Acronis

Acronis Cyber Infrastructure 3.5

Administrator's Guide

November 30, 2021

Copyright Statement

Copyright ©Acronis International GmbH, 2003-2021. All rights reserved.

"Acronis" and "Acronis Secure Zone" are registered trademarks of Acronis International GmbH.

"Acronis Compute with Confidence", "Acronis Startup Recovery Manager", "Acronis Instant Restore", and the Acronis logo are trademarks of Acronis International GmbH.

Linux is a registered trademark of Linus Torvalds.

VMware and VMware Ready are trademarks and/or registered trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Windows and MS-DOS are registered trademarks of Microsoft Corporation.

All other trademarks and copyrights referred to are the property of their respective owners.

Distribution of substantively modified versions of this document is prohibited without the explicit permission of the copyright holder.

Distribution of this work or derivative work in any standard (paper) book form for commercial purposes is prohibited unless prior permission is obtained from the copyright holder.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Third party code may be provided with the Software and/or Service. The license terms for such third-parties are detailed in the license.txt file located in the root installation directory. You can always find the latest up-to-date list of the third party code and the associated license terms used with the Software and/or Service at http://kb.acronis.com/content/7696.

Acronis patented technologies

Technologies, used in this product, are covered and protected by one or more U.S. Patent Numbers: 7,047,380; 7,246,211; 7,275,139; 7,281,104; 7,318,135; 7,353,355; 7,366,859; 7,383,327; 7,475,282; 7,603,533; 7,636,824; 7,650,473; 7,721,138; 7,779,221; 7,831,789; 7,836,053; 7,886,120; 7,895,403; 7,934,064; 7,937,612; 7,941,510; 7,949,635; 7,953,948; 7,979,690; 8,005,797; 8,051,044; 8,069,320; 8,073,815; 8,074,035; 8,074,276; 8,145,607; 8,180,984; 8,225,133; 8,261,035; 8,296,264; 8,312,259; 8,347,137; 8,484,427; 8,645,748; 8,732,121; 8,850,060; 8,856,927; 8,996,830; 9,213,697; 9,400,886; 9,424,678; 9,436,558; 9,471,441; 9,501,234; and patent pending applications.

Contents

1.	Abou	t This G	uide					
2.	Mana	naging the Storage Cluster						
	2.1	Manag	ing Networks and Traffic Types					
		2.1.1	Exclusive Traffic Types					
		2.1.2	Regular Traffic Types 5					
		2.1.3	Custom Traffic Types					
		2.1.4	Creating, Editing, and Deleting Networks					
		2.1.5	Creating, Editing, and Deleting Traffic Types 7					
	2.2	Configu	ring Node Network Interfaces					
		2.2.1	Setting Up Network Bonding 10					
		2.2.2	Creating VLAN Interfaces					
	2.3	Creating the Storage Cluster						
		2.3.1	Creating the Storage Cluster on the First Node					
		2.3.2	Adding Nodes to Storage Cluster					
		2.3.3	Assigning Disk Roles Manually 16					
	2.4	Connec	ting Remote iSCSI Devices to Storage Cluster Nodes					
		2.4.1	Assigning Disk Roles To Remote iSCSI Devices					
	2.5	Replaci	ng Node Disks					
		2.5.1	Replacing Storage Disks					
		2.5.2	Replacing Disks Manually 23					
	2.6	Perforr	ning Node Maintenance					
	2.7	Releasi	ng Nodes from the Storage Cluster					
	2.8	Remov	ing Nodes from the Unassigned List					
	2.9	Re-Add	ing Nodes to the Unassigned List					
3.	Moni	toring t	he Storage Cluster					

3.1 Monitoring the Entire Cluster			ring the Entire Cluster				
		3.1.1	I/O Activity Charts				
		3.1.2	Services Chart				
		3.1.3	Chunks Chart				
		3.1.4	Physical Space Chart				
			3.1.4.1 Understanding Physical Space				
		3.1.5	Logical Space Chart				
			3.1.5.1 Understanding Logical Space				
	3.2	Monito	ring Nodes				
		3.2.1	Understanding Node Role Icons 42				
		3.2.2	Monitoring Node Performance				
		3.2.3	Monitoring Node Disks				
			3.2.3.1 Monitoring the S.M.A.R.T. Status of Node Disks				
		3.2.4	Monitoring Node Network				
	3.3	Monito	ring Storage Cluster Objects via SNMP 45				
		3.3.1	Enabling SNMP Access				
		3.3.2	Accessing Storage Cluster Information Objects via SNMP				
			3.3.2.1 Listening to SNMP Traps 48				
		3.3.3	Monitoring the Storage Cluster with Zabbix				
		3.3.4	Storage Cluster Objects and Traps				
	3.4	Monito	ring Storage Cluster Remotely				
	3.5	Viewing	g Alerts and Audit Log and Sending E-mail Notifications				
		3.5.1	Viewing Alerts				
		3.5.2	Viewing Audit Log				
		3.5.3	Sending E-mail Notifications				
4	Mana	oging th	e Compute Cluster 69				
	4 1						
	4.2	Monito	ring the Compute Cluster 74				
	1,2	4.2.1	Used CPUs Chart 75				
		422	Reserved RAM Chart 76				
		4.2.3	Provisioned Storage Chart				
		4.2.4	VM Status Chart				
		4.2.5	Top VMs Chart				
		4.2.6	Alerts Chart				
	4.3	Managi	ng Compute Nodes				
	-						

	4.3.1	Adding No	odes to Compute Cluster			
	4.3.2	Releasing	Nodes from Compute Cluster			
4.4	Manag	ing Virtual I	Machines			
	4.4.1	Supported	d Guest Operating Systems			
	4.4.2	Creating \	/irtual Machines			
	4.4.3	Virtual Ma	achine Actions Overview			
	4.4.4	Enabling l	ogging inside Virtual Machines			
	4.4.5	Migrating	Virtual Machines			
	4.4.6	Reconfigu	ring and Monitoring Virtual Machines			
	4.4.7	Configuring Virtual Machine High Availability				
	4.4.8	Managing	Guest Tools			
		4.4.8.1	Installing Guest Tools			
		4.4.8.2	Uninstalling Guest Tools			
4.5	Manag	ing Images				
	4.5.1	Uploading	g, Editing, and Removing Images			
	4.5.2	Creating \	/olumes from Images			
	4.5.3	Mounting	ISO Images to Virtual Machines			
	4.5.4	Preparing	Templates			
		4.5.4.1	Preparing Linux Templates			
		4.5.4.2	Preparing Windows Templates			
4.6	Manag	ing Flavors				
4.7	Manag	ing SSH Key	/s			
4.8	Manag	ing Compu	te Network			
	4.8.1	Virtual Ne	twork Architecture			
		4.8.1.1	Private Network Connectivity			
		4.8.1.2	Public Network Connectivity			
	4.8.2	Creating, Editing, and Deleting Virtual Networks				
	4.8.3	Managing	Virtual Routers			
		4.8.3.1	Managing Router Interfaces			
		4.8.3.2	Managing Static Routes			
	4.8.4	Managing	Floating IP Addresses			
	4.8.5	Managing	Load Balancers			
		4.8.5.1	Managing Balancing Pools			
4.9	Manag	ing Compu [.]	te Storage			
	4.9.1	Managing	Compute Volumes			

		4.9.2	Creating,	, Editing, and Removing Volumes
		4.9.3	Cloning \	/olumes
		4.9.4	Attaching	g and Detaching Volumes
		4.9.5	Creating	Images from Volumes
		4.9.6	Managin	g Volume Snapshots
		4.9.7	Managin	g Storage Policies
	4.10	Manag	ing Placen	nents
		4.10.1	Creating	Placements
		4.10.2	Managin	g Nodes in Placements
		4.10.3	Managin	g Images in Placements
		4.10.4	Renamin	g Placements
		4.10.5	Deleting	Placements
	4.11	Manag	ng Kuberi	netes Clusters
	4.12	Destroy	ing the Co	ompute Cluster
5	Expo	rting Str	arago Sha	1/6
Э.	5 1	Evporti	ng Storag	
	5.1	5 1 1		rkflow Overview 148
		5.1.1	5 1 1 1	Managing Legacy iSCSI Targets 148
		512	Managin	g Target Groups
		5.1.2	5 1 2 1	Creating Target Groups 149
			5122	Adding Targets 152
			5123	Starting and Stopping Targets 154
			5124	Deleting Targets 155
			5.1.2.5	Deleting Target Groups
		5.1.3	Managin	g Volumes
			5.1.3.1	Creating Volumes
			5.1.3.2	Attaching Volumes to Target Groups
			5.1.3.3	Setting LUN Limits
			5.1.3.4	Detaching Volumes
			5.1.3.5	Deleting Volumes
		5.1.4	Restrictir	ng Access to Target Groups
			5.1.4.1	Managing Access Control Lists
			5.1.4.2	Managing CHAP Users
		5.1.5	Monitori	ng Target Groups
			5.1.5.1	Advanced iSCSI Monitoring via Grafana

5.2	Exporti	ing Storage	e via S3			
	5.2.1	S3 Storag	ge Infrastructure Overview			
	5.2.2	Planning the S3 Cluster				
	5.2.3	Sample S	33 Storage			
	5.2.4	Creating the S3 Cluster				
	5.2.5	Managing S3 Users				
		5.2.5.1	Adding S3 Users			
		5.2.5.2	Managing S3 Access Key Pairs			
	5.2.6	Managin	g S3 Buckets			
		5.2.6.1	Listing S3 Bucket Contents			
		5.2.6.2	Managing Acronis Notary in S3 Buckets			
	5.2.7	Best Prac	ctices for Using S3 in Acronis Cyber Infrastructure			
		5.2.7.1	S3 Bucket and Key Naming Policies			
		5.2.7.2	Improving Performance of PUT Operations			
	5.2.8	Replicati	ng S3 Data Between Datacenters			
	5.2.9	Monitoring S3 Access Points				
	5.2.10	Releasing Nodes from S3 Clusters				
	5.2.11	Supported Amazon S3 Features				
		5.2.11.1	Supported Amazon S3 REST Operations			
		5.2.11.2	Supported Amazon Request Headers			
		5.2.11.3	Supported Amazon Response Headers			
		5.2.11.4	Supported Amazon Error Response Headers			
		5.2.11.5	Supported Authentication Schemes			
5.3	Exporting Storage via NFS					
	5.3.1	1 Setting Up an NFS Cluster				
	5.3.2	Creating NFS Shares				
	5.3.3	Creating	NFS Exports			
		5.3.3.1	Creating the Root Export			
		5.3.3.2	Creating User Exports			
	5.3.4	Setting U	Ip User Authentication and Authorization			
		5.3.4.1	Authenticating NFS Share Users with Kerberos			
		5.3.4.2	Authorizing NFS Export Users with LDAP			
5.4	Connec	cting Acror	nis Backup Software to Storage Backends via Backup Gateway			
	5.4.1	Understa	anding the Infrastructure			
	5.4.2	Connecting to the Local Storage Cluster via Backup Gateway				

	5.4.3 Connecting to External NFS Shares via Backup Gateway		
	5.4.4	Connecti	ng to Public Cloud Storage via Backup Gateway
		5.4.4.1	Important Requirements and Restrictions
		5.4.4.2	Setting Up Backup Gateway
	5.4.5	Updating	g certificate for Backup Gateway
	5.4.6	Re-regist	ering Backup Gateway in a New Acronis Backup Advanced
	5.4.7	Migrating	g Backups from Older Acronis Solutions
		5.4.7.1	Migrating Backups from Acronis Storage 1.5
		5.4.7.2	Migrating Backups from Acronis Storage Gateway 1.6 and 1.7 (NFS)
	5.4.8	Managin	g Geo-Replication for Backup Gateway
		5.4.8.1	Enabling Geo-Replication
		5.4.8.2	Performing a Failover
		5.4.8.3	Updating the Geo-replication Configuration
		5.4.8.4	Disabling Geo-replication
	5.4.9	Monitori	ng Backup Gateway
		5.4.9.1	Advanced Backup Gateway Monitoring via Grafana
	5.4.10	Releasing	g Nodes from Backup Gateway
Man	nging Co	noral Sot	tings 229
6 1	Managi		as
0.1	6 1 1	Installing	z license Kevs 220
	6.1.2	Installing	SPLA Licenses 230
6.2	Managi		n Services 231
6.3	Managi	ing Tier Fr	233
6.4	Enablin		234
0.4	6.4.1	Configur	ing InfiniBand Devices 235
	6.4.2	Configur	ing RoCE and iWARP Devices
65	Enablin	o High Av	ailability
0.5	6 5 1	Fnahling	Management Node High Availability 238
6.6	Accessi	ng the Ad	min Panel via SSI 242
6.7	Backing	Un and F	Restoring Management Database 243
0.7	671	Restoring	g Management Database from Backun 245
	0.7.1	6711	Restoring with Existing Compute Cluster 245
6.8	Managi	ing Domai	ins Users and Projects 246
0.0	6.8.1	Managin	g Domains 247
	682	Managin	g Domain Lisers 249
	0.0.2	managin	g Domain OSCIS

6.

		6.8.2.1	Creating System Administrators
		6.8.2.2	Creating Domain Administrators
		6.8.2.3	Creating Project Members
	6.8.3	Managing	g Projects
		6.8.3.1	Assigning Members to Projects
		6.8.3.2	Editing Quotas for Projects
6.9	Managi	ng Update	s
6.10	Allowing	g root Acc	ess to Cluster Nodes over SSH
6.11	Adding	External D	NS Servers
6.12	Configu	iring the S	elf-Service Panel
6.13	Sending	g Problem	Reports

CHAPTER 1 About This Guide

This primary guide describes operations on Acronis Cyber Infrastructure that you can perform via the web-based admin panel. In particular, it explains how to:

- configure networking for both the storage and compute cluster;
- create and manage the storage cluster;
- set up and run storage services, including S3, iSCSI, NFS, and backup gateways;
- monitor the storage cluster;
- create and manage the compute cluster;
- create and manage virtual machines, volumes, images, and storage policies;
- perform auxiliary tasks: set up high availability, enable RDMA, manage licenses, send problem reports, and such.

CHAPTER 2

Managing the Storage Cluster

Before you create the storage cluster, you need to set up the networks and assign them to network interfaces as recommended in the Installation Guide. Next, configure an external DNS server as described in *Adding External DNS Servers* (page 263). Next, enable high availability of the management node (see *Enabling High Availability* (page 237)). Finally, make sure that storage nodes are shown on the **NODES** screen and proceed to create the storage cluster.

If you have remote iSCSI devices you wish to connect to cluster nodes, you can configure them prior to cluster creation as described in *Connecting Remote iSCSI Devices to Storage Cluster Nodes* (page 19).

2.1 Managing Networks and Traffic Types

To balance and optimize networking in Acronis Cyber Infrastructure, you can assign different types of traffic to separate networks. Assigning a traffic type to a network means that a firewall is configured on nodes connected to this network, specific ports are opened on node network interfaces, and the necessary iptables rules are set. For example, nodes connected to a network with only the **S3 public** traffic type will accept incoming connections only on ports 80 and 443.

As described in the Installation Guide, is it recommended to have these networks in Acronis Cyber Infrastructure:

• Internal network (traffic types: **Storage**, **Internal management**, **OSTOR private**, **ABGW private**), assigned to the first bonded connection

Note: If you plan to use RDMA over InfiniBand, move the traffic type **Storage** to a dedicated network and assign that network to the IB interface. See *Enabling RDMA* (page 234).

- Overlay network (traffic type **VM private**), assigned to a VLAN created on the second bonded connection
- Management network (traffic types: Admin panel, SSH, SNMP), assigned to a VLAN created on the second bonded connection
- External network (traffic types: **Self-service panel**, **Compute API**, **VM public**, **S3 public**, **ABGW public**, **iSCSI**, and **NFS**), assigned to a VLAN created on the second bonded connection
- Backup network (traffic types: **VM backups** and **VM public**), assigned to a VLAN created on the second bonded connection

You need to configure these networks on the **INFRASTRUCTURE** > **Networks** screen on the admin panel before you create the cluster (see *Creating, Editing, and Deleting Networks* (page 6)). By default, you have two preconfigured networks named **Public** and **Private**, according to their type. They can be considered templates that you can customize to create the desired (recommended) configuration.

Important: Some traffic types cannot be reassigned to a different network if they are in use. In addition, a public network cannot be renamed if it is used by a compute virtual network.

After you create the networks, proceed to create the remaining of the recommended VLAN interfaces on each node and assign them to networks as described in *Creating VLAN Interfaces* (page 12).

An example of recommended networks and their traffic types is:

Network	Traffic types
Internal	Storage, Internal management, OSTOR private, ABGW private
Overlay	VM private
Management	Admin panel, SSH, SNMP
External	Compute API, Self-service panel, S3 public, iSCSI, NFS, ABGW public, VM
	public

Table 2.1.1: Recommended network setup

Continued on next page

Network	Traffic types
Backup	VM backups, VM public

Table 2.1.1 – continued from previous page

The next three subsections describe all traffic types that can be assigned to networks.

2.1.1 Exclusive Traffic Types

Exclusivity means that such a traffic type can be added only to one network. Exclusive traffic types cannot be reassigned between networks if they are in use. To do that, you will have to delete the service that uses them first.

Internal management

Internal cluster management and transfers of node monitoring data to the admin panel. Without this traffic type, the administrator cannot control and monitor the cluster. The cluster, however, continues working. Uses any available port.

Storage

Internal transfers of data chunks, high availability service heartbeats, as well as data self-healing. This is the most critical traffic type that defines storage performance and enables cluster HA. Uses any available port.

OSTOR private

Internal data exchange between multiple S3/NFS services. Uses any available port.

ABGW private

Internal management of and data exchange between multiple Backup Gateway services. Uses any available port.

VM private

Network traffic between VMs in private virtual networks and VNC console traffic. Private virtual networks are implemented as VXLAN, overlay networking fully isolated on L2. Opens UDP port 4789 and TCP ports from 5900 to 5999.

Compute API

External access to standard OpenStack API endpoints. Opens the following ports:

- TCP 5000—Identity API v3
- TCP 6080—noVNC Websocket Proxy

- TCP 8004—Orchestration Service API v1
- TCP 8041—Gnocchi API (billing metering service)
- TCP 8774—Compute API
- TCP 8776—Block Storage API v3
- TCP 8780—Placement API
- TCP 9292—Image Service API v2
- TCP 9313—Key Manager API v1
- TCP 9513—Container Infrastructure Management API (Kubernetes service)
- TCP 9696—Networking API v2
- TCP 9888—Octavia API v2 (load balancer service)

VM backups

External access to NBD endpoints. Third-party backup management systems can pull VM backups using this traffic type. To be able to access backup agents installed in virtual machines, assign this traffic type along with **VM public**. Opens TCP ports from 49300 to 65535.

2.1.2 Regular Traffic Types

The traffic types listed further are not exclusive and can be added to multiple networks.

S3 public

External data exchange with the S3 access point. Uses TCP ports 80 and 443.

iSCSI

External data exchange with the iSCSI access point. Uses TCP port 3260.

NFS External data exchange with the NFS access point. Uses TCP/UDP ports 111, 892, and 2049.

ABGW public

External data exchange with Acronis Backup agents and Acronis Backup Cloud. Uses TCP port 44445.

Admin panel

External access to the admin panel. Uses TCP port 8888.

VM public

External data exchange between VMs and public networks (e.g., the Internet). When a node network

interface is assigned to a network with this traffic type, an Open vSwitch bridge is created on that network interface.

SSH Remote access to nodes via SSH. Uses TCP port 22.

SNMP

External access to storage cluster monitoring statistics via the SNMP protocol. Opens UDP port 161.

Self-service panel

External access to the self-service panel. Opens TCP port 8800.

2.1.3 Custom Traffic Types

You can create custom traffic types that will open desired TCP ports. Such traffic types can be added to multiple networks. See *Creating, Editing, and Deleting Traffic Types* (page 7).

