compute quotas with their usage for the project with the ID 6ef6f48f01b640ccb8ff53117b830fa3.

3.18.2 vinfra service compute quotas update

Update compute quotas:

```
usage: vinfra service compute quotas update [--cores <cores>] [--ram <ram>]
      [--floatingip <floatingip>]
      [--gigabytes <storage_policy>:<size>]
```



```
[--k8saas-cluster <cluster>]
<project_id>
```

```
--cores <cores>
    Number of cores

--ram <ram>
    Number of RAM, in megabytes

--floatingip <floatingip>
    Number of floating IP addresses

--gigabytes <storage_policy>:<size>
    Comma-separated list of <storage_policy>:<size>

--k8saas-cluster <cluster>
    Number of Kubernetes clusters

<project_id>
    Project ID
```

Example:

```
# vinfra service compute quotas update 6ef6f48f01b640ccb8ff53117b830fa3 --cores 10 \
--ram 20480 --gigabytes default:512
Operation successful.
```

This command updates compute quotas to 10 vCPUs, 20 GB of RAM, and 512 GB of disk space for the default storage policy.

3.19 Managing Updates

3.19.1 vinfra software-updates check-for-updates

Check for software updates:

```
usage: vinfra software-updates check-for-updates
```

Example:

```
# vinfra software-updates check-for-updates
+-----+-----+
| Field | Value |
```

```

+-----+-----+
| available_storage_release | | |
| last_check_datetime      | |
| nodes                    | - current_storage_release: |
|                          |   release: '756'           |
|                          |   version: 3.5.0          |
|                          |   downloaded_storage_release: null |
|                          |   host: node1.vstoragedomain |
|                          |   id: 51cc14d4-eec6-433e-a7b1-e1c5c7f9555e |
|                          |   orig_hostname: node1     |
|                          |   status: download_completed |
| services                 | []                          |
| status                   | uptodate                    |
| tasks                    | |                            |
+-----+-----+

```

3.19.2 vinfra software-updates eligibility-check

Check nodes' update eligibility:

```
usage: vinfra software-updates eligibility-check
```

Example:

```

# vinfra software-updates eligibility-check
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 0143aec7-f9ce-4654-ad48-edb6f4104e22 |
+-----+-----+

```

This command creates a task to check whether the nodes in the storage cluster are eligible for updates.

Task outcome:

```

# vinfra task show 0143aec7-f9ce-4654-ad48-edb6f4104e22
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.software_updates.tasks.EligibilityCheckTask |
| result | cluster_has_releasing_nodes: |
| | details: null |
| | exception: null |
| | message: null |
| | passed: true |
| | severity: critical |
| | cluster_unhealthy: |
| | details: null |
+-----+-----+

```

```

|         | exception: null
|         | message: null
|         | passed: true
|         | severity: critical
|         | not_enough_space_on_agents:
|         |   details: null
|         |   exception: null
|         |   message: null
|         |   passed: true
|         |   severity: critical
|         | not_enough_space_on_mn:
|         |   details: null
|         |   exception: null
|         |   message: null
|         |   passed: true
|         |   severity: critical
|         | postgres_not_running:
|         |   details: null
|         |   exception: null
|         |   message: null
|         |   passed: true
|         |   severity: critical
| state   | success
| task_id | 0143aec7-f9ce-4654-ad48-edb6f4104e22
+-----+

```

3.19.3 vinfra software-updates download

Download software updates:

```
usage: vinfra software-updates download
```

Example:

```

# vinfra software-updates download
+-----+
| Field | Value
+-----+
| task_id | 2f930030-22de-4ce5-bf00-05328ee672f0
+-----+

```

This command creates a task to download updates.

Task outcome:

```

# vinfra task show 2f930030-22de-4ce5-bf00-05328ee672f0
+-----+
| Field | Value
+-----+

```

```

| details |
| name    | backend.business.models.software_updates.tasks.DownloadSoftwareUpdatesTask |
| result  |
| state   | success
| task_id | 2f930030-22de-4ce5-bf00-05328ee672f0
+-----+

```

3.19.4 vinfra software-updates start

Start the software update procedure:

```
usage: vinfra software-updates start [--skip] [--force]
```

`--skip`

Skip and do not update nodes that cannot enter maintenance mode.

`--force`

Forcibly update and reboot (if needed) all nodes even if they cannot enter maintenance mode. Using this option may result in downtime.

If both parameters are omitted, the update will stop if a node cannot enter maintenance mode. Nodes that have already been updated will remain so.

Example:

```

# vinfra software-updates start
+-----+
| Field  | Value
+-----+
| task_id | 0eae9159-7595-42a7-8feb-d04df3e295c7 |
+-----+

```

This command creates a task to start updating.

Task outcome:

```

# vinfra task show 0eae9159-7595-42a7-8feb-d04df3e295c7
+-----+
| Field  | Value
+-----+
| details |
| name    | backend.business.models.software_updates.tasks.StartSoftwareUpdateTask |
| result  |
| state   | running
| task_id | 0eae9159-7595-42a7-8feb-d04df3e295c7
+-----+

```

3.19.5 vinfra software-updates pause

Show software updates status:

```
usage: vinfra software-updates pause
```

Example:

```
# vinfra software-updates pause
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | b02a686b-3214-447e-a9b4-43698aa9388b |
+-----+-----+
```

This command creates a task to pause updates.

Task outcome:

```
# vinfra task show b02a686b-3214-447e-a9b4-43698aa9388b
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                       |
| name    | backend.presentation.software_updates.tasks.PauseSoftwareUpdateTask |
| result  |                                       |
| state   | success                              |
| task_id | b02a686b-3214-447e-a9b4-43698aa9388b |
+-----+-----+
```

3.19.6 vinfra software-updates resume

Resume the software update procedure:

```
usage: vinfra software-updates resume [--skip] [--force]
```

`--skip`

Skip and do not update nodes that cannot enter maintenance mode.

`--force`

Forcibly update and reboot (if needed) all nodes even if they cannot enter maintenance mode. Using this option may result in downtime.

If both parameters are omitted, the update will stop if a node cannot enter maintenance mode. Nodes that have already been updated will remain so.

Example:

```
# vinfra software-updates resume
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 35323989-bdbc-4826-94c3-70ed7d06969d |
+-----+-----+
```

This command creates a task to resume the update.

Task outcome:

```
# vinfra task show 35323989-bdbc-4826-94c3-70ed7d06969d
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.software_updates.tasks.ResumeSoftwareUpdateTask |
| result | |
| state | success |
| task_id | 35323989-bdbc-4826-94c3-70ed7d06969d |
+-----+-----+
```

3.19.7 vinfra software-updates cancel

Cancel software updates:

```
usage: vinfra software-updates cancel
```

Example:

```
# vinfra software-updates cancel
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 7aeb20ba-1f9f-4f28-9790-086428d3e18e |
+-----+-----+
```

This command creates a task to cancel the update.

Task outcome:

```
# vinfra task show 7aeb20ba-1f9f-4f28-9790-086428d3e18e
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.software_updates.tasks.CancelSoftwareUpdateTask |
| result | |
| state | success |
+-----+-----+
```

```
| task_id | 7aeb20ba-1f9f-4f28-9790-086428d3e18e |
+-----+-----+
```

3.19.8 vinfra software-updates status

Check software update status:

```
usage: vinfra software-updates status
```

Example:

```
# vinfra software-updates status
+-----+-----+
| Field                | Value                |
+-----+-----+
| available_storage_release | release: '758'      |
|                        | version: 3.5.0      |
| last_check_datetime   | 2019-12-17T13:25:41.991763+00:00 |
| nodes                 | - current_storage_release: |
|                        |   release: '758'     |
|                        |   version: 3.5.0     |
|                        | downloaded_storage_release: null |
|                        | host: man-hci7-1.vstoragedomain |
|                        | id: 51cc14d4-eec6-433e-a7b1-e1c5c7f9555e |
|                        | orig_hostname: man-hci7-1 |
|                        | status: uptodate    |
| services               | []                  |
| status                  | uptodate            |
| tasks                   | - errors:           |
|                        |   message: None     |
|                        |   nodes: []         |
|                        | id: 8cf880d1-6648-450b-b96c-c87c27b9e181 |
|                        | name: StartSoftwareUpdateTask |
|                        | params:              |
|                        |   force: false      |
|                        |   skip: false       |
+-----+-----+
```

This command shows the node update status.

CHAPTER 4

Managing Backup Cluster

4.1 Creating, Showing, and Deleting Backup Cluster

4.1.1 vinfra service backup cluster create

Create the backup cluster:

```
usage: vinfra service backup cluster create --nodes <nodes> --domain <domain>
      --reg-account <reg-account>
      --reg-server <reg-server> --tier {0,1,2,3}
      --encoding <M>+<N> --failure-domain
      {disk,host,rack,row,room}
      --storage-type {local,nfs,s3,swift,azure,
      google} [--stdin]
      [--nfs-host <host>]
      [--nfs-export <export>]
      [--nfs-version <version>]
      [--s3-flavor <flavor>]
      [--s3-region <region>]
      [--s3-bucket <bucket>]
      [--s3-endpoint <endpoint>]
      [--s3-access-key-id <access-key-id>]
      [--s3-secret-key-id <secret-key-id>]
      [--s3-cert-verify <cert-verify>]
      [--swift-auth-url <auth-url>]
      [--swift-auth-version <auth-version>]
      [--swift-user-name <user-name>]
      [--swift-api-key <api-key>]
      [--swift-domain <domain>]
      [--swift-domain-id <domain-id>]
      [--swift-tenant <tenant>]
```



```

[--swift-tenant-id <tenant-id>]
[--swift-tenant-domain <tenant-domain>]
[--swift-tenant-domain-id <tenant-domain-id>]
[--swift-trust-id <trust-id>]
[--swift-region <region>]
[--swift-internal <internal>]
[--swift-container <container>]
[--swift-cert-verify <cert-verify>]
[--azure-endpoint <endpoint>]
[--azure-container <container>]
[--azure-account-name <account-nameE>]
[--azure-account-key <account-key>]
[--google-bucket <bucket>]
[--google-credentials <credentials>]

```

--nodes <nodes>

A comma-separated list of node hostnames or IDs

--domain <domain>

Domain name for the backup cluster

--reg-account <reg-account>

Partner account in the cloud or of an organization administrator on the local management server

--reg-server <reg-server>

URL of the cloud management portal or the hostname/IP address and port of the local management server

--tier {0,1,2,3}

Storage tier

--encoding <M>+<N>

Storage erasure coding mapping in the format:

- M: the number of data blocks
- N: the number of parity blocks

--failure-domain {disk,host,rack,row,room}

Storage failure domain

--storage-type {local,nfs,s3,swift,azure,google}

Storage type

--stdin

Ask to enter the registration password in the console.

Storage parameters for the nfs storage type:

- `--nfs-host <host>`
NFS hostname or IP address
- `--nfs-export <export>`
Full path to the NFS export
- `--nfs-version <version>`
NFS version (3 or 4)

Storage parameters for the s3 storage type:

- `--s3-flavor <flavor> (optional)`
Flavor name
- `--s3-region <region> (optional)`
Set region for Amazon S3.
- `--s3-bucket <bucket>`
Bucket name
- `--s3-endpoint <endpoint>`
Endpoint URL
- `--s3-access-key-id <access-key-id>`
Access key ID
- `--s3-secret-key-id <secret-key-id>`
Secret key ID
- `--s3-cert-verify <cert-verify> (optional)`
Allow self-signed certificate of the S3 endpoint

Storage parameters for the swift storage type:

- `--swift-auth-url <auth-url>`
Authentication (keystone) URL
- `--swift-auth-version <auth-version> (optional)`
Authentication protocol version
- `--swift-user-name <user-name>`
User name

--swift-api-key <api-key>
API key (password)

--swift-domain <domain> **(optional)**
Domain name

--swift-domain-id <domain-id> **(optional)**
Domain ID

--swift-tenant <tenant> **(optional)**
Tenant name

--swift-tenant-id <tenant-id> **(optional)**
Tenant ID

--swift-tenant-domain <tenant-domain> **(optional)**
Tenant domain name

--swift-tenant-domain-id <tenant-domain-id> **(optional)**
Tenant domain ID

--swift-trust-id <trust-id> **(optional)**
Trust ID

--swift-region <region> **(optional)**
Region name

--swift-container <container> **(optional)**
Container name

--swift-cert-verify <cert-verify> **(optional)**
Allow self-signed certificate of the Swift endpoint (true or false)

Storage parameters for the azure storage type:

--azure-endpoint <endpoint>
Endpoint URL

--azure-container <container>
Container name

--azure-account-name <account-name>
Account name

```
--azure-account-key <account-key>
```

Account key

Storage parameters for the google storage type:

```
--google-bucket <bucket>
```

Google bucket name

```
--google-credentials <credentials>
```

Path to the file with Google credentials

Example 1. Creating the backup cluster on the local storage:

```
# vinfra service backup cluster create --nodes 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb,\
74cbd22b-fb1b-4441-ae52-532078c54f9a,eeb06dce-4cfd-4c89-bc7f-4689ea5c7058 \
--storage-type local --domain dns.example.com --tier 0 --encoding 1+2 --failure-domain \
host --reg-account account@example.com --reg-server https://cloud.acronis.com/ --stdin
Password:
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | ee7e60c5-5447-4177-8581-26657ac380c0 |
+-----+-----+
```

This command creates a task to create the backup cluster from three nodes specified by ID on the local storage. It also specifies the domain name, tier, failure domain, registration account and server.

Task outcome:

```
# vinfra task show ee7e60c5-5447-4177-8581-26657ac380c0
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                     |
| name    | backend.presentation.abgw.tasks.RegisterAbgwTask |
| result  |                                     |
| state   | success                              |
| task_id | ee7e60c5-5447-4177-8581-26657ac380c0 |
+-----+-----+
```

Example 2. Creating the backup cluster on the S3 storage:

```
# vinfra service backup cluster create --nodes 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb,\
74cbd22b-fb1b-4441-ae52-532078c54f9a,eeb06dce-4cfd-4c89-bc7f-4689ea5c7058 \
--storage-type s3 --domain dns.example.com --tier 0 --encoding 1+2 --failure-domain \
host --s3-bucket mybucket --s3-endpoint s3.amazonaws.com --s3-access-key-id \
e302a06df8adbe9fAIF1 --s3-secret-key-id x1gXquRHQXuyiUJQoQMoAohA2TkYHer20o8tfPX7 \
--s3-cert-verify true --reg-account account@example.com --reg-server \
https://cloud.acronis.com/ --stdin
Password:
```

```

+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | 0fb53a6f-2bc4-410a-aa1c-b3cda6ca8570 |
+-----+-----+

```

This command creates a task to create the backup cluster from three nodes specified by ID on the S3 storage. It also specifies the domain name, tier, failure domain, registration account and server, as well as the required S3 parameters.

Task outcome:

```

# vinfra task show 0fb53a6f-2bc4-410a-aa1c-b3cda6ca8570
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                       |
| name    | backend.presentation.abgw.tasks.RegisterAbgwTask |
| result  |                                       |
| state   | success                               |
| task_id | 0fb53a6f-2bc4-410a-aa1c-b3cda6ca8570 |
+-----+-----+

```

Example 3. Creating the backup cluster on the NFS storage:

```

# vinfra service backup cluster create --nodes eeb06dce-4cfd-4c89-bc7f-4689ea5c7058 \
--storage-type nfs --domain dns.example.com --tier 0 --encoding 1+2 --failure-domain \
host --nfs-host nfs.example.com --nfs-export /myshare/myexport --nfs-version 4 \
--reg-account account@example.com --reg-server https://cloud.acronis.com/ --stdin
Password:
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | d76ceb22-48e7-4eac-b04f-03d3aa3377b7 |
+-----+-----+

```

This command creates a task to create the backup cluster from one node with the ID eeb06dce-4cfd-4c89-bc7f-4689ea5c7058 on the NFS storage. It also specifies the domain name, tier, failure domain, registration account and server, as well as the required NFS parameters.

Task outcome:

```

# vinfra task show d76ceb22-48e7-4eac-b04f-03d3aa3377b7
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                       |
| name    | backend.presentation.abgw.tasks.RegisterAbgwTask |
| result  |                                       |
| state   | success                               |
+-----+-----+

```

```
| task_id | d76ceb22-48e7-4eac-b04f-03d3aa3377b7 |
+-----+-----+
```

4.1.2 vinfra service backup cluster show

Display backup cluster details:

```
usage: vinfra service backup cluster show
```

Example:

```
# vinfra service backup cluster show
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| abgw_address   | dns.example.com                         |
| account_server | https://cloud.acronis.com               |
| dc_uid         | 44893a40296ecd9ae64567297a5b2b07-1577203369 |
| migration      | dns: null                               |
|                | ips: []                                  |
|                | running: false                          |
|                | time_left: 0.0                          |
| reg_type       | abc                                       |
| storage_params | access_key_id: e302a06df8adbe9fAIF1    |
|                | bucket: mybucket                        |
|                | cert_verify: true                       |
|                | endpoint: s3.amazonaws.com              |
|                | flavour: null                            |
|                | region: null                             |
|                | secret_key_id: x1gXquRHQXuyiUJQoQMoAohA2TkYHer20o8tfPX7 |
| storage_type   | s3                                       |
+-----+-----+
```

This command shows the domain name, registration details, and storage parameters of the backup cluster.

4.1.3 vinfra service backup cluster release

Delete the backup cluster and all its data:

```
usage: vinfra service backup cluster release [--reg-account <reg-account>]
                                           [--force] [--stdin]
```

`--reg-account <reg-account>`

Partner account in the cloud or of an organization administrator on the local management server

--force

Release the backup cluster but does not unregister it from your backup software.

Note: Choose this option only if you are sure that the cluster has already been unregistered from your backup software.

--stdin

Ask to enter the registration password in the console.

Example:

```
# vinfra service backup cluster release --reg-account account@example.com --stdin
Password:
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | cf270233-06d5-4a4a-8dea-443d6fb59b10 |
+-----+-----+
```

This command creates a task to delete the backup cluster with all its data and unregister it from your backup software.

Task outcome:

```
# vinfra task show cf270233-06d5-4a4a-8dea-443d6fb59b10
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.abgw.tasks.ReleaseNodesTask |
| result | |
| state | success |
| task_id | cf270233-06d5-4a4a-8dea-443d6fb59b10 |
+-----+-----+
```

4.2 Managing Backup Nodes

4.2.1 vinfra service backup node add

Add a list of nodes to the backup cluster:

```
usage: vinfra service backup node add --nodes <nodes>
```

--nodes <nodes>

A comma-separated list of node hostnames or IDs

Example:

```
# vinfra service backup node add --nodes 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | affe92f4-0c01-4a06-b91b-4ee0355d9a87 |
+-----+-----+
```

This command creates a task to add the node with the ID 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb to the backup cluster.

Task outcome:

```
# vinfra task show affe92f4-0c01-4a06-b91b-4ee0355d9a87
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.abgw.tasks.AssignNodesTask |
| result | |
| state | success |
| task_id | affe92f4-0c01-4a06-b91b-4ee0355d9a87 |
+-----+-----+
```

4.2.2 vinfra service backup node list

List nodes in the backup cluster:

```
usage: vinfra service backup node list
```

Example:

```
# vinfra service backup node list
+-----+-----+-----+
| id | host | is_online |
+-----+-----+-----+
| 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb | node003.vstoragedomain | True |
| 74cbd22b-fb1b-4441-ae52-532078c54f9a | node001.vstoragedomain | True |
| eeb06dce-4cf4-4c89-bc7f-4689ea5c7058 | node002.vstoragedomain | True |
+-----+-----+-----+
```

This command lists nodes in the backup cluster.

4.2.3 vinfra service backup node release

Release a list of nodes from the backup cluster:

```
usage: vinfra service backup node release --nodes <nodes>
```

--nodes <nodes>

A comma-separated list of node hostnames or IDs

Example:

```
# vinfra service backup node release --nodes 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | ea09642c-291c-4df8-87a5-a8958d6308c1 |
+-----+-----+
```

This command creates a task to release the node with the ID 2f3f6091-0d44-45aa-94e3-ebc2b65c0eeb from the backup cluster.

Task outcome:

```
# vinfra task show ea09642c-291c-4df8-87a5-a8958d6308c1
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.abgw.tasks.ReleaseNodesTask |
| result | |
| state | success |
| task_id | ea09642c-291c-4df8-87a5-a8958d6308c1 |
+-----+-----+
```

4.3 Updating Backup Cluster Certificates

Update certificates for the backup cluster:

```
usage: vinfra service backup cluster renew-certificates [--stdin]
                                     --reg-account <reg-account>
                                     --reg-server <reg-server>
```

--stdin

Ask to enter the registration password in the console.

```
--reg-account <reg-account>
```

Partner account in the cloud or of an organization administrator on the local management server

```
--reg-server <reg-server>
```

URL of the cloud management portal or the hostname/IP address and port of the local management server

Example:

```
# vinfra service backup cluster renew-certificates --reg-account account@example.com \
--reg-server https://cloud.acronis.com/ --stdin
Password:
+-----+
| Field | Value |
+-----+
| task_id | 7f1873a7-cd9b-49f3-ae17-fb14ff08ddf5 |
+-----+
```

This command creates a task to update certificates for the backup cluster.

Task outcome:

```
# vinfra task show 7f1873a7-cd9b-49f3-ae17-fb14ff08ddf5
+-----+
| Field | Value |
+-----+
| details | |
| name | backend.presentation.abgw.tasks.RenewRegistrationAbgwTask |
| result | |
| state | success |
| task_id | 7f1873a7-cd9b-49f3-ae17-fb14ff08ddf5 |
+-----+
```

4.4 Changing Storage Parameters

4.4.1 vinfra service backup storage-params show

Display storage parameters:

```
usage: vinfra service backup storage-params show
```

Example:

```
# vinfra service backup storage-params show
+-----+
| Field | Value |
+-----+
```

```

+-----+-----+
| abgw_address | dns.example.com |
| account_server | https://cloud.acronis.com |
| dc_uid | 44893a40296ecd9ae64567297a5b2b07-1577264050 |
| migration | dns: null |
| | ips: [] |
| | running: false |
| | time_left: 0.0 |
| reg_type | abc |
| storage_params | export: /myshare/myexport |
| | host: 10.94.129.70 |
| | version: 4 |
| storage_type | nfs |
+-----+-----+

```

This command shows the storage parameters of the backup cluster.

4.4.2 vinfra service backup storage-params change

Important: Change storage parameters with caution and only within the existing configuration. You can change the IP address of an external storage or its access credentials.

Modify storage parameters:

```

usage: vinfra service backup storage-params change --storage-type {local,nfs,s3,swift,
azure,google}
[--nfs-host <host>]
[--nfs-export <export>]
[--nfs-version <version>]
[--s3-flavor <flavor>]
[--s3-region <region>]
[--s3-bucket <bucket>]
[--s3-endpoint <endpoint>]
[--s3-access-key-id <access-key-id>]
[--s3-secret-key-id <secret-key-id>]
[--s3-cert-verify <cert-verify>]
[--swift-auth-url <auth-url>]
[--swift-auth-version <auth-version>]
[--swift-user-name <user-name>]
[--swift-api-key <api-key>]
[--swift-domain <domain>]
[--swift-domain-id <domain-id>]
[--swift-tenant <tenant>]
[--swift-tenant-id <tenant-id>]
[--swift-tenant-domain <tenant-domain>]
[--swift-tenant-domain-id <tenant-domain-id>]

```

```

[--swift-trust-id <trust-id>]
[--swift-region <region>]
[--swift-internal <internal>]
[--swift-container <container>]
[--swift-cert-verify <cert-verify>]
[--azure-endpoint <endpoint>]
[--azure-container <container>]
[--azure-account-name <account-name>]
[--azure-account-key <account-key>]
[--google-bucket <bucket>]
[--google-credentials <credentials>]

```

```
--storage-type {local,nfs,s3,swift,azure,google}
```

Storage type

Storage parameters for the nfs storage type:

```
--nfs-host <host>
```

NFS hostname or IP address

```
--nfs-export <export>
```

Full path to the NFS export

```
--nfs-version <version>
```

NFS version (3 or 4)

Storage parameters for the s3 storage type:

```
--s3-flavor <flavor> (optional)
```

Flavor name

```
--s3-region <region> (optional)
```

Set region for Amazon S3.

```
--s3-bucket <bucket>
```

Bucket name

```
--s3-endpoint <endpoint>
```

Endpoint URL

```
--s3-access-key-id <access-key-id>
```

Access key ID

```
--s3-secret-key-id <secret-key-id>
```

Secret key ID

--s3-cert-verify <cert-verify> **(optional)**

Allow self-signed certificate of the S3 endpoint

Storage parameters for the swift storage type:

--swift-auth-url <auth-url>

Authentication (keystone) URL

--swift-auth-version <auth-version> **(optional)**

Authentication protocol version

--swift-user-name <user-name>

User name

--swift-api-key <api-key>

API key (password)

--swift-domain <domain> **(optional)**

Domain name

--swift-domain-id <domain-id> **(optional)**

Domain ID

--swift-tenant <tenant> **(optional)**

Tenant name

--swift-tenant-id <tenant-id> **(optional)**

Tenant ID

--swift-tenant-domain <tenant-domain> **(optional)**

Tenant domain name

--swift-tenant-domain-id <tenant-domain-id> **(optional)**

Tenant domain ID

--swift-trust-id <trust-id> **(optional)**

Trust ID

--swift-region <region> **(optional)**

Region name

--swift-container <container> **(optional)**

Container name

`--swift-cert-verify <cert-verify> (optional)`

Allow self-signed certificate of the Swift endpoint (true or false)

Storage parameters for the azure storage type:

`--azure-endpoint <endpoint>`

Endpoint URL

`--azure-container <container>`

Container name

`--azure-account-name <account-name>`

Account name

`--azure-account-key <account-key>`

Account key

Storage parameters for the google storage type:

`--google-bucket <bucket>`

Google bucket name

`--google-credentials <credentials>`

Path to the file with Google credentials

Example:

```
# vinfra service backup storage-params change --storage-type nfs --nfs-host \
10.94.129.71 --nfs-export /myshare/myexport --nfs-version 4
Operation successful.
```

This command changes the NFS storage parameters for the backup cluster.