2.1.4 Creating, Editing, and Deleting Networks

If required, you can add a new network by doing as follows:

- 1. On the **INFRASTRUCTURE** > **Networks** screen, click **Edit** and then **Create network**.
- In the New network window, specify a network name. Network names must be alphanumerical and 2-32 characters long.

New network	×
Name net1	
	Cancel

3. Click Create.

- 4. Add the needed traffic types to the new network by ticking the corresponding checkboxes.
- 5. When finished, click **Save** to apply the changes.

To edit a network name or delete a custom network, click on the ellipsis icon next to it and select the action you want to perform.

You can only delete networks that are not assigned to any network adapters.

2.1.5 Creating, Editing, and Deleting Traffic Types

If required, you can add a new traffic type by doing as follows:

- 1. On the **INFRASTRUCTURE** > **Networks** screen, click **Edit** and then **Create traffic type**.
- 2. In the **Create traffic type** window, specify a traffic type name and port to open.

Traffic type names must be alphanumerical and 3-32 characters long.

Create traffic type	×
Name	
traffictype1	
Port 50000	
	Cancel
	Cancel

3. Click Create.

- 4. Add the newly created traffic type to one or more of your networks by ticking the corresponding checkboxes.
- 5. When finished, click **Save** to apply the changes.

To edit or delete a custom traffic type, make sure it is excluded from all networks, click the ellipsis icon next to it, and select the desired action.

2.2 Configuring Node Network Interfaces

After configuring the networks, you need to assign them to the network interfaces on each node. A network can only be assigned to one network interface per node.

To assign a network to a network interface, do the following:

- 1. On the **Infrastructure** > **Nodes** screen, click a node to configure.
- 2. On the node overview screen, click **NETWORK**.



- 3. Select a network interface and click **Configure**.
- 4. On the **Configure** screen, do one of the following:
 - To obtain the IP address, DNS, and routing settings from the DHCP server, select **Automatically** (DHCP).
 - To obtain just the IP address from the DHCP server, select **Automatically (DHCP address only)**.
 - To specify the IP address manually, select **Manually** and add the IP address.

Warning: Dynamic IP address allocation will cause network issues as soon as the IP addresses of cluster nodes will change. Configure static IP addresses from the start or as soon as possible.

Note: For information about configuring RDMA-enabled network interfaces, see *Enabling RDMA* (page 234).

× Configure
 Automatically (DHCP) Automatically (DHCP address only) Manually
✓ 10.37.130.250/24
Add Remove
Gateway
10.10.0.1
MTU
1500
DONE

- 5. If necessary, set up a gateway and a DNS server. The provided gateway will become node's default.
- 6. If you have set a custom maximum transmission unit (MTU) on the network hardware, set the same value in the corresponding field. See Step 2: Configuring the Network for more details.

Warning: Setting a custom MTU in admin panel prior to configuring it on the network hardware will result in network failure on the node and require manual resetting. Setting an MTU that differs from the one configured on the network hardware may result in network outage or poor performance.

- 7. Click **Done** to return to the list of network interfaces, do not change the selection, and click **Assign network**.
- 8. On the **Assign network** panel, select a network to connect the network interface to (for details, see *Managing Networks and Traffic Types* (page 2)) and click **Done**.

×	Assign network
 Pri No 	ivate one
D	ONE

2.2.1 Setting Up Network Bonding

Bonding multiple network interfaces is optional but provides the following benefits:

- High network availability. If one of the interfaces fails, the traffic will be automatically routed through the working interface(s).
- Higher network performance. For example, two bonded Gigabit interfaces will deliver the throughput of about 1.7 Gbit/s or up to 200 MB/s. For a storage node, the required number of network interfaces to bond may depend on the number of disks. For example, an HDD can deliver data at speeds of up to 1 Gbps.

To create a bond, do the following:

- 1. On the **Infrastructure** > **Nodes** screen, click the node to bond the network interfaces on.
- 2. On the node overview screen, click **NETWORK**.
- 3. In the **NETWORK** list, check network interfaces to bond, and click **Create bonding** in the menu to the right.
- 4. On the **Configure Bonding** panel, select the bonding type from the drop-down list. The balance-xor type is selected by default and recommended for both fault tolerance and good performance.

× Configure bonding

Туре	
balance-xor	~
 Automatically (DHCP) Automatically (DHCP address only) Manually 	
10.94.64.111/16	
10.94.57.114/16	
Add Remove	
Gateway	
MTU	
auto	
MAC	
auto	~

If you are bonding an Open vSwitch-based bridge used in the compute cluster with another network interface, two bonding types will be available: balance-tcp and active-backup. In the balance-tcp mode, load balancing is done based on L2-L4 data like the destination MAC address, IP address, and TCP port. This mode requires that Link Aggregation Control Protocol (LACP) be enabled on the physical switch the node is connected to. In the active-backup mode, one network interface is active and all other interfaces are passive. If the active interface fails, a passive one becomes active instead.

5. Set up network parameters as described in step 4 in *Configuring Node Network Interfaces* (page 8) and click **PROCEED**.

6. On the **Assign network** panel, select a network to connect the bonding network interface to (for details, see *Managing Networks and Traffic Types* (page 2)) and click **Done**.

×	Assign network
• P	rivate
() N	lone
	DONE

2.2.2 Creating VLAN Interfaces

To create a VLAN interface on a node, do the following:

- 1. On the **Infrastructure** > **Nodes** screen, click the node on which to configure VLAN.
- 2. On the node overview screen, click **NETWORK**.
- 3. Select a network interface and click **Create VLAN**.
- 4. On the **Configure VLAN** panel, specify a number for VLAN, add an IP address, and, if necessary, set up a gateway and a DNS server. The provided gateway will become node's default.

× Configure VLAN
VLAN #
100
 Automatically (DHCP) Automatically (DHCP address only) Manually
There are no items to show in this view.
Add Remove
Gateway
MTU
auto
PROCEED

5. On the **Assign network** panel, select a network to connect the VLAN network interface to (for details, see *Managing Networks and Traffic Types* (page 2)) and click **Done**.

2.3 Creating the Storage Cluster

Before you create the storage cluster, enable high availability of the management node as described in *Enabling High Availability* (page 237).

To create a storage cluster, you need to create a basic storage cluster on one (first) node, then populate it with more nodes.

If networks adapters on your nodes support RDMA (via RoCE, iWARP or IB) and you want to enable this functionality, you must do so before creating the storage cluster as explained in *Enabling RDMA* (page 234).

2.3.1 Creating the Storage Cluster on the First Node

- 1. Open the **INFRASTRUCTURE** > **Nodes** screen and click a node in the **UNASSIGNED** list.
- 2. On the node overview screen, click **Create cluster**.
- 3. In the **Cluster** field, type a name for the cluster. The name may only contain Latin letters (a-z, A-Z), numbers (0-9), underscores ("_") and hyphens ("-").

× New cluster

Create cluster on node node001		
Cluster		
cluster1		
Storage interface		
eth1 - 10.37.130.250	~	\$
Encryption		
NEW CLUSTER ADVANCED CONFIGURATION		

4. From the **Storage interface** drop-down list, select a node network interface connected to a network with the traffic type **Storage**.

If node network interfaces are not configured, click the cogwheel icon and assign a network with the traffic type **Storage** to a node's network interface.

- 5. If required, enable data encryption. To do this, check the Encryption box (see Managing Tier Encryption (page 233)) and proceed to create the cluster. Encryption will be enabled for all tiers by default. To enable encryption for particular tiers, click the cogwheel icon to open the Encryption Configuration panel, select tiers to encrypt, and click Done. You can later disable encryption for new chunk services (CS) on the SETTINGS > Advanced settings panel.
- Click New cluster to have Acronis Cyber Infrastructure assign the roles to disks automatically. Alternatively, click Advanced configuration to assign the roles to each drive manually and tweak other settings.

You can monitor cluster creation in the **HEALTHY** list of the **INFRASTRUCTURE** > **Nodes** screen. The creation might take some time depending on the number of disks to be configured. Once the automatic configuration is complete, the cluster is created.

2.3.2 Adding Nodes to Storage Cluster

To add an unassigned node to the storage cluster, do the following:

- 1. On the **INFRASTRUCTURE** > **Nodes** screen, click an unassigned node.
- 2. On the node overview screen, click **Join cluster**.
- 3. Make sure the network interface that is connected to a network with the traffic type **Storage** is selected from the **Storage interface** drop-down list.

If node network interfaces are not configured, click the cogwheel icon and assign a network with the traffic type **Storage** to a node's network interface.

× Join cluster		
Node to join: node002 Storage interface		
eth1 - 10.37.130.28	~	\$
JOIN CLUSTER ADVANCED CONFIGURATION		

4. Click **Join cluster** to have Acronis Cyber Infrastructure assign the roles to disks automatically and add the node to the current cluster. Alternatively, click **Advanced configuration** to assign the roles to each drive manually.

2.3.3 Assigning Disk Roles Manually

If you clicked **Advanced configuration** while creating a cluster or adding nodes to it, you will be taken to the list of drives on the node where you can manually assign roles to these drives. Do the following:

- On the Join cluster or New cluster panel, select a drive or check multiple drives in the list and click Configure.
- 2. On the **Choose role** screen, select a disk role:

Choose role Storage Metadata Cache Metadata+Cache Unassigned Cance Cache Cache

- **Storage**. Use the disk to store chunks and run a chunk service on the node. From the **Caching and checksumming** drop-down list, select one of the following:
 - Use SSD for caching and checksumming. Available and recommended only for nodes with SSDs.
 - **Enable checksumming** (default). Recommended for nodes with HDDs as it provides better reliability.
 - **Disable checksumming**. Not recommended for production. For an evaluation or testing environment, you can disable checksumming for nodes with HDDs, to provide better performance.

Data caching improves cluster performance by placing the frequently accessed data on an SSD.

Data checksumming generates checksums each time some data in the cluster is modified. When this data is then read, a new checksum is computed and compared with the old checksum. If the two are not identical, a read operation is performed again, thus providing better data reliability and integrity.

If a node has an SSD, it will be automatically configured to keep checksums when you add a node to a cluster. This is the recommended setup. However, if a node does not have an SSD drive, checksums will be stored on a rotational disk by default. It means that this disk will have to handle double the I/O, because for each data read/write operation there will be a corresponding checksum read/write operation. For this reason, you may want to disable checksumming on nodes without SSDs to gain performance at the expense of checksums. This can be especially useful for hot data storage.

To add an SSD to a node that is already in the cluster (or replace a broken SSD), you will need to release the node from the cluster, attach the SSD, choose to join the node to the cluster again, and, while doing so, select **Use SSD for caching and checksumming** for each disk with the role **Storage**.

With the **Storage** role, you can also select a tier from the **Tier** drop-down list. To make better use of data redundancy, do not assign all the disks on a node to the same tier. Instead, make sure that each tier is evenly distributed across the cluster with only one disk per node assigned to it. For more information, see the Installation Guide.

Note: If the disk contains old data that was not placed there by Acronis Cyber Infrastructure, the disk will not be considered suitable for use in Acronis Cyber Infrastructure.

- Metadata. Use the disk to store metadata and run a metadata service on the node.
- **Cache**. Use the disk to store write cache. This role is only for SSDs. To cache a specific storage tier, select it from the drop-down list. Otherwise, all tiers will be cached.
- **Metadata+Cache**. A combination of two roles described above.
- Unassigned. Remove the roles from the disk.

Take note of the following:

- If a physical server has a system disk with the capacity greater than 100GB, that disk can be additionally assigned the **Metadata** or **Storage** role. In this case, a physical server can have at least 2 disks.
- It is recommended to assign the **System+Metadata** role to an SSD. Assigning both these roles to an HDD will result in mediocre performance suitable only for cold data (e.g., archiving).
- The **System** role cannot be combined with the **Cache** and **Metadata+Cache** roles. The reason is that is I/O generated by the operating system and applications would contend with I/O generated by journaling, negating its performance benefits.
- 3. Click Done.
- 4. Repeat steps 1 to 3 for every disk you want to be used in the storage cluster.

5. Click **NEW CLUSTER** or **JOIN CLUSTER**. On the **Configuration summary** screen, check the number of disks per each configuration category.

METADATA SERVICE					
	Metadata			1	
STOR	AGE SERVICE	SSD cache	Same disk cache	No cache	
	Tier 0	1	0	0	
	Tier 1	3	0	0	
	Tier 2	0	1	0	
SSD (CACHE				
	Tier 0			1	
	Tier 1			1	
	PROCEEL	þ	CANCEL		

× Configuration summary

 Click PROCEED. You can monitor disk configuration progress in the HEALTHY list of the INFRASTRUCTURE > Nodes screen.

2.4 Connecting Remote iSCSI Devices to Storage Cluster Nodes

Acronis Cyber Infrastructure allows you to connect remote iSCSI devices to nodes and perceives their LUNs as storage disks. You can connect iSCSI devices to nodes at any time.

Note the following:

1. Remote iSCSI devices cannot be attached to hosts that belong to iSCSI target groups.

2. Only one IQN can be used per node. You can connect multiple iSCSI targets using other cluster nodes.

To connect a remote iSCSI device to a node, do the following:

1. On the **INFRASTRUCTURE** > **Nodes** screen, select a node, open its **DISKS** tab, and click **iSCSI target**.

Iqn.2014-06.com.vstorage:target1	
IP address: Port	+ Add portal
10.10.10.10	
20.20.20.20	
CHAP authentication (optional)	
Login	
user1	
Password	

× Remote iSCSI Target

- 2. In the **Remote iSCSI Target** window, do the following:
 - 2.1. Specify the IQN of the target.
 - 2.2. In the **Portal** and **Port** fields, specify the target's IP address and port (optional) and click the check icon.
 - 2.3. (Optional) If the target has multiple paths, click **Add portal** and configure it as in the previous step.
 - 2.4. (Optional) If necessary, check **CHAP authentication** and specify the credentials.
 - 2.5. Click Connect.

The target will be connected and all its LUNs will be initiated. Devices of the **iSCSI** type will appear in the node's **DISKS** list.

Note: If you add another LUN to a connected target, rescan it by running iscsiadm -m node -R in host's console.

To remove an iSCSI target, click **iSCSI Target**, **DELETE CONNECTION**, and **DELETE**.

2.4.1 Assigning Disk Roles To Remote iSCSI Devices

If the node had already been in the cluster before you connected the iSCSI device to it, assign disk roles to all its LUNs. To do this:

- 1. Select a disk with the **iSCSI** type and click **Assign**.
- 2. In the **Choose role** window, select **Storage** and click **Done**.
- 3. Repeat the above steps for every disk with the **iSCSI** type.

Even though you can assign **Metadata** or **Cache** roles to such disks, it is recommended only for single-node installations for Backup Gateway with SAN-provided redundancy. For more information on disk roles, see *Assigning Disk Roles Manually* (page 16).

2.5 Replacing Node Disks

When you replace a failed disk in a storage cluster node, you need to assign the same roles to it. Most roles need to be assigned by hand. For **Storage** disks (CS), however, this can be done automatically to speed up replacement.

2.5.1 Replacing Storage Disks

To have the **Storage** role assigned to replacement CS disks automatically, do the following:

- 1. Navigate to **SETTINGS** > **Advanced settings** > **DISK**.
- 2. Enable **Enable automatic configuration of new disks**.

- 3. Optionally, select tiers to scan for failed disks. If a disk on an unselected tier breaks down, you will have to assign roles to its replacement manually.
- 4. Click Save.



From now on, when you replace a failed **Storage** disk, the new disk will be detected, formatted, assigned the same role, and placed on the same tier (if applicable). You will see the result on node's **Disks** screen.

Take note of the following:

- If the new disk is larger than the old one, it will be assigned the **Storage** role.
- If the new disk is smaller than the old one, it will not be assigned the **Storage** role. Instead, you will be alerted about the size difference and will have to assign the role manually (or change the disk to a larger one).
- If the new disk is of a different type than the old one (e.g., if you replaced an SSD with an HDD or vice versa), it will not be assigned the **Storage** role. Instead, you will be alerted about the type difference and will have to assign the role manually (or change the disk to one of the needed type).
- If you enable this feature after a disk fails, its replacement will not be assigned the **Storage** role.
- If you accidentally remove and then re-attach a healthy **Storage** disk, its data will be reused.
- If you add a disk that does not replace any failed ones, it will not be assigned the **Storage** role.
- If you add a disk and one of the CSes is inactive or offline, the disk will be assigned the **Storage** role and a new CS will be created.
- If you attach multiple replacement disks at once, the **Storage** role will be assigned to them in no particular order, as long as their size and type fit. They will also be assigned to correct tiers, if applicable.

2.5.2 Replacing Disks Manually

To manually replace a disk and assign roles to the new one, do the following:

- 1. Open **INFRASTRUCTURE** > **Nodes** > <node> > **Disks**.
- 2. Select the failed disk, click **Release**.

DISKS NETWORK		
Q Search	Any st	iscsI target
Disk	S/N 👸	Details
sda	c13b77f6e4364141b989	Performance
sdc	6145366080be40328b40	郑 Release 🕁
✓ sdb	d99e21c7be6d4bdda0d0	ିଦ୍ରୁ [:] Blink

3. In the **Release disk** window, click **YES**.

× Release disks
Releasing disks triggers data migration from them as well as cluster replication and rebalancing to meet the configured redundancy levels. After the data has been migrated, the disks become unassigned and can be detached or re-used for other roles.
Forced release
Are you sure you want to release these disks?
YES NO

4. Replace the disk with a new one.

NETWORK

DISKS

5. Back on **INFRASTRUCTURE** > **Nodes** > <node> > **Disks**, select the unassigned disk and click **Assign**.

Q Search		Any status ~ Any rol	iSCSI target
Disk	Status	S/N 👸	Details
sda	🕑 ок	4e99968f275a4	Performance
sdc	🕑 ок	703075a61f9f4e	🗞 Assign 🖑
✓ sdb	🕑 ок	0f1dbadf0bf546	°Q⁼ Blink

6. On the **Choose role** screen, select a disk role:

Choose role Storage Metadata Cache Metadata+Cache Unassigned Cance Cache Cache

- **Storage**. Use the disk to store chunks and run a chunk service on the node. From the **Caching and checksumming** drop-down list, select one of the following:
 - Use SSD for caching and checksumming. Available and recommended only for nodes with SSDs.
 - **Enable checksumming** (default). Recommended for nodes with HDDs as it provides better reliability.
 - **Disable checksumming**. Not recommended for production. For an evaluation or testing environment, you can disable checksumming for nodes with HDDs, to provide better performance.

Data caching improves cluster performance by placing the frequently accessed data on an SSD.

Data checksumming generates checksums each time some data in the cluster is modified. When this data is then read, a new checksum is computed and compared with the old checksum. If the two are not identical, a read operation is performed again, thus providing better data reliability and integrity.

If a node has an SSD, it will be automatically configured to keep checksums when you add a node to a cluster. This is the recommended setup. However, if a node does not have an SSD drive, checksums will be stored on a rotational disk by default. It means that this disk will have to handle double the I/O, because for each data read/write operation there will be a corresponding checksum read/write operation. For this reason, you may want to disable checksumming on nodes without SSDs to gain performance at the expense of checksums. This can be especially useful for hot data storage.

To add an SSD to a node that is already in the cluster (or replace a broken SSD), you will need to release the node from the cluster, attach the SSD, choose to join the node to the cluster again, and, while doing so, select **Use SSD for caching and checksumming** for each disk with the role **Storage**.

With the **Storage** role, you can also select a tier from the **Tier** drop-down list. To make better use of data redundancy, do not assign all the disks on a node to the same tier. Instead, make sure that each tier is evenly distributed across the cluster with only one disk per node assigned to it. For more information, see the Installation Guide.

Note: If the disk contains old data that was not placed there by Acronis Cyber Infrastructure, the disk will not be considered suitable for use in Acronis Cyber Infrastructure.

- Metadata. Use the disk to store metadata and run a metadata service on the node.
- **Cache**. Use the disk to store write cache. This role is only for SSDs. To cache a specific storage tier, select it from the drop-down list. Otherwise, all tiers will be cached.
- **Metadata+Cache**. A combination of two roles described above.
- Unassigned. Remove the roles from the disk.