4.5 Changing Volume Parameters

4.5.1 vinfra service backup volume-params show

Display volume parameters:

```
usage: vinfra service backup volume-params show
```

Example:

```
# vinfra service backup volume-params show
+-----+-----+
| Field      | Value      |
+-----+-----+
| failure_domain | host      |
| redundancy    | m: 1      |
|               | n: 2      |
|               | type: raid6 |
| tier          | 0         |
+-----+-----+
```

This command shows the volume parameters of the backup cluster: failure domain, redundancy scheme, and tier.

4.5.2 vinfra service backup volume-params change

Important: Changing redundancy scheme is not recommended, because it may decrease cluster performance. The reason is that re-encoding demands a significant amount of cluster resources for a long period of time. If you still want to change the redundancy scheme, please contact the technical support.

Modify volume parameters:

```
usage: vinfra service backup volume-params change [--tier {0,1,2,3}] [--encoding <M>+<N>]
        [--failure-domain {disk,host,rack,
        row,room}]
```

`--tier {0,1,2,3}`

Storage tier

`--encoding <M>+<N>`

Storage erasure coding mapping in the format:

- M: the number of data blocks
- N: the number of parity blocks

`--failure-domain {disk,host,rack,row,room}`

Storage failure domain

Example:

```
# vinfra service backup volume-params change --tier 1 --encoding 1+0 \
--failure-domain host
```

```
+-----+
| Field | Value |
+-----+
| task_id | 28ae19dc-51c9-49bf-bd93-51a763fa181b |
+-----+
```

This command creates a task to change volume parameters of the backup cluster as follows:

- storage tier to 1
- erasure coding scheme to 1+0
- failure domain to host

Task outcome:

```
# vinfra task show 28ae19dc-51c9-49bf-bd93-51a763fa181b
+-----+
| Field | Value |
+-----+
| details | |
| name | backend.presentation.abgw.tasks.ChangeVolumeParamsTask |
| result | |
| state | success |
| task_id | 28ae19dc-51c9-49bf-bd93-51a763fa181b |
+-----+
```

4.6 Managing Backup Cluster Geo-replication

Important: To enable geo-replication, the backup clusters must be able to reach each other via domain names on TCP port 44445.

To enable geo-replication between two backup clusters, execute the following commands (see examples in the sections below):

1. On the cluster that will be configured as secondary, run `vinfra service backup geo-replication show` to learn its address and UID.
2. On the cluster that will be configured as primary, run `vinfra service backup geo-replication master setup`, using the address and UID of the secondary cluster.
3. On the primary cluster, run `vinfra service backup geo-replication master download-configs` to generate the configuration file of the primary cluster.

4. Move the configuration file of the primary cluster to the secondary cluster using the standard Linux command-line tool, for example, `scp`.
5. On the secondary cluster, run `vinfra service backup geo-replication slave setup` to upload the configuration file of the primary cluster.
6. On the primary cluster, run `vinfra service backup geo-replication master establish` to establish a connection between the primary and secondary clusters.
7. On one of the clusters, run `vinfra service backup geo-replication show` to check that geo-replication has been successfully enabled.

4.6.1 `vinfra service backup geo-replication show`

Display the geo-replication configuration:

```
usage: vinfra service backup geo-replication show
```

Example:

```
# vinfra service backup geo-replication show
+-----+-----+-----+-----+-----+
| Field | Value |
+-----+-----+-----+-----+
| self  | address: slave.example.com |
|       | datacenter_uid: e63a67388deb3c99d044eecbd7b79ad3-1577275849 |
+-----+-----+-----+-----+-----+-----+
```

This command shows the geo-replication configuration of the secondary cluster.

4.6.2 `vinfra service backup geo-replication master setup`

Configure geo-replication for the primary cluster:

```
usage: vinfra service backup geo-replication master setup --slave-cluster-address
<slave-cluster-address>
--slave-cluster-uid
<slave-cluster-uid>
```

```
--slave-cluster-address <slave-cluster-address>
```

Secondary cluster DNS name

```
--slave-cluster-uid <slave-cluster-uid>
```

Secondary cluster UID

Example:

```
# vinfra service backup geo-replication master setup --slave-cluster-address \
slave.example.com --slave-cluster-uid e63a67388deb3c99d044eecbd7b79ad3-1577275849
+-----+
| Field | Value |
+-----+
| task_id | 07df4a57-704e-47de-b681-615ee0c26a21 |
+-----+
```

This command creates a task to configure geo-replication for the primary backup cluster.

Task outcome:

```
# vinfra task show 07df4a57-704e-47de-b681-615ee0c26a21
+-----+
| Field | Value |
+-----+
| details | |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result | |
| state | success |
| task_id | 07df4a57-704e-47de-b681-615ee0c26a21 |
+-----+
```

4.6.3 vinfra service backup geo-replication master download-configs

Download the geo-replication configuration file of the primary cluster:

```
usage: vinfra service backup geo-replication master download-configs
      [--conf-file-path <conf-file-path>]
```

`--conf-file-path <conf-file-path>`

Path where the configuration file will be downloaded

Example:

```
# vinfra service backup geo-replication master download-configs \
--conf-file-path master_dc.conf
```

This command downloads the geo-replication configuration file of the primary cluster to the specified file.

4.6.4 vinfra service backup geo-replication slave setup

Configure geo-replication for the secondary cluster:

```
usage: vinfra service backup geo-replication slave setup
      --dc-config-file <dc-config-file>
```

--dc-config-file <dc-config-file>

Path to the configuration file of the primary cluster on the local server

Example:

```
# vinfra service backup geo-replication slave setup --dc-config-file master_dc.conf
+-----+-----+
| Field  | Value                                |
+-----+-----+
| task_id | d34b3a4f-6e16-4e60-b20a-844052945d3e |
+-----+-----+
```

This command creates a task to configure geo-replication for the secondary backup cluster.

Task outcome:

```
# vinfra task show d34b3a4f-6e16-4e60-b20a-844052945d3e
+-----+-----+
| Field  | Value                                |
+-----+-----+
| details |                                       |
| name    | backend.tasks.message_dispatcher.CommandDispatcher |
| result  |                                       |
| state   | success                              |
| task_id | d34b3a4f-6e16-4e60-b20a-844052945d3e |
+-----+-----+
```

4.6.5 vinfra service backup geo-replication master establish

Establish a connection between the primary and secondary clusters to enable geo-replication:

```
usage: vinfra service backup geo-replication master establish
```

Example:

```
# vinfra service backup geo-replication master establish
+-----+-----+
| Field  | Value                                |
+-----+-----+
```

```
| task_id | 014903e4-c2e6-4e03-b1af-06c28b672f6e |
+-----+-----+
```

This command creates a task to connect the primary and secondary clusters to enable geo-replication.

Task outcome:

```
# vinfra task show 014903e4-c2e6-4e03-b1af-06c28b672f6e
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result |
| state | success |
| task_id | 014903e4-c2e6-4e03-b1af-06c28b672f6e |
+-----+-----+
```

4.6.6 vinfra service backup geo-replication slave update-certificates

Update the primary cluster's configuration on the secondary cluster:

```
usage: vinfra service backup geo-replication slave update-certificates
       --dc-config-file <dc-config-file>
```

```
--dc-config-file <dc-config-file>
```

Path to the configuration file of the primary cluster

Example:

```
# vinfra service backup geo-replication slave update-certificates \
--dc-config-file primary_dc_updated.conf
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 0ab89de1-b02d-426b-a03c-b1922e610594 |
+-----+-----+
```

This command creates a task to update the configuration of the primary backup cluster.

Task outcome:

```
# vinfra task show 0ab89de1-b02d-426b-a03c-b1922e610594
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result |
```

```
| state | success |
| task_id | 0ab89de1-b02d-426b-a03c-b1922e610594 |
+-----+-----+
```

4.6.7 vinfra service backup geo-replication master disable

Disable geo-replication on the primary cluster:

```
usage: vinfra service backup geo-replication master disable
```

Example:

```
# vinfra service backup geo-replication master disable
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | dc2bb8ae-8e32-4d37-8d97-4c4c46189d27 |
+-----+-----+
```

This command creates a task to disable geo-replication on the primary cluster.

Task outcome:

```
# vinfra task show dc2bb8ae-8e32-4d37-8d97-4c4c46189d27
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result | |
| state | success |
| task_id | dc2bb8ae-8e32-4d37-8d97-4c4c46189d27 |
+-----+-----+
```

4.6.8 vinfra service backup geo-replication slave promote-to-master

Promote the secondary cluster to primary in the geo-replication configuration:

```
usage: vinfra service backup geo-replication slave promote-to-master
```

Example:

```
# vinfra service backup geo-replication slave promote-to-master
+-----+-----+
| Field | Value |
+-----+-----+
```

```
| task_id | 083a7d6e-3be8-490f-b468-a3f84abb3487 |
+-----+-----+
```

This command creates a task to promote the secondary cluster to primary in the geo-replication configuration.

Task outcome:

```
# vinfra task show 083a7d6e-3be8-490f-b468-a3f84abb3487
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result |
| state | success |
| task_id | 083a7d6e-3be8-490f-b468-a3f84abb3487 |
+-----+-----+
```

4.6.9 vinfra service backup geo-replication slave cancel

Cancel geo-replication for the secondary cluster:

```
usage: vinfra service backup geo-replication slave cancel
```

Example:

```
# vinfra service backup geo-replication slave cancel
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | ad977d03-995c-4677-9308-5e73ec8a2821 |
+-----+-----+
```

This command creates a task to cancel geo-replication for the secondary backup cluster.

Task outcome:

```
# vinfra task show ad977d03-995c-4677-9308-5e73ec8a2821
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result |
| state | success |
| task_id | ad977d03-995c-4677-9308-5e73ec8a2821 |
+-----+-----+
```

4.6.10 vinfra service backup geo-replication master cancel

Cancel geo-replication for the primary cluster:

```
usage: vinfra service backup geo-replication master cancel
```

Example:

```
# vinfra service backup geo-replication master cancel
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | e1931274-24a5-491e-a5f8-d24fdf4385f7 |
+-----+-----+
```

This command creates a task to cancel geo-replication for the primary backup cluster.

Task outcome:

```
# vinfra task show e1931274-24a5-491e-a5f8-d24fdf4385f7
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.tasks.message_dispatcher.CommandDispatcher |
| result | |
| state | success |
| task_id | e1931274-24a5-491e-a5f8-d24fdf4385f7 |
+-----+-----+
```

CHAPTER 5

Managing General Settings

5.1 Managing Licenses

5.1.1 vinfra cluster license load

Load a license from a key.

```
usage: vinfra cluster license load <license-key>
```

<license-key>

License key to register

Example:

```
# vinfra cluster license load A38600-3P6W74-RZSK58-Y9ZH05-2X7J48
+-----+
| Field          | Value                               |
+-----+
| capacity       | 7036767043584000                   |
| expiration_ts  | 1549583999                          |
| free_size      | 7036766991361913                   |
| keynumber      | VZSTOR.74418710.0000               |
| spla           | registered: false                   |
|                | registration_url: null              |
| status         | active                               |
| total_size     | 7036767043584000                   |
| used_size      | 52222087                            |
+-----+
+-----+
```

This command installs the license from the key A38600-3P6W74-RZSK58-Y9ZH05-2X7J48.

5.1.2 vinfra cluster license show

Show details of the installed license:

```
usage: vinfra cluster license show
```

Example:

```
# vinfra cluster license show
+-----+-----+
| Field      | Value                |
+-----+-----+
| capacity   | 7036767043584000    |
| expiration_ts | 1549583999          |
| free_size  | 7036766991361913    |
| keynumber  | VZSTOR.74418710.0000 |
| spla       | registered: false    |
|            | registration_url: null |
| status     | active               |
| total_size | 7036767043584000    |
| used_size  | 52222087             |
+-----+-----+
```

This command shows the details of the currently installed license.

5.1.3 vinfra cluster license update

Update the installed license:

```
usage: vinfra cluster license update [--server <ka-server>]
```

```
--server <ka-server>
```

Hostname[:port] of the key administration server (default: ka.parallels.com)

Example:

```
# vinfra cluster license update --server ka.parallels.com
+-----+-----+
| Field      | Value                |
+-----+-----+
| capacity   | 7036767043584000    |
| expiration_ts | 1549583999          |
| free_size  | 7036766991361913    |
| keynumber  | VZSTOR.74418710.0000 |
| spla       | registered: false    |
|            | registration_url: null |
| status     | active               |
| total_size | 7036767043584000    |
+-----+-----+
```

```
| used_size      | 52222087      |
+-----+-----+
```

This command updates the installed license.

5.2 Managing Domains

5.2.1 vinfra domain create

Create a new domain:

```
usage: vinfra domain create [--description <description>] [--enable | --disable] <name>
```

`--description <description>`

Domain description

`--enable`

Enable domain

`--disable`

Disable domain

`<name>`

Domain name

Example:

```
# vinfra domain create mydomain
+-----+-----+
| Field      | Value          |
+-----+-----+
| description |                |
| enabled     | True          |
| id          | ed408d00561c4a398f933c29e87cadab |
| name        | domain1       |
| projects_count | 0            |
+-----+-----+
```

This command creates and enables the domain `mydomain`.

5.2.2 vinfra domain list

List all available domains:

```
usage: vinfra domain list
```

Example:

```
# vinfra domain list
+-----+-----+-----+-----+
| id                | name    | enabled | description          |
+-----+-----+-----+-----+
| default           | Default | True    | The default domain  |
| 24986479ee3246048d3ef2a065ea99f5 | mydomain | True    |                      |
+-----+-----+-----+-----+
```

This command lists domains used in the compute cluster.

5.2.3 vinfra domain show

Display information about a domain:

```
usage: vinfra domain show <domain>
```

<domain>

Domain ID or name

Example:

```
# vinfra domain show mydomain
+-----+-----+
| Field          | Value          |
+-----+-----+
| description    |                |
| enabled        | True           |
| id             | 24986479ee3246048d3ef2a065ea99f5 |
| name           | mydomain       |
| projects_count | 0              |
+-----+-----+
```

This command shows the details of the domain mydomain.

5.2.4 vinfra domain set

Modify an existing domain:

```
usage: vinfra domain set [--description <description>] [--enable | --disable]
                        [--name <name>] <domain>
```

--description <description>

Domain description

--enable

Enable domain

--disable

Disable domain

--name <name>

Domain name

<domain>

Domain ID or name

Example:

```
# vinfra domain set mydomain --description "A custom domain"
+-----+-----+
| Field      | Value                |
+-----+-----+
| description | A custom domain     |
| enabled    | True                 |
| id         | 24986479ee3246048d3ef2a065ea99f5 |
| name       | mydomain             |
| projects_count | 0                   |
+-----+-----+
```

This command adds the description for the domain mydomain.

5.2.5 vinfra domain delete

Delete a domain:

```
usage: vinfra domain delete <domain>
```

<domain>

Domain ID or name

Example:

```
# vinfra domain delete mydomain
Operation successful
```

This command deletes the domain `mydomain`.

5.3 Managing Domain Users

5.3.1 vinfra domain user list-available-roles

List available user roles:

```
usage: vinfra domain user list-available-roles
```

Example:

```
# vinfra domain user list-available-roles
+-----+-----+-----+-----+
| id      | name      | description                                     | scope  |
+-----+-----+-----+-----+
| abgw    | ABGW     | Can create and manage Acronis Backup          | - system |
|         |          | Gateway.                                     |         |
| admin   | Administrator | Can perform all management operations.      | - system |
| cluster | Cluster   | Can create cluster, join nodes to cluster,   | - system |
|         |          | and manage (assign and release) disks.      |         |
| compute | Compute   | Can create and manage compute cluster.      | - system |
| domain_admin | Domain Admin | Can manage users, projects and all          | - domain |
|         |          | resources in a domain.                     |         |
| image_upload | Image Upload | Can manage compute images.                  | - domain |
| iscsi   | Block Storage | Can create and manage iSCSI targets, LUNs  | - system |
|         |          | and CHAP users.                            |         |
| login   | Login     | Can login in web UI.                        | []     |
| network | Network   | Can modify network settings and roles.      | - system |
| nfs     | NFS       | Can create and manage NFS.                  | - system |
| project_admin | Project Admin | Can manage virtual objects inside a        | - project |
|         |          | project.                                    |         |
| s3      | S3        | Can create and manage S3 cluster.           | - system |
| ssh     | SSH       | Can add and remove SSH keys for cluster     | - system |
|         |          | nodes access.                               |         |
| updates | Updates   | Can install updates.                        | - system |
| viewer  | Viewer    | Viewer role (read only)                     | - system |
+-----+-----+-----+-----+
```

This command lists all available user roles.

5.3.2 vinfra domain user create

Create a new domain user:

```
usage: vinfra domain user create [--email <email>] [--description <description>]
                                [--assign <project> <role>]
                                [--domain-permissions <domain_permissions>]
                                [--system-permissions <system_permissions>]
                                [--enable | --disable] --domain <domain> <name>
```

--email <email>

User email

--description <description>

User description

--assign <project> <role>

Assign a user to a project with one or more permission sets. Specify this option multiple times to assign the user to multiple projects.

- <project>: project ID or name
- <role>: user role in the project (project_admin)

--domain-permissions <domain_permissions>

A comma-separated list of domain permissions. View the list of available domain permissions using `vinfra domain user list-available-roles | grep domain`.

--system-permissions <system_permissions>

A comma-separated list of system permissions. View the list of available system permissions using `vinfra domain user list-available-roles | grep system`.

--enable

Enable user

--disable

Disable user

--domain <domain>

Domain name or ID

<name>

User name

Example:

```
# vinfra domain user create --domain mydomain --name myuser \
--domain-permissions domain_admin
Password:
+-----+-----+
| Field          | Value          |
+-----+-----+
| assigned_projects | []             |
| description      |                |
| domain_permissions | - domain_admin |
| email           |                |
| enabled         | True          |
| id              | a9c67c6acf1f4df1818fdeeee0b4bd5e |
| name            | myuser        |
| role            | domain_admin  |
| system_permissions | []            |
+-----+-----+
```

This command creates and enables a new administrator account `myuser` within the domain `mydomain`. It also sets password for the new user.

5.3.3 vinfra domain user list

List all users in a domain:

```
usage: vinfra domain user list --domain <domain>
```

```
--domain <domain>
    Domain name or ID
```

Example:

```
# vinfra domain user list --domain mydomain -c id -c name -c enabled \
-c domain_permissions -c assigned_projects
+-----+-----+-----+-----+-----+
| id      | name  | enabled | domain_permissions | assigned_projects |
+-----+-----+-----+-----+-----+
| a9c6<...> | myuser | True   | - domain_admin    | []                |
+-----+-----+-----+-----+-----+
```

This command lists users in the domain `mydomain`. (The output is abridged to fit on page.)

5.3.4 vinfra domain user show

Display information about a domain user:

```
usage: vinfra domain user show --domain <domain> <user>
```

```
--domain <domain>
```

Domain ID or name

```
<user>
```

User ID or name

Example:

```
# vinfra domain user show myuser --domain mydomain
+-----+-----+
| Field          | Value                               |
+-----+-----+
| assigned_projects | []                                   |
| description      |                                       |
| domain_permissions | - domain_admin                     |
| email           |                                       |
| enabled         | True                                |
| id              | a9c67c6acf1f4df1818fdeeee0b4bd5e  |
| name            | myuser                              |
| role            | domain_admin                        |
| system_permissions | []                                   |
+-----+-----+
```

This command shows the details of the user `myuser` from the domain `mydomain`.

5.3.5 vinfra domain user set

Modify the parameters of a domain user:

```
usage: vinfra domain user set [--password] [--email <email>]
                             [--description <description>]
                             [--assign <project> <role>]
                             [--domain-permissions <domain_permissions>]
                             [--system-permissions <system_permissions>]
                             [--enable | --disable] [--name <name>]
                             --domain <domain> <user>
```

```
--password
```

Request the password from stdin


```
--email <email>
    User email

--description <description>
    User description

--assign <project> <role>
    Assign a user to a project with one or more permission sets. Specify this option multiple times to
    assign the user to multiple projects.

    • <project>: project ID or name

    • <role>: user role in the project (project_admin)

--domain-permissions <domain_permissions>
    A comma-separated list of domain permissions. View the list of available domain permissions using
    vinfra domain user list-available-roles | grep domain.

--system-permissions <system_permissions>
    A comma-separated list of system permissions. View the list of available system permissions using
    vinfra domain user list-available-roles | grep system.

--enable
    Enable user

--disable
    Disable user

--name <name>
    User name

--domain <domain>
    Domain name or ID

<user>
    User ID or name
```

Example:

```
# vinfra domain user set myuser --domain mydomain \
--assign myproject project_admin
+-----+-----+
| Field          | Value          |
+-----+-----+
```

```

| assigned_projects | [] |
| description      |   |
| domain_permissions | - domain_admin |
| email           |   |
| enabled          | True |
| id              | a9c67c6acf1f4df1818fdeeee0b4bd5e |
| name            | myuser |
| role            | domain_admin |
| system_permissions | [] |
+-----+-----+

```

This command assigns the user `myuser` from the domain `mydomain` to the project `myproject` as a project administrator.

5.3.6 `vinfra domain user delete`

Remove a domain user:

```
usage: vinfra domain user delete --domain <domain> <user>
```

```
--domain <domain>
    Domain ID or name
```

```
<user>
    User ID or name
```

Example:

```
# vinfra domain user delete myuser --domain mydomain
Operation successful
```

This command deletes the user `myuser` from the domain `mydomain`.

5.4 Managing Domain Projects

5.4.1 `vinfra domain project create`

Create a new domain project:

```
usage: vinfra domain project create [--description <description>] [--enable | --disable]
      --name <name> --domain <domain>
```

```
--description <description>
```

Project description

```
--enable
```

Enable project

```
--disable
```

Disable project

```
--name <name>
```

Project name

```
--domain <domain>
```

Domain name or ID

Example:

```
# vinfra domain project create --domain mydomain --name myproject \
--description "A custom project"
+-----+-----+
| Field      | Value                                |
+-----+-----+
| description | A custom project                    |
| domain_id  | 9f7e68938fe946a2a862e360bbe40d98  |
| enabled    | True                                 |
| id         | d1c4d6198fb940e6b971cf306571ebbd  |
| name       | myproject                           |
+-----+-----+
```

This command creates and enables the project `myproject` within the domain `mydomain` and adds a description to it.

5.4.2 vinfra domain project list

List all projects in a domain:

```
usage: vinfra domain project list --domain <domain>
```

```
--domain <domain>
```

Domain name or ID

Example:

```
# vinfra domain project list --domain mydomain
+-----+-----+-----+-----+-----+-----+
```

```
| id          | name      | enabled | description      | domain_id |
+-----+-----+-----+-----+-----+
| d1c4<...> | myproject | True    | A custom project | 9f7e<...> |
+-----+-----+-----+-----+-----+
```

This command lists projects in the domain `mydomain`. (The output is abridged to fit on page.)

5.4.3 `vinfra domain project show`

Show details of a domain project:

```
usage: vinfra domain project show --domain <domain> <project>
```

```
--domain <domain>
    Domain name or ID
```

```
<project>
    Project ID or name
```

Example:

```
# vinfra domain project show myproject --domain mydomain
+-----+-----+
| Field          | Value                                     |
+-----+-----+
| description    | A custom project                         |
| domain_id     | 9f7e68938fe946a2a862e360bbe40d98       |
| enabled       | True                                     |
| id            | d1c4d6198fb940e6b971cf306571ebbd      |
| members_count | 0                                        |
| name         | myproject                               |
+-----+-----+
```

This command shows the details of the project `myproject` from the domain `mydomain`.

5.4.4 `vinfra domain project set`

Modify an existing project:

```
usage: vinfra domain project set [--description <description>] [--enable | --disable]
    [--name <name>] --domain <domain> <project>
```

```
--description <description>
    Project description
```

```
--enable
    Enable project

--disable
    Disable project

--name <name>
    Project name

--domain <domain>
    Domain name or ID

<project>
    Project ID or name
```

Example:

```
# vinfra cluster domain project set myproject --domain mydomain --disable
+-----+-----+
| Field      | Value                |
+-----+-----+
| description | A custom project     |
| domain_id  | 9f7e68938fe946a2a862e360bbe40d98 |
| enabled    | False                |
| id         | d1c4d6198fb940e6b971cf306571ebbd |
| name       | myproject            |
+-----+-----+
```

This command disables the project `myproject` from the domain `mydomain`.