Take note of the following:

- If a physical server has a system disk with the capacity greater than 100GB, that disk can be additionally assigned the **Metadata** or **Storage** role. In this case, a physical server can have at least 2 disks.
- It is recommended to assign the **System+Metadata** role to an SSD. Assigning both these roles to an HDD will result in mediocre performance suitable only for cold data (e.g., archiving).
- The **System** role cannot be combined with the **Cache** and **Metadata+Cache** roles. The reason is that is I/O generated by the operating system and applications would contend with I/O generated by journaling, negating its performance benefits.
- 7. Click Done.

The disk will be assigned the chosen role and added to the cluster.

2.6 Performing Node Maintenance

Whenever you need to perform service operations on a cluster node, place it in the maintenance mode. When you do so, the node stops allocating new chunks of storage data, but continues to handle I/O operations for the core storage services such as MDS, CS, and cache. Other node's services (compute, Backup Gateway, iSCSI, S3, and NFS) can either be relocated or left as is during maintenance. Once the node is in the maintenance mode, you can shut it down and perform service operations on it. Once you are done, power on the node and return it to operation in the admin panel.

Important: It is recommended to have five MDS services in the storage cluster. In this case, when a node running MDS service is shut down during maintenance, the cluster can survive failure of another node.

Before placing a node in the maintenance mode, do the following:

- If the node hosts virtual machines, they will be relocated. Make sure that other compute nodes have enough resources to accommodate these VMs.
- If the node hosts iSCSI targets, make sure that iSCSI initiators are configured to use multiple IP addresses from the same target group.
- If the node runs an S3 gateway, remove its IP addresses from DNS records of S3 access points. Otherwise, some of S3 clients may experience connection timeouts.

To place a node in the maintenance mode, do the following:

- 1. On the **INFRASTRUCTURE** > **Nodes** screen, click the desired node.
- 2. On the node overview screen, click Enter maintenance.
- 3. In the **Enter maintenance** window, choose to **Evacuate** or **Ignore** the following workloads during maintenance:
 - **Block storage**. iSCSI target groups are highly available, with multiple targets running on different nodes. When the node enters maintenance, the target it hosts is stopped and the preferred path is moved to another node in the target group within 60 seconds. Thus the service is not interrupted during maintenance.
- **Block storage (version 2.4 or earlier)**. Old iSCSI targets created in version 2.4 or earlier will be evacuated from the node and migrated back after maintenance. To avoid this, it is recommended to convert old targets to new target groups as described in *Managing Legacy iSCSI Targets* (page 148).
- **Compute**. Evacuating virtual machines from the node means migrating them live one by one to other compute nodes. If you choose to ignore them, they will continue running until you reboot or shut down the node. In this case, they will be stopped and shelved, resulting in downtime. They also will not be started automatically once the node is up again.

Important: Suspended VMs cannot be evacuated from the node and will be ignored.

- **S3**. You can evacuate S3 services from this node to other nodes in the S3 cluster or ignore them. In the latter case, they will continue running until you reboot or shut down the node, resulting in downtime. They will be started automatically once the node is up again.
- **NFS**. You can evacuate NFS services from this node to other nodes in the NFS cluster or ignore them. In the latter case, they will continue running until you reboot or shut down the node, resulting in downtime. They will be started automatically once the node is up again.
- **ABGW**. This service is highly available, with multiple instances spread across different nodes. Placing this node in maintenance mode will stop one of the instances but the others will continue working, so the service will not be interrupted.

Cluster self-healing is automatic restoration of storage cluster data that becomes unavailable when a storage node (or a disk) goes offline. If this happens during maintenance, self-healing is delayed (for 30 minutes by default) to save cluster resources. If the node goes back online before the delay ends, self-healing is not necessary.

You can manually configure the replication timeout by setting the mds.wd.offline_tout_mnt parameter, in milliseconds, with the vstorage -c <cluster_name> set-config command.

In addition, any non-redundant chunks of data on the node will become unavailable if the node goes offline. They will, however, be moved to other storage nodes if you check the box **Relocate non-redundant data**. They may also be temporarily moved to another tier if the current one is full.

In general, all CSes on the node will continue serving data even in the maintenance mode unless the node goes offline. They will not, however, be used to allocate new data, so placing the node in maintenance may reduce the free space in the storage cluster.

Enter maintenance	×
Choose how the storage and compute workloads will be managed while the node " node003 " in maintenance.	is
Block storage O Evacuate O Ignore	
All iSCSI target groups on this node are highly available and have multiple targets spread across different nodes. Putting this node to maintenance will stop one of the targets but the others will continue working, so the block storage service will not be interrupted. Make sure that iSCSI initiators are configured to use multiple IP addresses from the same target group.	
Compute O Evacuate O Ignore	
All virtual machines on this node will continue running until you reboot or shut down the node. In this case, they will be stopped, resulting in downtime. They also will not be started automatically once the node is up again.	
Start data healing After 30 minutes ()	
Self healing for storage data is enabled for this node. Configure it before putting the node to maintenance.	
Cancel	

4. Click Enter.

If a service cannot be evacuated from the node for some reason, entering maintenance will be halted. You will need to decide on how to proceed: exit maintenance so that all services on the node are returned to their normal state; or force maintenance so that the services that could not be evacuated are stopped during node reboot or shutdown. On the node overview screen, click **Enter maintenance**, choose the desired action, and click **Continue**.

Entering maintenance halted

X

Failed to evacuate NFS services

🛕 🔹 Show resources (1) 👗

🗄 share1

The node "**node003**" cannot enter maintenance due to the issue shown above. Choose how to proceed and click "**Continue**". Alternatively, close this window, solve the issue manually, and retry.



Any services that have been prepared for maintenance will return to their normal state.

Force maintenance

The services will continue running until you reboot or shut down the node. In this case, they will be stopped, resulting in downtime. They will be started automatically once the node is up again.



Nodes in maintenance can be returned to operation or released.

To return a node to operation, click **Exit maintenance** on its overview screen.

N	odes > node003 MAINTENANCE			ADD NODE
				Exit maintenance
	CPU cores 8	CPU usage 100% 75%	Memory usage 4.7 3.7 2.8	Release
	RAM 15.51 GB	25%6 0%6 66.00 66.00 72.00 12.00	1.9 0.9 0.6B 0.6D 65 ^D 55 ^D 72 ^D 4.0 ^D	

2.7 Releasing Nodes from the Storage Cluster

To release a node means to remove it from the cluster (e.g., for maintenance). As the node may be running services needed by the cluster, do the following prior to releasing it to avoid cluster degradation:

- If the node runs one of the three required metadata services, add a metadata role to another node.
 You need to make sure that the cluster has at least three metadata services running at any time.
- 2. If the node has any access points, make sure that the same access points are configured on other nodes in the cluster as well.
- 3. If the node is in an iSCSI target group, remove it from the target group first.
- 4. If the node has an S3 or backup gateway, reconfigure DNS for S3 and Backup Gateway access points to remove the node from DNS records. Next, release the node from the S3 and Backup Gateway clusters in the corresponding sections of the **STORAGE SERVICES** screen.
- 5. If the node is in the compute cluster, remove it from the compute cluster first.
- 6. Make sure the cluster has enough storage space to accommodate the data from the released node.

Once you initiate the release, the cluster will start replicating data chunks that were stored on the released node and distributing them among other storage nodes in the cluster. Depending on the amount of data to replicate, the process may take as much as several hours.

If necessary, you can also release a node forcibly, that is, without replication.

Warning: Releasing nodes forcibly may result in data loss.

To release a node from a cluster, do the following:

- 1. On the **INFRASTRUCTURE** > **Nodes** screen, click the node to release.
- 2. On the node overview screen, click **Release**.
- 3. If necessary, in the **Release** node window, check force to release the node forcibly (highly not recommended).
- Click Yes. The released node will return to the UNASSIGNED list on the INFRASTRUCTURE > Nodes screen.

× Release node
Releasing a node triggers data migration from the node as well as cluster replication and rebalancing to meet the configured redundancy levels. After the data has been migrated, the node becomes "Unassigned" and can be powered off or completely removed from the system.
Force release
Are you sure you want to release this node?
YES NO

2.8 Removing Nodes from the Unassigned List

Nodes in the **UNASSIGNED** list can be completely removed from Acronis Cyber Infrastructure if they are not in the high availability cluster.

To completely remove a node from the admin panel, do the following:

- 1. Select it in the **UNASSIGNED** list on the **INFRASTRUCTURE** > **Nodes** screen and click **Remove (forget)**.
- 2. For security purposes, clean up node certificates and identity by deleting the following from the node:

```
# rm -rf /usr/libexec/vstorage-ui-backend/ca
# rm -rf /etc/nginx/ssl
# rm -f /etc/vstorage/host_id
# rm -f /etc/vstorage/vstorage-ui-agent.conf
```

Note: After such a cleanup, the only way to add the node back to the cluster is by re-installing Acronis Cyber Infrastructure on it from scratch.

2.9 Re-Adding Nodes to the Unassigned List

Nodes removed from Acronis Cyber Infrastructure can be re-added to the **UNASSIGNED** list in two ways:

• By logging in to the node via SSH and running

/usr/libexec/vstorage-ui-agent/bin/register-storage-node.sh -m MN_ADDRESS -t TOKEN in the node's console (MN_ADDRESS is the management node IP address and TOKEN is the token obtained in the admin panel).

Note: You can only do this if you have not cleaned up the node as described in *Removing Nodes from the Unassigned List* (page 32).

• By reinstalling Acronis Cyber Infrastructure on the node from scratch.

CHAPTER 3

Monitoring the Storage Cluster

Acronis Cyber Infrastructure uses the Prometheus monitoring system to monitor performance and availability of both the entire cluster and its components. It also generates alerts, which you can configure to be sent as notifications via e-mail.

3.1 Monitoring the Entire Cluster

The overall storage cluster statistics are available on the **MONITORING** > **Dashboard** screen. Pay attention to the storage cluster status that can be one of the following:

HEALTHY

All cluster components are active and operate normally.

UNAVAILABLE

Not enough information about the cluster state (e.g., because the cluster is inaccessible).

DEGRADED

Some of the cluster components are inactive or inaccessible. The cluster is trying to heal itself, data replication is scheduled or in progress.

ERROR

The cluster has too many inactive services, automatic replication is disabled. If the cluster enters this state, troubleshoot the nodes or contact the support team.

To view the storage cluster statistics in full screen, click **Fullscreen mode**. To exit the fullscreen mode, press **Esc** or **Exit fullscreen mode**.

For advanced monitoring, click **Grafana dashboard**. A separate browser tab will open with preconfigured Grafana dashboards where you can manage existing dashboards, create new ones, share them between users, configure alerting, etc. The dashboards use the Prometheus data source. Its metrics are stored for 7 days. If you want to increase this retention period, you can configure it manually as described in Configuring Retention Policy for Prometheus Metrics. For more information, refer to Grafana documentation .



The default time interval for the charts is 12 hours. To zoom into a particular time interval, select the internal with the mouse; to reset zoom, double click any chart.

3.1.1 I/O Activity Charts

The **Read** and **Write** charts show the history of the cluster I/O activity as the speed of read and write I/O operations in megabytes per second and the number of read and write I/O operations per second (IOPS). For example:



3.1.2 Services Chart

On the **Services** chart, you can monitor two types of services:

- Metadata services (MDS). The number all disks with the metadata role. Ensure that at least three MDSes are running at all times.
- Chunk services (CS). The number of all disks with the storage role.

Typical statistics may look like this:



If some of the services were not in the healthy state for some time, these time periods will be highlighted in red on the chart.

3.1.3 Chunks Chart

You can monitor the state of all chunks in the cluster on the **Chunks** chart. Chunks can be in the following 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.

Healthy chunks are highlighted on the scale in green, offline in red, blocked in yellow, and degraded in grey. For example:

Chunks		
100 % Heal	thy	
Offline 0	Blocked 0	Degraded 0
Replicatior	1	
Read 0 B/s	Write 0 B/s	ETA O

The **Replication** section shows the information about replication activity in the cluster.

3.1.4 Physical Space Chart

The **Physical space** chart shows the current usage of physical space in the entire storage cluster and on each particular tier. The used space includes the space occupied by all data chunks and their replicas plus the space occupied by any other data.

Physical space Total: 3.21 TiB free of 3.6 TiB	
1.15 TiB free of 1.44 TiB	Tier 0
1.41 TIB free of 1.44 TIB	Tier 1
664.67 GiB free of 737.8 GiB	Tier 2

3.1.4.1 Understanding Physical Space

The total physical disk space is a total of all the disk space on all storage disks on the same tier. The used physical space is a total of all the user data on the storage disks of the same tier, considering the redundancy mode. The free disk space is the total physical space minus the used physical space.

To better understand how physical disk space is calculated, consider the following example:

	Used/Total (Free), GiB		
	Tier 0, 3+2 encodingTier 1, 2 replicasTier 2, no redunda		Tier 2, no redundancy
	(67% overhead)	(100% overhead)	
Node 1	334/1024 (690)	134/512 (378)	50/256 (206)
Node 2	334/1024 (690)	133/512 (379)	50/256 (206)
Node 3	334/1024 (690)	133/512 (379)	
Node 4	334/1024 (690)		
Node 5	334/1024 (690)		
Reported	1670/5120 (3450)	400/1536 (1136)	100/512 (412)
summary			

Table 3.1.4.1.1: Physical space example

The cluster has ten disks with the storage role: five 1024 GiB disks are assigned to tier 0, three 512 GiB disks to tier 1, and two 256 GiB disk to tier 2. There is no other data on the disks (like system files, for example). Tier 0 stores 1000 GiB of user data in the 3+2 encoding mode. Tier 1 stores 200 GiB of user data in the 2 replicas mode. Tier 2 stores 100 GB of user data with no redundancy.

No matter what redundancy mode is used, the cluster attempts to spread data chunks evenly across disks of the same tier.

In this example, the physical disk space on each tier is reported as follows:

- On tier 0, the total disk space is 5120 GiB, the used disk space is 1670 GiB, and the free disk space is 3450 GiB;
- On tier 1, the total disk space is 1536 GiB, the used disk space is 400 GiB, and the free disk space is 1136 GiB;
- On tier 2, the total disk space is 512 GiB, the used disk space is 100 GiB, and the free disk space is 456 GiB.

3.1.5 Logical Space Chart

The **Logical space** chart represents all the space allocated to different services for storing user data. This includes the space occupied exclusively by user data. Replicas and erasure coding metadata are not taken into account.



3.1.5.1 Understanding Logical Space

When monitoring disk space information in the cluster, keep in mind that logical space is the amount of free disk space that can be used for storing user data in the form of data chunks and all their replicas. Once this space runs out, no data can be written to the cluster.

To better understand how logical disk space is calculated, consider the following example:

- The cluster has three disks with the storage role. The first disk has 200 GB of space, the second one has 500 GB, and the third one has 1 TB.
- If the redundancy mode is set to three replicas, each data chunk must be stored as three replicas on three different disks with the storage role.

In this example, the available logical disk space will be 200 GB, that is, equal to the capacity of the smallest disk with the storage role. The reason is that each replica must be stored on a different disk. So once the space on the smallest disk (i.e. 200 GB) runs out, no new chunk replicas can be created unless a new disk with the storage role is added or the redundancy mode is changed to two replicas.

With the two replicas redundancy mode, the available logical disk space would be 700 GB, because the two smallest disks combined can hold 700 GB of data.

3.2 Monitoring Nodes

Nodes added to the infrastructure are listed on the **INFRASTRUCTURE** > **Nodes** screen, grouped by status. If the storage cluster has not been created yet, you will only see nodes in the **UNASSIGNED** list. If the storage cluster exists, its nodes will be listed on the screen.

A node can have one of the following statuses:

UNASSIGNED

The node is not assigned to a cluster.

HEALTHY

All the storage services on the node are running.

ENTERING MAINTENANCE...

The node is entering maintenance. The services it hosts are either being evacuated or stopped.

ENTERING MAINTENANCE HALTED

The node cannot enter maintenance, because some of its services cannot be evacuated.

IN MAINTENANCE

The node is in maintenance mode. It does not participate in new chunk allocation.

EXITING MAINTENANCE...

The node is exiting maintenance. Nodes exiting maintenance cannot be managed.

OFFLINE

The node cannot be reached from the admin panel, although it may still be up and its services may be running.

FAILED

One or more storage services on the node have failed.

The default time interval for the charts is 12 hours. To zoom into a particular time interval, select the internal with the mouse; to reset zoom, double click any chart.

3.2.1 Understanding Node Role Icons

On the **Nodes** screen, nodes included in the storage cluster are shown with small icons that represent their roles. Icons provide an overview of cluster infrastructure and the status of some services on each node. All existing node icons are listed below with their description.

lcon	Role	Description	
	Management node	The node hosts cluster management services and the	
		admin panel. The primary node in the infrastructure.	
	Storage	The node has disks with the storage role. It runs chunk services, stores all data, and provides access to it. In case a CS fails, the icon changes its color to red.	
	Metadata	The node has disks with the metadata role. It runs metadata services, stores cluster metadata, controls the amount of chunk replicas, and logs important cluster events. In case an MDS fails, the icon changes its color to red.	
	Master Metadata	The master node in the metadata quorum. If the master MDS fails, another available MDS is selected as master.	
0	Backup Gateway	The node runs the Backup Gateway service and participates in the Backup Gateway cluster.	
iscsi	iSCSI	The node hosts iSCSI targets and exports storage space over the iSCSI protocol.	
53	53	The node participates in the S3 cluster and exports storage space over the S3 protocol.	

Table 3.2.1.1: List of existing node icons

3.2.2 Monitoring Node Performance

To monitor the performance of a cluster node, open the **Nodes** screen and click the node. On the node overview screen, you will see performance statistics described below.

The overall statistics include:

- the number of CPUs and the amount of RAM,
- CPU usage, in percent over time,
- RAM usage, in megabytes or gigabytes over time.



The **DISKS** section shows:

- the number of HDD and SSD drives and their statuses,
- node I/O activity over time on the read and write charts.



The **NETWORK** section shows:

· the list of network interfaces and their statuses,

• the amount of transmitted (TX) and received (RX) traffic over time.



The following sections provide more information on disk and network usage.

3.2.3 Monitoring Node Disks

To monitor the usage and status of node disks, click the **DISKS** link on the node overview screen. You will see a list of all disks on the node and their status icons.

A disk status icon shows the combined status of S.M.A.R.T. and the service corresponding to the disk role. It can be one of the following:

OK The disk and service are healthy.

Failed

The service has failed or S.M.A.R.T. reported an error.

Releasing

The service is being released. When the process finishes, the disk status will change to OK.

To monitor performance of a particular disk, select it and click **Performance**. The **Drive performance** panel will display the I/O activity of the disk.

To view information about the disk, including its S.M.A.R.T. status, click **Details**.

To have the disk blink its activity LED, select the disk, and click **Blink**. To have the disk stop blinking, click **Unblink**.

3.2.3.1 Monitoring the S.M.A.R.T. Status of Node Disks

The S.M.A.R.T. status of all disks is monitored by a tool installed along with Acronis Cyber Infrastructure. Run every 10 minutes, the tool polls all disks attached to nodes, including journaling SSDs and system disks, and reports the results to the management node.

For the tool to work, make sure the S.M.A.R.T. functionality is enabled in node's BIOS.

If a S.M.A.R.T. warning message is shown in the node status, one of that node's disks is in pre-failure condition and should be replaced. If you continue using the disk, keep in mind that it may fail or cause performance issues.

Pre-failure condition means that at least one of these S.M.A.R.T. counters is not zero:

- Reallocated Sector Count
- Reallocated Event Count
- Current Pending Sector Count
- Offline Uncorrectable

3.2.4 Monitoring Node Network

To monitor the node's network usage, click **NETWORK** on the node overview screen.

To display the performance charts of a specific network interface, select it in the list and click **Performance**. When monitoring network performance, keep in mind that if the **Receive and transmit errors** chart is not empty, the network is experiencing issues and requires attention.