5.4.5 vinfra domain project user list

List users of a project:

```
usage: vinfra domain project user list --domain <domain> <project>
```

```
--domain <domain>
    Domain name or ID

<project>
    Project ID or name
```

Example:

```
# vinfra domain project user list myproject --domain mydomain
+-----+-----+-----+-----+
| id           | name   | description | role           |
+-----+-----+-----+-----+
| eb0203e6b8a641d8be5b54b2f3fc9f47 | myuser |              | project_admin |
+-----+-----+-----+-----+
```

This command lists users of the project `myproject` within the domain `mydomain`.

5.4.6 vinfra domain project user remove

Remove a user from a project:

```
usage: vinfra domain project user remove --user <user> --domain <domain> <project>
```

`--user <user>`

User name or ID

`--domain <domain>`

Domain name or ID

`<project>`

Project ID or name

Example:

```
# vinfra domain project user remove myproject --domain mydomain --user myuser
Operation successful
```

This command removes the user `myuser` from the project `myproject` within the domain `mydomain`.

5.4.7 vinfra domain project delete

Delete a domain project:

```
usage: vinfra domain project delete --domain <domain> <project>
```

`--domain <domain>`

Domain name or ID

`<project>`

Project ID or name

Example:

```
# vinfra domain project delete myproject --domain mydomain
Operation successful
```

This command deletes the project `myproject` from the domain `mydomain`.

5.5 Managing SSH Keys

5.5.1 vinfra cluster sshkey add

Add an SSH public key from a file:

```
usage: vinfra cluster sshkey add <file>
```

<file>

SSH public key file

Example:

```
# vinfra cluster sshkey add id_rsa.pub
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 100a54ce-0bf5-4bc0-8e46-2e8b952343e6 |
+-----+-----+
```

This command creates a task to add a public SSH key from the file `mykey.pub` to the list of trusted keys.

Task outcome:

```
# vinfra task show 100a54ce-0bf5-4bc0-8e46-2e8b952343e6
+-----+-----+
| Field | Value |
+-----+-----+
| args | - admin |
| | - 1 |
| kwargs | key: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACueW0956J/u5kjWnia7zePChoTMVBtsh1TDNg0skM |
| | shfHWUzzfydi3/4sTrJ++6dtIoS1D3VVHvHBvp456PT5e/eVy7u0SipOPPoDY2vS2IEY+zjT6MYABi6oEYom |
| | Dbi7CsRL02HcTWzAkooZNlimWPggYaMT10BZOKAvNB+Ctpkw8JaT5PRve8UVfjxIQIzL6pQ0f0CDeCHgDsvw |
| | xK7SrQ0vBzTlF9mWkGdTGy+R0JrgGk+v9PvDXZweK+qS54uaGmpB6ZRkKMroIk3h+nZ4y/1eQ6m1C8Aspa0 |
| | nnaMaNK0tw0ibrd3MDroMcqkJWTTH/cukD3sB+MjL6nmFlrrAFRU6PBkwysIio6/XHS9jG+TI7NeRApkHnwi |
| | vwIWEKSg6ppaiLUsMi/46KCHzde20zgO8Hd0R5d7hNN/80mhD7b+bY9wig+VTMoQFQYSWrIy/qLL95ws4amg |
| | nXOIksNFjffEE/+1McZxt3j5kqnjW7OT2/xkqqWoumaM+FEPLNijL18yb29/XJr/cQZX5R9iXSk33DVjhln/ |
| | HG7xpHqAtrXbvKY8zI8t23otGT/rSvWRWV/wgPBZVSWtsE99FEMmwmk/b3KuPhi0jK0IUkcv5UBL+NLHw4 |
| | rZRiYgw/fWXPO3f6ZSLLJXtW4iW+BQL60qQWUNQ== |
| | user@example.com |
| name | backend.presentation.nodes.ssh.tasks.CreateSshKeyTask |
+-----+-----+
```

```

| result | id: 6a2fb834-4bc6-4597-ae74-7cacf96b7c75 |
|         | key: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACueW0956J/u5kjWnia7zePChoTMVBtsh1TDNg0skM |
|         | shfHWUzzfydi3/4sTrJ++6dtIoS1D3VVHvHBvp456PT5e/eVy7u0Sip0PPoDY2vS2IEY+zjT6MYABi6oEYom |
|         | Dbi7CsRL02HcTWzAkooZNIimWPggYaMT10BZOKAvNB+Ctpkw8JaT5PRve8UVfjxIQIzL6pQ0f0CDeCHgDsvw |
|         | xK7SrqOvBzTlF9mWkGdTGy+R0JrgGk+v9PvDXZweK+qS54uaGmpB6ZRkKMroIk3h+nZ4y/1eQ6m1C8Aspa0 |
|         | nnaMaNK0tw0ibrd3MDroMcqkJWTTH/cukD3sB+MjL6nmFlrrAFRU6PBkwysIio6/XHS9jG+TI7NeRApkHnwi |
|         | vwIWEKSg6ppaiLUsMi/46KCHzde20zg08Hd0R5d7hNN/80mhD7b+bY9wig+VTMoQFQYSWrIy/qLL95ws4amg |
|         | nX0IksNFjffEE/+lMcZxt3j5kqnjW70T2/xkqqWoumaM+FEPLNijL18yb29/XJr/cQZX5R9iXSk33DVjhln/ |
|         | HG7xpHqAtrXbvKY8zI8t23otGT/rSvWRWV/wgPBZVSWtsE99FEMmwxk/b3KuPhi0jK0IUKcv5UBL+NLHw4 |
|         | rZRiYgw/fWXPO3f6ZSLLJXtW4iW+BQL60qQWUNQ== |
|         | user@example.com |
|         | label: user@example.com |
| state  | success |
| task_id | 100a54ce-0bf5-4bc0-8e46-2e8b952343e6 |
+-----+

```

5.5.2 vinfra cluster sshkey list

Show the list of added SSH public keys:

```
usage: vinfra cluster sshkey list
```

Example:

```

# vinfra cluster sshkey list
+-----+-----+-----+
| id | key | label |
+-----+-----+-----+
| 8ccf7f1b-6a53-4d74-99ce-c410d51a9921 | ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQACueW0956J/u5kjWnia7zePChoTMVBtsh1TDNg0skMg5shfHWUzzfydi3/4sTrJ++6dtIoS1D3VVHvHBvp456PT5e/eVy7u0Sip0PPoDY2vS2IEY+zjT6MYABi6oEYomIIDbi7CsRL02HcTWzAkooZNIimWPggYaMT10BZOKAvNB+Ctpkw8JaT5PRve8UVfjxIQIzL6pQ0f0CDeCHgDsvwcmxK7SrqOvBzTlF9mWkGdTGy+R0JrgGk+v9PvDXZweK+qS54uaGmpB6ZRkKMroIk3h+nZ4y/1eQ6m1C8Aspa0f5nnaMaNK0tw0ibrd3MDroMcqkJWTTH/cukD3sB+MjL6nmFlrrAFRU6PBkwysIio6/XHS9jG+TI7NeRApkHnwi0vwIWEKSg6ppaiLUsMi/46KCHzde20zg08Hd0R5d7hNN/80mhD7b+bY9wig+VTMoQFQYSWrIy/qLL95ws4amgAqnX0IksNFjffEE/+lMcZxt3j5kqnjW70T2/xkqqWoumaM+FEPLNijL18yb29/XJr/cQZX5R9iXSk33DVjhln/EyHG7xpHqAtrXbvKY8zI8t23otGT/rSvWRWV/wgPBZVSWtsE99FEMmwxk/b3KuPhi0jK0IUKcv5UBL+NLHw4gQrZRiYgw/fWXPO3f6ZSLLJXtW4iW+BQL60qQWUNQ== | user@example.com |
+-----+-----+-----+

```



```
| user@example.com |
+-----+-----+-----+
```

This command lists trusted SSH keys.

5.5.3 vinfra cluster sshkey delete

Remove an SSH public key from storage cluster nodes:

```
usage: vinfra cluster sshkey delete <sshkey>
```

<sshkey>

SSH key value

Example:

```
# vinfra cluster sshkey delete 8ccf7f1b-6a53-4d74-99ce-c410d51a9921
+-----+-----+-----+
| Field | Value |
+-----+-----+-----+
| task_id | 053802b2-b4c3-454d-89e2-6d6d312dd2ed |
+-----+-----+-----+
```

This command creates a task to delete the SSH key with the ID 8ccf7f1b-6a53-4d74-99ce-c410d51a9921.

Task outcome:

```
# vinfra task show 053802b2-b4c3-454d-89e2-6d6d312dd2ed
+-----+-----+-----+
| Field | Value |
+-----+-----+-----+
| args | - admin |
| | - 1 |
| | - 8ccf7f1b-6a53-4d74-99ce-c410d51a9921 |
| kwargs | {} |
| name | backend.presentation.nodes.ssh.tasks.RemoveSshKeyTask |
| state | success |
| task_id | 053802b2-b4c3-454d-89e2-6d6d312dd2ed |
+-----+-----+-----+
```

5.6 Managing External DNS Servers

5.6.1 vinfra cluster settings dns show

Display DNS servers:

```
usage: vinfra cluster settings dns show
```

Example:

```
# vinfra cluster settings dns show
+-----+-----+
| Field          | Value                                |
+-----+-----+
| dhcp_nameservers | 10.10.0.10,10.10.0.11,10.37.130.2 |
| nameservers     | 10.10.0.11,10.10.0.10              |
+-----+-----+
```

This command lists the currently used DNS servers: both internal (obtained via DHCP) and external (static set by the user).

5.6.2 vinfra cluster settings dns set

Set DNS servers:

```
usage: vinfra cluster settings dns set --nameservers <nameservers>
```

```
--nameservers <nameservers>
```

A comma-separated list of DNS servers

Example:

```
# vinfra cluster settings dns set --nameservers 8.8.8.8
+-----+-----+
| Field          | Value                                |
+-----+-----+
| dhcp_nameservers | - 10.10.0.10 |
|                 | - 10.10.0.11 |
|                 | - 10.37.130.2 |
| nameservers     | - 8.8.8.8    |
+-----+-----+
```

This command sets the external DNS server to 8.8.8.8.

5.7 Configuring Management Node High Availability

5.7.1 `vinfra cluster ha create`

Create a HA configuration:

```
usage: vinfra cluster ha create --virtual-ip <network:ip> --nodes <nodes> [--force]
```

`--virtual-ip <network:ip>`

HA configuration mapping in the format:

- `network`: network to include in the HA configuration (must include at least one of these traffic types: Internal management, Admin panel, Self-service panel, or Compute API).
- `ip`: virtual IP address that will be used in the HA configuration.

Specify this option multiple times to create a HA configuration for multiple networks.

`--nodes <nodes>`

A comma-separated list of node IDs or hostnames

`--force`

Skip checks for minimal hardware requirements

Example:

```
# vinfra cluster ha create --virtual-ip Private:10.37.130.200 \
--virtual-ip Public:10.94.129.79 --nodes 94d58604-6f30-4339-8578-adb7903b7277,\
f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4,7d7d37b8-4c06-4f1a-b3a6-4b54257d70ce
+-----+-----+
| Field   | Value                                     |
+-----+-----+
| task_id | 80a00e55-335d-4d41-bac4-5fee4791d423 |
+-----+-----+
```

This command creates a task to create a management node HA cluster from nodes with the IDs 94d58604-6f30-4339-8578-adb7903b7277, f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4, and 7d7d37b8-4c06-4f1a-b3a6-4b54257d70ce.

The command must specify the network with the traffic type `Internal management` as well as one with the traffic type `Admin panel`.

Important: After the HA cluster has been created, the admin panel will only be accessible at the provided public IP address. Log in to said address via SSH to continue managing Virtuozzo Hybrid Infrastructure with the `vinfra` CLI tool. You may also need to set the `VINFRA_PASSWORD` environment variable again, because you will access different HA cluster nodes on each log in where it may not have been set.

Task outcome:

```
# vinfra task show 80a00e55-335d-4d41-bac4-5fee4791d423
+-----+-----+
| Field | Value |
+-----+-----+
| details |
| name | backend.presentation.ha.tasks.CreateHaConfigTask |
| result | compute_task_id: c5125024-5472-4420-b8b6-e03971ab952c |
| | ha_cluster_location: |
| | - https://10.94.129.79:8888 |
| | nodes: |
| | - id: 94d58604-6f30-4339-8578-adb7903b7277 |
| | ipaddr: 10.37.130.118 |
| | is_primary: false |
| | - id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
| | ipaddr: 10.37.130.134 |
| | is_primary: true |
| | - id: 7d7d37b8-4c06-4f1a-b3a6-4b54257d70ce |
| | ipaddr: 10.37.130.246 |
| | is_primary: false |
| | primary_node_location: https://10.94.62.243:8888 |
| | virtual_ips: |
| | - ip: 10.94.129.79 |
| | roles_set: 5f0adc1d-c10f-46c1-b7b8-dd1aacab613b |
| | - ip: 10.37.130.200 |
| | roles_set: 5a0401b5-9b42-4d8b-8372-71c747230033 |
| state | success |
| task_id | 80a00e55-335d-4d41-bac4-5fee4791d423 |
+-----+-----+
```

5.7.2 vinfra cluster ha update

Update the HA configuration:

```
usage: vinfra cluster ha update [--virtual-ip <network:ip>]
                                [--nodes <nodes>] [--force]
```

`--virtual-ip <network:ip>`

HA configuration mapping in the format:

- `network`: network to include in the HA configuration (must include at least one of these traffic types: **Internal management**, **Admin panel**, **Self-service panel**, or **Compute API**).
- `ip`: virtual IP address that will be used in the HA configuration.

Specify this option multiple times to create an HA configuration for multiple networks.

`--nodes <nodes>`

A comma-separated list of node IDs or hostnames

`--force`

Skip checks for minimal hardware requirements

Example:

```
# vinfra cluster ha update --nodes 94d58604-6f30-4339-8578-adb7903b7277,\
f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4,4b83a87d-9adf-472c-91f0-782c47b2d5f1
+-----+
| Field | Value |
+-----+
| task_id | 565e9146-254b-4f7a-a2ff-b7119c95baa9 |
+-----+
```

This command creates a task to update the management node HA configuration, that is, include the nodes with the IDs `94d58604-6f30-4339-8578-adb7903b7277`, `f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4`, and `4b83a87d-9adf-472c-91f0-782c47b2d5f1`.

Task outcome:

```
# vinfra task show 565e9146-254b-4f7a-a2ff-b7119c95baa9
+-----+
| Field | Value |
+-----+
| details |
| name | backend.presentation.ha.tasks.UpdateHaConfigTask |
| result | compute_task_id: 84994caf-3a02-43ea-b904-48632f0379c7 |
| | ha_cluster_location: |
| | - https://10.94.129.79:8888 |
| | nodes: |
| | - id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
| | ipaddr: 10.37.130.134 |
| | is_primary: true |
| | - id: 4b83a87d-9adf-472c-91f0-782c47b2d5f1 |
| | ipaddr: 10.37.130.127 |
| | is_primary: false |
| | - id: 94d58604-6f30-4339-8578-adb7903b7277 |
| | ipaddr: 10.37.130.118 |
| | is_primary: false |
| | primary_node_location: https://10.94.62.243:8888 |
+-----+
```

```

|         | virtual_ips: |
|         | - ip: 10.94.129.79 |
|         |   roles_set: 5f0adc1d-c10f-46c1-b7b8-dd1aacab613b |
|         | - ip: 10.37.130.200 |
|         |   roles_set: 5a0401b5-9b42-4d8b-8372-71c747230033 |
| state   | success |
| task_id | 565e9146-254b-4f7a-a2ff-b7119c95baa9 |
+-----+-----+

```

5.7.3 vinfra cluster ha show

Display the HA configuration:

```
usage: vinfra cluster ha show
```

Example:

```

# vinfra cluster ha show
+-----+-----+
| Field          | Value |
+-----+-----+
| ha_cluster_location | - https://10.94.129.79:8888 |
| nodes          | - id: 94d58604-6f30-4339-8578-adb7903b7277 |
|                | ipaddr: 10.37.130.118 |
|                | is_primary: false |
|                | - id: f59dabdb-bd1c-4944-8af2-26b8fe9ff8d4 |
|                | ipaddr: 10.37.130.134 |
|                | is_primary: true |
|                | - id: 4b83a87d-9adf-472c-91f0-782c47b2d5f1 |
|                | ipaddr: 10.37.130.127 |
|                | is_primary: false |
| primary_node_location | https://10.94.62.243:8888 |
| virtual_ips      | - ip: 10.37.130.200 |
|                  | roles_set: 5a0401b5-9b42-4d8b-8372-71c747230033 |
|                  | - ip: 10.94.129.79 |
|                  | roles_set: 5f0adc1d-c10f-46c1-b7b8-dd1aacab613b |
+-----+-----+

```

This command shows the management node HA cluster configuration.

5.7.4 vinfra cluster ha delete

Delete the HA configuration:

```
usage: vinfra cluster ha delete
```

Example:

```
# vinfra cluster ha delete
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | c1f3e9c3-0a7b-455a-96d4-cef3b7e86e62 |
+-----+-----+
```

This command creates a task to delete the management node HA cluster.

Task outcome:

```
# vinfra task show c1f3e9c3-0a7b-455a-96d4-cef3b7e86e62
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.ha.tasks.DeleteHaConfigTask |
| result | |
| state | success |
| task_id | c1f3e9c3-0a7b-455a-96d4-cef3b7e86e62 |
+-----+-----+
```

5.8 Managing Cluster Backups

5.8.1 vinfra cluster backup create

Create a backup:

```
usage: vinfra cluster backup create
```

Example:

```
# vinfra cluster backup create
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | e4b4f891-cc3a-4308-a321-d76265ef7b5b |
+-----+-----+
```

This command creates a task to back up the storage cluster.

Task outcome:

```
# vinfra task show e4b4f891-cc3a-4308-a321-d76265ef7b5b
+-----+-----+
| Field | Value |
+-----+-----+
```

```

| details |
| name    | backend.presentation.backups.tasks.BackupManagedNodeTask |
| result  | status: finished |
| state   | success |
| task_id | e4b4f891-cc3a-4308-a321-d76265ef7b5b |
+-----+

```

5.8.2 vinfra cluster backup show

Show backup information:

```
usage: vinfra cluster backup show
```

Example:

```

# vinfra cluster backup show
+-----+-----+
| Field          | Value |
+-----+-----+
| last_backup_date | 2019-08-21T15:41:24+00:00 |
| last_backup_location | /mnt/vstorage/webcp/backup/ |
| ready          | True |
| tasks          | [] |
+-----+-----+

```

This command shows the details of the last cluster backup and the ID of the ongoing backup task, if any.

5.9 Managing Storage Tier Encryption

5.9.1 vinfra cluster settings encryption show

Display storage tier encryption:

```
usage: vinfra cluster settings encryption show
```

Example:

```

# vinfra cluster settings encryption show
+-----+-----+
| Field | Value |
+-----+-----+
| tier0  | False |
| tier1  | False |
| tier2  | False |

```



```
| tier3 | False |
+-----+-----+
```

This command shows encryption status of each storage tier.

5.9.2 vinfra cluster settings encryption set

Set storage tier encryption:

```
usage: vinfra cluster settings encryption set [--tier-enable {0,1,2,3}]
      [--tier-disable {0,1,2,3}]
```

`--tier-enable {0,1,2,3}`

Enable encryption for storage tiers. This option can be used multiple times.

`--tier-disable {0,1,2,3}`

Disable encryption for storage tiers. This option can be used multiple times.

Example:

```
# vinfra cluster settings encryption set --tier-enable 2
+-----+-----+
| Field | Value |
+-----+-----+
| tier0  | False |
| tier1  | False |
| tier2  | True  |
| tier3  | False |
+-----+-----+
```

This command enables encryption for the storage tier 2.

5.10 Managing Automatic Storage Disk Configuration

5.10.1 vinfra cluster settings automatic-disk-replacement show

Show automatic storage disk configuration:

```
usage: vinfra cluster settings automatic-disk-replacement show
```

Example:

```
# vinfra cluster settings automatic-disk-replacement show
+-----+-----+
| Field | Value |
+-----+-----+
| tier0  | True  |
| tier1  | False |
| tier2  | True  |
| tier3  | False |
+-----+-----+
```

This command shows the status of automatic disk configuration for each storage tier.

5.10.2 vinfra cluster settings automatic-disk-replacement set

Change automatic storage disk configuration:

```
usage: vinfra cluster settings automatic-disk-replacement set
       [--tier0 {on,off}] [--tier1 {on,off}] [--tier2 {on,off}]
       [--tier3 {on,off}]
```

--tier0 {on,off}

Enable or disable automatic storage disk configuration for tier 0

--tier1 {on,off}

Enable or disable automatic storage disk configuration for tier 1

--tier2 {on,off}

Enable or disable automatic storage disk configuration for tier 2

--tier3 {on,off}

Enable or disable automatic storage disk configuration for tier 3

Example:

```
# vinfra cluster settings automatic-disk-replacement set \
  --tier0 on --tier1 on --tier2 on --tier3 on
+-----+-----+
| Field | Value |
+-----+-----+
| tier0  | True  |
| tier1  | True  |
| tier2  | True  |
| tier3  | True  |
+-----+-----+
```

This command enables automatic storage disk configuration for all storage tiers.

5.11 Managing Alerts

5.11.1 vinfra cluster alert list

List alert log entries:

```
usage: vinfra cluster alert list [--all]
```

--all

Show both enabled and disabled alerts

Example:

```
# vinfra cluster alert list --all
+-----+-----+-----+-----+
| id | type           | datetime           | severity | enabled |
+-----+-----+-----+-----+
| 1  | Network warning | 2018-08-30T18:02:14 | warning  | True    |
| 2  | Network warning | 2018-08-30T18:02:14 | warning  | True    |
| 3  | Network warning | 2018-08-30T18:02:14 | warning  | True    |
| 4  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 5  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 6  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 7  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 8  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 9  | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 10 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 11 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 12 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 13 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 14 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
| 15 | Network warning | 2018-08-31T13:02:15 | warning  | True    |
+-----+-----+-----+-----+
```

This command lists all alerts in the log and shows whether they are enabled or disabled.

5.11.2 vinfra cluster alert show

Show details of the specified alert log entry:

```
usage: vinfra cluster alert show <alert>
```

<alert>

Alert ID

Example:

```
# vinfra cluster alert show 1
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| _type      | undefined_speed                          |
| cluster_id |                                           |
| cluster_name |                                           |
| datetime   | 2018-08-30T18:02:14.855302+00:00        |
| details    | host: stor-1.example.com.vstoragedomain. |
| enabled    | True                                      |
| group      | node                                      |
| host       | stor-1.example.com.vstoragedomain.      |
| id         | 1                                         |
| message    | Network interface "eth1" on node "stor-1.example.com.vstoragedomain." has |
|            | an undefined speed                       |
| node_id    | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55    |
| object_id  | eth1                                      |
| severity   | warning                                   |
| suspended  |                                           |
| type       | Network warning                          |
+-----+-----+
```

This command shows the details of alert with ID 1.

5.11.3 vinfra cluster alert delete

Remove an entry from the alert log:

```
usage: vinfra cluster alert delete <alert>
```

<alert>

Alert ID

Example:

```
# vinfra cluster alert delete 1
+-----+-----+
| Field      | Value                                     |
+-----+-----+
| _type      | undefined_speed                          |
| cluster_id |                                           |
| cluster_name |                                           |
| datetime   | 2018-08-30T18:02:14.855302+00:00        |
| details    | host: stor-1.example.com.vstoragedomain. |
| enabled    | True                                      |
| group      | node                                      |
| host       | stor-1.example.com.vstoragedomain.      |
+-----+-----+
```

```

| id          | 1 |
| message     | Network interface "eth1" on node "stor-1.example.com.vstoragedomain." has an |
|             | undefined speed |
| node_id     | 4f96acf5-3bc8-4094-bcb6-4d1953be7b55 |
| object_id   | eth1 |
| severity    | warning |
| suspended   | |
| type        | Network warning |
+-----+-----+

```

This command deletes the alert with the ID 1 from the log.