To display the details of a network interface, click **Details**. The **Network details** panel shows the interface state, bandwidth, MTU, MAC address, and all IP addresses.

3.3 Monitoring Storage Cluster Objects via SNMP

You can monitor cluster objects via the Simple Network Management Protocol (SNMP). The implementation conforms to the same Structure of Management Information (SMI) rules as the data in the standard SNMP

context: all objects are organized in a tree; each object identifier (OID) is a series of integers corresponding to tree nodes and separated by dots.

General information:

- The OID of the root subtree with all the objects you can monitor is 1.3.6.1.4.1.8072.161.1.
- The VSTORAGE-MIB.txt information base file is required to monitor the objects. You can download the file at https://<admin_panel_IP>:8888/api/v2/snmp/mibs/.

The following subsections describe ways to enable and use SNMP to monitor cluster objects.

3.3.1 Enabling SNMP Access

To monitor cluster objects, enable the SNMP access on the node. Do the following in the admin panel:

- 1. Open UDP port 161 on the management node as follows:
 - 1.1. On the **INFRASTRUCTURE** > **Networks** screen, click **Edit**.
 - 1.2. Add the **SNMP** traffic type to your public network by ticking the corresponding checkbox.
 - 1.3. Click **Save** to apply changes.
- On the SETTINGS > Advanced settings > SNMP tab, check Enable SNMP on the management node. The network management system (SNMP monitor) will be enabled, giving you access to the cluster via the SNMP protocol.

Encryption	SNMP	Notifications	RDMA
Enable SNMF	on the m	anagement node	
You can acces Download MI Download Za	ss SNMP or B file bbix templ	n the management ate	node on port 161
Send SNM	1P traps to	this network mana	agement system
Save			

- 3. Click the provided link to download the MIB file and set it up in your SNMP monitor.
- 4. If required, have Acronis Cyber Infrastructure send SNMP traps to your SNMP monitor. Do the following:
 - 4.1. Check Send SNMP traps to this network management system.
 - 4.2. Specify the **IP address**, **Port**, and **Community** of the network management system.

By default, the snmptrapd daemon uses port 162. The default community is public.

- 4.3. If required, click **Send test trap** to test the service.
- 5. Click **Save** to apply changes.

3.3.2 Accessing Storage Cluster Information Objects via SNMP

You can access storage cluster information objects with SNMP tools of your choice, e.g., the free Net-SNMP suite for Linux.

To obtain storage cluster information on a node with the admin panel, place the MIB file to /usr/share/snmp/mibs and run the snmpwalk command. For example:

```
# snmpwalk -M /usr/share/snmp/mibs -m VSTORAGE-MIB -v 2c -c public \
localhost:161 VSTORAGE-MIB:cluster
```

Typical output may be the following:

```
VSTORAGE-MIB::clusterName.0 = STRING: "cluster1"
VSTORAGE-MIB::healthStatus.0 = STRING: "healthy"
VSTORAGE-MIB::usedLogicalSpace.0 = Counter64: 173732322
VSTORAGE-MIB::totalLogicalSpace.0 = Counter64: 1337665179648
VSTORAGE-MIB::freeLogicalSpace.0 = Counter64: 1318963253248
VSTORAGE-MIB::licenseStatus.0 = STRING: "unknown"
VSTORAGE-MIB::licenseCapacity.0 = Counter64: 1099511627776
VSTORAGE-MIB::licenseExpirationStatus.0 = STRING: "None"
VSTORAGE-MIB::ioReadOpS.0 = Counter64: 0
VSTORAGE-MIB::ioWriteOpS.0 = Counter64: 0
VSTORAGE-MIB::ioReads.0 = Counter64: 0
VSTORAGE-MIB::ioWrites.0 = Counter64: 0
VSTORAGE-MIB::csActive.0 = Counter64: 11
VSTORAGE-MIB::csTotal.0 = Counter64: 11
VSTORAGE-MIB::mdsAvail.0 = Counter64: 4
VSTORAGE-MIB::mdsTotal.0 = Counter64: 4
```

3.3.2.1 Listening to SNMP Traps

To start listening to SNMP traps, do the following:

 Configure the snmptrapd daemon to log SNMP traps, allow them to trigger executable actions, and resend data to the network. To do this, uncomment the following public community string in the /etc/snmp/snmptrapd.conf file:

authCommunity log,execute,net public

- 2. Configure the firewall to allow inbound traffic on UDP port 162.
- 3. Download the VSTORAGE-MIB.txt file and place it to the /usr/share/snmp/mibs directory.
- 4. Start the daemon and specify the MIB file:

snmptrapd -M /usr/share/snmp/mibs -m VSTORAGE-MIB -n -f

By default, traps will be logged to /var/log/messages. You can redirect them to a custom log file with the -Lf <path> option. For example:

snmptrapd -M /usr/share/snmp/mibs -m VSTORAGE-MIB -n -f -Lf /tmp/traps.log

- 5. Send a test trap from the **SETTINGS** > **Advanced settings** > **SNMP** tab in the admin panel.
- 6. View the log file:

tail -f /tmp/traps.log 2019-10-14 12:51:50 node001.vstoragedomain [UDP: [10.94.80.22]:40029->\ [10.94.80.22]:162]:#012DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: \ (111150521) 12 days, 20:45:05.21#011SNMPv2-MIB::snmpTrapOID.0 = OID: \ NET-SNMP-MIB::netSnmp.161.3.100#011NET-SNMP-MIB::netSnmp.161.2.1 = STRING: "TestTrap"\ #011NET-SNMP-MIB::netSnmp.161.2.2 = STRING: "It is the test trap from VStorage"\ #011NET-SNMP-MIB::netSnmp.161.2.3 = Counter64: 0

3.3.3 Monitoring the Storage Cluster with Zabbix

To configure cluster monitoring in Zabbix, do the following:

 On the SETTINGS > Advanced settings > SNMP tab, click the provided link to download a template for Zabbix.

Note: The template is compatible with Zabbix 3.x.

2. In Zabbix, click **Configuration** > **Templates** > **Import** and **Browse**.

Import file	Browse vstorage.	xml		
Rules	Upda	te existing Cre	ate new Delei	te missing
	Groups			
	Hosts			
	Templates			
	Template screens			
	Template linkage			
	Applications			
	Items			
	Discovery rules			
	Triggers			
	Graphs			
	Web scenarios			
	Screens			
	Maps			
	Images			
	Value mappings			
	Import Cancel			

- 3. Navigate to the template, select it, and click **Import**.
- 4. Click **Configuration** > **Hosts** > **Create host**.

Hostname	management.node.hostname]
Visible name	Storage Cluster]
Groups	In groups Other groups	
	Discovered hosts Hypervisors Linux servers Templates Virtual machines Zabbix servers	
New group	vstorage]
Agent interfaces	IP address DNS name Connect to Port Default Add	
SNMP interfaces	ID.250.14.15 Use bulk requests Add	DNS 161
JMX interfaces	Add	
IPMI interfaces	Add	
Description		
Monitored by proxy	(no proxy) -	
Enabled		
	Add Cancel	

- 5. On the **Host** tab, do the following:
 - 5.1. Specify the **Host name** of the management node and its **Visible name** in Zabbix.
 - 5.2. Specify vstorage in the **New group** field.
 - 5.3. **Remove** the **Agent Interfaces** section.
 - 5.4. **Add** an **SNMP interfaces** section and specify the management node IP address.
- 6. On the **Templates** tab, click **Select** next to the **Link new templates** field.
- 7. In the **Zabbix Server: Templates** window, check the Template VStorageSNMP template and click **Select**.



8. Back on the **Templates** tab, click the **Add** link in the **Link new templates** section. The VStorageSNMP template will appear in the **Linked templates** group.

Linked templates	Name	Action	
Link new templates	Template VStorageSNMP × type here to search		Select
	Add Add Cancel		

9. Having configured the host and added its template, click the **Add** button.

Linked templates	Name Template VStorageSNMP	Action Unlink
Link new templates	type here to search Add	Select
	Add	

In a few minutes, the cluster's SNMP label in the **Availability** column on the **Configuration** > **Hosts** screen will turn green.

Name 🔺	Applications	Items	Triggers	Graphs	Discovery	Web	Interface	Templates	Status	Availability	Agent encryption	Info
Cluster	Applications 2	ltems 32	Triggers 7	Graphs 3	Discovery 2	Web	10.250.14.15: 161	VStorageSNMP	Enabled	ZBX SNMP JMX IPMI	NONE	

To monitor cluster's parameters, open the **Monitoring** > **Latest data** screen, set the filter's **Host groups** to vstorage and click **Apply**.

You can create performance charts on the **Configuration** > **Hosts** > **<cluster>** > **Graphs** tab and a workplace for them on the **Monitoring** > **Screens** tab.

3.3.4 Storage Cluster Objects and Traps

Cluster-related objects that you can monitor:

VSTORAGE-MIB:cluster

General cluster information.

VSTORAGE-MIB:csStatTable

Chunk server statistics table.

VSTORAGE-MIB:mdsStatTable

Metadata server statistics table.

VSTORAGE-MIB::clusterName

Cluster name.

VSTORAGE-MIB::healthStatus

Cluster health status.

VSTORAGE-MIB::usedLogicalSpace

The space occupied by all data chunks and their replicas plus the space occupied by any other data stored on cluster nodes' disks.

VSTORAGE-MIB::totalLogicalSpace

The total space on all cluster nodes' disks.

VSTORAGE-MIB::freeLogicalSpace

The unused space on all cluster nodes' disks.

VSTORAGE-MIB::licenseStatus

License status.

VSTORAGE-MIB::licenseCapacity

The maximum disk space available as defined by license.

VSTORAGE-MIB::licenseExpirationStatus

License expiration status.

VSTORAGE-MIB::ioReadOpS

Current read speed in operations per second.

VSTORAGE-MIB::ioWriteOpS

Current write speed in operations per second.

VSTORAGE-MIB::ioReads

Current read speed in bytes per second.

VSTORAGE-MIB::ioWrites

Current read write in bytes per second.

VSTORAGE-MIB::csActive

The number of active chunk servers.

VSTORAGE-MIB::csTotal

The total number of chunk servers.

VSTORAGE-MIB::mdsAvail

The number of running metadata servers.

VSTORAGE-MIB::mdsTotal

The total number of metadata servers.

VSTORAGE-MIB::s3OsAvail

The number of running S3 object servers.

VSTORAGE-MIB::s3OsTotal

The total number of S3 object servers.

VSTORAGE-MIB::s3NsAvail

The number of running S3 name servers.

VSTORAGE-MIB::s3NsTotal

The total number of S3 name servers.

VSTORAGE-MIB::s3GwAvail

The number of running S3 gateways.

VSTORAGE-MIB::s3GwTotal

The total number of S3 gateways.

CS-related objects that you can monitor:

VSTORAGE-MIB::csId

Chunk server identifier.

VSTORAGE-MIB::csStatus

Current chunk server status.

VSTORAGE-MIB::csIoReadOpS

Current read speed of a chunk server in operations per second.

VSTORAGE-MIB::csIoWriteOpS

Current write speed of a chunk server in operations per second.

VSTORAGE-MIB::csloWait

The percentage of time spent waiting for I/O operations. Includes time spent waiting for synchronization.

VSTORAGE-MIB::csloReadS

Current read speed of a chunk server in bytes per second.

VSTORAGE-MIB::csloWriteS

Current write speed of a chunk server in bytes per second.

MDS-related objects you can monitor:

VSTORAGE-MIB::mdsId

Metadata server identifier.

VSTORAGE-MIB::mdsStatus

Current metadata server status.

VSTORAGE-MIB::mdsMemUsage

The amount of memory used by a metadata server.

VSTORAGE-MIB::mdsCpuUsage

The percentage of the CPU's capacity used by a metadata server.

VSTORAGE-MIB::mdsUpTime

Time since the startup of a metadata server.

SNMP traps triggered by the specified alerts:

license expired

The license has expired.

license_isnot_loaded

The license is not loaded.

too few free space

The cluster is running out of logical space.

too_few_free_phys_space

The cluster is running out of physical space.

offline node

A cluster node is offline.

too few nodes

Too few cluster nodes are left.

too few mdses

Too few MDSes are left.

too_much_mdses

More than one MDS is on a node.

too few cses

Too few CSes are left.

failed mds

The MDS service has failed.

failed cs

The CS service has failed.

cses_on_single_tier_have_different_journalling_settings

A CS has incorrect journalling settings.

cses_on_single_tier_have_different_encryption_settings

A CS has incorrect encryption settings.

smart_failed

A disk has failed a S.M.A.R.T. check.

disk_failed

A disk has failed.

too_few_root_space

The root partition on a node is out of space.

too_few_space_on_metadata_disk

An MDS disk is out of space.

low_level_network_settings

A network interface is missing important features.

half_duplex

A network interface is not in the full duplex mode.

low_speed

A network interface has speed lower than 1 Gbps.

undefined_speed

A network interface has an undefined speed.

network link

A network interface is misconfigured.

abgw_cert_expired

Backup Gateway certificate has expired or will expire soon.

iscsi_redundancy_disk

The failure domain set for an iSCSI LUN does not make it highly available.

s3_redundancy_disk

The failure domain set for an S3 cluster does not make it highly available.

software_updates

Software updates exist for a node.

no_internet_connection

No internet connection on a node.

disk_write_cache_disabled

Disk write cache is disabled.

disk_write_cache_status_unknown

Disk write cache has an unknown status.

compute_unavailable

The compute cluster has failed.

oom_happened

OOM killer has been triggered.

kernel_not_current

The kernel is outdated on a node.

no_ha

High availability for the admin panel is not configured.

time_not_synced

Time is not synced on a node.

iscsi_upgrade_failed

iSCSI major upgrade has failed.

backend_backup_is_too_old

The last management node backup has failed, does not exist, or is too old.

other

Other alerts.

spla_push_stats_failed

Unable to push space usage statistics.

spla_license_load_failed

Unable to apply a SPLA license.

spla_get_space_usage_failed

Unable to get space usage.

To see the full list of generated alerts with their descriptions, refer to Viewing Alerts (page 62).

3.4 Monitoring Storage Cluster Remotely

You can monitor your storage cluster via Prometheus remotely. To do this, you need to open a TCP port for Prometheus API to be accessible from the outside.

To open a port, do the following:

- 1. On the **INFRASTRUCTURE** > **Networks** screen, click **Edit** and then **Create traffic type**.
- 2. In the Create traffic type window, specify a custom name in the Name field and 9090 in the Port field.

Name		
Prometheus		
Port		
9090		

- 3. Click Create.
- 4. Add the newly created traffic type to your public network by ticking the corresponding checkbox.
- 5. Click **Save** to apply the changes.

You can now access the built-in Prometheus web-based user interface at http://<admin_panel_IP_address>:9090. For more information on using Prometheus, refer to its documentation.

If you have an external Grafana account and want to use it for monitoring Acronis Cyber Infrastructure, you can add Prometheus as a data source as follows:

- 1. Log in into your Grafana user interface.
- 2. Click the cogwheel icon in the left menu and select **Data Sources**.
- 3. On the **Data Sources** tab, click **Add data source**.
- 4. On the **Data Sources / New** screen, specify the following parameters:
 - 4.1. Enter a custom data source name in the **Name** field.
 - 4.2. Set **Type** to Prometheus.
 - 4.3. Enter http://<admin_panel_IP_address>:9090 in the URL field.
- 5. Click Save & Test.

If the specified parameters are correct, the Data source is working message will appear.

Promethe Promethe http://172 Server (D	RUS RUS 2.16.25.3:9090 efault) With Credentials With CA Cert	8		0 0	Default Help >	
Promethe Promethe http://172 Server (D	RUS 2.16.25.3:9090 efault) With Credentials With CA Cert	6		0 • •	Default Help >	
Promethe http://172 Server (Do	EUS 2.16.25.3:9090 efault) With Credentials With CA Cert	6	0	•	Help >	
http://172 Server (D	2.16.25.3:9090 efault) With Credentials With CA Cert	0	0	0	Help +	
http://172 Server (D	2.16.25.3:9090 efault) With Credentials With CA Cert	0	0	0	Help +	
Server (D	efault) With Credentials With CA Cert	0	0	•	Help +	
	With Credentials	0	0			
0	With Credentials With CA Cert	0 0				
0	With CA Cert	0				
			0			
ation (Insecu ITP Setti kies	ings Id Name (9					
	0					
	0					
GET	• 0					
urce is wor	king					
	TTP Setti dies Ad 15s 60s GET Nurce is wor	ation (Insecure)	ation (Insecure)	ation (Insecure)	ation (Insecure)	ation (Insecure)

Using the newly added Prometheus data source, you can import the default Grafana dashboards from Acronis Cyber Infrastructure or create new ones.

3.5 Viewing Alerts and Audit Log and Sending E-mail Notifications

This section describes Acronis Cyber Infrastructure alerts and audit log and how to send out e-mail notifications about alerts, warnings, and errors.

3.5.1 Viewing Alerts

The **Alerts** tab lists all the alerts logged by Acronis Cyber Infrastructure. An alert is generated and logged each time one of the following conditions is met or events happen:

- a critical issue has happened with a cluster, its components (CS, MDS), disks, nodes, or services;
- cluster requires configuration or more resources to build or restore its health;
- network requires configuration or is experiencing issues that may affect performance;
- · license has expired;
- cluster is about to or has run out of available space.

Q Search			
Туре	Message V	Date and time	Resource
	No internet connection on the node node003	Nov 23, 2018, 4:21 PM	node
	No internet connection on the node node002	Nov 23, 2018, 4:21 PM	node
	No internet connection on the node node001	Nov 23, 2018, 4:16 PM	node
	Network interface "eth1" on node "node003.v	Nov 23, 2018, 4:46 PM	node

To view the details of an alert, select it on the **MONITORING** > **Alerts** screen and click **Details** in the menu on the right.

Alerts can be ignored (deleted from the alerts list) or postponed for several hours. Postponed alerts reappear in the list after some time.

To ignore or postpone an alert, select it and click the corresponding button.

The following alerts are generated and displayed in the admin panel:

Title	Message	Severity					
License alerts							
License is not loaded	License is not installed.	warning					
License expired	The license of cluster " <cluster_name>" has expired.</cluster_name>	critical					
	Contact your reseller to update your license						
	immediately!						
Cluster alerts							
Cluster is out of space	Cluster has just <free_space> TB</free_space>	warning					
	(<free_space_in_percent>%) of physical storage</free_space_in_percent>						
	space left. You may want to free some space or add						
	more storage capacity.						
	Cluster " <cluster_name>" has run out of storage</cluster_name>	warning					
	space allowed by license. No more data can be						
	written. Please contact your reseller to update your						
	license immediately!						
Not enough cluster	Cluster " <cluster_name>" has only {1,2} node(s)</cluster_name>	warning					
nodes	instead of the recommended minimum of 3. Add						
	{2,1} or more nodes to the cluster.						
High availability for the	Configure high availability for the admin panel in	critical					
admin panel must be	SETTINGS > Management node. Otherwise the						
configured	admin panel will be a single point of failure.						
The last management	Management node backup is older than	critical					
node backup has failed,	<number_of_days> days.</number_of_days>						
does not exist, or is too	Management node backup does not exist.	critical					
old							
Metadata service alerts							
Not enough metadata	Cluster " <cluster_name>" has only one MDS. There</cluster_name>	warning					
disks	is only one disk with the metadata role at the						
	moment. Losing this disk will completely destroy all						
	cluster data irrespective of the redundancy schema.						
	Cluster " <cluster_name>" requires more disks with</cluster_name>	warning					
	the metadata role. Losing one more MDS will halt						
	cluster operation.						

Continued on next page
Title	Message	Severity
Configuration warning	Node " <hostname>" has more than one metadata</hostname>	warning
	service located on it. It is recommended to have	
	only one metadata service per node. Delete the	
	extra metadata service(s) from this node and create	
	them on other nodes instead.	
Service failed	Metadata service # <id> is in the "<status>" state.</status></id>	warning
	Node: <hostname>. Disk: <disk_name>. Disk serial:</disk_name></hostname>	
	<disk_serial>.</disk_serial>	
Metadata disk is out of	Metadata disk on node " <hostname>" is running</hostname>	warning
space	out of space.	
Chunk service alerts		
Not enough disks with	Cluster " <cluster_name>" has no disks with the</cluster_name>	warning
storage role	storage role.	
	Cluster " <cluster_name>" has too few available</cluster_name>	warning
	CSes.	
Service failed	Storage service # <id> is in the "<status>" state.</status></id>	warning
	Node: <hostname>. Disk: <disk_name>. Disk serial:</disk_name></hostname>	
	<disk_serial>.</disk_serial>	
CS configuration is not	CS# <cs_id> on tier <tier> has incorrect journalling</tier></cs_id>	warning
optimal	settings.	
	Encryption is disabled for CS# <cs_id> on tier <tier></tier></cs_id>	warning
	but is enabled for other CSes on the same tier.	
Node alerts		
Node is offline	Node " <hostname>" is offline.</hostname>	warning
Software updates exist	Software updates are available for the node	warning
	" <hostname>".</hostname>	
Kernel is outdated	Node " <hostname>" is not running the latest kernel.</hostname>	warning
OOM killer triggered	OOM killer has been triggered on node	warning
	" <hostname>".</hostname>	
Time is not synced	Time on node " <hostname>" is out of sync with the</hostname>	warning
	management node.	
Disk alerts		

I able 2.5.1.1 continued from providue p	
ab = 3.3.1.1 = CONTINUED II ON DIEVIOUS D	age

Continued on next page

Title	Message	Severity
S.M.A.R.T. warning	Disk " <disk_name>"(<serial>) on node</serial></disk_name>	critical
	" <hostname>" has failed a S.M.A.R.T. check.</hostname>	
Disk error	Disk " <disk_name>" (<serial>) failed on node</serial></disk_name>	critical
	" <hostname>".</hostname>	
Disk is out of space	Root partition on node " <hostname>" is running out</hostname>	warning
	of space.	
Disk write cache	Disk write cache is disabled for disk " <disk_name>"</disk_name>	warning
disabled	on node " <hostname>".</hostname>	
Disk write cache status	Cannot determine the status of write cache for disk	warning
unknown	" <disk_name>" on node "<hostname>".</hostname></disk_name>	
Network alerts		
No internet connection	No Internet connection on the node " <hostname>".</hostname>	warning
Network warning	Network interface " <iface_name>" has incorrect</iface_name>	warning
	settings: <duplex> duplex and <speed> speed.</speed></duplex>	
	Network interface " <iface_name>" on node</iface_name>	warning
	" <hostname>" is missing important features (or has</hostname>	
	them disabled): " <feature_name>".</feature_name>	
	Network interface " <iface_name>" on node</iface_name>	warning
	" <hostname>" is not in the full duplex mode.</hostname>	
	Network interface " <iface_name>" on node</iface_name>	warning
	" <hostname>" has speed lower than the minimally</hostname>	
	required 1 Gbps.	
	Network interface " <iface_name>" on node</iface_name>	warning
	" <hostname>" has an undefined speed.</hostname>	
Other alerts		
Compute cluster is	Compute cluster has failed. Management of virtual	critical
failed	machines may be unavailable.	

Table 3.5.1.1 – continued from previous page	j
--	---

Continued on next page

Title	Message	Severity
Redundancy warning	iSCSI LUN <lun_id> of target group</lun_id>	warning
	" <target_group>" is set to failure domain "disk" even</target_group>	
	though <number_of_nodes> nodes are available. It</number_of_nodes>	
	is recommended to set the failure domain to "host"	
	so that the LUN can survive host failures in addition	
	to disk failures.	
	S3 is set to failure domain "disk" even though	warning
	<number_of_nodes> nodes are available. It is</number_of_nodes>	
	recommended to set the failure domain to "host"	
	so that S3 can survive host failures in addition to	
	disk failures.	
Certificate expiration	Acronis Backup Gateway certificate has expired or	warning
	will expire soon.	
iSCSI major upgrade	iSCSI major upgrade failed. Will be retried	critical
failed		
SPLA license alerts		
Unable to apply SPLA	Unable to apply SPLA license to cluster	critical
license	" <cluster_name>". Contact your reseller to solve</cluster_name>	
	this issue!	
Unable to push space	Unable to push space usage statistics for cluster	warning
usage statistics	" <cluster_name>". Make sure that the management</cluster_name>	
	node is connected to the Internet.	
Unable to get space	Unable to get space usage for cluster	warning
usage	" <cluster_name>".</cluster_name>	

Table	3.5.1.1	_	continued	from	previous page	2
10010	0.0.1.1		contaca		previous puge	-

3.5.2 Viewing Audit Log

The **MONITORING** > **Audit log** screen lists all management operations performed by users and their activity events.

Q Search			Show more details
Date and time 🗸	User	Activity	Ś
Nov 23, 2018, 4:36 PM	admin	Create virtual machine	Start creation of vir
Nov 23, 2018, 4:34 PM	admin	Create virtual machine	Start creation of vir
Nov 23, 2018, 4:19 PM	admin	Deploy compute cluster	Start deployment o
Nov 23, 2018, 4:18 PM	admin	Join node	Join node "node003

To view detailed information on a log entry, select it and click **Show more details**.

3.5.3 Sending E-mail Notifications

Acronis Cyber Infrastructure can send automatic e-mail notifications about errors, warnings, and alerts.

To set up e-mail notifications, do the following:

- 1. On the **SETTINGS** > **Advanced settings** > **NOTIFICATIONS** tab, specify the following information:
 - 1.1. In the **From** and **Sender name** fields, the notification sender's e-mail and name.
 - 1.2. In the **To** field, one or more notification recipient e-mails, one per line.
 - 1.3. In the **User account** and **User password** fields, the credentials of the notification sender registered on the SMTP server.
 - 1.4. In the **SMTP server** field, the DNS name of the SMTP server, either public (e.g., smtp.gmail.com) or the one in your organization.

The management node must be able to access the SMTP server.

- 1.5. If required, a custom **SMTP port** the server uses.
- 1.6. In the **Security** field, the security protocol of the SMTP server.

From potifior@oxample.com			To
notinei @example.com	1		user2@example.com
Sender name			user3@example.com
Event Notifier			
Licor account			Test
notifier			
			Send notifications about
User password			Errors
•••••			✓ Warnings
			✓ Information
SMTP server			
smp.example.com			
SMTP port	Security		
465	SSL	ř	

2. Tick the checkboxes for alerts you want to be notified about.

3. Click Save.

To send a test e-mail, specify your e-mail registered on the SMTP server in both the **From** and **To** fields and click **Test**.

CHAPTER 4

Managing the Compute Cluster

4.1 Creating the Compute Cluster

Before creating a compute cluster, make sure the network is set up according to recommendations in *Managing Networks and Traffic Types* (page 2). The basic requirements are: (a) the traffic types **VM private**, **VM public**, **Compute API**, and **VM backups** must be assigned to networks; (b) the nodes to be added to the compute cluster must be connected to these networks and to the same network with the **VM public** traffic type.

Warning: The **Compute API** and **VM private** traffic types cannot be reassigned after the compute cluster deployment.

Important: The **VM public** traffic type cannot be removed from a network that has a public virtual network created on top of it.

Besides, high availability for the management node should also be enabled (see *Enabling High Availability* (page 237)).

Also take note of the following:

1. Creating the compute cluster prevents (and replaces) the use of the management node backup and restore feature.

2. If nodes to be added to the compute cluster have different CPU models, consult Setting Virtual Machines CPU Model.

To create the compute cluster, open the **COMPUTE** screen, click **Create compute cluster** and do the following in the **Configure compute cluster** window:

Configure compute cluster

1. In the **Nodes** section, select nodes to add to the compute cluster, make sure the network state of each selected node is **Configured**, and click **Next**.

Nodes in the management node high availability cluster are automatically selected to join the compute cluster.

8						
Nodes	Select nodes to add to the compute cluster.					
Public network	Searc	h	Q			
Add-on services		Name 🕇	Node status	IP address	Network state	
		node001 🛈	Healthy	10.37.130.60	🥝 Configured	¢
• Summary	Z	node002	Healthy	10.37.130.67	🥝 Configured	¢
	~	node003	Healthy	10.37.130.190	Configured	¢
					Ne	xt

If node network interfaces are not configured, click the cogwheel icon, select networks as required, and click **Apply**.

Note: The compute cluster must have at least three nodes to allow self-service users to enable high availability for Kubernetes master nodes.

2. In the **Public network** section, enable IP address management if needed and provide the required details for the public network.

With IP address management enabled, Acronis Cyber Infrastructure will handle virtual machine IP addresses and provide the following features:

• Allocation pools. You can specify ranges of IP addresses that will be automatically assigned to VMs.

 \times

- Built-in DHCP server. Assigns IP addresses to virtual machines. With the DHCP server enabled, VM
 network interfaces will automatically be assigned IP addresses: either from allocation pools or, if
 there are no pools, from network's entire IP range. With the DHCP server disabled, VM network
 interfaces will still get IP addresses, but you will have to manually assign them inside VMs.
- Custom DNS servers. You can specify DNS servers that will be used by VMs. These servers will be delivered to virtual machines via the built-in DHCP server.

With IP address management disabled:

- VMs connected to a network will be able to obtain IP addresses from DHCP servers in that network.
- Spoofing protection will be disabled for all VM network ports. Each VM network interface will accept all traffic, even frames addressed to other network interfaces.

In any case, you will be able to manually assign static IP addresses from inside VMs.

If you choose to enable IP address management, select a physical network to connect the public virtual network to and optionally specify its gateway. The subnet IP range in the CIDR format will be filled in automatically. If you choose to leave IP address management disabled, select a physical network to connect the public virtual network to.

By default, the public network will be shared between all future projects. You can disable this option on the network panel after the compute cluster is created.

Nodes	Specify the subnet CIDR and gateway for the public virtual network.
Public network	IP address management
DHCP and DNS	Physical network Public
• Add-on services	Subnet CIDR 10.94.0.0/16
• Summary	Gateway (optional) 10.94.0.1
	Back Next

Configure compute cluster

The selected public network will appear in the list of virtual networks on compute cluster's **NETWORKS** tab.

×

Click Next.

 If you enabled IP address management on the previous step, you will move on to the DHCP and DNS section. In it, enable or disable the built-in DHCP server and specify one or more allocation pools and DNS servers. Click Next.

Configure compute cluster		×
Nodes	Set DHCP and specify one or more allocation pools for the public virtual network.	
Public network	Enable the built-in DHCP server.	
DHCP and DNS	Allocation pools	+ Add pool
Add-on services	10.94.129.128 — 10.94.129.255 128 addresses available	<i>0</i> ū
• Summary	DNS servers	+ Add server
	10.94.0.10	Ø Ū
		Back Next

4. In the **Add-on services** section, enable additional services that will be installed during the compute cluster deployment. You can also install these services later (see *Managing Add-On Services* (page 231)).

Х

Configure compute cluster

• Nodes Add-on services • Public network Image: Constant additional services for your compute cluster: • DHCP and DNS Image: Constant additional services service • Add-on services Image: Constant additional service allows you to deploy scalable and production-ready Kubernetes clusters with pre-integrated persistent storage. • Summary Image: Constant additional service analyses workload scaling and improves application availability and security. Image: Constant additional service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gonocchi API. Image: Constant additional service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gonocchi API. Image: Constant additional service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gonocchi API.		
 Public network DHCP and DNS Add-on services Summary Load balancer service The load balancer service enables workload scaling and improves application availability and security. Billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. 	Nodes	Add-on services You can install additional services for your compute cluster.
 DHCP and DNS Add-on services Summary Load balancer service enables workload scaling and improves application availability and security. Billing metering service The builling metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. 	Public network	
• Add-on services The Kubernetes service allows you to deploy scalable and production-ready Kubernetes clusters with pre-integrated persistent storage. • Summary	DHCP and DNS	✤ Kubernetes service
Summary	Add-on services	The Kubernetes service allows you to deploy scalable and production-ready Kubernetes clusters with pre-integrated persistent storage.
Summary Load balancer service The load balancer service enables workload scaling and improves application availability and security. Security. Set Billing metering service The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. Back Next	e Cummon	
The load balancer service enables workload scaling and improves application availability and security. Signature Billing metering service The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. Back	• Summary	Load balancer service
Billing metering service The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. Back Next		The load balancer service enables workload scaling and improves application availability and security.
Billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API.		
The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API. Back Next		Billing metering service
Back Next		The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The meters can be accessed via the Gnocchi API.
Back Next		
		Back Next

Important: To be able to deploy and work with Kubernetes clusters, make the following services accessible:

- the etcd discovery service at https://discovery.etcd.io—from all management nodes and the public network with the VM public traffic type
- the public Docker Hub repository at https://registry-1.docker.io—from the public network with the **VM public** traffic type
- the compute API—from the public network with the **VM public** traffic type
- the Kubernetes API at the public or floating IP address of the Kubernetes load balancer or master VM on port 6443—from all management nodes

If the **Compute API** traffic type is added to a private network that is inaccessible directly from the network with the **VM public** traffic type but exposed to public networks via NAT and available publicly

via the DNS name, you need to set the DNS name for the compute API as described in Setting a DNS Name for the Compute API.

Note: Installing Kubernetes automatically installs the load balancer service as well.

5. In the **Summary** section, review the configuration and click **Create cluster**.

Configure compute cluster

-			
Nodes	Review the compute cluster detail:	s and go back to change them if nece	essary.
Public network	Nodes	node001 (10.37.130.60) node003 (10.37.130.190) node002 (10.37.130.67)	
DHCP and DNS	Subnet CIDR	10.94.0.0/16	
Add-on services	Gateway	10.94.0.1	
	Physical network	Public	
• Summary	DHCP	Enabled	
	Allocation pools	10.94.129.128 — 10.94.129.255	128 addresses available
	DNS servers	10.94.0.10	
	Add-on services	Kubernetes service Load balancer service Billing metering service	
			Back Create cluster

You can monitor compute cluster deployment on the **Compute** screen.

4.2 Monitoring the Compute Cluster

After you create the compute cluster, you can monitor it on the **COMPUTE** > **Overview** screen.

The compute cluster status is displayed on top of the screen and can be one of the following:

HEALTHY

All compute cluster components and nodes operate normally.

 \times

CONFIGURING

The compute cluster configuration (the default CPU model for VMs or the number of compute nodes) is changing.

WARNING

The compute cluster operates normally but some issues have been detected.

CRITICAL

The compute cluster has encountered a critical problem and is not operatonal.

The charts show the information on CPU, RAM, and storage usage; the number of virtual machines grouped by status and resource consumption; and compute-related alerts.

4.2.1 Used CPUs Chart

This chart displays CPU utilization of the compute cluster. The following statistics are available:

System

The number of logical cores used by system and storage services on all nodes in the compute cluster.

VMs The number of logical cores used by virtual machines on all nodes in the compute cluster.

Free The number of unused logical cores on all nodes in the compute cluster.

Fenced

The number of CPUs on all fenced nodes in the compute cluster.

Total

The total number of logical cores on all nodes in the compute cluster.

Provisioned vCPUs

The number of vCPUs provisioned for all VMs in the compute cluster.



A similar chart is available for each individual node in the compute cluster.

4.2.2 Reserved RAM Chart

This chart displays RAM utilization of the compute cluster. The following statistics are available:

System

The amount of RAM reserved for system and storage services on all nodes in the compute cluster.

VMs The amount of RAM provisioned for all VMs in the compute cluster.

Free The amount of free RAM on all nodes in the compute cluster.

Fenced

The amount of RAM on all fenced nodes in the compute cluster.

Total

The total amount of RAM on all nodes in the compute cluster.

Used by VMs

The amount of RAM actually used by all VMs in the compute cluster.



A similar chart is available for each individual node in the compute cluster.

4.2.3 Provisioned Storage Chart

This chart shows usage of storage space by the compute cluster. The following statistics are available:

Used

The amount of storage space actually occupied by data in all volumes provisioned in the compute cluster.

Free The amount of unused space in all volumes provisioned in the compute cluster.

Total

The total size of volumes provisioned in the compute cluster.

Free physical space

The amount of physical space available in the storage cluster.



4.2.4 VM Status Chart

The **VMs status** chart shows the total number of virtual machines in the compute cluster and groups them by status, which can be the following:

Running

The number of virtual machines that are up and running.

In progress

The number of virtual machines that are in a transitional state: building, restarting, migrating, etc.

Stopped

The number of virtual machines that are suspended or powered off.

Error

The number of virtual machines that have failed. You can reset state for such VMs to their last stable state.

VMs status 2 Total VMs	
Running 2	In progress • 0
Stopped	Error 0

To see a full list of virtual machines filtered by the chosen status, click the number next to the status icon.

4.2.5 Top VMs Chart

The **Top VMs** chart lists virtual machines with the highest resource consumption sorted by **CPU**, **RAM**, or **Storage** in descending order. To switch between lists, click the desired resource.

Top VMs		
vCPU RAM	Storage	
🕅 vm1		1
₩ vm2		1
	Show all	

To see a full list of virtual machines in the compute cluster, click **Show all**.

4.2.6 Alerts Chart

The **Alerts** chart lists all the alerts related to the compute cluster sorted by severity. Alerts include the following:

Critical

The compute cluster has encountered a critical problem. For example, one or more of its components have been unavailable for more than 10 seconds or some resource has exceeded its soft limit.

Warning

The compute cluster is experiencing issues that may affect its performance. For example, one or more of its components operate slowly or some resource is approaching its soft limit.

Other

Some other issue has happened with the compute cluster. For example, its license is about to expire or has expired.

To see a full list of compute-related alerts, click **Show all**.

4.3 Managing Compute Nodes

To make your infrastructure more resilient and redundant, you can create a high availability configuration of three nodes.

Management node HA and compute cluster are tightly coupled, so changing nodes in one usually affects the other. Take note of the following:

- 1. Each node in the HA configuration must meet the requirements to the management node listed in the Installation Guide. If the compute cluster is to be created, its hardware requirements must be added as well.
- 2. If the HA configuration has been created before the compute cluster, all nodes in it will be added to the compute cluster.
- 3. If the compute cluster has been created before HA configuration, only nodes in the compute cluster can be added to the HA configuration. For this reason, to add a node to HA configuration, add it to the compute cluster first.
- 4. If both the HA configuration and compute cluster include the same three nodes, single nodes cannot be removed from the compute cluster. In such a case, the compute cluster can be destroyed completely, but the HA configuration will remain. This is also true vice versa, the HA configuration can be deleted, but the compute cluster will continue working.

Note: The compute cluster must have at least three nodes to allow self-service users to enable high availability for Kubernetes master nodes.

Nodes in the compute cluster are shown on the **Nodes** screen.

Clicking a node, you can see the following information about it:

- node CPU and RAM usage,
- node name, status, and IP address,
- hosted virtual machines and their resource consumption.



The next subsections describe how to add nodes to and remove nodes from the compute cluster.

4.3.1 Adding Nodes to Compute Cluster

Note: Before changing nodes in the compute cluster, see limitations in *Managing Compute Nodes* (page 80).

To add one or more nodes to your compute cluster, do the following:

- 1. Click Add node on the Nodes screen. The Add node window will open.
- 2. If required, configure network on each node not marked green: click the cogwheel icon, assign networks with the compute-related traffic types to node NICs, and click **Apply**.
- 3. Select nodes and click Add.

Add	node			×
Choos	e one or more nodes to add to the compute cluster			
	Name ↑	IP address	Network state 🕇	
	node002.vstoragedomain.	10.37.130.28	🕑 Configured	٥
			Cancel	bb

The added nodes will appear on the **Nodes** screen.

If several nodes are in the management node HA group, they all must be added to the compute cluster.

4.3.2 Releasing Nodes from Compute Cluster

Note: Before changing nodes in the compute cluster, see limitations in *Managing Compute Nodes* (page 80).