5.12 Managing Audit Log

5.12.1 vinfra cluster auditlog list

List all audit log entries:

```
usage: vinfra cluster auditlog list
```

Example:

```

# vinfra cluster auditlog list
+-----+-----+-----+-----+-----+
| id | username | type | activity | timestamp |
+-----+-----+-----+-----+-----+
| 1 | admin | LoginUser | User login | 2018-09-07T08:33:44 |
| 2 | admin | ChangeNetwokInterface | Configure network | 2018-09-07T09:53:58 |
| 3 | admin | UpInterface | Bring up interface | 2018-09-07T09:54:44 |
| 4 | admin | ChangeNetwokInterface | Configure network | 2018-09-07T09:54:54 |
| 5 | admin | CreateBonding | Create bonding | 2018-09-07T09:57:24 |
| 17 | admin | RemoveNode | Forget node | 2018-09-07T12:21:59 |
| 14 | admin | RemoveNetworkIface | Delete interface | 2018-09-07T12:17:14 |
| 15 | admin | RemoveNode | Forget node | 2018-09-07T12:17:49 |
| 6 | admin | UpInterface | Bring up interface | 2018-09-07T10:59:28 |
| 7 | admin | ChangeNetwokInterface | Configure network | 2018-09-07T10:59:46 |
| 9 | admin | UpInterface | Bring up interface | 2018-09-07T11:42:29 |
| 10 | admin | UpInterface | Bring up interface | 2018-09-07T11:42:42 |
| 11 | admin | CreateBonding | Create bonding | 2018-09-07T11:43:46 |
| 12 | admin | ChangeNetwokInterface | Configure network | 2018-09-07T11:52:17 |
| 13 | admin | ChangeNetwokInterface | Configure network | 2018-09-07T11:52:44 |
| 16 | admin | RemoveNode | Forget node | 2018-09-07T12:21:51 |
| 8 | admin | CreateBonding | Create bonding | 2018-09-07T11:00:39 |
| 18 | admin | RemoveNode | Forget node | 2018-09-07T12:22:08 |
| 19 | admin | UpInterface | Bring up interface | 2018-09-07T12:33:16 |
| 20 | admin | CreateVLAN | Create VLAN | 2018-09-07T12:34:18 |
| 21 | admin | RemoveNetworkIface | Delete interface | 2018-09-07T13:26:40 |

```

```

| 22 | admin | LoginUser | User login | 2018-09-07T14:50:06 |
| 23 | admin | LoginUser | User login | 2018-09-07T14:51:34 |
| 24 | admin | CreateNetworkRolesSet | Create custom role set | 2018-09-07T15:06:03 |
| 25 | admin | ChangeNetworkRolesSet | Configure custom role set | 2018-09-07T15:37:50 |
| 26 | admin | RemoveNetworkRolesSet | Remove custom role set | 2018-09-07T15:39:31 |
| 27 | admin | CreateNetworkRole | Create custom role | 2018-09-07T15:58:50 |
| 28 | admin | RemoveNetworkRole | Remove custom role | 2018-09-07T16:20:22 |
+-----+-----+-----+-----+-----+

```

This command lists the audit log entries.

5.12.2 vinfra cluster auditlog show

Show details of an audit log entry:

```
usage: vinfra cluster auditlog show <auditlog>
```

<auditlog>

Audit log ID

Example:

```

# vinfra cluster auditlog show 1
+-----+-----+
| Field | Value |
+-----+-----+
| activity | User login |
| cluster_id | |
| cluster_name | |
| component | Users |
| details | [] |
| id | 1 |
| message | User "admin" login |
| node_id | |
| result | success |
| session_id | 817a19beaf244f92604fbf4b40af2c29 |
| task_id | 5686556295049300 |
| timestamp | 2018-09-07T08:33:44.175797+00:00 |
| type | LoginUser |
| username | admin |
+-----+-----+

```

This command shows the details of the audit log entry with the ID 1.

5.13 Sending Problem Reports

Generate and send a problem report:

```
usage: vinfra cluster problem-report [--email <email>]
                                     [--description <description>] [--send]
```

--email <email>

Contact email address

--description <description>

Problem description

--send

Generate the problem report archive and send it to the technical support team

Example:

```
# vinfra cluster problem-report --email test@example.com --description "Test report" --send
+-----+-----+
| Field | Value |
+-----+-----+
| task_id | 8bcfb92f-f02b-4de8-8e44-3426047630e3 |
+-----+-----+
```

This command creates a task to send a problem report with the description "Test report" to the technical support team and use `test@example.com` as a reply-to address. Note the problem report ID in the task details. You will need to mention it in the support ticket.

Task outcome:

```
+-----+-----+
| Field | Value |
+-----+-----+
| details | |
| name | backend.presentation.reports.tasks.ReportProblemTask |
| result | id: '1001923113' |
| | path: /var/cache/problem-reports/report-2018-12-10T15:33:23.391329.tar.gz |
| state | success |
| task_id | 37d5c13a-001c-4789-8242-96825a17deda |
+-----+-----+
```

CHAPTER 6

Monitoring Storage Cluster

Monitoring the storage cluster is very important because it allows you to check the status and health of all computers in the cluster and react as necessary.

The main command for monitoring is `vstorage -c <cluster_name> top`. It invokes a text user interface that you can control with keys (press **h** for help).

6.1 Monitoring General Storage Cluster Parameters

By monitoring general parameters, you can get detailed information about all components of the storage cluster, its overall status and health. To display this information, use the `vstorage -c <cluster_name> top` command. For example:


```

Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
IO:      read   0B/s ( 0ops/s), write  0B/s ( 0ops/s)

```

MDSID	STATUS	%CTIME	COMMITTS	%CPU	MEM	UPTIME	HOST
M 3	avail	0.0%	0/s	1.1%	192m	19d 23h	management.655c19da7e854d6f.nodes.svc.vstoragedomain:2510
1	avail	0.0%	0/s	0.2%	192m	20d 0h	management.b2823b72aeff4ddb.nodes.svc.vstoragedomain:2510
2	avail	0.0%	0/s	0.0%	192m	19d 23h	management.bda1f22b3a854b6c.nodes.svc.vstoragedomain:2510

CSID	STATUS	SPACE	AVAIL	REPLICAS	UNIQUE	IOWAIT	IOLAT (ms)	QDEPTH	HOST
1027	active	492.0G	451.4G	295	12	0%	0/0	0.0	management.655c19da7e854d6f.nodes.svc.v
1025	active	492.0G	449.5G	305	22	0%	0/0	0.0	management.b2823b72aeff4ddb.nodes.svc.v
1026	active	492.0G	453.0G	289	6	0%	0/0	0.0	management.bda1f22b3a854b6c.nodes.svc.v

CLID	LEASES	READ	WRITE	RD OPS	WR OPS	FSYNCS	IOLAT (ms)	HOST
2050	1/222	6B/s	6B/s	0ops/s	0ops/s	0ops/s	0.13/1	management.b2823b72aeff4ddb.nodes.
2226	1/2	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.bda1f22b3a854b6c.nodes.
2142	0/0	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.655c19da7e854d6f.nodes.

TIME	SYS	SEV	MESSAGE
21-12-18 12:06:24	MDS	INF	Add new MDS#3 at 10.37.130.79:2510 by request from 10.37.130.79:45672
21-12-18 12:06:24	MON	INF	MDS#3 was started
21-12-18 12:06:35	MON	INF	MDS#3 was stopped
21-12-18 12:06:35	MON	INF	CS#1027 was started
21-12-18 12:06:36	MDS	INF	New CS#1027 at 10.37.130.79:45742 (0.0.0.655c19da7e854d6f), tier=0
21-12-18 12:06:36	MON	INF	MDS#3 was started
21-12-18 12:06:38	MDS	INF	CS#1027 is active
21-12-18 12:06:45	MDS	INF	The cluster physical free space: 1.4Tb (99%), total 1.4Tb

The command above shows detailed information about the stor1 cluster. The general parameters (highlighted in red) are as follows.

Cluster

Overall status of the cluster:

Healthy

All chunk servers in the cluster are active.

Unknown

There is not enough information about the cluster state (e.g., because the master MDS server was elected a while ago).

Degraded

Some of the chunk servers in the cluster are inactive.

Failure

The cluster has too many inactive chunk servers; the automatic replication is disabled.

SMART warning

One or more physical disks attached to cluster nodes are in pre-failure condition. For details, see *Monitoring Physical Disks* (page 208).

Space

Amount of disk space in the cluster:

Free Free physical disk space in the cluster.

Allocatable

Amount of logical disk space available to clients. Allocatable disk space is calculated on the basis of the current replication parameters and free disk space on chunk servers. It may also be limited by license.

Note: For more information on monitoring and understanding disk space usage in clusters, see *Understanding Disk Space Usage* (page 202).

MDS nodes

Number of active MDS servers as compared to the total number of MDS servers configured for the cluster.

Epoch time

Time elapsed since the MDS master server election.

CS nodes

Number of active chunk servers as compared to the total number of chunk servers configured for the cluster.

In parentheses, you can see the additional information on these chunk servers:

- Active chunk servers (avail.) that are currently up and running in the cluster.
- Inactive chunk servers (inactive) that are temporarily unavailable. A chunk server is marked as inactive during its first 5 minutes of inactivity.
- Offline chunk servers (offline) that have been inactive for more than 5 minutes. A chunk server changes its state to offline after 5 minutes of inactivity. Once the state is changed to offline, the cluster starts replicating data to restore the chunks that were stored on the offline chunk server.

License

Key number under which the license is registered on the Key Authentication server and license state.

Replication

Replication settings. The normal number of chunk replicas and the limit after which a chunk gets blocked until recovered.

IO Disk IO activity in the cluster:

- Speed of read and write I/O operations, in bytes per second.
- Number of read and write I/O operations per second.

6.2 Monitoring Metadata Servers

MDS servers are a critical component of any storage cluster, and monitoring the health and state of MDS servers is a crucial task. To monitor MDS servers, use the `vstorage -c <cluster_name> top` command. For example:

```
Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
IO:      read      0B/s ( 0ops/s), write      0B/s ( 0ops/s)
```

MDSID	STATUS	%CTIME	COMMITTS	%CPU	MEM	UPTIME	HOST
M 3	avail	0.0%	0/s	1.1%	192m	19d 23h	management.655c19da7e854d6f.nodes.svc.vstoragedomain:2510
1	avail	0.0%	0/s	0.2%	192m	20d 0h	management.b2823b72aeff4ddb.nodes.svc.vstoragedomain:2510
2	avail	0.0%	0/s	0.0%	192m	19d 23h	management.bda1f22b3a854b6c.nodes.svc.vstoragedomain:2510

CSID	STATUS	SPACE	AVAIL	REPLICAS	UNIQUE	IOWAIT	IOLAT (ms)	QDEPTH	HOST
1027	active	492.0G	451.4G	295	12	0%	0/0	0.0	management.655c19da7e854d6f.nodes.svc.v
1025	active	492.0G	449.5G	305	22	0%	0/0	0.0	management.b2823b72aeff4ddb.nodes.svc.v
1026	active	492.0G	453.0G	289	6	0%	0/0	0.0	management.bda1f22b3a854b6c.nodes.svc.v

CLID	LEASES	READ	WRITE	RD OPS	WR OPS	FSYNCS	IOLAT (ms)	HOST
2050	1/222	6B/s	6B/s	0ops/s	0ops/s	0ops/s	0.13/1	management.b2823b72aeff4ddb.nodes.svc.v
2226	1/2	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.bda1f22b3a854b6c.nodes.svc.v
2142	0/0	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.655c19da7e854d6f.nodes.svc.v

TIME	SYS	SEV	MESSAGE
21-12-18 12:06:24	MDS	INF	Add new MDS#3 at 10.37.130.79:2510 by request from 10.37.130.79:45672
21-12-18 12:06:24	MON	INF	MDS#3 was started
21-12-18 12:06:35	MON	INF	MDS#3 was stopped
21-12-18 12:06:35	MON	INF	CS#1027 was started
21-12-18 12:06:36	MDS	INF	New CS#1027 at 10.37.130.79:45742 (0.0.0.655c19da7e854d6f), tier=0
21-12-18 12:06:36	MON	INF	MDS#3 was started
21-12-18 12:06:38	MDS	INF	CS#1027 is active
21-12-18 12:06:45	MDS	INF	The cluster physical free space: 1.4Tb (99%), total 1.4Tb

The command above shows detailed information about the `stor1` cluster. The monitoring parameters for MDS servers (highlighted in red) are as follows:

MDSID

MDS server identifier (ID).

The letter “M” before ID, if present, means that the given server is the master MDS server.

STATUS

MDS server status.

%CTIME

Total time the MDS server spent writing to the local journal.

COMMITTS

Local journal commit rate.

%CPU

MDS server activity time.

MEM

Amount of physical memory the MDS server uses.

UPTIME

Time elapsed since the last MDS server start.

HOST

MDS server hostname or IP address.

6.3 Monitoring Chunk Servers

By monitoring chunk servers, you can keep track of the disk space available in the storage cluster. To monitor chunk servers, use the `vstorage -c <cluster_name> top` command. For example:

```

Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
IO:      read    0B/s ( 0ops/s), write    0B/s ( 0ops/s)

```

MDSID	STATUS	%TIME	COMMITTS	%CPU	MEM	UPTIME	HOST
M 3	avail	0.0%	0/s	1.1%	192m	19d 23h	management.655c19da7e854d6f.nodes.svc.vstoragedomain:2510
1	avail	0.0%	0/s	0.2%	192m	20d 0h	management.b2823b72aeff4ddb.nodes.svc.vstoragedomain:2510
2	avail	0.0%	0/s	0.0%	192m	19d 23h	management.bda1f22b3a854b6c.nodes.svc.vstoragedomain:2510

CSID	STATUS	SPACE	AVAIL	REPLICAS	UNIQUE	IOWAIT	IOLAT (ms)	QDEPTH	HOST
1027	active	492.0G	451.4G	295	12	0%	0/0	0.0	management.655c19da7e854d6f.nodes.svc.v
1025	active	492.0G	449.5G	305	22	0%	0/0	0.0	management.b2823b72aeff4ddb.nodes.svc.v
1026	active	492.0G	453.0G	289	6	0%	0/0	0.0	management.bda1f22b3a854b6c.nodes.svc.v

CLID	LEASES	READ	WRITE	RD OPS	WR OPS	FSYNCS	IOLAT (ms)	HOST
2050	1/222	6B/s	6B/s	0ops/s	0ops/s	0ops/s	0.13/1	management.b2823b72aeff4ddb.nodes.svc.v
2226	1/2	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.bda1f22b3a854b6c.nodes.svc.v
2142	0/0	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.655c19da7e854d6f.nodes.svc.v

TIME	SYS	SEV	MESSAGE
21-12-18 12:06:24	MDS	INF	Add new MDS#3 at 10.37.130.79:2510 by request from 10.37.130.79:45672
21-12-18 12:06:24	MON	INF	MDS#3 was started
21-12-18 12:06:35	MON	INF	MDS#3 was stopped
21-12-18 12:06:35	MON	INF	CS#1027 was started
21-12-18 12:06:36	MDS	INF	New CS#1027 at 10.37.130.79:45742 (0.0.0.655c19da7e854d6f), tier=0
21-12-18 12:06:36	MON	INF	MDS#3 was started
21-12-18 12:06:38	MDS	INF	CS#1027 is active
21-12-18 12:06:45	MDS	INF	The cluster physical free space: 1.4Tb (99%), total 1.4Tb

The command above shows detailed information about the stor1 cluster. The monitoring parameters for chunk servers (highlighted in red) are as follows:

CSID Chunk server identifier (ID).

STATUS

Chunk server status:

active

The chunk server is up and running.

failed

The chunk server process is running but a problem has occurred with the CS disk.

inactive

The chunk server is temporarily unavailable. A chunk server is marked as inactive during its first 5 minutes of inactivity.

offline

The chunk server is inactive for more than 5 minutes. After the chunk server goes offline, the cluster starts replicating data to restore the chunks that were stored on the affected chunk server.

dropped

The chunk server was removed by the administrator.

maintenance

The node that the chunk server is located on is in maintenance.

SPACE

Total amount of disk space on the chunk server.

AVAIL

Available disk space on the chunk server.

REPLICAS

Number of replicas stored on the chunk server.

UNIQUE

Number of chunks that do not have replicas.

IOWAIT

Percentage of time spent waiting for I/O operations being served.

IOLAT

Average/maximum time, in milliseconds, the client needed to complete a single IO operation during the last 20 seconds.

QDEPTH

Average chunk server I/O queue depth.

HOST

Chunk server hostname or IP address.

FLAGS

The following flags may be shown for active chunk servers:

- J** The CS uses a write journal.
- C** Checksumming is enabled for the CS. Checksumming lets you know when a third party changes the data on the disk.
- D** Direct I/O, the normal state for a CS without a write journal.
- c** The chunk server's write journal is clean, there is nothing to commit from the write journaling SSD to the HDD where the CS is located.

6.3.1 Understanding Disk Space Usage

Usually, you get the information on how disk space is used in your cluster with the `vstorage top` command. This command displays the following disk-related information: total space, free space, and allocatable space. For example:

```
# vstorage -c stor1 top
connected to MDS#1
Cluster 'stor1': healthy
Space: [OK] allocatable 180GB of 200GB, free 1.6TB of 1.7TB
...
```

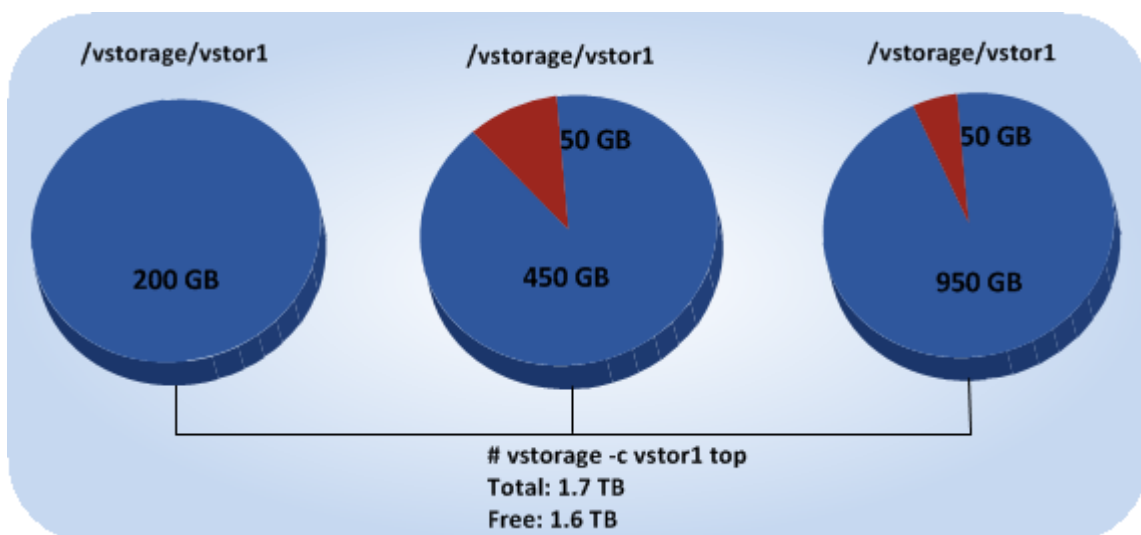
In this command output:

- 1.7TB is the total disk space in the `stor1` cluster. The total disk space is calculated on the basis of used and free disk space on all partitions in the cluster. Used disk space includes the space occupied by all data chunks and their replicas plus the space occupied by any other files stored on the cluster partitions.

Let us assume that you have a 100 GB partition and 20 GB on this partition are occupied by some files. Now if you set up a chunk server on this partition, this will add 100 GB to the total disk space of the cluster, though only 80 GB of this disk space will be free and available for storing data chunks.

- 1.6TB is the free disk space in the `stor1` cluster. Free disk space is calculated by subtracting the disk space occupied by data chunks and any other files on the cluster partitions from the total disk space.

For example, if the amount of free disk space is 1.6 TB and the total disk space is 1.7 TB, this means that about 100 GB on the cluster partitions are already occupied by some files.




- `allocatable 180GB of 200GB` is the amount of free disk space that can be used for storing data chunks. See **Understanding allocatable disk space** below for details.

6.3.1.1 Understanding Allocatable Disk Space

When monitoring disk space information in the cluster, you also need to pay attention to the space reported by the `vstorage top` utility as *allocatable*. Allocatable space is the amount of disk space that is free and can be used for storing user data. Once this space runs out, no data can be written to the cluster.

Calculation of allocatable disk space is illustrated on the following example:

- The cluster has 3 chunk servers. The first chunk server has 200 GB of disk space, the second one — 500 GB, and the third one — 1 TB.
- The default replication factor of 3 is used in the cluster, meaning that each data chunk must have 3 replicas stored on three different chunk servers.



In this example, the available disk space is 200 GB, which equals the amount of disk space on the smallest chunk server:

```
# vstorage -c stor1 top
connected to MDS#1
Cluster 'stor1': healthy
```



```
Space: [OK] allocatable 180GB of 200GB, free 1.6TB of 1.7TB
...
```

In this cluster configuration each server is set to store one replica for each data chunk. So once the disk space on the smallest chunk server (200 GB) runs out, no more chunks in the cluster can be created until a new chunk server is added or the replication factor is decreased.

If the replication factor changes to 2, the `vstorage top` command will report the available disk space as 700 GB:

```
# vstorage set-attr -R /mnt/vstorage replicas=2:1
# vstorage -c stor1 top
connected to MDS#1
Cluster 'stor1': healthy
Space: [OK] allocatable 680GB of 700GB, free 1.6TB of 1.7TB
...
```

The available disk space has increased because now only 2 replicas are created for each data chunk and new chunks can be made even if the smallest chunk server runs out of space (in this case, replicas will be stored on a bigger chunk server).

Allocatable disk space may also be limited by license.

6.3.1.2 Viewing Space Occupied by Data Chunks

To view the total amount of disk space occupied by all user data in the cluster, run the `vstorage top` command and press the V key on your keyboard. Once you do this, your command output should look like the following:

```
# vstorage -c stor1 top
Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h, cluster version: 128
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline), storage version: 128
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
Chunks: [OK] 323 (100%) healthy, 0 (0%) standby, 0 (0%) degraded, 0 (0%) urgent,
        0 (0%) blocked, 0 (0%) pending, 0 (0%) offline, 0 (0%) replicating,
        0 (0%) overcommitted, 0 (0%) deleting, 0 (0%) void
FS: 20.3GB in 757 files, 757 inodes, 244 file maps, 323 chunks, 889 chunk replicas
IO:      read   0B/s ( 0ops/s), write  0B/s ( 0ops/s)
IO total: read 37.1GB ( 473Kops), write 133.7GB ( 4.7Mops)
Repl IO: read   0B/s, write:  0B/s
Sync rate: 0ops/s, datasync rate: 0ops/s
IO QDEPTH: 0.0 aver, 0.0 max
...
```

The **FS** field shows the size of all user data in the cluster without consideration for replicas.

6.3.2 Exploring Chunk States

The following is a list of all possible chunk states.

Healthy

Number and percentage of chunks that have enough active replicas. The normal state of chunks.

Offline

Number and percentage of chunks all replicas of which are offline. Such chunks are completely inaccessible for the cluster and cannot be replicated, read from or written to. All requests to an offline chunk are frozen until a CS that stores that chunk's replica goes online.

Get offline chunk servers back online as fast as possible to avoid losing data.

Blocked

Number and percentage of chunks which have fewer active replicas than the set minimum amount. Write requests to a blocked chunk are frozen until it has at least the set minimum amount of replicas. Read requests to blocked chunks are allowed, however, as they still have some active replicas left. Blocked chunks have a higher replication priority than degraded chunks.

Having blocked chunks in the cluster increases the risk of losing data, so postpone any maintenance on working cluster nodes and get offline chunk servers back online as fast as possible.

Degraded

Number and percentage of chunks whose active replicas are few but not below the set minimum. Such chunks can be read from and written to. However, in the latter case a degraded chunk becomes urgent.

Replicating

Number and percentage of chunks which are being replicated. Write operations on such chunks are frozen until replication ends.

Void Number and percentage of chunks that have been allocated but never used yet. Such chunks contain no data. It is normal to have some void chunks in the cluster.

Pending

Number and percentage of chunks that must be replicated immediately. For a write request from client to a chunk to complete, the chunk must have at least the set minimum amount of replicas. If it does not, the chunk is blocked and the write request cannot be completed. As blocked chunks must be

replicated as soon as possible, the cluster places them in a special high-priority replication queue and reports them as pending.

Urgent

Number and percentage of chunks which are degraded and have non-identical replicas. Replicas of a degraded chunk may become non-identical if some of them are not accessible during a write operation. As a result, some replicas happen to have the new data while some still have the old data. The latter are dropped by the cluster as fast as possible. Urgent chunks do not affect information integrity as the actual data is stored in at least the set minimum amount of replicas.

Overcommitted

Number and percentage of chunks that have more replicas than normal. Usually these chunks appear after the normal number of replicas has been lowered or a lot of data has been deleted. Extra replicas are eventually dropped, however, this process may slow down during replication.

Deleting

Number and percentage of chunks queued for deletion.

6.4 Monitoring Clients

By monitoring clients, you can check the status and health of servers that you use to access virtual machines. To monitor clients, use the `vstorage -c <cluster_name> top` command. For example:

```

Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
IO:      read    0B/s ( 0ops/s), write    0B/s ( 0ops/s)

```

MDSID	STATUS	%CTIME	COMMITTS	%CPU	MEM	UPTIME	HOST
M 3	avail	0.0%	0/s	1.1%	192m	19d 23h	management.655c19da7e854d6f.nodes.svc.vstoragedomain:2510
1	avail	0.0%	0/s	0.2%	192m	20d 0h	management.b2823b72aeff4ddb.nodes.svc.vstoragedomain:2510
2	avail	0.0%	0/s	0.0%	192m	19d 23h	management.bda1f22b3a854b6c.nodes.svc.vstoragedomain:2510

CSID	STATUS	SPACE	AVAIL	REPLICAS	UNIQUE	IOWAIT	IOLAT (ms)	QDEPTH	HOST
1027	active	492.0G	451.4G	295	12	0%	0/0	0.0	management.655c19da7e854d6f.nodes.svc.v
1025	active	492.0G	449.5G	305	22	0%	0/0	0.0	management.b2823b72aeff4ddb.nodes.svc.v
1026	active	492.0G	453.0G	289	6	0%	0/0	0.0	management.bda1f22b3a854b6c.nodes.svc.v

CLID	LEASES	READ	WRITE	RD OPS	WR OPS	FSYNCS	IOLAT (ms)	HOST
2050	1/222	6B/s	6B/s	0ops/s	0ops/s	0ops/s	0.13/1	management.b2823b72aeff4ddb.nodes.
2226	1/2	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.bda1f22b3a854b6c.nodes.
2142	0/0	0B/s	0B/s	0ops/s	0ops/s	0ops/s	0/0	management.655c19da7e854d6f.nodes.