To release one or more nodes from the compute cluster, do the following:

- 1. On the **Nodes** screen, either
 - select the nodes and click Release nodes above the list, or
 - click the ellipsis icon next to a node and select Release, or
 - click a node to open its details, then click **Release node** on the top toolbar.
- 2. In the **Release node** window, confirm the action by clicking **Release**.

The selected nodes will disappear from the **Nodes** screen.

If the node to be released has VMs on it, they must be migrated to other nodes first.

4.4 Managing Virtual Machines

Each virtual machine (VM) is an independent system with an independent set of virtual hardware. Its main features are the following:

- A virtual machine resembles and works like a regular computer. It has its own virtual hardware. Software applications can run in virtual machines without any modifications or adjustment.
- Virtual machine configuration can be changed easily, e.g., by adding new virtual disks or memory.
- Although virtual machines share physical hardware resources, they are fully isolated from each other (file system, processes, sysctl variables) and the compute node.
- A virtual machine can run any supported guest operating system.

The following table lists the current virtual machine configuration limits:

Resource	Limit
RAM	1 TiB
CPU	48 logical CPUs
Storage	15 volumes, 512 TiB each
Network	15 NICs

	Table 4.4.1:	Virtual	machine	hardware
--	--------------	---------	---------	----------

A logical CPU is a core (thread) in a multicore (multithreading) processor.

4.4.1 Supported Guest Operating Systems

The following guest operating systems have been tested and are supported in virtual machines:

Operating System	Edition	Architecture
Windows Server 2019	Essentials, Standard, Datacenter	x64
Windows Server 2016	Essentials, Standard, Datacenter	x64

Table 4.4.1.1: Windows guest operating systems

Continued on next page

Operating System	Edition	Architecture
Windows Server 2012 R2	Essentials, Standard, Datacenter	x64
Windows Server 2012	Standard, Datacenter	x64
Windows Server 2008 R2	Standard, Datacenter	x64
Windows Server 2008	Standard, Datacenter	x64
Windows 10	Home, Professional, Enterprise, Enterprise 2016	x64
	LTSB	
Windows 8.1	Home, Professional, Enterprise	x64
Windows 7	Home, Professional, Enterprise	x64

Table 4.4.1.1 – continued from previous page

Table 4.4.1.2: Linux guest operating systems

Operating System	Architecture
CentOS 8.x	x64
CentOS 7.x	x64
CentOS 6.x	x64
RHEL 8.x	x64
RHEL 7.x	x64
Debian 9.x	x64
Ubuntu 20.04.x	x64
Ubuntu 18.04.x	x64
Ubuntu 16.04.x	x64

4.4.2 Creating Virtual Machines

Before you proceed to creating VMs, check that you have these:

- A guest OS source (see *Managing Images* (page 104)):
 - a distribution ISO image of a guest OS to install in the VM, or
 - a template that is a boot volume in the QCOW2 format, or
 - a boot volume

Note: To obtain a boot volume, create a volume as described in *Managing Compute Volumes* (page 131), attach it to a VM, install an operating system in it, then delete the VM.

- A storage policy for volumes (see *Managing Storage Policies* (page 138))
- A flavor (see *Managing Flavors* (page 110))
- One or more virtual networks (see Managing Compute Network (page 114))
- An SSH key (see *Managing SSH Keys* (page 112))

Note: You can specify an SSH key only when creating VMs from a template or boot volume.

Note: Virtual machines are created with the host CPU model by default. Having compute nodes with different CPUs may lead to live migration issues. To avoid them, you can manually set CPU model for all new VMs as described in Setting Virtual Machines CPU Model.

To create a VM, do the following:

 On the COMPUTE > Virtual machines > VIRTUAL MACHINES tab, click Create virtual machine. A window will open where you will need to specify VM parameters.

Create virtual machine			
Review the virtual machine details and go	back to change them if necessary.		
Name vm1			
Deploy from: O Image O Volum	e		
Image	Specify	Ø	
Columes	Specify	0	
Flavor	Specify	0	
X Networks	Specify	Ø	
		Deploy	

- 2. Specify a name for the new VM.
- 3. In **Deploy from**, choose **Volume** if you have a boot volume or want to create one. Otherwise, choose **Image**.
- 4. Depending on your choice, click the pencil icon in the **Volumes** or **Image** section and do one of the following:
 - In the **Images** window, select the ISO image or template and click **Done**.

lma	ges					×
Searc	ch		Q			
		Name 🕇	Туре	Min. volume size	OS Type	Size
0		cirros	Template	1 GB	linux	13 MB
You ca	an add in	nages to this list	on the Images tab	. Then reload the page.	Cancel	Done

- In the **Volumes** window, do one of the following:
 - If you have prepared a volume with an installed guest OS, click **Attach**, find and select the volume, and click **Done**.

	Attach volume	×
	Volume vol1 (f71f6053-5b9b-4e33-8046-80b11139ab07), 1	~
	Cancel	ach
Crea	te volume	×
Nam vol1	5	
Size 1	GiB) Min. 1 GiB, Max. 512 TiB	
Stora def	ge policy ault	~
	Delete on termination	
	Cancel	Add

- 5. Optionally, in the **Volumes** window, click **Add** or **Attach** to create or attach any other volumes you need. To select a volume as bootable, place it first in the list by clicking the up arrow button next to it.
- After you select an image or a volume, the **Placement** drop-down list is displayed. Placements are created by the administrator to group nodes or VMs sharing a distinctive feature, like a special license. Select the placement corresponding to the VM characteristics. For more information, see *Managing Placements* (page 140).

7. In the **Flavor** window, choose a flavor and click **Done**.

Flav	or		×
	Name 🕆	vCPU 🕇	Memory
0	😚 tiny	1	512 MiB
	😚 small	1	2 GiB
	😚 medium	2	4 GiB
	😯 large	4	8 GiB
	😯 xlarge	8	16 GiB
You ca	an add flavors to this list on the Flavors tab. Then reload the page.	Cancel	Done

8. In the network window, click **Add**, select a virtual network interface and click **Add**. It will appear in the **Network interfaces** list.

Add network interface		×
Network public		~
	Cancel	Add

You can edit additional parameters of newly added network interfaces, like IP and MAC addresses and spoofing protection. To do this, click interface's ellipsis icon, then **Edit**, and set parameters in the **Edit network interface** window.

Network in	terfaces			×
Search Q				+ Add
Name	IP address	MAC address	Spoofing prot	
private	Auto	Auto	Yes	
			<pre> P E T T R </pre>	dit 👆
			Cancel	Done

You will not be able to edit these parameters later. Instead, you will be able to delete the old network interface and replace it with a new one.

Edit network interface	×
Name	
private	
Subnet CIDR 192.168.128.0/24	
IP address	
MAC address	
Spoofing protection	
	Cancel Save

Click **Done**.

9. (Optional) If you are deploying the VM from a template or boot volume (not an ISO image), you can specify the following:

• An SSH key to be injected into the VM. To do it, select an SSH key in the **Select an SSH key** window, and click **Done**.



Note: To be able to connect to the VM via SSH, make sure the VM template or boot volume has cloud-init and OpenSSH installed (see *Preparing Templates* (page 107)).

 User data to customize the VM after launch. You can specify user data in one of two formats: cloud-config or shell script. To do it, write a script in the **Customization script** field or browse a file on your local server to load the script from.

Provide a customization script	×
Provide user data to customize the VM after launch. User d in one of two formats: cloud-config or shell script. For the g be customizable, the template must have cloud-init installe	lata can be guest OS to ed.
Customization script #cloud-config user: myuser password: password chpasswd: {expire: False} ssh_pwauth: True	
Load from file user-data	Browse
Cancel	Save

Note: For the guest OS to be customizable, make sure the VM template or boot volume has cloud-init installed (see *Preparing Templates* (page 107)).

To inject a script in a Windows VM, refer to the Cloudbase-Init documentation. For example, you can set a new password for the account using the following script:



- 10. Back in the **Create virtual machine** window, click **Deploy** to create and boot the VM.
- 11. If you are deploying the VM from an ISO image (not a boot volume template or a volume with a pre-installed guest OS), select the VM, click **Console**, and install the guest OS using the built-in VNC console.

- 12. (Optional) If you are deploying the VM from a prepared template with an injected SSH key, you can connect to it via SSH using the username and the VM IP address:
 - For Linux templates, enter the username that is default for the cloud image OS (for example, for a CentOS cloud image, the default login is centos).
 - For Windows templates, enter the username that you specified during Cloudbase-Init installation.

For example:

ssh myuser@10.10.10.10

4.4.3 Virtual Machine Actions Overview

After you create a virtual machine, you can manage it using the actions available for its current state. To see the full list of available actions, click the ellipsis button next to a VM or on top of its panel. Actions include:

- Run powers up a VM.
- **Console** connects to running VMs via the built-in VNC console. In the console browser window, you can send a key combination to a VM, take a screenshot of the console window, and download the console log.
- **Reboot** soft-reboots a running VM.
- Shut down gracefully shuts down a running VM.
- Hard reboot cuts off and restores power, then starts a VM.
- **Power off** forcibly cuts off power from a VM.
- **Shelve** unbinds a stopped VM from the node it is hosted on and releases its reserved resources such as CPU and RAM. A shelved VM remains bootable and retains its configuration, including the IP addresses.

Virtual machines in other states can be shelved by clicking **Shut down** or **Power off** and selecting the checkbox **Shelve virtual machine** in the confirmation window.

- Unshelve spawns a shelved VM on a node with enough resources to host it.
- **Suspend** saves the current VM state to a file.

This may prove useful, for example, if you need to restart the host but do not want to quit the applications currently running in the VM or restart its guest OS.

• **Resume** restores a VM from suspended state.

• **Download console log** downloads the console log. Make sure logging is enabled inside the VM, otherwise the log will be empty (for more information, see *Enabling Logging inside Virtual Machines* (page 94)).

Examining console logs may be useful in troubleshooting failed virtual machines.

- **Reset state** resets the VM stuck in a failed or transitional state to its last stable state: active, shut down or shelved.
- Delete removes a VM from the compute cluster.
- **Migrate** moves a VM to another node in the compute cluster (for more information, see *Migrating Virtual Machines* (page 95)).

4.4.4 Enabling Logging inside Virtual Machines

VM's console log will contain log messages only if the TTY1 and TTYS0 logging levels are enabled inside the VM. For example, you can enable them as follows in Linux VMs:

- 1. Add the line GRUB_CMDLINE_LINUX_DEFAULT="console=tty1 console=ttyS0" to the file /etc/default/grub.
- 2. Depending on the boot loader, run either

grub-mkconfig -o /boot/grub/grub.cfg

or

grub2-mkconfig -o /boot/grub2/grub.cfg

3. Reboot the VM.

In Windows VMs, you can enable Emergency Management Services (EMS) console redirection for this purpose. Do the following:

- 1. Start Windows PowerShell with administrator privileges.
- In the PowerShell console, set the COM port and baud rate for EMS console redirection. As Windows VMs have only the COM1 port with the transmission rate of 9600 bps, run:

bcdedit /emssettings EMSPORT:1

3. Enable EMS for the current boot entry:

bcdedit /ems on

You may also enable driver status logging to see the list of loaded drivers. This can be useful for troubleshooting a faulty driver or long boot process. You can do this as follows:

- 1. Start System Configuration with administrator privileges.
- 2. In the **System Configuration** windows, open the **Boot** tab, select the checkboxes **OS boot information** and **Make all boot settings permanent**.
- 3. Confirm the changes and restart the system.

4.4.5 Migrating Virtual Machines

VM migration helps facilitate cluster upgrades and workload balancing between compute nodes. Acronis Cyber Infrastructure allows you to perform two types of migration:

- · Cold migration for stopped and suspended virtual machines
- Hot migration for running virtual machines (allows you to avoid VM downtime)

For both migration types, a virtual machine is migrated between compute nodes using shared storage, so no block device migration takes place.

Hot migration consists of the following steps:

- 1. All VM memory is copied to the destination node while the virtual machine keeps running on the source node. If a VM memory page changes, it is copied again.
- 2. When only a few memory pages are left to copy, the VM is stopped on the source node, the remaining pages are transferred, and the VM is restarted on the destination node.

Large virtual machines with write-intensive workloads write to memory faster than memory changes can be transferred to the destination node, thus preventing migration from converging. For such VMs, the auto-converge mechanism is used. When a lack of convergence is detected during live migration, VM's vCPU execution speed is throttled down, which also slows down writing to VM memory. Initially, virtual machine's vCPU is throttled by 20% and then by 10% during each iteration. This process continues until writing to VM memory slows down enough for migration to complete or the VM vCPU is throttled by 99%.

Note: Virtual machines are created with the host CPU model by default. Having compute nodes with different CPUs may lead to live migration issues. To avoid them, you can manually set CPU model for all new VMs as described in Setting Virtual Machines CPU Model.

To migrate a VM, do the following:

1. On the **COMPUTE** > **Virtual machines** > **VIRTUAL MACHINES** tab, click a VM to migrate, click the ellipsis button and choose **Migrate**.

OVERVIEW NODES	VIRTUAL MACHINES	VOI 〈 〉
vm1	▶ Console 💥	Reboot
a un tru		🕁 Shut down
Overview		औද Hard reboot
Details		O Power off
Status	Active	00 Suspend
Node	node003.vstorage	🕩 Migrate 🔥
Creation time	November 8, 2018	■ Download console log
SSH key	Not specified	The Delete

- 2. In the new window, specify the destination node:
 - **Auto**. Automatically select the optimal destination among cluster nodes based on available CPU and RAM resources.
 - Select the destination node manually from the drop-down list.

Migrate virtual machine	×
Migrate virtual machine "vm1 " to the following node	
Node node001.vstoragedomain.	~
Cold migration	
Cancel	Migrate

3. By default, running VMs are migrated live. You can change the migration mode to offline by ticking the **Cold migration** checkbox. A VM will be stopped and restarted on the destination node after migration.

4. Click **Migrate** to reserve resources on the destination node and start migration.

The admin panel will show the migration progress.

4.4.6 Reconfiguring and Monitoring Virtual Machines

To monitor virtual machine's CPU, storage, and network usage, select the VM and open the **Monitoring** tab.

The default time interval for the charts is 12 hours. To zoom into a particular time interval, select the internal with the mouse; to reset zoom, double click any chart.

The following performance charts are available:

CPU / RAM

CPU and RAM usage by the VM.

Network

Incoming and outgoing network traffic.

Storage read/write

Amount of data read and written by the VM.

Read/write latency

Read and write latency. Hovering the mouse cursor over a point on the chart, you can also see the average and maximum latency for that moment as well as the 95 and 99 percentiles.

To reconfigure a VM, select it and, on the **Overview** tab, click the pencil icon next to a parameter you need to change. You cannot do the following:

- · Change, detach, or delete the boot volume
- Manage non-boot volumes except attaching and detaching
- · Modify previously added network interfaces
- Attach and detach network interfaces to and from shelved VMs
- · Change the flavor for running and shelved VMs

4.4.7 Configuring Virtual Machine High Availability

High availability keeps virtual machines operational if the node they are located on fails due to kernel crash, power outage and such or becomes unreachable over the network. Graceful shutdown is not considered a failure event.

Important: The compute cluster can survive the failure of only one node.

In the event of failure, the system will attempt to evacuate affected VMs automatically, that is, migrate them offline with auto-scheduling to other healthy compute nodes in the following order:

- VMs with the "Active" status are evacuated first and automatically started.
- VMs with the "Shut down" status are evacuated next and remain stopped.
- All other VMs are ignored and left on the failed node.

If something blocks the evacuation, for example, destination compute nodes lack resources to host the affected VMs, these VMs remain on the failed node and receive the "Error" status. You can evacuate them manually after solving the issue (providing sufficient resources, joining new nodes to the cluster, etc.). To do this, click the ellipsis button next to such a VM or open its panel and click **Evacuate**.

vm2				
S Evacuate 2 Reset state	S Evacuate			
Overview	Monitoring			
Details				
Status	× Error			
Fault	No valid host was found.			
Node	node002			
Creation time	May 28, 2019 5:57 PM			
SSH key	Not specified			
VM ID	3a4d5d86-5cfa-4ac7-8bfa-d06cf2411bca			
Project ID	bbf0bad8210547b58365fc743548deb6			

When the failed node becomes available again, it is fenced from scheduling new VMs on it and can be returned to operation manually. To do it, click the ellipsis button next to the fenced node or open its panel and then click **Return to operation**.
node002	×
Return to operation	
Used CPUs	Reserved RAM
8 • System: 0.12 Total • Free: 7.88	15.51 GiB Total • System: 3.75 GiB • VMs: 4 GiB Free: 7.76 GiB
Provisioned vCPUs: 2	Used by VMs: 0 bytes
Details	
Name	node002
Status	A Fenced
IP address	10.37.130.118
Node ID	40f769cc-7d32-4bf1-9c16-fa50966f26b0

By default, high availability for virtual machines is enabled automatically after creating the compute cluster. If required, you can disable it manually as follows:

- 1. Click the VM for which you wish to disable HA.
- 2. On the VM panel, click the pencil icon next to the **High availability** parameter.
- 3. In the **High availability** window, disable HA for the VM and click **Save**.



Virtual machines with disabled HA will not be evacuated to healthy nodes in case of failover.

4.4.8 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, as well as for creating consistent snapshots of a running VM's disks (refer to *Managing Volume Snapshots* (page 136)).

4.4.8.1 Installing Guest Tools

Installation of the guest tools inside a virtual machine includes steps performed by users with different roles: a system administrator and a VM user.

As a system administrator, you need to do the following:

- 1. Upload the guest tools ISO files located in the /usr/share/vz-guest-tools/ directory on any compute node to a network share or FTP server:
 - for a Windows guest, upload vz-guest-tools-win.iso
 - for a Linux guest, upload vz-guest-tools-lin.iso
- 2. Provide access to the uploaded ISO file to a VM user.

As a VM user, log in to the virtual machine and do the following:

- Inside a Windows VM:
 - 1. Download the Windows guest tools ISO image provided by your system administrator.

- 2. Mount the image inside the VM.
 - On Windows 8 or Windows Server 2012 or newer, you can natively mount an ISO image. To do this, right-click the guest tools ISO image and select **Mount**.
 - On Windows 7 and Windows Server 2008, you need a third-party application to mount ISO images.
- 3. Go to the mounted optical drive in Explorer and install the guest tools by running setup.exe.
- 4. After the installation is complete, restart the VM.
- Inside a Linux VM:
 - 1. Download the Linux guest tools ISO image provided by your system administrator.
 - 2. Create a mount point for the optical drive with the guest tools image and run the installer:

#	mkdir	/mnt/cdrom
#	mount	<path_to_guest_tools_iso> /mnt/cdrom</path_to_guest_tools_iso>
#	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.

4.4.8.2 Uninstalling Guest Tools

If you find out that the guest tools are incompatible with some software inside a virtual machine, you can uninstall them as follows:

- Inside a Windows VM:
 - 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.

- Inside a Linux VM:
 - 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.

4.5 Managing Images

Acronis Cyber Infrastructure allows you to upload ISO images and templates that can be used to create VM volumes. An ISO image is a typical OS distribution that needs to be installed on disk. In turn, a template is a ready volume in the QCOW2 format with an installed operating system and applications and a set minimum size. Many OS vendors offer templates of their operating systems under the name "cloud images". For a list of guest OSes supported in virtual machines, see *Supported Guest Operating Systems* (page 84).

Note: Images are stored according to the default storage policy.

4.5.1 Uploading, Editing, and Removing Images

To add an image, do the following:

- 1. On the **COMPUTE** > **Virtual machines** > **IMAGES** tab, click **Add image**.
- 2. In the **Add image** window, do the following:
 - 2.1. Click **Browse** and select a template or ISO file.
 - 2.2. Specify an image name to be shown in the admin panel.
 - 2.3. Select a correct OS type from the drop-down list.

Important: OS type affects VM parameters like hypervisor settings. VMs created from an image with a wrong OS type may not work correctly, e.g., crash.