TIME	SYS	SEV	MESSAGE
21-12-18 12:06:24	MDS	INF	Add new MDS#3 at 10.37.130.79:2510 by request from 10.37.130.79:45672
21-12-18 12:06:24	MON	INF	MDS#3 was started
21-12-18 12:06:35	MON	INF	MDS#3 was stopped
21-12-18 12:06:35	MON	INF	CS#1027 was started
21-12-18 12:06:36	MDS	INF	New CS#1027 at 10.37.130.79:45742 (0.0.0.655c19da7e854d6f), tier=0
21-12-18 12:06:36	MON	INF	MDS#3 was started
21-12-18 12:06:38	MDS	INF	CS#1027 is active
21-12-18 12:06:45	MDS	INF	The cluster physical free space: 1.4Tb (99%), total 1.4Tb

The command above shows detailed information about the stor1 cluster. The monitoring parameters for clients (highlighted in red) are as follows.

CLID Client identifier (ID).

LEASES

Average number of files opened for reading/writing by the client and not yet closed, for the last 20 seconds.

READ

Average rate, in bytes per second, at which the client reads data, for the last 20 seconds.

WRITE

Average rate, in bytes per second, at which the client writes data, for the last 20 seconds.

RD_OPS

Average number of read operations the client made per second, for the last 20 seconds.

WR_OPS

Average number of write operations the client made per second, for the last 20 seconds.

FSYNCS

Average number of sync operations the client made per second, for the last 20 seconds.

IOLAT

Average/maximum time, in milliseconds, the client needed to complete a single IO operation, for the last 20 seconds.

HOST

Client hostname or IP address.

6.5 Monitoring Physical Disks

The S.M.A.R.T. status of physical disks is monitored by the `smartctl` tool installed along with Virtuozzo Hybrid Infrastructure. For it to work, S.M.A.R.T. functionality must be enabled in the node's BIOS. The tool is run every 10 minutes as a cron job also added during installation. The `smartctl` tool polls all physical disks attached to nodes in the cluster, including caching and journaling SSDs, and reports the results to the MDS server.

You can view disk poll results for the last 10 minutes in the output of the `vstorage top` command. For example:

```
Cluster 'stor1': healthy, SMART warning
Space: [OK] allocatable 100GB (+778GB unlicensed) of 926GB, free 924GB of 926GB
MDS nodes: 1 of 1, epoch uptime: 7d 22h
CS nodes: 2 of 2 (2 avail, 0 inactive, 0 offline)
Replication: 1 norm, 1 limit
IO:      read      0B/s ( 0ops/s), write      0B/s ( 0ops/s)

MDSID STATUS  %CTIME  COMMITS  %CPU  MEM  UPTIME  HOST
M   1 avail    0.0%    0/s     0.0%  48m   7d 22h pcs36.qa.sw.ru:2510

CSID STATUS  SPACE  AVAIL  REPLICAS  UNIQUE  IOWAIT  IOLAT(ms)  QDEPTH  HOST
1025 active   9.1GB  7.1GB    0         0       0%       0/0       0.0 pcs36.q
1026 active  916GB  870GB    0         0       0%       0/0       0.0 pcs36.q

CLID  LEASES  READ  WRITE  RD_OPS  WR_OPS  FSYNCS  IOLAT(ms)  HOST

TIME  SYS SEV MESSAGE
01-07-14 16:42:19 MON WRN CS#1026 was stopped
01-07-14 16:42:26 JRN INF MDS#1 at 10.29.2.16:2510 became master
01-07-14 16:42:26 MDS WRN License not installed, please add license using comma
01-07-14 16:42:29 MON WRN MDS#1 was stopped
01-07-14 16:42:44 MDS INF CS#1025, CS#1026 are active
01-07-14 16:42:53 MDS INF The cluster is healthy with 2 active CS
01-07-14 16:42:53 MDS INF The cluster physical free space: 925.0Gb (99%), total
```

If the **SMART warning** message is shown in the main table, one of the physical disks is in pre-failure condition according to S.M.A.R.T. Press **d** to switch to the disks table to see more details. For example:

```
Cluster 'stor1': healthy, SMART warning
Space: [OK] allocatable 100GB (+778GB unlicensed) of 926GB, free 924GB of 926GB
MDS nodes: 1 of 1, epoch uptime: 7d 22h
CS nodes: 2 of 2 (2 avail, 0 inactive, 0 offline)
Replication: 1 norm, 1 limit
IO:      read      0B/s ( 0ops/s), write      0B/s ( 0ops/s)
```

DISK	SMART	TEMP	CAPACITY	SERIAL	MODEL	HOST
sdc	OK	27C	931GB	1374X80PS	TOSHIBA DT01ACA100	pcs36.qa
sde	warn	31C	931GB	MSE5235U36ZHJU	Hitachi HDS721010DLE630	pcs36.qa

The disks table shows the following parameters:

DISK Disk name assigned by operating system.

SMART

Disk's S.M.A.R.T. status:

OK The disk is healthy.

Warn

The disk is in pre-failure condition. Pre-failure condition means that at least one of these S.M.A.R.T. counters is nonzero:

- Reallocated Sector Count
- Reallocated Event Count
- Current Pending Sector Count
- Offline Uncorrectable

TEMP

Disk temperature in Celsius.

CAPACITY

Disk capacity.

SERIAL

Disk serial number.

MODEL

Disk model.

HOST

Disk's host address.

To disable S.M.A.R.T. disk monitoring, delete the corresponding cron job.

6.6 Monitoring Event Logs

You can use the `vstorage -c <cluster_name> top` utility to monitor significant events happening in the storage cluster. For example:

```
Cluster 'stor1': healthy
Space: [OK] allocatable 1.32TB of 1.44TB, free 1.39TB of 1.44TB
MDS nodes: 3 of 3, epoch uptime: 19d 23h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 1 norm, 1 limit
IO:      read    0B/s ( 0ops/s), write    0B/s ( 0ops/s)

MDSID STATUS  %CTIME  COMMITS  %CPU  MEM  UPTIME  HOST
M   3 avail    0.0%    0/s    1.1% 192m 19d 23h management.655c19da7e854d6f.nodes.svc.vstoragedomain:2510
   1 avail    0.0%    0/s    0.2% 192m 20d 0h management.b2823b72aeff4ddb.nodes.svc.vstoragedomain:2510
   2 avail    0.0%    0/s    0.0% 192m 19d 23h management.bda1f22b3a854b6c.nodes.svc.vstoragedomain:2510

CSID STATUS  SPACE  AVAIL  REPLICAS  UNIQUE  IOWAIT  IOLAT (ms)  QDEPTH  HOST
1027 active  492.0G 451.4G   295      12     0%      0/0     0.0 management.655c19da7e854d6f.nodes.svc.v
1025 active  492.0G 449.5G   305      22     0%      0/0     0.0 management.b2823b72aeff4ddb.nodes.svc.v
1026 active  492.0G 453.0G   289       6     0%      0/0     0.0 management.bda1f22b3a854b6c.nodes.svc.v

CLID  LEASES  READ  WRITE  RD OPS  WR OPS  FSYNCS  IOLAT (ms)  HOST
2050  1/222  6B/s  6B/s  0ops/s  0ops/s  0ops/s  0.13/1 management.b2823b72aeff4ddb.nodes.v
2226  1/2    0B/s  0B/s  0ops/s  0ops/s  0ops/s  0/0 management.bda1f22b3a854b6c.nodes.v
2142  0/0    0B/s  0B/s  0ops/s  0ops/s  0ops/s  0/0 management.655c19da7e854d6f.nodes.v

TIME          SYS SEV MESSAGE
21-12-18 12:06:24 MDS INF Add new MDS#3 at 10.37.130.79:2510 by request from 10.37.130.79:45672
21-12-18 12:06:24 MON INF MDS#3 was started
21-12-18 12:06:35 MON INF MDS#3 was stopped
21-12-18 12:06:35 MON INF CS#1027 was started
21-12-18 12:06:36 MDS INF New CS#1027 at 10.37.130.79:45742 (0.0.0.655c19da7e854d6f), tier=0
21-12-18 12:06:36 MON INF MDS#3 was started
21-12-18 12:06:38 MDS INF CS#1027 is active
21-12-18 12:06:45 MDS INF The cluster physical free space: 1.4Tb (99%), total 1.4Tb
```

The command above shows the latest events in the `stor1` cluster. The information on events (highlighted in red) is given in a table with the following columns:

TIME

Time of event.

SYS Component of the cluster where the event happened (e.g., MDS for an MDS server or JRN for local journal).

SEV Event severity.

MESSAGE

Event description.

The following table lists basic events displayed when you run the `vstorage top` utility.

Table 6.6.1: Basic events

Event	Severity	Description
MDS#<N> (<addr>:<port>) lags behind for more than 1000 rounds	JRN err	Generated by the MDS master server when it detects that MDS#<N> is stale. This message may indicate that some MDS server is very slow and lags behind.
MDS#<N> (<addr>:<port>) didn't accept commits for <i>M</i> sec	JRN err	Generated by the MDS master server if MDS#<N> did not accept commits for <i>M</i> seconds. MDS#<N> gets marked as stale. This message may indicate that the MDS service on MDS#<N> is experiencing a problem. The problem may be critical and should be resolved as soon as possible.
MDS#<N> (<addr>:<port>) state is outdated and will do a full resync	JRN err	Generated by the MDS master server when MDS#<N> will do a full resync. MDS#<N> gets marked as stale. This message may indicate that some MDS server was too slow or disconnected for such a long time that it is not really managing the state of metadata and has to be resynchronized. The problem may be critical and should be resolved as soon as possible.
MDS#<N> at <addr>:<port> became master	JRN info	Generated every time a new MDS master server is elected in the cluster. Frequent changes of MDS masters may indicate poor network connectivity and may affect the cluster operation.
The cluster is healthy with <i>N</i> active CS	MDS info	Generated when the cluster status changes to healthy or when a new MDS master server is elected. This message indicates that all chunk servers in the cluster are active and the number of replicas meets the set cluster requirements.

Continued on next page

Table 6.6.1 – continued from previous page

Event	Severity	Description
The cluster is degraded with N active, M inactive, K offline CS	MDS warn	Generated when the cluster status changes to degraded or when a new MDS master server is elected. This message indicates that some chunk servers in the cluster are <ul style="list-style-type: none"> • inactive, i.e. do not send any registration messages, or • offline, i.e. have been inactive for longer than <code>mds.wd.offline_tout</code>, which is 5 min by default.
The cluster failed with N active, M inactive, K offline CS (<code>mds.wd.max_offline_cs=<n></code>)	MDS err	Generated when the cluster status changes to failed or when a new MDS master server is elected. This message indicates that the number of offline chunk servers exceeds <code>mds.wd.max_offline_cs</code> , which is 2 by default. When the cluster fails, the automatic replication is not scheduled any more. So the cluster administrator must take action to either repair failed chunk servers or increase <code>mds.wd.max_offline_cs</code> . Setting this value to 0 disables the failed mode completely.
The cluster is filled up to $<N>\%$	MDS info/warn	Shows the current space usage in the cluster. A warning is generated if the disk space consumption equals or exceeds 80%. It is important to have spare disk space for data replicas if one of the chunk servers fails.
Replication started, N chunks are queued	MDS info	Generated when the cluster starts automatic data replication to recover the missing replicas.
Replication completed	MDS info	Generated when the cluster finishes automatic data replication.

Continued on next page

Table 6.6.1 – continued from previous page

Event	Severity	Description
CS#<N> has reported hard error on <i>path</i>	MDS warn	Generated when the chunk server CS#<N> detects disk data corruption. You are recommended to check the hardware for errors and replace corrupted disks as soon as possible.
CS#<N> has not registered during the last <i>T</i> sec and is marked as inactive/offline	MDS warn	Generated when the chunk server CS#<N> has been unavailable for a while. In this case, the chunk server first gets marked as inactive. After 5 minutes, the state is changed to offline, which starts automatic replication of data to restore the replicas that were stored on the offline chunk server.
Failed to allocate <i>N</i> replicas for ' <i>path</i> ' by request from <addr>:<port> - <i>K</i> out of <i>M</i> chunks servers are available	MDS warn	Generated when the cluster cannot allocate chunk replicas, for example, when it runs out of disk space.
Failed to allocate <i>N</i> replicas for ' <i>path</i> ' by request from <addr>:<port> since only <i>K</i> chunk servers are registered	MDS warn	Generated when the cluster cannot allocate chunk replicas because not enough chunk servers are registered in the cluster.

6.7 Monitoring Replication Parameters

When you configure replication parameters, keep in mind that the new settings do not come into effect immediately. For example, increasing the default replication parameter for data chunks may take some time to complete, depending on the new value of this parameter and the number of data chunks in the cluster.

To check that the new replication parameters have been successfully applied to your cluster:

1. Run the `vstorage -c <cluster_name> top` command.
2. Press **V** to display additional information about the cluster. Typical command output may look like this:

```
# vstorage -c stor1 top
Cluster 'stor1': healthy
Space: [OK] allocatable 448.6GB of 492.0GB, free 1.39TB of 1.44TB
```

```
MDS nodes: 3 of 3, epoch uptime: 20d 0h
CS nodes: 3 of 3 (3 avail, 0 inactive, 0 offline)
License: ACTIVE (expiration: 01/10/2021, capacity: 10TB, used: 20.3GB)
Replication: 3 norm, 2 limit
Chunks: [Warning] 187 (57%) healthy, 0 (0%) standby, 0 (0%) degraded, 135 (41%) urgent,
        0 (0%) blocked, 0 (0%) pending, 0 (0%) offline, 1 (0%) replicating,
        0 (0%) overcommitted, 0 (0%) deleting, 0 (0%) void
IO:      read    0B/s ( 0ops/s), write 106KB/s ( 7ops/s)
...
```

3. Check the **Chunks** field for the following:

- When decreasing the replication parameters, look for chunks that are in the overcommitted or deleting state. If the replication process is complete, no chunks with these states should be present in the output.
- When increasing the replication parameters, look for chunks that are in the blocked or urgent state. If the replication process is complete, no chunks with these states should be present in the output. Besides, when the process is still in progress, the value of the healthy parameter is less than 100%.

For more information on available chunk statuses, see *Exploring Chunk States* (page 205).

CHAPTER 7

Accessing Storage Clusters via iSCSI

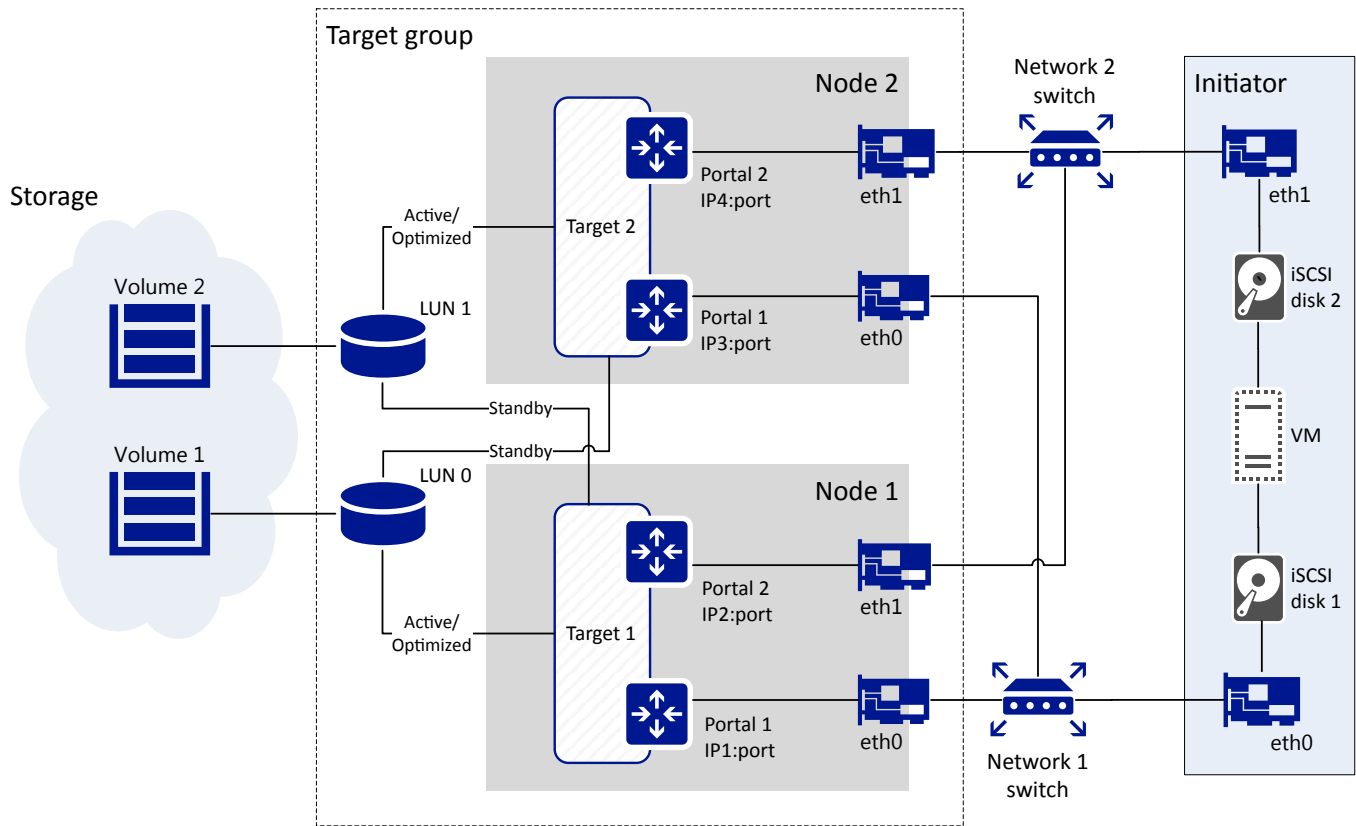
Virtuozzo Hybrid Infrastructure allows you to export cluster disk space to external operating systems and third-party virtualization solutions in the form of LUN block devices over iSCSI in a SAN-like manner.

In Virtuozzo Hybrid Infrastructure, you can create groups of redundant targets running on different storage nodes. To each target group you can attach multiple storage volumes with their own redundancy provided by the storage layer. These volumes are exported by targets as LUNs.

Each node in a target group can host a single target for that group if Ethernet is used, or one target per FC port if Fibre Channel is used. If one of the nodes in a target group fails along with its target(s), healthy targets from the same group continue to provide access to the LUNs previously serviced by the failed target(s).

You can create multiple target groups on same nodes. A volume, however, may only be attached to one target group at any moment of time.

The figure below shows a typical setup for exporting Virtuozzo Hybrid Infrastructure disk space via iSCSI.



The figure shows two volumes located on redundant storage provided by Virtuozzo Hybrid Infrastructure. The volumes are attached as LUNs to a group of two targets running on Virtuozzo Hybrid Infrastructure nodes. Each target has two portals, one per network interface with the iSCSI traffic type. This makes a total of four discoverable endpoints with different IP addresses. Each target provides access to all LUNs attached to the group.

Targets work in the ALUA mode, so one path to the volume is preferred and considered Active/Optimized while the other is Standby. The Active/Optimized path is normally chosen by the initiator (Explicit ALUA). If the initiator cannot do so (either does not support it or times out), the path is chosen by the storage itself (Implicit ALUA).

Network interfaces `eth0` and `eth1` on each node are connected to different switches for redundancy. The initiator, e.g., VMware ESXi, is connected to both switches as well and provides volumes as iSCSI disks 1 and 2 to a VM via different network paths.

If the Active/Optimized path becomes unavailable for some reason (e.g., the node with the target or network switch fails), the Standby path through the other target will be used instead to connect to the volume. When the Active/Optimized path is restored, it will be used again.

7.1 iSCSI Workflow Overview

The typical workflow of exporting volumes via iSCSI is as follows:

1. Assign the network with the traffic type **iSCSI** to a network interface on each node that you will add to a target group. See *Managing Traffic Types and Networks* (page 5).
2. Create a target group on chosen nodes, providing details for target WWNs and portals. Targets will be created automatically and added to the group. Target portals will be created on specified network interfaces and ports. See *Creating Target Groups* (page 218).
3. Create volumes and attach them to the target group. See *Managing Volumes* (page 111).
4. Optionally, enable CHAP authorization for the target group, create CHAP accounts, and assign them to the target group. See *Managing CHAP Accounts* (page 229).
5. Optionally, enable ACL authorization for the target group, create a list of initiators that will be allowed to access only specific LUNs. Initiators not on the list will be able to access all LUNs in the target group. See *Managing LUN Views* (page 230).
6. Start the target group. See *Starting and Stopping Target Groups* (page 220).
7. Connect initiators to targets using standard tools of your operating system or product, e.g., `iscsiadm`. Use the `vstorage-target session-list` command to view iSCSI sessions active on a node in a target group.

7.2 Configuring CLI Tool

Before you can use the `vstorage-target` CLI tool to export volumes via iSCSI, set it up as described further. Perform these steps on each node where you plan to run iSCSI targets.

1. Create a configuration file `/etc/vstorage/iscsi/config.json` with at least these mandatory parameters:

```
{
  "ClusterName": "cluster1",
  "VolumesRoot": "/vols/iscsi/vols",
}
```

Where `ClusterName` is the name of your storage cluster and `VolumesRoot` is the path to the directory for iSCSI volumes.

You can also set these optional parameters:

- "PcsLogLevel", log level, ranges from 1 (log errors only) to 7 (log all, including debug messages).
- "LogPath", path to log files, the default is "/var/log/vstorage" (the log will be saved to vstorage-target.log).
- "GetTimeout", the timeout for the initiator's command to read target port group status, the default is 3000 ms.

1. Enable the target monitor service:

```
# systemctl start vstorage-target-monitor.service
# systemctl enable vstorage-target-monitor.service
```

2. Create the iSCSI volume directory if it does not exist:

```
# mkdir -p /mnt/vstorage/vols/iscsi/
```

If you modify the configuration file later, restart the TCM monitor service to apply changes:

```
# systemctl restart vstorage-target-monitor.service
```

7.3 Managing Target Groups

This section explains how to create and manage groups of iSCSI targets.

7.3.1 Creating Target Groups

Before you create any target groups, assign the network with the iSCSI traffic type to a network interface on each node that you will add to a target group.

To create a target group, you will need a configuration file with a list of nodes to add to the group as well as target WWNs and portals. For example:

```
[
  {
    "NodeId": "01baeabee73e4a0d",
    "WWN": "iqn.2013-10.com.vstorage:test1",
    "Portals": [
      {
        "Addr": "192.168.10.11",
        "Port": 3025
      }
    ]
  }
],
```

```

{
  "NodeId": "0d90158e9d2444e1",
  "WWN": "iqn.2013-10.com.vstorage:test2",
  "Portals": [
    {
      "Addr": "192.168.10.12",
      "Port": 3025
    }
  ]
},
{
  "NodeId": "a9eca47661a64031",
  "WWN": "iqn.2013-10.com.vstorage:test3",
  "Portals": [
    {
      "Addr": "192.168.10.13",
      "Port": 3025
    }
  ]
}
]

```

In this configuration file:

- `NodeId` is a node identifier that you can obtain from `/etc/vstorage/host_id` on a node.
- `WWN` is a target world wide name:
 - an IQN if iSCSI protocol is used, e.g., `iqn.2013-10.com.vstorage:test1` (you can only customize the last part after the colon), or
 - a WWPN in NAA format if Fibre Channel protocol is used, e.g., `naa.21000024ff586d3b` (you can obtain the port number from `/sys/class/fc_host/host6/port_name`).
- `Portals` is one or more target portals, IP address and port combinations that the target will be accessible at. The IP address `Addr` belongs to a public network interface on the node that handles the iSCSI traffic type. The port `Port` is optional and defaults to 3260 if omitted.

Once you have the configuration file, e.g., `tg1.json`, you can create the target group with the `vstorage-target tg-create` command. For example, to create an iSCSI target group, run:

```

# vstorage-target tg-create -name tg1 -targets tg1.json -type ISCSI
{
  "Id": "3d8364f5-b830-4211-85af-3a19d30ebac4"
}

```

When you run the command, targets are created on the nodes specified in the configuration file and joined to the target group, target portals are created on the specified network interfaces and ports.

7.3.2 Starting and Stopping Target Groups

When you create a target group, its targets are initially stopped. You can start them with the `vstorage-target tg-start` command. For example:

```
# vstorage-target tg-start -id 3d8364f5-b830-4211-85af-3a19d30ebac4
```

This command starts all targets in the group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`.

All targets in a group can either be running or stopped. So if you add targets to a group of running targets, the new targets will be started automatically.

To stop a target group, use the `vstorage-target tg-stop` command. For example:

```
# vstorage-target tg-stop -id 3d8364f5-b830-4211-85af-3a19d30ebac4
```

7.3.3 Listing Target Groups

You can list target groups with the `vstorage-target tg-list` command that displays basic information about groups. For example:

```
# vstorage-target tg-list
[
  {
    "Id": "3d8364f5-b830-4211-85af-3a19d30ebac4",
    "Name": "tg1",
    "Type": "ISCSI",
    "Running": true,
    "ACL": false,
    "ChapAuth": false,
    "CHAP": {},
    "Mode": 0
  },
  {
    "Id": "78c3b51e-fd9a-485b-91ce-bc0a8171c89d",
    "Name": "tg2",
    "Type": "ISCSI",
    "Running": false,
    "ACL": false,
    "ChapAuth": false,
    "CHAP": {},
    "Mode": 0
  }
]
```

To print complete information about all target groups, use `vstorage-target tg-list -all`.

7.3.4 Printing Details of Target Groups

To print the details of a specific target group, use the `vstorage-target tg-status` command. For example:

```
# vstorage-target tg-status -id faeacacd-eba6-416c-9a7f-b5ba9e372e16
```

This command prints the complete details of the target group with the ID `faeacacd-eba6-416c-9a7f-b5ba9e372e16`. One parameter to pay attention to is `NodeState`. It indicates whether a node is in sync with the target group, i.e. aware of its current configuration. The following states can be shown:

- `synced`, node is in sync with the target group.
- `syncing`, node is syncing with the target group.
- `failed`, node failed to sync with the target group (see the `Error` parameter for details).
- `offline`, node is offline.
- `disabled`, node is disabled and its target is offline.

7.3.5 Managing Persistent Reservations of Target Groups

SCSI-2 reservations allow initiators to gain exclusive access to a LUN and prevent other initiators from making changes to that LUN at the same time. Such reservations are typically released by the initiator after changes have been made to the LUN. They are, however, also released on initiator failures or logical unit resets. SCSI-3 introduces persistent reservations that remain in case of failures or resets and are released by the initiator when needed. They also allow multiple initiators to communicate with the LUN in a controlled manner.

In Virtuozzo Hybrid Infrastructure, persistent reservations are used mostly to support Microsoft Hyper-V nodes working in Failover Clusters.

SCSI persistent reservations are enabled by default. You can enable and disable them—for all volumes in the target group—as follows:

```
# vstorage-target tg-pr -id <tg_ID> -enable  
# vstorage-target tg-pr -id <tg_ID> -disable
```

Where `<tg_ID>` is the ID of the target group for which persistent reservations are set.

Note: For persistent reservations to work, the `vstorage-target-manager` service must be running on all MDS nodes.

7.3.6 Deleting Target Groups

To delete a target group, use the `vstorage-target tg-delete` command. For example:

```
# vstorage-target tg-delete -id 3d8364f5-b830-4211-85af-3a19d30ebac4
```

This command deletes the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`.

7.4 Managing iSCSI Volumes

This section describes how to create and manage volumes to be exported via iSCSI.

7.4.1 Creating iSCSI Volumes

To create a volume, use the `vstorage-target vol-create` command. For example:

```
# vstorage-target vol-create -name vol1 -size 1G \  
-vstorage-attr "replicas=3:2 failure-domain=host tier=0"  
{  
  "Id": "3277153b-5296-49c5-9b66-4c200ddb343d"  
}
```

This command creates a 1 GB volume named `vol1` on storage tier 0 with 3:2 replication and host as failure domain.

7.4.2 Listing and Printing Details of iSCSI Volumes

To list volumes, use the `vstorage-target vol-list` command. For example:

```
# vstorage-target vol-list  
[  
  "3277153b-5296-49c5-9b66-4c200ddb343d",  
  "a12110d5-cbbc-498a-acdd-a8567286f927",
```

```
"d5cc3c13-cfb4-4890-a20d-fb80e2a56278"
]
```

Use `vstorage-target vol-stat -all` to print details of all volumes. To print details of a specific volume, run `vstorage-target vol-stat -id <vol_ID>`.

7.4.3 Attaching iSCSI Volumes to Target Groups

To attach a volume to a target group, use the `vstorage-target tg-attach` command. A volume cannot be attached to multiple target groups at the same time. For example:

```
# vstorage-target tg-attach -id 3d8364f5-b830-4211-85af-3a19d30ebac4 \
-volume 3277153b-5296-49c5-9b66-4c200ddb343d -lun 0 -node bbfd0e7a26b1406d
```

This command attaches the volume with the ID `3277153b-5296-49c5-9b66-4c200ddb343d` to a target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4` as LUN 0. LUN ID numbering must start with 0. The same command sets the PREFERRED bit to the node with the ID `bbfd0e7a26b1406d`. The default Active/Optimized path will go via this node.

7.4.4 Setting the Active/Optimized Path for iSCSI Volumes

To set an Active/Optimized path for an iSCSI volume, use the `vstorage-target vol-set` command. It will only work if the specified node is STABLE.

Note: Make sure the new preferred node is reachable by the initiator.

```
# vstorage-target vol-set -id 3d8364f5-b830-4211-85af-3a19d30ebac4 \
-pref-node bbfd0e7a26b1406d
```

This command sets the Active/Optimized path for the volume with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4` to the node with the ID `bbfd0e7a26b1406d`.

7.4.5 Viewing iSCSI Volume ALUA Information

To view the ALUA information for an iSCSI volume, use the `vstorage-target vol-info` command. For example:

```
# vstorage-target vol-info -id 3d8364f5-b830-4211-85af-3a19d30ebac4
Volume ID:      3d8364f5-b830-4211-85af-3a19d30ebac4
```

```
Name:          vol1
Size:          1073741824
Used:          1073152
Serial:        d2be0e84fd7f
Attrs:         map[]
TG:            4708b908-8c2d-444c-91b1-a1e18a96d4fc
LUN:           0
```

```

*** Node #0 ***
-----
NodeId:        bbfd0e7a26b1406d
State:         synced
TPGs:          vstorage_tpg_0
ALUA:          active
Preferred:     1
WWNs:         iqn.2014-06.com.vstorage:target1 [2]
Portals:       10.37.130.61
```

This command shows the ALUA details for the volume with the ID 3d8364f5-b830-4211-85af-3a19d30ebac4.

7.4.6 Viewing and Setting iSCSI Volume Parameters

To view and set volume parameters, e.g. redundancy mode, failure domain, or tier, use the commands `vstorage-target vol-attr get` and `vstorage-target vol-attr set`, respectively. For example:

```
# vstorage-target vol-attr get -id d5cc3c13-cfb4-4890-a20d-fb80e2a56278
{
  "chunk-size": "268435456",
  "client-ssd-cache": "1",
  "failure-domain": "host",
  "replicas": "3:2",
  "tier": "0"
}
# vstorage-target vol-attr set -id d5cc3c13-cfb4-4890-a20d-fb80e2a56278 \
-vstorage-attr "replicas=2:1 tier=1"
```

7.4.7 Increasing iSCSI Volume Size

To increase the size of a volume, use the `vstorage-target vol-grow` command. For example:

```
# vstorage-target vol-grow -id d5cc3c13-cfb4-4890-a20d-fb80e2a56278 -size 2G
```

7.4.8 Setting iSCSI Volume Limits

To set read/write limits for a volume, use the `vstorage-target vol-limits` command. For example:

```
# vstorage-target vol-limits -id d5cc3c13-cfb4-4890-a20d-fb80e2a56278 -read-bps 10485760 \  
-write-bps 10485760
```

This command sets read/write speed for the volume with the ID `d5cc3c13-cfb4-4890-a20d-fb80e2a56278` to 10485760 bytes per second.

7.4.9 Detaching iSCSI Volumes from Target Groups

To detach a volume from a target group, use the `vstorage-target tg-detach` command. LUN 0 must be detached last. For example:

```
# vstorage-target tg-detach -id 3d8364f5-b830-4211-85af-3a19d30ebac4 \  
-volume d5cc3c13-cfb4-4890-a20d-fb80e2a56278
```

This command detaches the volume with the ID `d5cc3c13-cfb4-4890-a20d-fb80e2a56278` from the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`.

7.4.10 Deleting iSCSI Volumes

To delete a volume, use the `vstorage-target vol-delete` command. You cannot delete volumes attached to target groups. For example:

```
# vstorage-target vol-delete -id d5cc3c13-cfb4-4890-a20d-fb80e2a56278
```

This command deletes the volume with the ID `d5cc3c13-cfb4-4890-a20d-fb80e2a56278`.

7.5 Managing Nodes in Target Groups

This section describes how to manage nodes in relation to target groups.

7.5.1 Adding Nodes to Target Groups

To add a node to a target group, create a configuration file with details about target WWN and portal. The target will be created automatically on the added node. One node can be added to multiple target groups, and the same network interfaces on it can be used simultaneously by multiple targets from different groups.

For example:

```
# vstorage-target node-add -node bbfd0e7a26b1406d -tg 3d8364f5-b830-4211-85af-3a19d30ebac4 \
-targets target.json
```

This command adds the node with the ID `bbfd0e7a26b1406d` to the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`. It also creates a target on it according to the `target.json` configuration file that looks as follows:

```
[
  {
    "NodeId": "bbfd0e7a26b1406d",
    "WWN": "iqn.2013-10.com.vstorage:test2",
    "Portals": [
      {
        "Addr": "10.94.104.89",
        "Port": 3260
      }
    ]
  }
]
```

7.5.2 Listing Nodes in Target Groups

You can list all nodes in all target groups and their detailed information using the `vstorage-target node-list` command. For example:

```
# vstorage-target node-list
[
  {
    "ID": "bbfd0e7a26b1406d",
    "Status": "STABLE",
    "Enabled": true,
    "MonitorOnline": true,
    "Version": "7.10.32",
    "Address": "10.94.104.89:40135",
    "ActiveVolumes": [
      "0937f0e3-91a9-4dfc-8c10-6202bdc792c8"
    ]
  }
]
```

```
},  
]
```

7.5.3 Enabling and Disabling Nodes in Target Groups

To enable or disable a node in all target groups it belongs to, use the `vstorage-target node-set` command. Enabling a node starts its targets, while disabling a node stops its targets and moves the active path to another node. These operations are also performed when the node exits and enters maintenance, respectively.

For example, to enable a node with the ID `bbfd0e7a26b1406d`, run

```
# vstorage-target node-set -node bbfd0e7a26b1406d -enable
```

Before disabling a node, make sure there are other STABLE nodes where the Active/Optimized path can be moved. Otherwise, an I/O error will occur.

To disable a node with the ID `bbfd0e7a26b1406d`, run

```
# vstorage-target node-set -node bbfd0e7a26b1406d -disable
```

You can check the node status with the `vstorage-target node-list` command, see [Listing Nodes in Target Groups](#) (page 226).

7.5.4 Deleting Nodes from Target Groups

To delete a node from a target group, use the `vstorage-target node-del` command. You can only delete a node if it is not on the Active/Optimized path. Otherwise, you need to move the A/O path to another node either by disabling the node (see [Enabling and Disabling Nodes in Target Groups](#) (page 227)) or manually (see [Setting the Active/Optimized Path for iSCSI Volumes](#) (page 223)).

```
# vstorage-target node-del -tg 3d8364f5-b830-4211-85af-3a19d30ebac4 -node bbfd0e7a26b1406d
```

This command deletes the node with the ID `bbfd0e7a26b1406d` from the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`.

7.6 Managing Targets and Portals

This section describes how to create and manage targets.

The optimal way is to create a single target per node if you use the iSCSI protocol and one target per FC port if you use the FC protocol.

7.6.1 Creating Targets

Typically, targets are created automatically when you create target groups or add nodes to them. However, as you can delete target(s) from a node without removing the node from a target group, you can also create target(s) on such a node again. Use the `vstorage-target target-create` command. For example:

```
# vstorage-target target-create -tg 3d8364f5-b830-4211-85af-3a19d30ebac4 -json target.json
```

This command creates a target based on the `target.json` configuration file in the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`. The configuration file lists target details like the node to create the target on, WWN, and portal. For example:

```
{
  "NodeId": "bbfd0e7a26b1406d",
  "WWN": "iqn.2013-10.com.vstorage:test22",
  "Portals": [
    {
      "Addr": "10.94.104.90",
      "Port": 3260
    }
  ]
}
```

7.6.2 Adding and Removing Target Portals

To add a portal to a target, use the `vstorage-target target-portal add` command. For example:

```
# vstorage-target target-portal add -wwn iqn.2013-10.com.vstorage:test2 -addr 10.94.104.90 \
-tg 3d8364f5-b830-4211-85af-3a19d30ebac4
```

This command adds a portal with the IP address `10.94.104.90` and default port `3260` to the target with the IQN `iqn.2013-10.com.vstorage:test2` in the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4`.

To delete a portal from a target, use the `vstorage-target target-portal del` command. For example:

```
# vstorage-target target-portal del -wwn iqn.2013-10.com.vstorage:test2 -addr 10.94.104.90 \
-tg 3d8364f5-b830-4211-85af-3a19d30ebac4
```

This command deletes the portal created before.

7.6.3 Deleting Targets

To delete a target from a target group (as well as the node it is on), use the `vstorage-target target-delete` command. For example:

```
# vstorage-target target-delete -tg 3d8364f5-b830-4211-85af-3a19d30ebac4 \
-wwn iqn.2013-10.com.vstorage:test22
```

This command deletes the target with the IQN `iqn.2013-10.com.vstorage:test22` from the target group with the ID `3d8364f5-b830-4211-85af-3a19d30ebac4` and from the node it is located on.

A node that has no targets left on it is deleted from the target group.

7.7 Managing CHAP Accounts

The Challenge-Handshake Authentication Protocol (CHAP) provides a way to restrict access to targets and their LUNs by requiring a user name and a password from the initiator. CHAP accounts apply to entire target groups. Fibre Channel target groups do not use CHAP.

To use CHAP, enable it for the target group:

```
# vstorage-target tg-auth -enable-chap -id <tg_ID>
```

7.7.1 Creating and Listing CHAP Accounts

To create a CHAP account, use the `vstorage-target account-create` command. For example:

```
# vstorage-target account-create -user user1 -desc "User for TG1"
Enter Password:
```

The password must be 12 to 16 characters long.

To list existing CHAP accounts and their details, use the `vstorage-target account-list` command.

7.7.2 Changing CHAP Account Details

To change the password or description of a CHAP account, use the `vstorage-target account-set` command. For example:

```
# vstorage-target account-set description -user user1 -desc "A new description"
# vstorage-target account-set password -user user1
Enter Password:
```

7.7.3 Assigning CHAP Accounts to Target Groups

To assign a CHAP account to a target group, use the `vstorage-target tg-chap` command. For example:

```
# vstorage-target tg-chap set -id faeacacd-eba6-416c-9a7f-b5ba9e372e16 -user user1
```

To remove an assignment, run

```
# vstorage-target tg-chap del -id faeacacd-eba6-416c-9a7f-b5ba9e372e16 -user user1
```

7.7.4 Deleting CHAP Accounts

To delete an unused CHAP account, use the `vstorage-target account-delete` command. For example:

```
# vstorage-target account-delete -user user1
```

7.8 Managing LUN Views

LUN views provide a way to create and manage an access control list (ACL) that limits access to chosen LUNs for specific initiators. Initiators not on the list have access to all LUNs in iSCSI target groups. Volumes exported via Fibre Channel target groups, however, can only be accessed by initiators that are added to group ACL.

To use ACL-based authorization, enable it for the target group:

```
# vstorage-target tg-auth -enable-acl -id <tg_ID>
```

7.8.1 Creating LUN Views

To create a LUN view for an initiator, use the commands `vstorage-target tg-initiator add` or `vstorage-target view-add`. The former command adds an initiator to target group's ACL and creates a view for it. The latter command is used to add views to initiators that are already on the ACL.

For example:

```
# vstorage-target tg-initiator add -alias initiator2 -luns 0,1 \  
-tg ee764519-80e3-406e-b637-8d63712badf1 -wwn iqn.1994-05.com.redhat:1535946874d
```

This command adds the initiator with the IQN `iqn.1994-05.com.redhat:1535946874d` to the ACL of the target group with the ID `ee764519-80e3-406e-b637-8d63712badf1` and creates a view allowing it to access the LUNs with the IDs `0` and `1`.

Another example:

```
# vstorage-target view-add -tg faeacacd-eba6-416c-9a7f-b5ba9e372e16 -lun 2 -map 2 \  
-wwn iqn.1994-05.com.redhat:1535946874d
```

This command adds a view for the same initiator allowing it to access LUN 2 as well.

7.8.2 Listing LUN Views

To list LUN views for an initiator, use the `vstorage-target view-list` command. For example:

```
# vstorage-target view-list -tg ee764519-80e3-406e-b637-8d63712badf1 \  
-wwn iqn.1994-05.com.redhat:1535946874d
```

This command lists views for the initiator with the IQN `iqn.1994-05.com.redhat:1535946874d`.

7.8.3 Changing LUN View Details

To change LUN views for an initiator, use the `vstorage-target view-set` command. For example:

```
# vstorage-target view-set -luns 1 -tg ee764519-80e3-406e-b637-8d63712badf1 \  
-wwn iqn.1994-05.com.redhat:1535946874d
```

This command allows the initiator with the IQN `iqn.1994-05.com.redhat:1535946874d` to access only LUN 1. Essentially, it deletes all LUN views for it excluding specified.

7.8.4 Deleting LUN Views

To delete a LUN view for an initiator, use the `vstorage-target view-del` command.

```
# vstorage-target view-del -lun 1 -tg ee764519-80e3-406e-b637-8d63712badf1 \  
-wwn iqn.1994-05.com.redhat:1535946874d
```

This command deletes the view for LUN 1 for the initiator with the IQN `iqn.1994-05.com.redhat:1535946874d`.

CHAPTER 8

Advanced Tasks

This chapter describes miscellaneous configuration and management tasks.

8.1 Updating Kernel with ReadyKernel

ReadyKernel is a kpatch-based service shipped with Virtuozzo Hybrid Infrastructure and available out-of-the-box on physical servers with active licenses. ReadyKernel offers a more convenient, rebootless alternative to updating the kernel the usual way and allows you not to wait for scheduled server downtime to apply critical security updates. ReadyKernel enables you to receive cumulative kernel patches that fix critical security issues and apply these patches without having to reboot the server. ReadyKernel updates are released for kernels younger than 18 months. When a kernel becomes older than 18 months, you need to switch to a newer kernel to keep receiving ReadyKernel updates.

Upon installation, the patches are loaded into server RAM and immediately applied to the kernel. If the server reboots, these patches are reapplied to the kernel on boot. You can check the details of the applied ReadyKernel patch at any time with `readykernel info`.

If later you install a new kernel or a major kernel update that requires a reboot, the downloaded patches will remain on the server but will not be applied.

In Virtuozzo Hybrid Infrastructure, ReadyKernel is set to automatically download and apply updates. Checks for new patches are added to each yum transaction that takes place on any node in the infrastructure.

Even though ReadyKernel requires no user interaction by default, you can read the following subsections to understand how this tool works and manage it if needed.

8.1.1 Installing ReadyKernel Patches Automatically

ReadyKernel is enabled by default and checks for new patches daily at 12:00 server time by means of a `cron.d` script. If a patch is available, ReadyKernel will download, install, and load it for the current kernel.

To disable automatic updating, run

```
# readykernel autoupdate disable
```

You can re-enable automatic updating later with the following command:

```
# readykernel autoupdate enable <hour>
```

The service will check for patches daily at the specified `<hour>` (set in 24-hour format, server time).

8.1.2 Managing ReadyKernel Patches Manually

8.1.2.1 Downloading, Installing, and Loading ReadyKernel Patches

To download, install, and instantly load the latest ReadyKernel patch for the current kernel, do the following:

1. Check for new ReadyKernel patches:

```
# readykernel check-update
```

2. If a new patch is available, download, install, and instantly load it for the current kernel by running:

```
# readykernel update
```

ReadyKernel patches are cumulative, i.e. the latest patch includes all the previous ones. To keep the kernel secure, you only need to install and load the latest patch.

8.1.2.2 Loading and Unloading ReadyKernel Patches

To manually load the latest installed ReadyKernel patch to the kernel, do one of the following:

- If an older patch is already loaded, unload it first, then load the latest patch by running:

```
# readykernel load-replace
```

- If no older patches are loaded, load the latest patch by running:

```
# readykernel load
```

To unload the patch from the current kernel, run

```
# readykernel unload
```

8.1.2.3 Installing and Removing ReadyKernel Patches for Specific Kernels

If multiple kernels are installed on the server, you can install a ReadyKernel patch for a specific kernel:

```
# yum install readykernel-patch-<kernel_version>
```

To remove a specific ReadyKernel patch from the server, run

```
# yum remove readykernel-patch-<kernel_version>
```

8.1.2.4 Downgrading ReadyKernel Patches

If you experience problems with the latest ReadyKernel patch, you can downgrade it to an older version if one is available.

To downgrade a patch for the current kernel to the previous version, run

```
# yum downgrade readykernel-patch-$(uname -r)
```

To downgrade a patch for a specific kernel to the previous version, run

```
# yum downgrade readykernel-patch-<kernel_version>
```

You can run these commands multiple times to downgrade to the patch version you need. Alternatively, you can downgrade a patch to a specific version by specifying the desired patch version. For example:

```
# yum downgrade readykernel-patch-12.7-0.4-17.v17
```

8.1.2.5 Disabling Loading of ReadyKernel Patches on Boot

If for some reason you do not want ReadyKernel patches to be applied at boot time, run the following command:

```
# readykernel autoload disable
```

To re-enable automatic loading of ReadyKernel patches on boot, run


```
# readykernel autoload enable
```

8.1.2.6 Managing ReadyKernel Logs

ReadyKernel logs event information in `/var/log/messages` and `/var/log/kpatch.log`. You can specify logging parameters for the latter in the configuration file `/etc/logrotate.d/kpatch`. For more information on parameters you can use, see the `logrotate` man page.

8.2 Managing Guest Tools

This section explains how to install and uninstall the guest tools. This functionality is required for *Running Commands in Virtual Machines without Network Connectivity* (page 241).

8.2.1 Installing Guest Tools

To be able to install the guest tools in virtual machines, you first need to create and upload compute images from the supplied guest tools ISO files located in `/usr/share/vz-guest-tools/`. Execute the following commands on one of the compute nodes:

- For Linux guest tools:

```
# vinfra service compute image create vz-guest-tools-lin \  
--file /usr/share/vz-guest-tools/vz-guest-tools-lin.iso --os-distro linux  
Uploading image to server [Elapsed Time: 0:00:05] ...
```

- For Windows guest tools:

```
# vinfra service compute image create vz-guest-tools-win \  
--file /usr/share/vz-guest-tools/vz-guest-tools-win.iso --os-distro windows  
Uploading image to server [Elapsed Time: 0:00:09] ...
```

Next, you need to attach the created image to a VM and run the guest tools installer. The steps differ for new and already existing VMs and are described in the following subsections.

8.2.1.1 Installing Guest Tools in New VMs

When you create a new VM, you can attach the guest tools image to it and install the guest tools after the operating system. To do this, perform the following steps on a compute node:

1. Create a new VM with the guest tools image. For example, to create a Linux VM `centos`, run:

```
# vinfra service compute server create centos --network id=private --flavor medium \  
--volume source=blank,size=64,boot-index=0,type=disk \  
--volume source=image,id=centos7,size=3,boot-index=1,type=cdrom \  
--volume source=image,id=vz-guest-tools-lin,size=1,boot-index=2,type=cdrom
```

Note: Round up the size of volumes to be created from images. E.g., if the OS distribution image is 2.6 GB, use `size=3`.

In this example, the first volume is a blank virtual HDD, the second volume is the OS distribution image `centos7`, and the third volume is the guest tools image `vz-guest-tools-lin`. Make sure to specify the correct boot order by means of the `boot-index` parameter.

2. Log in to the virtual machine and install an operating system in it.
3. Run guest tools installer inside the VM:
 - Inside a Linux VM, create a mount point for the optical drive with the guest tools image and run the installer:

```
# mkdir /mnt/cdrom  
# mount /dev/sr1 /mnt/cdrom  
# bash /mnt/cdrom/install
```

- Inside a Windows VM, launch the installer in the AutoPlay window if `autorun` is enabled. Otherwise open the optical drive in Explorer and run `setup.exe`.

After installing guest tools, restart the VM.

Note: Guest tools rely on the QEMU guest agent that is installed alongside the tools. The agent service must be running for the tools to work.

8.2.1.2 Installing Guest Tools in Existing VMs

The steps you need to perform to install the guest tools in existing VMs depend on the guest OS type. They are described in the following subsections.

8.2.1.2.1 Installing Guest Tools in Existing Linux VMs

To install the guest tools in an existing Linux virtual machine, do the following:

1. Create a volume from the uploaded guest tools image. For example:

```
# vinfra service compute volume create vz-guest-tools-lin-vol --storage-policy default \
--size 1 --image vz-guest-tools-lin
```

2. Attach the guest tools volume to the virtual machine. For example:

```
# vinfra service compute server volume attach \
--server centos vz-guest-tools-lin-vol
+-----+-----+
| Field | Value |
+-----+-----+
| device | /dev/sr1 |
| id     | 1a40012a-7976-47a1-81f1-ff498cba90af |
+-----+-----+
```

3. Log in to the virtual machine, create a mount point for the optical drive with the guest tools image and run the installer:

```
# mkdir /mnt/cdrom
# mount /dev/sr1 /mnt/cdrom
# bash /mnt/cdrom/install
```

Note: Guest tools rely on the QEMU guest agent that is installed alongside the tools. The agent service must be running for the tools to work.