Add image	×
Image file Fedora-LXDE-Live-x86_64-27-1.6.iso	Browse
Name Fedora-LXDE-Live-x86_64-27-1.6.iso	
Select OS distribution Generic Linux	~
Share between all projects	
Ca	ancel Add

- 3. Optionally, select the **Share between all projects** checkbox. With the option disabled, the image will only be available in the **admin** project of the **Default** domain.
- 4. Click **Add** to upload the image.

The admin panel will show the upload progress.

Important: Do not reload the page while the image is being uploaded or the process will fail.

To edit an image, select it and click the pencil icon next to a parameter you need to change.

To remove an image, click the ellipsis button next to it and **Delete**.

4.5.2 Creating Volumes from Images

You can create volumes from both ISO images and templates. Do the following:

1. On the image panel, click **Create volume**.

2. In the **Create volume** window, specify the volume name, size, and choose a storage policy.

Create volume	×	
Name vol1		
Size (GiB) 10 Min. 1 GiB, Max. 512 TiB		
Storage policy default	¥	
Image: centos7-minimal		
	Cancel	

3. Click Create.

The new volume will appear on the **COMPUTE** > **Storage** > **VOLUMES** tab.

4.5.3 Mounting ISO Images to Virtual Machines

Note: This feature is supported only for Linux virtual machines. You can, however, mount an ISO image to a Windows VM via the vinfra tool. Use the insructions provided in Installing Guest Tools in Existing VMs. Replace the guest tools image in the examples with your image.

To mount an ISO image to a Linux VM, do as follows:

- 1. Create a volume from the ISO image as described in *Creating Volumes from Images* (page 105).
- 2. Attach the resulting volume to the desired VM as described in *Attaching and Detaching Volumes* (page 134).

The mounted disk will appear inside the Linux VM.

4.5.4 Preparing Templates

To be able to connect to a virtual machine via SSH or perform some of the initial configuration tasks, you need to prepare a VM template (or a boot volume for Windows VMs) before creating a VM from it. The steps you need to perform to prepare the VM template differ depending on the guest operating system and are described in the sections below.

4.5.4.1 Preparing Linux Templates

As all Linux guests have OpenSSH Server pre-installed by default, you only need to make sure a Linux template has cloud-init installed.

The easiest way to get a Linux template with cloud-init installed is to obtain it from its official repository or build one with the diskimage-builder tool. (For more information, refer to Creating Linux Templates.)

4.5.4.2 Preparing Windows Templates

Windows guests have neither Cloudbase-Init nor OpenSSH Server pre-installed by default. You need to install and configure them manually as follows:

Note: If you only need to configure Cloudbase-Init, skip step 4.

- 1. Create a VM from a Windows image as described in Creating Virtual Machines (page 85).
- 2. Log in to the VM and install the guest OS using the built-in VNC console.
- 3. Create a new administrator account that will be used for SSH connections and log in with it.
- 4. Install and configure OpenSSH Server as follows:
 - 4.1. Run Windows PowerShell with administrator privileges and set the execution policy to unrestricted to be able to run scripts:

> Set-ExecutionPolicy Unrestricted

4.2. Download OpenSSH Server (for example, from the GitHub repository), extract the archive into the C:\Program Files directory, and install it by running:

> & 'C:\Program Files\OpenSSH-Win64\install-sshd.ps1'

4.3. Start the sshd service and set its startup type to "Automatic":

```
> net start sshd
> Set-Service sshd -StartupType Automatic
```

- 4.4. Open TCP port 22 for the OpenSSH service in the Windows Firewall:
 - on Windows 8.1, Windows Server 2012, and newer versions, run

```
> New-NetFirewallRule -Protocol TCP -LocalPort 22 -Direction Inbound \
-Action Allow -DisplayName OpenSSH
```

on Windows 7, Windows Server 2008, and Windows Server 2008 R2, run

```
> netsh advfirewall firewall add rule name=sshd dir=in action=allow \
protocol=TCP localport=22
```

4.5. Open the C:\ProgramData\ssh\sshd_config file:

> notepad 'C:\ProgramData\ssh\sshd_config'

Comment out the following lines at the end of the file:

```
#Match Group administrators
#AuthorizedKeysFile __PROGRAMDATA__/ssh/administrators_authorized_keys
```

And save the changes.

4.6. Create the .ssh directory in C:\Users\<current_user> and an empty authorized_keys file inside it:

```
> cd C:\Users\<current_user>
> mkdir .ssh
> notepad .\.ssh\authorized_keys
```

Remove the .txt extension from the created file:

> move .\.ssh\authorized_keys.txt .\.ssh\authorized_keys

4.7. Modify the permissions for the created file to disable inheritance as follows:

```
> icacls .\.ssh\authorized_keys /inheritance:r
```

- 5. Download Cloudbase-Init (for example, from the official site), launch the installation, and follow the on-screen instructions:
 - 5.1. In the **Configuration options** window, enter the current username in the **Username** field:

Important: The user account password will be reset on the next VM startup. You will be able to log in with this account using the key authentication method or you can set a new password with a customization script (see *Creating Virtual Machines* (page 85)).

🗑 Cloudbase-Init 0.9.11 Setup	-		\times
Configuration options Options for guest startup initialization	cloud		-
Username:			
Use metadata password			
User's local groups (comma separated list): Administrators			
Serial port for logging:			
Run Cloudbase-Init service as LocalSystem			
Back	Next	Cano	el

5.2. When the installation is complete, do not run Sysprep and click **Finish**:



5.3. Run Windows PowerShell with administrator privileges and open the file C:\Program

Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf:

> notepad 'C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf'

Add metadata_services and plugins in two lines:

metadata_services=cloudbaseinit.metadata.services.configdrive.ConfigDriveService,\
cloudbaseinit.metadata.services.httpservice.HttpService
plugins=cloudbaseinit.plugins.common.mtu.MTUPlugin,\
cloudbaseinit.plugins.windows.ntpclient.NTPClientPlugin,\
cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin,\
cloudbaseinit.plugins.common.networkconfig.NetworkConfigPlugin,\
cloudbaseinit.plugins.windows.licensing.WindowsLicensingPlugin,\
cloudbaseinit.plugins.common.sethostname.ExtendVolumesPlugin,\
cloudbaseinit.plugins.common.setuserpassword.SetUserPasswordPlugin,\
cloudbaseinit.plugins.common.userdata.UserDataPlugin,\
cloudbaseinit.plugins.windows.winrmlistener.ConfigWinRMCertificateAuthPlugin,\
cloudbaseinit.plugins.common.localscripts.LocalScriptsPlugin

Note: Make sure to remove all backslashes in the lines above.

And save the changes.

6. Stop the VM.

After the VM shuts down, you can either

- · delete it to make its boot volume available for creating new VMs, or
- convert the VM boot volume to a template (see Creating Images from Volumes (page 135)).

4.6 Managing Flavors

A flavor in Acronis Cyber Infrastructure is a configuration template that simplifies VM deployment. It allows you to set the number of virtual CPU cores and the amount of RAM a virtual machine will use. By default, five predefined flavors are created with the following parameters:

Name	vCPUs	Memory
tiny	1	512 MiB
small	1	2 GiB
medium	2	4 GiB
large	4	8 GiB
xlarge	8	16 GiB

Table 4	.6.1:	Default	flavors
---------	-------	---------	---------

Important: When choosing a flavor for a VM, make sure it satisfies the hardware requirements of the guest OS. For more information on VM hardware, see *Managing Virtual Machines* (page 84).

To create a custom flavor, do the following:

- 1. On the **Compute > Virtual machines > FLAVORS** tab, click **Create flavor**.
- 2. In the **Create flavor** window, specify a flavor name, a number of virtual CPU cores, an amount of RAM and click **Create**.

Create flavor			×
Specify flavor paramete	rs		
Name flavor1			
Memory 8	Units GiB	~	
vCPU 4			
		Cancel	Create

To delete one or more flavors, select them and click **Delete flavors**.

4.7 Managing SSH Keys

Use of SSH keys allows you to secure SSH access to virtual machines. You can generate a key pair on a client from which you will connect to VMs via SSH. The private key will be stored on the client and you will be able to copy it to other nodes. The public key will need to be uploaded to Acronis Cyber Infrastructure and specified during VM creation. It will be injected into the VM by cloud-init and used for OpenSSH authentication. Keys injection is supported for both Linux and Windows virtual machines.

Note: You can specify an SSH key only if you deploy a VM from a template or boot volume (not an ISO image).

Before using the SSH keys feature, make sure the following requirements are met:

- The cloud-init utility is installed in a VM template or boot volume.
- OpenSSH Server is installed in a Windows template or boot volume.

For the instructions on preparing templates or boot volumes, see *Preparing Templates* (page 107).

To add a public key, do the following:

1. Generate an SSH key pair on a client using the ssh-keygen utility:

ssh-keygen −t rsa

- 2. On the Compute > Virtual machines > SSH KEYS tab, click Add key.
- 3. In the **Add SSH key** window, specify a key name and copy the key value from the generated public key located in /root/.ssh/id_rsa.pub. Optionally, you can add a key description.

Х

Add SSH key

For the key to be successfully injected into the VM, the template must contain the cloud-init package.

Name root_node001vstoragedomain

Description (optional) My public key

Key value

9MANMUTVzgDu/xFh0Nm2HKNV4GWGVAGGbGNqBfkjDBOq/wfJ OrrwXQXghgmvD+FCeGlEh3YCxeVlMS6/PgnbZEfOG9o4QianAGs8 kMrrF8zL6svL8qOvlWUxsGoJT+3WmXT+fF5OExm01XDau0vhmhT 6Vl6KDON2Y14YthzBQxGheUEhjUC45xvklQXi0oYxa0eGi1Ed3s3bX ICWbDQsJSvaluRviqMKE7x6M+iWSgm9wuzBwM1+SKHtiaKsDKyQ zPqpmGVkl4tj7X9gWRhM2trKqd0CkKkd2lgezDReTgQOerJ5+YTPg qlKnbNPAMSn root@node001.vstoragedomain

Cancel Add

To delete one or more keys, select them and click **Delete**.

Note: If a key has been injected into one or more VMs, it will remain inside those VMs even if you delete it from the admin panel.

4.8 Managing Compute Network

In the compute cluster, you can create and manage two types of virtual networks:

Private

VXLAN-based overlay virtual networks that can be used for intercommunication between VMs. Each private network is isolated from other private networks as well as public networks.

Public

Virtual networks that use IP address ranges of public physical networks. Such networks can be used to provide Internet access to VMs.

Each public virtual network can use IP addresses of only one physical network.

In Acronis Cyber Infrastructure, compute networking also includes virtual routers, floating public IP addresses, and network load balancers.

The next subsections explain the virtual network architecture and describe how to add, edit, and delete virtual networks as well as manage virtual routers, floating IP addresses, and load balancers.

4.8.1 Virtual Network Architecture

Acronis Cyber Infrastructure supports distributed virtual switching on the basis of Open vSwitch. The latter runs on every compute node and forwards network traffic between virtual machines on the same node and between virtual machines and physical networks. Distributed virtual switching provides centralized management and monitoring of virtual network configuration across all nodes in the compute cluster.

4.8.1.1 Private Network Connectivity

VXLAN technology used for private virtual networks allows creating logical L2 networks in L3 networks by encapsulating (tunneling) Ethernet frames over UDP packets.

The physical representation of private network connectivity can be shown as follows:



On the figure above:

- Three virtual machines are distributed across the compute cluster and connected to two private virtual networks via two virtual switches: VM1 and VM2 belong to one private virtual network, VM3 belongs to the other one.
- For each virtual network, the DHCP server runs on the management node.
- The virtual router that runs on the management node connects the two private virtual networks and the public virtual network created on top of the physical one, thus enabling connectivity between the VMs from different private virtual networks.
- The compute nodes are connected to the physical switch via the eth0 network interfaces and reside in one L2 segment.
- The eth0 network interfaces are connected to the physical network with the VM private and VM public traffic types.
- The physical router provides access to public networks, such as the Internet.



Logically the private networking scheme can be represented as follows:

4.8.1.2 Public Network Connectivity

Public virtual networks are connected to physical networks on Layer 2.

The physical representation of public network connectivity can be shown as follows:



On the figure above:

- Five virtual machines are distributed across the compute cluster and connected to two public virtual networks via two physical switches: VM1 and VM2 belong to one public virtual network, while VM3, VM4, and VM5 belong to the other one.
- For each virtual network, the DHCP server runs on the management node.
- The compute nodes are connected to one physical switch via the eth0 network interfaces and to the other physical switch via eth1 and reside in two separate L2 segments.
- The eth0 and eth1 network interfaces are connected to the physical networks with the VM public traffic type.
- The physical router interconnects two public virtual networks created on top of the physical ones and provides access to public networks, such as the Internet.

Logically the public networking scheme can be represented as follows:



4.8.2 Creating, Editing, and Deleting Virtual Networks

To add a new virtual network, do the following:

- 1. On the **COMPUTE > Network > NETWORKS** tab, click **Create virtual network**.
- 2. In the **Network configuration** section, configure the network parameters:
 - 2.1. Enable or disable IP address management.

With IP address management enabled, Acronis Cyber Infrastructure will handle virtual machine IP addresses and provide the following features:

- Allocation pools. You can specify ranges of IP addresses that will be automatically assigned to VMs.
- Built-in DHCP server. Assigns IP addresses to virtual machines. With the DHCP server enabled, VM network interfaces will automatically be assigned IP addresses: either from allocation pools or, if there are no pools, from network's entire IP range. With the DHCP server disabled, VM network interfaces will still get IP addresses, but you will have to manually assign them inside VMs.
- Custom DNS servers. You can specify DNS servers that will be used by VMs. These servers will be delivered to virtual machines via the built-in DHCP server.

With IP address management disabled:

• VMs connected to a network will be able to obtain IP addresses from DHCP servers in that network.

• Spoofing protection will be disabled for all VM network ports. Each VM network interface will accept all traffic, even frames addressed to other network interfaces.

In any case, you will be able to manually assign static IP addresses from inside VMs.

- 2.2. Choose network type.
- 2.3. Provide network details depending on type:
 - For a private network, specify a name. If IP address management is enabled, specify network's IPv4 address range in **Subnet CIDR**. Optionally specify a gateway. If you leave the **Gateway** field blank, the gateway will be omitted from network settings.
 - For a public network, specify a name and choose a physical network with the VM public traffic type (that is not already used by a public network). If IP address management is enabled, optionally specify a gateway. If you leave the Gateway field blank, the gateway will be omitted from network settings. The Subnet CIDR field will be filled in automatically. Optionally, select the Share between all projects checkbox. With the enabled option, self-service users can use this public network to connect virtual machines directly to it, connect private networks to it via virtual routers, and allocate floating IP addresses from it. With the disabled option, the public network will be available in the self-service panel only for connecting private networks via virtual routers and floating IPs allocation.

Click Next.

 \sim

Cr					\sim
•	Network configuration	Configure network settings.			
0		IP address management			
0		Type O Private O Public			
		Name public			
		Public	~	() Only networks with the VM public traffic type can be selected.	
		Subnet CIDR 10.94.0.0/16			
		Gateway (optional) 10.94.0.1			
		Share between all projects			
				Ne	ext

Create virtual network

3. If you enabled IP address management on the previous step, you will move on to the **DHCP and DNS** section. In it, enable or disable the built-in DHCP server and specify one or more allocation pools and DNS servers. Click **Next**.

Create virtual network	×
Network configuration	Set DHCP and specify one or more allocation pools for the public virtual network.
DHCP and DNS	Enable the built-in DHCP server.
• Summary	Allocation pools + Add pool
	10.94.10.2 — 10.94.10.128 126 addresses available 🧷 🛅
	DNS servers + Add server
	20.20.20.20
	10.10.10 / 🔟
	Back Next

este virtual network Cr

4. In the **Summary** section, review the configuration and click **Add virtual network**.

Х

•	Network configuration	Review the virtual network details and go back to change them if necessary	
•	DHCP and DNS	Туре	Public
		Name	public
•	Summary	Physical network	Public
		Subnet CIDR	10.94.0.0/16
		Gateway	10.94.0.1
		DHCP	Enabled
		Allocation pools	10.94.10.2 — 10.94.10.128 126 addresses available
		DNS servers	20.20.20.20 10.10.10.10
			Back Create virtual network

Create virtual network

To view and edit parameters of a virtual network, click it on the **NETWORKS** tab. On the virtual network panel, you can change the virtual network name, gateway, DHCP settings, allocation pools, and DNS servers. To do this, click the pencil icon, enter a new value, and click the check mark icon to confirm.

To delete a virtual network, click the ellipsis icon next to it and **Delete**. To remove multiple virtual networks at once, select them and click **Delete**. Before deleting a virtual network, make sure no VMs are connected to it.

4.8.3 Managing Virtual Routers

Virtual routers provide L3 services such as routing and Source Network Address Translation (SNAT) between private and public networks or different private networks:

- a virtual router between private and public networks provides access to public networks, such as the Internet, for VMs connected to this private network;
- a virtual router between different private networks provides network communication for VMs connected to these private networks.

A virtual router has two types of ports:

- an external gateway that is connected to a public network,
- an internal port that is connected to a private network.

Note: A router can only connect networks with enabled IP management.

To create a virtual router, do the following:

- 1. On the **COMPUTE** > **Network** > **NETWORKS** tab, make sure the virtual networks that are to be connected to a router have a gateway specified.
- 2. Navigate to the **ROUTERS** tab and click **Add router**.
- 3. In the Add router window:
 - 3.1. Specify a router name.
 - 3.2. From the **Network** drop-down menu, select a public network through which external access will be provided via an external gateway. The new external gateway will pick an unused IP address from the selected public network.
 - 3.3. In the **Add internal interfaces** section, select one or more private networks to connect to a router via internal interfaces. The new internal interfaces will attempt to use the gateway IP address of the selected private networks by default.
 - 3.4. Optionally, select or deselect the **SNAT** checkbox to enable or disable SNAT, respectively, on the external gateway of the router. With SNAT enabled, the router replaces VM private IP addresses with the public IP address of its external gateway.

Add virtual router	×
Name router1	
Specify a network through which public netwo	rks will be accessed.
Network public: 10.94.0.0/16	~
SNAT ()	
Add internal interfaces	+ Add
private: 192.168.128.0/24	× ū
_	

4. Click Create.

To edit a router name, click the ellipsis icon next to it and **Rename**.

To remove a virtual router, click the ellipsis icon next to it and **Delete**. To remove multiple virtual networks at once, select them and click **Delete**. Before deleting a virtual router, make sure no floating IP addresses are associated with any network it is connected to.

4.8.3.1 Managing Router Interfaces

You can add an external router interface as follows:

Note: To change an external gateway, remove the existing one first.

1. On **Routers** screen, click the router name to open the list of its interfaces.

- 2. Click **Add** on the toolbar, or click **Add interface** if there are no interfaces to show.
- 3. In the Add interface window, do the following:
 - 3.1. Choose **External gateway**.
 - 3.2. From the **Network** drop-down menu, select a public network to connect to the router. The new interface will pick an unused IP address from the selected public network. You can also provide a specific IP address from the selected public network to assign to the interface in the **IP address** field.
 - 3.3. Optionally, select or deselect the **SNAT** checkbox to enable or disable SNAT, respectively, on the external gateway of the router. With SNAT enabled, the router replaces VM private IP addresses with the public IP address of its external gateway.

Add interface	×
 External gateway Internal interface Specify new interface parameters 	
Network public: 10.94.0.0/16	~
IP address (optional)	
By adding a router interface you connect the selected network router. The new interface will pick an unused IP address from selected public network. You can also provide a specific IP add from the selected public network to assign to the interface.	to the the ress
Cancel	Add

4. Click Add.

To edit the external gateway parameters, click the ellipsis icon next to it and **Edit**. In the **Edit interface** window, you can change the external gateway IP address and enable or disable SNAT on it. To save your changes, click **Save**.

You can add an internal router interface as follows:

- 1. On **Routers** screen, click the router name to open the list of its interfaces.
- 2. Click **Add**.
- 3. In the Add interface window, select a network to connect to the router from the Network drop-down menu. The new interface will attempt to use the gateway IP address of the selected private network by default. If it is in use, specify an unused IP address from the selected private network to assign to the interface in the IP address field.