8.2.1.2.2 Installing Guest Tools in Existing Windows VMs

To install the guest tools in an existing Windows virtual machine, do the following:

1. Power off the Windows VM. For example, to stop the win10 VM, run:

```
# vinfra service compute server stop win10
```

- Convert its system volume to a template image. You will need the volume ID that you can obtain with `vinfra service compute volume list`. For example, to use the `win10` VM boot volume, run:

```
# vinfra service compute volume list | grep win10
| 7116d747-a1e1-4200-bd4a-25cc51ef006c | win10/windows_10_pro_x64.iso/Boot volume | <...> |
| ef2f1979-7811-4df6-9955-07e2fc942858 | win10/windows_10_pro_x64.iso/CD/DVD volume | <...> |
# vinfra service compute volume upload-to-image 7116d747-a1e1-4200-bd4a-25cc51ef006c | grep id
| id | 79da5239-b2bb-4779-ada2-46cb8da8ba0e
```

- Create a new Windows VM from the template, attaching the guest tools image to it during creation. For example:

```
# vinfra service compute server create newvm --network id=private --flavor medium \
--volume source=image,id=79da5239-b2bb-4779-ada2-46cb8da8ba0e,size=64,boot-index=0,type=disk \
--volume source=image,id=vz-guest-tools-win,size=1,boot-index=1,type=cdrom
```

Note: The size of volume to be created from a template image must be equal to or greater than the minimum volume size specified in the image metadata. You can learn the minimum volume size by using `vinfra service compute image show <image_id> | grep min_disk`.

In this example, the first volume is the template of the original VM's system disk and the second volume is the guest tools image. Make sure to specify the correct boot order by means of the `boot-index` parameter.

- Once the image is mounted inside the Windows VM, launch the installer in the AutoPlay window if `autorun` is enabled. Otherwise open the optical drive in Explorer and run `setup.exe`.

After installing guest tools, restart the VM.

Note: Guest tools rely on the QEMU guest agent that is installed alongside the tools. The agent service must be running for the tools to work.

8.2.2 Uninstalling Guest Tools

The steps you need to perform to remove guest tools depend on the guest OS and are described in the following sections.

8.2.2.1 Uninstalling Guest Tools from Linux VMs

To uninstall the guest tools from a Linux guest, log in to the virtual machine and do as follows:

1. Remove the packages:

- 1.1. On RPM-based systems (CentOS and other):

```
# yum remove dkms-vzvirtio_balloon prl_nettool qemu-guest-agent-vz vz-guest-udev
```

- 1.2. On DEB-based systems (Debian and Ubuntu):

```
# apt-get remove vzvirtio-balloon-dkms prl-nettool qemu-guest-agent-vz vz-guest-udev
```

If any of the packages listed above are not installed on your system, the command will fail. In this case, exclude these packages from the command and run it again.

2. Remove the files:

```
# rm -f /usr/bin/prl_backup /usr/share/qemu-ga/VERSION /usr/bin/install-tools \  
/etc/udev/rules.d/90-guest_iso.rules /usr/local/bin/fstrim-static /etc/cron.weekly/fstrim
```

3. Reload the udev rules:

```
# udevadm control --reload
```

After removing guest tools, restart the virtual machine.

8.2.2.2 Uninstalling Guest Tools from Windows VMs

To uninstall the guest tools for Windows, log in to the virtual machine and do as follows:

1. Remove the QEMU device drivers from the device manager.

Important: Do not remove the VirtIO/SCSI hard disk driver and NetKVM network driver. Without the former, the VM will not boot; without the latter, the VM will lose network connectivity.

2. Uninstall the QEMU guest agent and guest tools from the list of installed applications.
3. Stop and delete Guest Tools Monitor:

```
> sc stop VzGuestToolsMonitor
> sc delete VzGuestToolsMonitor
```

4. Unregister Guest Tools Monitor from Event Log:

```
> reg delete HKLM\SYSTEM\CurrentControlSet\services\eventlog\Application\
VzGuestToolsMonitor
```

5. Delete the autorun registry key for RebootNotifier:

```
> reg delete HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run /v \
VzRebootNotifier
```

6. Delete the C:\Program Files\Qemu-ga\ directory.

If `VzGuestToolsMonitor.exe` is locked, close all the Event Viewer windows. If it remains locked, restart the `eventlog` service:

```
> sc stop eventlog
> sc start eventlog
```

After removing the guest tools, restart the virtual machine.

8.3 Running Commands in Virtual Machines without Network Connectivity

If a VM cannot access a network for some reason, you can still run commands in it from the node the VM resides on. The VM in question must have the guest tools installed in it (see *Managing Guest Tools* (page 236)).

You will need the VM ID that you can obtain with `vinfra service compute server list`. You can also use a `virsh` domain name that you can get using `virsh list`.

8.3.1 Running Commands in Linux Virtual Machines

To run an arbitrary command inside a Linux VM and receive the output to your console, use the `virsh x-exec` command. For example:

```
# virsh x-exec 1d45a54b-0e20-4d5e-8f11-12c8b4f300db /usr/bin/bash -c 'lsblk'
NAME          MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
```

```

loop0      7:0    0 945.9M 1 loop
loop1      7:1    0    5G 1 loop
  live-rw  253:0  0    5G 0 dm  /
  live-base 253:1  0    5G 1 dm
loop2      7:2    0   32G 0 loop
  live-rw  253:0  0    5G 0 dm  /
sda        8:0    0   64G 0 disk
sdc        8:32  0    1G 1 disk
sr0        11:0   1    2G 0 rom  /run/initramfs/live

```

To copy a file to a Linux VM, use the `virsh x-exec` and `cat` commands. For example:

```

# virsh x-exec 1d45a54b-0e20-4d5e-8f11-12c8b4f300db \
--shell 'cat > test.file' < /home/test.file

```

To get a file from a Linux VM, use the `virsh x-exec` and `cat` commands as well. For example:

```

# virsh x-exec 1d45a54b-0e20-4d5e-8f11-12c8b4f300db \
--shell 'cat /home/test.file' > test.file

```

8.3.2 Running Commands in Windows Virtual Machines

To run an arbitrary command inside a Windows VM and receive the output to your console, use the `virsh x-exec` command. For example:

```

# virsh x-exec bbf4a6ec-865f-4e2c-ac21-8639d1bfb85c --shell dir c:\
Volume in drive C has no label.
Volume Serial Number is D0BE-A8D1

Directory of c:\

06/10/2009  01:42 PM                24 autoexec.bat
06/10/2009  01:42 PM                 10 config.sys
07/13/2009  06:37 PM          <DIR>      PerfLogs
11/12/2018  07:45 AM          <DIR>      Program Files
11/12/2018  07:55 AM          <DIR>      test
11/12/2018  06:23 AM          <DIR>      Users
11/12/2018  07:53 AM          <DIR>      Windows
                2 File(s)                34 bytes
                5 Dir(s)  59,329,495,040 bytes free

```

To copy a file to a Windows VM, use the `virsh x-exec` and `prl_cat` commands. For example:

```

# virsh x-exec bbf4a6ec-865f-4e2c-ac21-8639d1bfb85c \
--shell '%programfiles%\qemu-ga\prl_cat' 'c:\test\test.file' < /home/test.file

```

To get a file from a Windows VM, use the `virsh x-exec` and `type` commands. For example:

```
# virsh x-exec bbf4a6ec-865f-4e2c-ac21-8639d1bfb85c \
--shell type 'c:\test\test.file' > test.file
```

8.4 Setting Virtual Machines CPU Model

Virtual machines are created with the host CPU model by default. If nodes in the compute cluster have different CPUs, live migration of VMs between compute nodes may not work or applications inside VMs that depend on particular CPUs may not function properly. To avoid this, you can find out which CPU model offers compatibility across all nodes in the compute cluster and manually set it as the compute cluster default. In this case, however, the compute cluster CPU model will be the least advanced one and compute nodes will lose CPU capabilities of a more advanced processor.

To set the compute cluster CPU model, do the following:

1. Run `virsh capabilities` on each node to print an XML document with information on node's CPU. Join the `<cpu>` sections from all XML outputs to a single XML file, e.g., `cpu-compare.xml`.
2. Compare the CPU features using `virsh cpu-baseline`. For example:

```
# virsh cpu-baseline cpu-compare.xml | grep model
<model fallback='allow'>IvyBridge</model>
```

The command will print the most compatible CPU model across all nodes.

3. Set this CPU model for the compute cluster. For example:

```
# vinfra service compute set --cpu-model IvyBridge
```

Take note of the following:

- For the list of supported CPU models, run `vinfra service compute show`.
- Changing CPU model affects only new VMs (i.e. those created after the change).

8.5 Creating Linux Templates

If you do not have a ready Linux template, you can build one with the `diskimage-builder` tool. The disk image is created with only the root user that has neither password nor SSH keys. You can use the `user data` and `cloud-init` methods to perform initial configuration tasks on VMs that will be deployed from the disk image,

for example, create custom user accounts. For more options to customize a VM during boot, refer to the [cloud-init documentation](#).

To create a template and deploy a VM from it, do as follows:

1. Install the `diskimage-builder` package:

```
# yum install diskimage-builder
```

2. For the RHEL 7 guest OS, download the cloud image from the [Red Hat Customer Portal](#) (login required) and execute:

```
# export DIB_LOCAL_IMAGE=<path_to_rhel7_image>
```

3. Execute the following command to build a disk image with installed `cloud-init` for the desired Linux guest. For example:

```
# disk-image-create vm centos7 -t qcow2 -o centos7
```

where

- `centos7` is the name of a guest OS. Can be one of the following: `centos6`, `centos7`, `debian`, `rhel7`, or `ubuntu`.

By default, using the `ubuntu` element will create a disk image for Ubuntu 16.04. To build the Ubuntu 18.04 disk image, add the `DIB_RELEASE=bionic` to the command as follows: `DIB_RELEASE=bionic disk-image-create vm ubuntu -t qcow2 -o ubuntu18`.

- `-o` sets the name for the resulting disk image file.

4. Upload the created disk image using the `vinfra` tool to the compute cluster:

```
# vinfra service compute image create centos7-image --os-distro centos7 \  
--disk-format qcow2 --file centos7.qcow2
```

where

- `centos7-image` is the name of a new image.
- `centos7` is the OS distribution. Can be one of the following: `centos6`, `centos7`, `debian9`, `rhel7`, `ubuntu16.04`, and `ubuntu18.04`.
- `centos7.qcow2` is the QCOW2-image created on step 3.

5. Create the user-data configuration file with a custom user account:

```
# cat <<EOF > user-data
#cloud-config
user: myuser
password: password
chpasswd: {expire: False}
ssh_pwauth: True
EOF
```

where `myuser` is the name of a custom user and `password` is a password for the account.

6. Launch the deployment of a VM from the disk image using the configuration file as user data:

```
# vinfra service compute server create centos7-vm --flavor medium --network public \
--user-data user-data --volume source=image,id=centos7-image,size=10
```

where


- `centos7-vm` is the name of a new VM,
- `user-data` is the configuration file created in step 5,
- `centos7-image` is the image added to the compute cluster in step 4.

For more information on managing compute objects via the `vinfra` tool, see *Managing Compute Cluster* (page 46).

8.6 Connecting to OpenStack Command-Line Interface

For managing the compute cluster, you can also use the OpenStack command-line client, which is automatically installed along with the Virtuozzo Hybrid Infrastructure.

To connect to and be able to use the OpenStack CLI, do the following:

1. Locate the node with the management role in the admin panel. On the **INFRASTRUCTURE > Nodes** screen, the management node has the icon . For example:



2. Access the management node via SSH and log in as the service user. For example:

```
# ssh node001.vstoragedomain
# su - vstoradmin
```

3. Generate the admin OpenRC script that sets environment variables:

```
# kolla-ansible post-deploy
```

The command will create the `/etc/kolla/admin-openrc.sh` bash script:

```
export OS_PROJECT_DOMAIN_NAME=Default
export OS_USER_DOMAIN_NAME=Default
export OS_PROJECT_NAME=admin
export OS_USERNAME=vstorage-service-user
export OS_PASSWORD=<password>
export OS_AUTH_URL=https://<MN_IP_address>:5000/v3
export OS_IDENTITY_API_VERSION=3
export OS_AUTH_TYPE=password
export OS_INSECURE=true
export PYTHONWARNINGS="ignore:Unverified HTTPS request is being made"
export NOVACLIENT_INSECURE=true
export NEUTRONCLIENT_INSECURE=true
export CINDERCLIENT_INSECURE=true
export OS_PLACEMENT_API_VERSION=1.22
```

By default, the script is created to authorize in the `admin` project under the `vstorage-service-user` user for managing the compute cluster with administrative privileges.

4. To perform administrative actions, run this script:

Important: You need to run the script each session.

```
# source /etc/kolla/admin-openrc.sh
```

If you want to work in another project under another user, you need to make changes to the `admin-openrc.sh` script. For example, to authorize in the `myproject` project under the `myuser` user within the `mydomain` domain, do the following:

1. Copy the script to the chosen directory with a new name. For example:

```
# cp /etc/kolla/admin-openrc.sh /root/myscript.sh
```

2. Open the copied script for editing and change the first five variables as follows:

```
export OS_PROJECT_DOMAIN_NAME=mydomain
export OS_USER_DOMAIN_NAME=mydomain
export OS_PROJECT_NAME=myproject
export OS_USERNAME=myuser
export OS_PASSWORD=<myuser_password>
```

Leave other variables as is and save your changes.

3. Run the modified script:

Important: You need to run the script each session.

```
# source /root/myscript.sh
```

Now you can work in the project you have authorized in by executing OpenStack commands with the `--insecure` option. For example:

```
# openstack --insecure server list
+-----+-----+-----+-----+-----+-----+
| ID           | Name | Status | Networks          | Image | Flavor |
+-----+-----+-----+-----+-----+-----+
| 32b0f95d-477f-46b5-<...> | vm1  | ACTIVE | private=192.168.128.87 |      | tiny   |
+-----+-----+-----+-----+-----+-----+
```

To learn how to secure OpenStack API traffic, refer to *Securing OpenStack API Traffic with SSL* (page 248).

8.7 Setting a DNS Name for the Compute API

By means of the **Compute API** traffic type, Virtuozzo Hybrid Infrastructure exposes a public endpoint that listens to OpenStack API requests. By default, it points to the IP address of the management node (or to its virtual IP address if high availability is enabled).

In some cases, you need to modify all public endpoints to use the domain name resolvable to the management node IP address (or its virtual IP), for example, to secure OpenStack API traffic with an SSL certificate without the `subjectAltName` field or to make the Kubernetes service access the compute API via the DNS name.

To modify all public endpoints to use the domain name, execute:

```
# openstack --insecure endpoint list | grep public
| 44aa0f53a40e4e52b1c7eeeb20c7811e | <...> | https://10.94.16.12:8774/v2.1/(tenant_id)s |
| 5a845b4b813047c292db73c42dad5efd | <...> | https://10.94.16.12:8780 |
| 0b906e518b1041c8b94af7f410403369 | <...> | https://10.94.16.12:9696 |
| d80af756adf1449f9237c3aeebc9206a | <...> | https://10.94.16.12:8004/v1/(tenant_id)s |
| d0e8c7da7d174e1f9aa4efbc6dff2113 | <...> | https://10.94.16.12:5000/v3 |
| 0e6d3a39d6c44aa883984a35dde434bb | <...> | https://10.94.16.12:9292 |
| 7d901686bca549f9b294e572f046f634 | <...> | https://10.94.16.12:8776/v2/(tenant_id)s |
| 1b68ac7c3f7949fbaeef4a815fe6f3b1 | <...> | https://10.94.16.12:8776/v3/(tenant_id)s |

# openstack --insecure endpoint set \
--url https://<DNS_name>:8774/v2.1/(tenant_id)s 44aa0f53a40e4e52b1c7eeeb20c7811e
# openstack --insecure endpoint set \
--url https://<DNS_name>:8780 5a845b4b813047c292db73c42dad5efd
# openstack --insecure endpoint set \
--url https://<DNS_name>:9696 0b906e518b1041c8b94af7f410403369
# openstack --insecure endpoint set \
--url https://<DNS_name>:8004/v1/(tenant_id)s d80af756adf1449f9237c3aeebc9206a
# openstack --insecure endpoint set \
--url https://<DNS_name>:5000/v3 d0e8c7da7d174e1f9aa4efbc6dff2113
# openstack --insecure endpoint set \
--url https://<DNS_name>:9292 0e6d3a39d6c44aa883984a35dde434bb
# openstack --insecure endpoint set \
--url https://<DNS_name>:8776/v2/(tenant_id)s 7d901686bca549f9b294e572f046f634
# openstack --insecure endpoint set \
--url https://<DNS_name>:8776/v3/(tenant_id)s 1b68ac7c3f7949fbaeef4a815fe6f3b1
```

8.8 Securing OpenStack API Traffic with SSL

Traffic to and from the public endpoint that listens to OpenStack API requests can be secured with an SSL certificate. However, as domain names are not used by default, the certificate will need a `subjectAltName` field containing the aforementioned management node IP address. If it does not have such a field, you will need to modify the public endpoint to use a domain name that you have a certificate for.

To secure public OpenStack API traffic with SSL, do the following:

1. Upload the certificate and then private key in the admin panel, on the **SETTINGS > Management node > SSL ACCESS** screen.

- Place the CA certificate file to operating system's trusted bundle:

```
# cp ca.pem /etc/pki/ca-trust/source/anchors/
# update-ca-trust extract
```

Alternatively, you can append the `--os-cacert ca.pem` option to each OpenStack client call.

- If your certificate does not have the `subjectAltName` field, modify all public endpoints to use the domain name for which you have the certificate for as described in *Setting a DNS Name for the Compute API* (page 247). This domain name must resolve to the management node IP address (or to its virtual IP address if high availability is enabled).
- In your OpenRC script, change `OS_AUTH_URL` to the same domain name and remove all parameters related to insecure access. For example:

```
export OS_PROJECT_DOMAIN_NAME=Default
export OS_USER_DOMAIN_NAME=Default
export OS_PROJECT_NAME=admin
export OS_USERNAME=admin
export OS_PASSWORD=<ADMIN_PASSWORD>
export OS_AUTH_URL=https://<DOMAIN_NAME>:5000/v3
export OS_IDENTITY_API_VERSION=3
```

Now you can run OpenStack commands without the `--insecure` option.

8.9 Using Metering for Compute Resources

You can collect usage data of compute resources using [Gnocchi](#). This time series database processes and stores measurement data of compute resources and provides access to it via REST API or the command-line tool.

Measurements can be sampled from such compute resources as virtual machines, VM disks and interfaces, compute networks, volumes, etc. All resources are being revised: if any attribute of a resource changes, this change is recorded in the history of the resource. For a VM, for example, you can measure the amount of allocated memory and virtual CPUs as well as the memory and CPU usage.

An entity storing aggregates composed of a timestamp and value is called a metric. A metric is attached to a specific resource and associated to an archive policy. A policy defines how long aggregates are kept in a metric and how they are computed (minimum, maximum, average, etc.).

The following metrics are available for aggregation:

Table 8.9.1: Compute resource metrics

Metric	Type	Resource type	Description
memory	gauge	instance	Amount of RAM allocated to the VM, in megabytes
memory.usage	gauge	instance	Percentage of RAM used by the VM
vcpus	gauge	instance	Number of virtual CPUs allocated to the VM
cpu	cumulative	instance	Amount of CPU time used by the VM, in nanoseconds
disk.device.read.requests	cumulative	instance_disk	Number of read requests
disk.device.write.requests	cumulative	instance_disk	Number of write requests
disk.device.read.bytes	cumulative	instance_disk	Amount of data read, in bytes
disk.device.write.bytes	cumulative	instance_disk	Amount of data written, in bytes
network.incoming.bytes	cumulative	instance_network_interface	Incoming network traffic, in bytes
network.outgoing.bytes	cumulative	instance_network_interface	Outgoing network traffic, in bytes
network.incoming.packets	cumulative	instance_network_interface	Incoming network traffic, in packets
network.outgoing.packets	cumulative	instance_network_interface	Outgoing network traffic, in packets
image.size	gauge	image	Size of the uploaded image, in bytes
volume.size	gauge	volume	Size of the volume, in gigabytes
snapshot.size	gauge	volume	Size of the volume snapshot, in gigabytes

Continued on next page

Table 8.9.1 – continued from previous page

Metric	Type	Resource type	Description
magnum.cluster	gauge	coe_cluster	Number of Container Orchestration Engine (COE), i.e. Kubernetes, clusters

Cumulative metrics are polled every 5 minutes and increase over time, while gauge metrics are updated on events and show fluctuating values.

For gauge metrics, the default archive policy is `ceilometer-low`, and for cumulative metrics, it is `ceilometer-low-rate`. These policies imply that all computed aggregates are kept for one day with 5-minute granularity and for one month with 1-hour granularity. The difference between them is in aggregation methods they use: the `ceilometer-low` policy stores only mean values, while `ceilometer-low-rate` stores mean values as well as the average of delta values per interval.

8.9.1 Enabling Metering Service

Note: If you have already installed metering services via the admin panel, skip this section.

You can enable metering services in your compute cluster by doing one of the following:

- If you have no compute cluster yet, deploy it and enable metering by adding the `--enable-metering` option to the `vinfra service compute cluster create` command. For example:

```
# vinfra service compute create --nodes <node1_id>[,<node2_id>,...] --enable-metering
```

- If you have already created the compute cluster, use the following command:

```
# vinfra service compute cluster set --enable-metering
```

Note: The metering service will only take into account compute objects created after it has been enabled.

These commands open port 8041 and enable two Gnocchi services: `gnocchi-api`, an HTTP server, and `gnocchi-metricd`, a metric daemon.

8.9.2 Using Gnocchi Command-line Tool

After enabling metering, you can get access to your compute resource metrics either via REST API or using the Gnocchi command-line tool. To be able to use the tool, do as follows:

1. Install the Gnocchi client:

```
# yum install python-gnocchiclient
```

2. Generate the admin OpenRC file:

```
# kolla-ansible post-deploy
# source /etc/kolla/admin-openrc.sh
```

Now you can use the `gnocchi` command with the `--insecure` option. See also *Securing OpenStack API Traffic with SSL* (page 248).

For example, you can do as follows:

1. View the existing resources:

```
# gnocchi --insecure resource list -c id -c type -c project_id
+-----+-----+-----+
| id          | type                | project_id |
+-----+-----+-----+
| 238597c7<...> | volume              | c1bf1<...> |
| 3c78558f<...> | instance            | c1bf1<...> |
| 44f1896f<...> | instance_network_interface | c1bf1<...> |
| 880e9efc<...> | instance_disk       | c1bf1<...> |
+-----+-----+-----+
```

The output shows that the compute cluster hosts one virtual machine with one NIC and one disk that is also present as a volume.

2. Check all the available metrics for the resources:

```
# gnocchi --insecure metric list
+-----+-----+-----+-----+-----+
| id          | archive_policy/name | name                | unit   | resource_id |
+-----+-----+-----+-----+-----+
| 243c7ac<...> | ceilometer-low     | disk.root.size     | GB     | 3c7855<...> |
| 365e454<...> | ceilometer-low-rate | network.outgoing.packets | packet | 44f189<...> |
| 4fbd3e8<...> | ceilometer-low-rate | disk.device.read.requests | request | 880e9e<...> |
| 54519fa<...> | ceilometer-low     | compute.instance.booting.time | sec    | 3c7855<...> |
| 5e1406f<...> | ceilometer-low-rate | disk.device.write.bytes | B      | 880e9e<...> |
| 66a96c2<...> | ceilometer-low     | vcpus              | vcpu   | 3c7855<...> |
| 722ea97<...> | ceilometer-low     | memory             | MB     | 3c7855<...> |
| 7c961ab<...> | ceilometer-low-rate | disk.device.write.requests | request | 880e9e<...> |
| 87e9fb7<...> | ceilometer-low-rate | network.incoming.packets | packet | 44f189<...> |
+-----+-----+-----+-----+-----+
```

```

| 9d56321<...> | ceilometer-low-rate | disk.device.read.bytes | B | 880e9e<...> |
| b8be8f7<...> | ceilometer-low-rate | cpu | ns | 3c7855<...> |
| c1961bb<...> | ceilometer-low | disk.ephemeral.size | GB | 3c7855<...> |
| c9b61e0<...> | ceilometer-low | volume.size | GB | 238597<...> |
| d06a58c<...> | ceilometer-low-rate | network.outgoing.bytes | B | 44f189<...> |
| e2d9981<...> | ceilometer-low-rate | network.incoming.bytes | B | 44f189<...> |
| eaac2b5<...> | ceilometer-low | memory.usage | MB | 3c7855<...> |
+-----+-----+-----+-----+-----+

```

3. View measures for a particular metric. For example:

```

# gnocchi --insecure measures show cpu --resource-id 3c78558f-08bf-47e2-ba3e-bdb13e7b25bb
+-----+-----+-----+
| timestamp | granularity | value |
+-----+-----+-----+
| 2019-12-11T17:05:00+03:00 | 300.0 | 2.2756e+11 |
| 2019-12-11T17:10:00+03:00 | 300.0 | 2.8897e+11 |
| 2019-12-11T17:15:00+03:00 | 300.0 | 3.7367e+11 |
| 2019-12-11T17:20:00+03:00 | 300.0 | 4.64e+11 |
| 2019-12-11T17:25:00+03:00 | 300.0 | 7.6104e+11 |
+-----+-----+-----+

```

By default, the mean aggregation method is used. To obtain how much CPU time is consumed per interval, use the `--aggregation rate:mean` option:

```

# gnocchi --insecure measures show cpu --aggregation rate:mean \
--resource-id 3c78558f-08bf-47e2-ba3e-bdb13e7b25bb
+-----+-----+-----+
| timestamp | granularity | value |
+-----+-----+-----+
| 2019-12-11T17:10:00+03:00 | 300.0 | 61410000000.0 |
| 2019-12-11T17:15:00+03:00 | 300.0 | 84700000000.0 |
| 2019-12-11T17:20:00+03:00 | 300.0 | 90330000000.0 |
| 2019-12-11T17:25:00+03:00 | 300.0 | 2.9704e+11 |
| 2019-12-11T17:30:00+03:00 | 300.0 | 3.64e+11 |
+-----+-----+-----+

```

You can see the full list of `gnocchi` commands in the [OpenStack documentation](#).

8.9.3 Viewing Resource Usage Per Project

To get usage of compute resources allocated to all virtual machines that belong to a particular project, use the command `vinfra service compute quotas show --usage <project_id>`. For example:

```

# vinfra service compute quotas show 6ef6f48f01b640ccb8ff53117b830fa3 --usage
+-----+-----+
| Field | Value |
+-----+-----+

```

```

| compute.cores.limit           | 20 |
| compute.cores.used           | 2  |
| compute.ram.limit            | 40960 |
| compute.ram.used             | 4096 |
| k8saas.cluster.limit         | 10 |
| k8saas.cluster.used          | 0  |
| lbaas.loadbalancer.limit     | 10 |
| lbaas.loadbalancer.used      | 0  |
| network.floatingip.limit     | 10 |
| network.floatingip.used      | 0  |
| storage.gigabytes.default.limit | 1024 |
| storage.gigabytes.default.used | 66  |
+-----+-----+

```

The output shows that VMs included in the project with the ID 62af79f31ae5488aa33077d02af48282 were allocated 2 vCPUs, 4 GB of RAM, and 66 GB of disk space.

8.10 Configuring Memory Policy for Storage Services

You can configure memory limits and guarantees for storage services at runtime using the `vinfra memory-policy vstorage-services` commands. You can do this for the entire cluster or a specific node.

The following memory parameters can be configured manually:

- memory guarantee
- swap size
- page cache (which, in turn, is set using `cache ratio`, `minimum`, and `maximum`)

Page cache is calculated according to the following formula:

```
$PAGE_CACHE = minimum <= ratio * $TOTAL_MEMORY <= maximum
```

The `minimum` and `maximum` values are hard limits that are applied if the `ratio * $TOTAL_MEMORY` value is outside these limits.

To better understand how page cache size is calculated, consider the following examples:

Table 8.10.1: Page cache examples

	Example 1 (cache size is within limits)	Example 2 (cache size equals minimum)	Example 3 (cache size equals maximum)
Total memory	4 GiB	4 GiB	4 GiB
Cache ratio	0.5	0.1	0.9
Cache minimum	1 GiB	2 GiB	1 GiB
Cache maximum	3 GiB	3 GiB	3 GiB
Cache size	2 GiB	2 GiB	3 GiB

If memory parameters are set both per node and per cluster, the per-node ones are applied. If no memory parameters are configured manually, the memory management is performed automatically by the `vcmmmd` daemon as follows:

- Each CS (e.g., storage disk) requires 512 MiB of RAM for page cache.
- The page cache minimum is 1 GiB.
- If the total memory is less than 48 GiB, the page cache maximum is calculated as two-thirds of it.
- If the total memory is greater than 48 GiB, the page cache maximum is 32 GiB.

To check the current memory parameters for storage services set by `vcmmmd`, run

```
# vcmmddctl list
name                type active guarantee    limit  swap  cache
<...>
vstorage.slice/vstorage-services.slice SRVC   yes  1310720 24522132    0  1048576
```

8.10.1 vinfra memory-policy vstorage-services per-cluster change

Change per-cluster memory parameters:

```
usage: vinfra memory-policy vstorage-services per-cluster change
       [--guarantee <guarantee>] [--swap <swap>] [--cache-ratio <cache-ratio>]
       --cache-minimum <cache-minimum> --cache-maximum <cache-maximum>]
```

`--guarantee <guarantee>`

Guarantee, in bytes

`--swap <swap>`

Swap size, in bytes, or -1 if unlimited

```
--cache-ratio <cache-ratio>
    Cache ratio from 0 to 1 inclusive

--cache-minimum <cache-minimum>
    Minimum cache, in bytes

--cache-maximum <cache-maximum>
    Maximum cache, in bytes
```

Example:

```
# vinfra memory-policy vstorage-services per-cluster change --guarantee 8796093022208 \
--swap 1099511627776 --cache-ratio 0.5 --cache-minimum 1099511627776 \
--cache-maximum 3298534883328
+-----+-----+
| Field   | Value                               |
+-----+-----+
| cache   | maximum: 3298534883328 |
|         | minimum: 1099511627776 |
|         | ratio: 0.5              |
| guarantee | 8796093022208         |
| swap     | 1099511627776         |
+-----+-----+
```

This command sets the storage services memory parameters for all nodes in the storage cluster as follows:

- the memory guarantee to 8 GB
- the swap size to 1 GB
- the page cache limits: the minimum to 1 GB, the maximum to 3 GB, and the cache ratio to 0.5

8.10.2 vinfra memory-policy vstorage-services per-cluster show

Show per-cluster memory parameters:

```
usage: vinfra memory-policy vstorage-services per-cluster show
```

Example:

```
# vinfra memory-policy vstorage-services per-cluster show
+-----+-----+
| Field   | Value                               |
+-----+-----+
| cache   | maximum: 3298534883328 |
|         | minimum: 1099511627776 |
|         | ratio: 0.5              |
+-----+-----+
```

```
| guarantee | 8796093022208 |
| swap      | 1099511627776 |
+-----+-----+
```

This command lists the storage services memory parameters for all nodes in the storage cluster.

8.10.3 vinfra memory-policy vstorage-services per-cluster reset

Reset per-cluster parameters to default:

```
usage: vinfra memory-policy vstorage-services per-cluster reset [--guarantee]
      [--swap] [--cache]
```

`--guarantee`

Reset only the guarantee.

`--swap`

Reset only the swap size.

`--cache`

Reset only cache values.

Example:

```
# vinfra memory-policy vstorage-services per-cluster reset --cache
+-----+-----+
| Field  | Value          |
+-----+-----+
| cache  |                 |
| guarantee | 8796093022208 |
| swap   | 1099511627776 |
+-----+-----+
```

This command resets the manually configured page cache limits to default for all nodes in the storage cluster.

8.10.4 vinfra memory-policy vstorage-services per-node change

Change per-node memory parameters:

```
usage: vinfra memory-policy vstorage-services per-node change [--guarantee <guarantee>]
      [--swap <swap>] [--cache-ratio <cache-ratio> --cache-minimum <cache-minimum>
      --cache-maximum <cache-maximum>] --node <node>
```

`--guarantee <guarantee>`

Guarantee, in bytes

`--swap <swap>`
 Swap size, in bytes, or -1 if unlimited

`--cache-ratio <cache-ratio>`
 Cache ratio from 0 to 1 inclusive

`--cache-minimum <cache-minimum>`
 Minimum cache, in bytes

`--cache-maximum <cache-maximum>`
 Maximum cache, in bytes

`--node <node>`
 Node ID or hostname

Example:

```
# vinfra memory-policy vstorage-services per-node change --guarantee 8796093022208 \
--swap 1099511627776 --cache-ratio 0.5 --cache-minimum 1099511627776 \
--cache-maximum 3298534883328 --node 7ffa9540-5a20-41d1-b203-e3f349d62565
+-----+-----+
| Field   | Value                               |
+-----+-----+
| cache   | maximum: 3298534883328 |
|         | minimum: 1099511627776 |
|         | ratio: 0.5              |
| guarantee | 8796093022208         |
| swap     | 1099511627776         |
+-----+-----+
```

This command sets the storage services memory parameters for the node with the ID `7ffa9540-5a20-41d1-b203-e3f349d62565` as follows:

- the memory guarantee to 8 GB
- the swap size to 1 GB
- the page cache limits: the minimum to 1 GB, the maximum to 3 GB, and the cache ratio to 0.5

8.10.5 vinfra memory-policy vstorage-services per-node show

Show per-node memory parameters:

```
usage: vinfra memory-policy vstorage-services per-node show --node <node>
```

```
--node <node>
    Node ID or hostname
```

Example:

```
# vinfra memory-policy vstorage-services per-node show \
--node 7ffa9540-5a20-41d1-b203-e3f349d62565
+-----+-----+
| Field  | Value                |
+-----+-----+
| cache  | maximum: 13194139533312 |
|        | minimum: 8796093022208  |
|        | ratio: 0.7             |
| guarantee | 8796093022208         |
| swap   | 1099511627776         |
+-----+-----+
```

This command lists the storage services memory parameters set for the node with the ID 7ffa9540-5a20-41d1-b203-e3f349d62565.

8.10.6 vinfra memory-policy vstorage-services per-node reset

Reset per-node memory parameters to defaults:

```
usage: vinfra memory-policy vstorage-services per-node reset [--guarantee] [--swap]
      [--cache] --node <node>
```

```
--guarantee
    Reset only the guarantee.
```

```
--swap
    Reset only the swap size.
```

```
--cache
    Reset only the cache values.
```

```
--node <node>
    Node ID or hostname
```

Example:

```
# vinfra memory-policy vstorage-services per-node reset --cache \
--node 7ffa9540-5a20-41d1-b203-e3f349d62565
+-----+-----+
| Field  | Value                |
+-----+-----+
```




```
| cache | |
| guarantee | 8796093022208 |
| swap | 1099511627776 |
+-----+-----+
```

This command resets the manually configured page cache limits to default for the node with the ID 7ffa9540-5a20-41d1-b203-e3f349d62565.

8.11 Migrating Virtual Machines from VMware vCenter

Starting from Virtuozzo Hybrid Infrastructure 3.5, you can migrate virtual machines from VMware vCenter 5.0 or newer to Virtuozzo Hybrid Infrastructure using the `virt-v2v` tool. You will need to create a `virt-v2v` appliance virtual machine to transfer and convert the data from.



8.11.1 Deploying the Appliance VM

To create a `virt-v2v` appliance VM, follow these steps:

1. Download the `virt-v2v` appliance image from the [official repository](#).
2. Upload the image to Virtuozzo Hybrid Infrastructure. For example:

```
# vinfra service compute image create virt-v2v-img \
--file vmware_to_vip.qcow2
```

3. Create an SSH key for the appliance if you do not have one. For example:

```
# vinfra service compute key create publickey \
--public-key virt-v2v-app-key.pub
```

4. Create a virtual machine and deploy the uploaded image in it. The VM needs at least two CPUs, 4 GiB RAM, and enough storage space to accommodate the largest VM to be migrated to Virtuozzo Hybrid

Infrastructure. It must also be connected to the network that handles the **Compute API** traffic type and the network with access to VMware vCenter API. For example:

```
# vinfra service compute server create virt-v2v-appliance \
--flavor medium --key-name <key>
--network id=<compute_API> --network id=<vcenter_API> \
--volume source=image,id=virt-v2v-img,size=<size>
```

Where:

- <key> is the SSH key to authorize in the appliance VM.
- <compute_API> is the network that handles the traffic type **Compute API**.
- <vcenter_API> is the network that can access the VMware vCenter API.
- <size> is the disk size. For online migration, it must be enough to accommodate the largest VM of all you are going to migrate. For offline migration, it must be enough to accommodate twice as much.

8.11.2 Setting Up Authentication in the Appliance VM

1. Log in to the appliance VM as the admin user with the SSH key.
2. Get root privileges, for example, with `sudo -i`.
3. Create a bash script that will export OpenStack credentials:

```
# cat > user-openrc.sh << EOF
export OS_PROJECT_DOMAIN_NAME=Domain_name
export OS_USER_DOMAIN_NAME=Domain_name
export OS_PROJECT_NAME=Project_name
export OS_USERNAME=user_name
export OS_PASSWORD=Password
export OS_AUTH_URL=https://<admin_panel_IP_addr>:5000/v3
export OS_IDENTITY_API_VERSION=3
export OS_INSECURE=true
export NOVACLIENT_INSECURE=true
export NEUTRONCLIENT_INSECURE=true
export CINDERCLIENT_INSECURE=true
export LIBGUESTFS_BACKEND=direct
EOF
```

Note: You will need the administrator credentials for the project that the appliance VM belongs to.

- Copy the OpenStack root CA certificate and CA keys from the Virtuozzo Hybrid Infrastructure management node:

```
# scp root@<MN_IP>:/usr/libexec/vstorage-ui-backend/ca/ca.* /etc/pki/ca-trust/source/anchors/
# update-ca-trust extract
```

Where <MN_IP> is the management node IP address. For more information, see *Securing OpenStack API Traffic with SSL* (page 248).

- Create a file with the VMware vCenter password to supply to virt-v2v. For example:

```
# echo $vCenterPass > password.txt
```

Alternatively, you can enter the password during migration or supply it to virt-v2v with the `--password-file` option.

8.11.3 Migrating Virtual Machines to Virtuozzo Hybrid Infrastructure Online

Take note of the following before you proceed:

- You can migrate VMs created on vCenter 5.0 or newer.
- Remove VMware tools from Windows VMs before the migration to avoid issues on boot afterwards. You will not be able to do this after the migration. VMware tools will be removed from Linux guests automatically.

- Log in to the appliance VM as the admin user with the SSH key.
- Get root privileges, for example, with `sudo -i`.
- Set OpenStack credentials:

```
# source user-openrc.sh
```

- Test the connection between libvirt and VMware vCenter. For example:

```
# virsh -c 'vpx://<domain>%5c<user>@<hostname>?no_verify=1' list --all
Enter root's password for vcenter.example.com: ***
+-----+-----+-----+
| Id | Name          | <...> |
+-----+-----+-----+
| - | Fedora 20     | <...> |
| - | Windows 2008 | <...> |
+-----+-----+-----+
```

Where <hostname> is the name of the VMware ESXi host that runs virtual machines. Its full path looks like <vCenter_hostname>/<datacenter_name>/<cluster_name>/<server_hostname> and can be found in VMware vCenter.

If the VPX username contains a backslash (e.g., <domain>\<user>), escape it with %5c: <domain>%5c<user>. Similarly, escape spaces with %20.

5. Check the OpenStack connection and find out the virt-v2v appliance ID. For example:

```
# openstack --insecure server list
+-----+-----+-----+-----+
| ID                | Name                | Status | <...> |
+-----+-----+-----+-----+
| 635ae4cc-4c01-461a-ae63-91ca4187a7b1 | virt-v2v-appliance | ACTIVE | <...> |
+-----+-----+-----+-----+
```

6. Shut down the VM. Windows VMs must be shut down gracefully for the migration to be successful.
7. Migrate the VM to a volume in Virtuozzo Hybrid Infrastructure. For example:

```
# virt-v2v -ip password.txt -ic 'vpx://<domain>%5c<user>@<hostname>?no_verify=1' \
'Windows 2008' -o openstack -oo server-id=635ae4cc-4c01-461a-ae63-91ca4187a7b1
```

8. Find out the new volume's ID or name. For example:

```
# openstack --insecure volume list
+-----+-----+-----+-----+-----+
| ID                | Name | Status  | Size | Attached to |
+-----+-----+-----+-----+-----+
| 024b6843-2de3-4e25-a6e1-2b6ea2d601cf | sda1 | available | 64 |              |
+-----+-----+-----+-----+-----+
```

9. Create a virtual machine in Virtuozzo Hybrid Infrastructure based on the new volume. For example:

```
# vinfra service compute server create migratedvm \
--network id=private --network id=public \
--volume source=volume,id=sda1,size=64 --flavor medium
```

8.11.4 Migrating Virtual Machines to Virtuozzo Hybrid Infrastructure

Offline

If the network connection between the virt-v2v appliance VM and VMware vCenter is inferior, you can manually copy the VMs to a USB drive, connect it to the virt-v2v appliance VM, and convert them to Virtuozzo Hybrid Infrastructure. Do the following:

1. Copy all of the VM files, including vmdk and vmx, to a USB drive.

2. Attach the USB drive to a host in the same local network as the appliance VM.
3. Log in to the appliance VM as the `admin` user with the SSH key.
4. Get root privileges, for example, with `sudo -i`.
5. Copy VM files to the appliance VM, for example, using `rsync` or `scp`.
6. Set OpenStack credentials:

```
# source user-openrc.sh
```

7. Migrate the VM to a volume in Virtuozzo Hybrid Infrastructure. For example:

```
# virt-v2v -i vmx <VM_config> -o openstack \
-oo server-id=635ae4cc-4c01-461a-ae63-91ca4187a7b1
```

Where `<VM_config>` is the VM configuration file in the `vmx` format.

8. Find out the new volume's ID or name. For example:

```
# openstack --insecure volume list
+-----+-----+-----+-----+-----+
| ID                | Name | Status  | Size | Attached to |
+-----+-----+-----+-----+-----+
| 024b6843-2de3-4e25-a6e1-2b6ea2d601cf | sda1 | available | 64 |              |
+-----+-----+-----+-----+-----+
```

9. Create a virtual machine in Virtuozzo Hybrid Infrastructure based on the new volume. For example:

```
# vinfra service compute server create migratedvm \
--network id=private --network id=public \
--volume source=volume,id=sda1,size=64 --flavor medium
```

8.12 Changing the Default Load Balancer Flavor

By default, a load balancer is created using the private `amphora` flavor that cannot be managed via `vinfra`. You can, however, change it using the OpenStack command-line tool as follows:

1. Generate the admin OpenRC file:

```
# kolla-ansible post-deploy
# source /etc/kolla/admin-openrc.sh
```

2. Check if the default `amphora` flavor exists:

```
# openstack --insecure flavor list --all
```

ID	Name	RAM	Disk	Ephemeral	VCPUs	Is Public
100	tiny	512	0	0	1	True
101	small	2048	0	0	1	True
102	medium	4096	0	0	2	True
103	large	8192	0	0	4	True
104	xlarge	16384	0	0	8	True
amphora	amphora	1024	10	0	2	False

3. Delete this flavor:

```
# openstack --insecure flavor delete amphora
```

4. Create a new amphora flavor with custom parameters. For example:

```
# openstack --insecure flavor create amphora --id amphora --ram 4096 \
--vcpus 4 --disk 15 --private
```

Field	Value
OS-FLV-DISABLED:disabled	False
OS-FLV-EXT-DATA:ephemeral	0
disk	15
id	amphora
name	amphora
os-flavor-access:is_public	False
properties	
ram	4096
rxtx_factor	1.0
swap	
vcpus	4

5. Change the load balancer flavor by performing its failover. For example:

```
# openstack --insecure loadbalancer failover mylbaas
```

The load balancer `mylbaas` will be recreated with 4 vCPUs, 4 GB of RAM, and 15 GB of disk space.

8.13 Configuring Retention Policy for Prometheus Metrics

The Prometheus service used for monitoring the cluster runs and stores its data on the management node. By default, Prometheus metrics are stored for 7 days. This retention period can be insufficient for troubleshooting purposes. You can increase this period in the Prometheus configuration file as follows:

1. On the management node, open the `/etc/sysconfig/prometheus` file to edit, set the needed retention period for the `STORAGE_RETENTION` option, and save your changes. For example:

```
STORAGE_RETENTION="--storage.tsdb.retention.time=30d"
```

2. Restart the Prometheus service:

```
systemctl restart prometheus.service
```

In case high availability is enabled in the storage cluster, repeat these steps for the other two management nodes.

However, with a long retention period, the root partition where the data is stored may run out of free space. To prevent this, you can define the maximum size for the Prometheus metrics. The oldest data will be removed first. To change the time retention policy to the size retention policy, do as follows:

1. On the management node, open the `/etc/sysconfig/prometheus` file to edit, change the flag for the `STORAGE_RETENTION` option, and save your changes. For example:

```
STORAGE_RETENTION="--storage.tsdb.retention.size=10GB"
```

2. Restart the Prometheus service:

```
systemctl restart prometheus.service
```

In case high availability is enabled in the storage cluster, repeat these steps for the other two management nodes.

8.14 Viewing Cluster Logs

When you encounter a problem in Virtuozzo Hybrid Infrastructure, you can send a problem report as described in [Sending Problem Reports](#). The report will gather all the logs needed to troubleshoot the problem and send them to the technical support team.

Alternatively, you can investigate the root cause of your problem using the logs listed in table below.

Table 8.14.1: Cluster log location

Service	Log location	Description
Metadata	/vstorage/mds/logs/mds.log.zst on the storage node hosting an MDS service	Storage metadata service events
Storage	/vstorage/<id>/cs/logs/cs.log.zst on the storage node hosting a CS service <hr/> Note: To find out the log location of a particular CS on the node, run <code>vstorage -c <cluster_name> list-services -C</code> . <hr/>	Chunk service events
Storage mountpoint	/var/log/vstorage/<cluster_name>/vstorage-mount.*.blog on any storage node	Software-defined storage mounting on each node
Management node	/var/log/vstorage-ui-backend/messages.log and /var/log/vstorage-ui-backend/celery*.log on the management node	Management node and admin panel events
	/var/log/vstorage-ui-agent/* on any storage node	Agent controller component events
Backup Gateway	/var/log/vstorage/abgw.log*zst on any node in the Backup Gateway cluster <hr/> Note: The latest log is <code>abgw.log.zst</code> , older ones are renamed to <code>abgw.log.0.zst</code> , <code>abgw.log.1.zst</code> , etc. <hr/>	Backup Gateway cluster deployment and management
iSCSI	/var/log/vstorage/iscsi/vstorage-target.log on any node in an iSCSI target group	iSCSI target management
	/var/log/vstorage/iscsi/vstorage-target-monitor.log on any node in an iSCSI target group	iSCSI target monitoring

Continued on next page

Table 8.14.1 – continued from previous page

Service	Log location	Description
	/var/log/vstorage/iscsi/scst.log.zst on any node in an iSCSI target group	SCST service logs
S3	/var/log/ostor/NS-* on the S3 node with NS services	S3 name server events
	/var/log/ostor/OS-* on the S3 node with OS services	S3 object server events
	/var/log/ostor/S3GW-* on the S3 node with GW services	S3 gateway events
	/var/log/nginx/* on any node in the S3 cluster	nginx service logs
	/var/log/ostor/GR-* on the S3 node with GR services	S3 geo-replicator service events
NFS	/var/log/ganesha/ganesha.log and /var/log/ostor/ostorfs.log.gz on any node in the NFS cluster	NFS server events
	/var/log/vstorage/vstorage-nfsd.log on any node in the NFS cluster	NFS service events
	/var/log/ostor/FS-* on the node hosting an NFS share	FS service events
	/var/log/ostor/OS-* on the node hosting an NFS share	OS service events
Compute	/var/log/vstorage-ui-backend/ansible.log on the controller node	Compute cluster and add-on deployment
	/var/log/hci/beholder/beholder.log on the controller node	Notifications about all compute events, including VM placement
	/var/log/hci/nova/* on the compute node hosting a VM <hr/> Note: In case of problems during VM migration, view /var/log/hci/nova/nova-compute.log on the source and destination compute nodes. <hr/>	Virtual machine management
	/var/log/hci/neutron/neutron-l3-agent.log on any compute node	Virtual routing events
	/var/log/hci/neutron/neutron-openvswitch-agent.log on the compute node hosting a VM	VM network interface management

Continued on next page

Table 8.14.1 – continued from previous page

Service	Log location	Description
	/var/log/hci/cinder/* on the controller node	Compute volume management
	/var/log/hci/glance/glance-api.log on the controller node	Image service API requests
	/var/log/hci/octavia/octavia-worker.log and /var/log/hci/octavia/octavia-api.log on the controller node	Load balancing service management
	/var/log/hci/magnum/magnum-conductor.log, /var/log/hci/magnum/magnum-api.log, and /var/log/hci/heat/heat-engine.log on the controller node	Kubernetes service and VM stack deployment and management
	/var/log/hci/gnocchi/* and /var/log/hci/ceilometer/* on any compute node	Billing metering service management
High availability	/var/log/vstorage-ui-backend/ha.log on all management nodes	High availability management
Updates	/var/log/vstorage-ui-backend/software-updates.log on the management node	Software update orchestration
	/var/log/vstorage-ui-agent/software-updates.log on any storage node	Software update downloading and installation on each node

To open log files, use the following commands:

- for log files:

```
# less <log_file>.log
```

- for blog files:

```
# blogcat <log_file>.blog | less
```

- for gz files:

```
# zless <log_file>.gz
```

- for zst files:

```
# zstdless <log_file>.zst
```