Add interface	\times
Specify new interface parameters	
Network Select	~
IP address (optional)	
By adding a router interface you connect the selected network to the router. The new interface will attempt to use the gateway IP address the selected private network by default. If it is in use, specify an unit IP address from the selected private network to assign to the interf	ie is of used ace.
Cancel	ld

4. Click **Add**.

To remove a router interface, click the ellipsis icon next to it and **Delete**. To remove multiple interfaces at once, select them and click **Delete**.

4.8.3.2 Managing Static Routes

You can also configure static routes of a router by manually adding entries into its routing table. This can be useful, for example, if you do not need a mutual connection between two private networks and want only one private network to be accessible from the other.

Consider the following example:

- the virtual machine vm1 is connected to the private network private1 (192.168.128.0/24) via the network interface with IP address 192.168.128.10,
- the virtual machine vm2 is connected to the private network private2 (192.168.30.0/24) via the network interface with IP address 192.168.30.10,
- the router router1 connects the network private1 to the public network via the external gateway with the IP address 10.94.129.73,
- the router router2 connects the network private2 to the public network via the external gateway with the IP address 10.94.129.74.

To be able to access vm2 from vm1, you need to add a static route for router1, specifying the CIDR of private2, that is 192.168.30.0/24, as the destination subnet and the external gateway IP address of router2, that is 10.94.129.74, as the next hop IP address. In this case, when an IP packet for 192.168.30.10 reaches router1, it will be forwarded to router2 and then to vm2.

To create a static route for a router, do the following:

- 1. On the **STATIC ROUTES** tab of a virtual router, click **Add static route**.
- 2. In the **Add static route** window, specify the destination subnet range and mask in CIDR notation and the next hop's IP address. The next hop's IP address must belong to one of the networks that the router is connected to.

Add static route	×
Specify static route parameters	
Destination subnet and mask 192.168.30.0/24	
Next hop 10.94.129.74	
The next hop's IP address must belong to one of the networks t router is connected to.	hat the
Cancel	Add

3. Click Add.

To edit a static route, click the ellipsis icon next to it and **Edit**. In the **Edit static route** window, change the desired parameters and click **Save**.

To remove a static route, click the ellipsis icon next to it and **Delete**. To remove multiple routes at once, select them and click **Delete**.

4.8.4 Managing Floating IP Addresses

A virtual machine connected to a virtual private network can be accessed from public networks, such as the Internet, by means of a floating IP address. Such an address is picked from a public network and mapped to VM's private IP address. The floating and private IP addresses are used at the same time on the VM's network interface. The private IP address is used to communicate with other VMs on the private network. The floating IP address is used to access the VM from public networks. The VM guest operating system is unaware of the assigned floating IP address.

Note the following prerequisites:

- 1. A VM must have a fixed private IP address.
- 2. A virtual router must connect the public network from which a floating IP will be picked with VM's private network.

You can create a floating IP address and assign it to a VM as follows:

- 1. On the **COMPUTE > Network > FLOATING IPS** tab, click **Add floating IP**.
- 2. In the **Add floating IP address**, select a public network from which a floating IP will be picked and a VM network interface with a fixed private IP address.

public: 10.94.0.0/16	
	Ť
Virtual machine	
myvm — private: 192.168.128.5	~ >

3. Click **Add**.

A floating IP address can be re-assigned to another virtual machine. Do the following:

- 1. Click the ellipsis icon next to the floating IP address and then click **Unassign**.
- Once the VM name disappears in the Assigned to column, click the ellipsis icon again and choose Assign.
- 3. In the Assign floating IP address window, select a VM network interface with a fixed private IP address.
- 4. Click Assign.

To remove a floating IP address, unassign it from a VM as described above, then click the ellipsis icon again and choose **Delete**.

4.8.5 Managing Load Balancers

Acronis Cyber Infrastructure offers load balancing as a service for the compute infrastructure. Load balancing ensures fault tolerance and improves performance of web applications by distributing incoming network traffic across virtual machines from a balancing pool. A load balancer receives and then routes incoming requests to a suitable VM based on a configured balancing algorithm and VM health.

Load balancers can be created and edited by self-service users. Administrators can only monitor, disable/enable, and delete load balancers. In the admin panel, load balancers are shown on the **COMPUTE** > **Network** > **LOAD BALANCERS** tab.

Note: For self-service users to be able to create highly available load balancers, the compute cluster must have at least two nodes.

NETW	ORKS FLOATING IPS	ROUTERS	LOAD BALANCERS						
-• F	lters Search		۹						
	Name 🕇	St	atus 🧅	IP address 👃	Project \downarrow	Floating IP 🔱	Members state	Members total 🔱	¢
	പ്പെ LBaaS1	o	Active	192.168.30.159	project1	10.94.129.75		6	

To monitor a load balancer's performance and health, select it and open the **Overview** tab.

The following charts are available:

Members state

The total number of members in the balancing pools grouped by status: "Healthy", "Unhealthy", "Error", and "Disabled".

CPU/RAM

CPU and RAM usage by the load balancer.

Network

Incoming and outgoing network traffic.

Active connections

The number of active connections.

Error requests

The number of error requests.

By default, a load balancer instance is created with 2 vCPUs and 1 GB of memory. Check the **CPU/RAM** chart to see how the load balancer handles its workload. If the chart shows high values (about 80%), the load balancer needs more resources. You can add resources by changing the load balancer flavor as described in Changing the Default Load Balancer Flavor.

You can see the load balancer parameters on its **Properties** tab.

In the **Virtual machines** field, you can see the name of load balancer instances. Click it to open the VM's panel.

To disable/enable or remove a load balancer, click the ellipsis icon next to it and the desired action. To remove multiple load balancers at once, select them and click **Delete**.

4.8.5.1 Managing Balancing Pools

To see a list of balancing pools in a load balancer, click its name.

Netwo	ork > Load balancers > LBaaS1				0
Sear	ch Q				
	Balancing pool	Status	Members state	Members total	¢
	\bigotimes HTTP on port 80 \rightarrow HTTP on port 80	Active		3	
	\bigotimes HTTPS on port 443 \rightarrow HTTPS on port 443	Active		3	

You can open the pool's panel to monitor its performance and health on the **Overview** tab, see its parameters on the **Properties** tab, and manage its members on the **Members** tab.

To remove a balancing pool, click the ellipsis icon next to it and click **Delete**. To remove multiple balancing pools at once, select them and click **Delete**.

4.9 Managing Compute Storage

In Acronis Cyber Infrastructure, compute storage comprises volumes provisioned to virtual machines. The provisioned storage space can exceed available physical space. Rules for storing volumes can be set via storage policies.

The next subsections describe how to manage volumes and create storage policies for them.

4.9.1 Managing Compute Volumes

A volume in Acronis Cyber Infrastructure is a virtual disk drive that can be attached to a VM. The integrity of data in volumes is protected by the redundancy mode specified in the storage policy.

Note: Additional virtual disks attached to VMs need to be initialized inside the guest OS by standard means before they can be used.

4.9.2 Creating, Editing, and Removing Volumes

To create a volume, do the following:

1. On the **COMPUTE** > **Storage** > **VOLUMES** tab, click **Create volume**.

Create volume	×
Name vol1	
Size (GiB) Min. 1 GiB, Max. 512 TiB	
Storage policy default	~
	Cancel Create

2. In the **Create volume** window, specify a volume name and size in gigabytes, select a storage policy, and click **Add**.

To edit a volume, select it and click the pencil icon next to a parameter you need to change. Note the following restrictions:

- You cannot shrink volumes.
- To extend volumes that are in use, stop the VM first.
- You cannot change the volume redundancy type.

To remove a volume, click its ellipsis button then click **Delete**. To remove multiple volumes at once, select them and click **Delete**. To remove a volume that is in use, detach it first.

Note: A volume is removed along with all its snapshots.

4.9.3 Cloning Volumes

You can clone volumes that are not attached to VMs or attached to stopped VMs. To clone a volume, do the following:

- 1. On the **COMPUTE** > **Storage** > **VOLUMES** tab, click a volume.
- 2. In volume details that opens, click **Clone**.

vol1	×
→I Attach 🗗 Clone 🔀 Create snapshot	🖾 Create image 🛅 Delete
Overview	Snapshots
Details	
Status 🥝 Availa	ble

3. In the **Clone volume** window that opens, specify a volume name, size, and storage policy. Click **Clone**.

Clone volur	ne		×
Name Clone_vol1			
Size (GiB) 1	Min. 1 GiB, Max. 512 TiB		
Storage policy default			~
		Cancel	Clone

4.9.4 Attaching and Detaching Volumes

To add a writable virtual disk drive to a VM, attach a volume to it. To do this:

- On the COMPUTE > Storage > VOLUMES tab, click the ellipsis button next to an unused volume and click Attach in the context menu.
- 2. In the **Attach volume** window, select the VM from the drop-down list and click **Done**.

Attach volume	×
Choose a volume to attach	
Volume vol1 Virtual machine	~
vm1	Cancel Done

To detach a volume, do the following:

- 1. Click the ellipsis button next to the volume that is in use.
- 2. If the VM is not running, click **Detach**. If the VM is running, you can only click **Force detach** to immediately detach the volume with a risk of data loss.

4.9.5 Creating Images from Volumes

To create multiple VMs with the same boot volume, you can create an image from an existing boot volume and deploy VMs from it. Make sure to install cloud-init in the volume before creating the image.

Do the following:

- 1. Power off the VM that the original volume is attached to.
- Switch to the COMPUTE > Storage > VOLUMES tab, click volume's ellipsis button and choose Create image.

V	m1/cirros/Boot volume			×
	Details		ତ Create	image
	Status	In use		
	Volume ID	6e4edfa2-727f-466e-8519-482ae4	1422b56	
	Bootable	true		

3. In the **Create image** window, enter an image name and click **Create**.

Create image	×
Name vol1-image Volume: vm1/cirros/Boot volume	
	Cancel
The new image will appear on the **IMAGES** tab.

4.9.6 Managing Volume Snapshots

You can save the current state of a VM file system or user data by creating a snapshot of a volume. A snapshot of a boot volume may be useful, for example, before updating VM software. If anything goes wrong, you will be able to revert the VM to a working state at any time. A snapshot of a data volume can be used for backing up user data and testing purposes.

To create a snapshot of a volume, do the following:

- 1. On the **COMPUTE** > **Storage** > **VOLUMES** tab, click a volume.
- 2. In the volume panel that opens, switch to **Snapshots** and click **Create snapshot**.

vm1/cirros/Boot volume			X
Create snapshot			
Overview		Snapshots	
Details			
Status	🖸 In use		
Volume ID	f9a98c59-	3025-4634-ace0-f4ed0b5667f0	

Note: To create a consistent snapshot of a running VM's volume, make sure the guest tools are installed in the VM. QEMU guest agent included in the guest tools image automatically quiesces the filesystem during snapshotting. For the instructions on installing the guest tools, see *Installing Guest Tools* (page 101).

Once the snapshot is created, you can see and manage it on the **Snapshots** tab on the volume panel.

vm1/cirros/Boot volume

 \times

🔁 Create snapshot

	Overview	Snapshots	
a few s	seconds ago		
~ 오	Snapshot-vm1/cirros/Boot volume Created on : April 30, 2019 5:40 PM Status: Available Description: N/A Size: 1 GiB Storage policy: default Snapshot ID: e14d3256-4284-402f-9785-5c76db	5485407	

To see the full list of available actions, click the ellipsis button next to a snapshot. Actions include:

- Create volume creates a new volume from the snapshot.
- Create image creates a template image from the snapshot.
- **Revert to snapshot** discards all changes that have been made to the volume since the snapshot was taken. This action is available only for VMs with the "Shut down" and "Shelved offloaded" statuses.

Warning: As each volume has only one snapshot branch, all snapshots created after the snapshot you are reverting to will be deleted. If you want to save a subsequent snapshot before reverting, create a volume or an image from it first.

- Edit changes the snapshot name and description.
- **Reset** resets the snapshot stuck in the "Error" state or one of transitional states to the "Available" state.
- **Delete** removes the snapshot.

4.9.7 Managing Storage Policies

A storage policy is a group of parameters that define how to store VM volumes: how redundant they must be and on what storage tier they need to be located.

When you deploy the compute cluster, a default storage policy is created that enforces the best replication scheme allowed by the number of nodes in the storage cluster. The default policy cannot be deleted or renamed and is always applied to uploaded images and base volumes created from these images.

Note: A base volume is created from a source image when you deploy a VM. It is not used directly by a VM, but all volumes that a VM actually uses (which are listed on the **VOLUMES** tab) are in fact deltas (differences) from the base volume. It is important to keep base volumes available as VM volumes depend on them. For that, you need the default storage policy to enforce multiple replicas.

If the storage cluster does not have enough nodes to enable multiple replicas (not recommended), you can adjust the default storage policy once you add more nodes to the storage cluster. It will be applied to images and base volumes that already exist in the compute cluster.

To apply custom redundancy schemes to VM volumes, you can create custom storage policies in addition to the default one. To create a custom storage policy, do the following:

- 1. On the **COMPUTE** > **Storage** > **STORAGE POLICIES** tab, click **Create storage policy**.
- 2. In the Create storage policy window, specify a name and select the following:

Create storag	e policy		×
Name policy1			
poncy			
Tier Tier 0	~	Failure domain Host	~
 No redundancy 			
 2 replicas 	100% ov	erhead	
3 replicas	200% ov	erhead	
		Cancel	reate

- 2.1. In **Tier**, a tier to store volumes on.
- 2.2. In **Failure domain**, a placement policy for data pieces or replicas.
- 2.3. In **Type**, a data redundancy type and mode.
- 3. Click Create.

To edit the policy name, select the policy and click the pencil icon on its panel.

To remove a policy, select it and click **Delete policy**. A policy cannot be removed if it governs existing volumes.

4.10 Managing Placements

A placement is a group of nodes in the compute cluster that share a distinctive feature. It can be a special license for software to run in VMs or an advanced CPU model. You can create placements to make sure that VMs that need a specific feature are located on nodes that have it. The same node can be included into several placements.

Take note of the following:

- A virtual machine that is assigned a placement can only be migrated between nodes in this placement.
 Similarly, a VM without a placement can only be migrated between nodes that are not in placements.
 When adding nodes to placements, make sure to provide migration options for various scenarios, including high availability and maintenance. Avoid situations when VMs cannot migrate because of limitations imposed by placements.
- Self-service users can only use placements by creating VMs from images added to those placements. If all the nodes are in placements, make sure that at least one image is added to a placement. Otherwise self-service users will not be able to create virtual machines.

4.10.1 Creating Placements

To create a placement, follow these steps:

- In the admin panel, open COMPUTE > Nodes and switch to the PLACEMENTS tab. Click Create placement.
- 2. Specify a name for the new placement. The name should clearly state the distinctive feature of nodes in the placement. For example, Microsoft Windows Server license.

Create placement							
 Self-service users can choose placements for virtual machines they create. A placement is a group of nodes sharing a distinctive feature, e.g., a special license for software to run in VMs or an advanced CPU model. 							
The name of a placement should clearly describe the distinctive feature of nodes in it. For example, "Microsoft Windows Server license".							
Name MS Windows Server license							
Description (optional)							

3. Select nodes to be included in the placement.

Add nodes to place virtual machines on.	
Nodes	+ Add

4. Optionally, add images to the placement. VMs created from such images will automatically be added to the placement.



5. Click Create.

The new placement will appear in the list.

4.10.2 Managing Nodes in Placements

To add a node to a placement, do as follows:

- 1. In the admin panel, open **COMPUTE** > **Nodes** and switch to the **PLACEMENTS** tab. Click the required placement.
- 2. In the **Nodes** pane, click **Add**.
- 3. Select the nodes to be included in the placement and click **Add**.

To remove nodes from a placement, select them on the placement's nodes pane and click **Remove**.

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.

4.10.3 Managing Images in Placements

To have VMs deployed from an image added to a placement automatically, add that image to a placement. You can do this with images that are shared between all projects (see *Managing Images* (page 104)).

Do the following:

- 1. In the admin panel, open **COMPUTE** > **Nodes** and switch to the **PLACEMENTS** tab. Click the required placement.
- 2. Click **Add images** on the toolbar.
- 3. Select one or more images and click **Add**.

When you choose this image while creating a VM, the corresponding placement will be selected automatically.

You can also use an image added to a placement to create a volume added to the same placement (see *Creating, Editing, and Removing Volumes* (page 131)). Snapshots and virtual machines created from this volume will inherit the placement too.

To remove an image from a placement, open the placement details, switch to the **Images** tab and click the bin icon next to the image you want to remove.

4.10.4 Renaming Placements

To rename a placement, do the following:

- 1. In the admin panel, open **COMPUTE** > **Nodes** and switch to the **PLACEMENTS** tab.
- 2. Select the required placement and click **Edit** on the toolbar.
- 3. Enter a new name and click **Save**.

4.10.5 Deleting Placements

Prior to deleting a placement, make sure there are no nodes in it. A placement with nodes cannot be deleted. Remove nodes from the placement, if any, and proceed as follows:

- 1. In the admin panel, open **COMPUTE** > **Nodes** and switch to the **PLACEMENTS** tab.
- 2. Select the required placement and click **Delete** on the toolbar.
- 3. In the confirmation message, click **Delete placement**.

4.11 Managing Kubernetes Clusters

In Acronis Cyber Infrastructure, Kubernetes clusters are created by self-service users. To be able to use the service, however, you need to install it in the admin panel. The Kubernetes service can be deployed along with the compute cluster or later (see *Creating the Compute Cluster* (page 69) or *Managing Add-On Services* (page 231)). The network requirements for Kubernetes clusters are described in Kubernetes Network Requirements.

You can also perform the following actions in the admin panel:

• view Kubernetes cluster details

Cancel

- change Kubernetes service parameters
- delete Kubernetes clusters

0

To view the details of a Kubernetes cluster, open **COMPUTE** > **Kubernetes** and click a cluster.

You can change such parameters as the size and storage policy of the system volume that is used by the Kubernetes management services and etcd. The new parameters will be applied to system volumes on master nodes in new Kubernetes clusters. System volumes in existing Kubernetes clusters will retain their previous parameters. To improve the stability of Kubernetes clusters, it is highly recommended to select a storage policy with an SSD-based tier.

To change the Kubernetes service parameters, click **Settings** on the **Kubernetes clusters** screen, set the desired storage policy and size for the system volume, and click **Done**.

Kubernetes service parameters

The system volume is used by the Kubernetes management services and etcd. To improve the stability of Kubernetes clusters, it is highly recommended to select a storage policy with an SSD-based tier.

Storage policy	Disk size (GiB)	Min. 10 GiB,
default	10	Max. 512 TiB

Creation of Kubernetes clusters in a project will only be available if a project quota is set for the selected storage policy.

To delete a Kubernetes cluster, open **COMPUTE** > **Kubernetes**, click a cluster, and click **Delete** on the toolbar.

To learn how to create and manage Kubernetes clusters in the self-service panel, refer to Managing Kubernetes Clusters.

Х

Done

4.12 Destroying the Compute Cluster

To destroy the compute cluster, do the following:

- 1. Delete all virtual machines from all nodes.
- Release the compute nodes. To do this, select all nodes on the COMPUTE > Nodes screen and click Release nodes. Regular (non-management) nodes will be released first. Management nodes will follow. Releasing the management nodes will destroy the compute cluster.

CHAPTER 5

Exporting Storage Space

Acronis Cyber Infrastructure allows you to export storage space as:

- Block storage via iSCSI for virtualization, databases and other needs.
- Object storage for storing unlimited number of files via an Amazon S3 compatible protocol. You can store data like media files, backups, Open Xchange files and access the storage using Dropbox-like applications. You can build your own Amazon S3 compatible object storage services as a part of your cloud offering or for internal needs.
- A back-end for Acronis Backup Cloud and Acronis Backup Advanced backups.
- NFS exports.

5.1 Exporting Storage via iSCSI

Acronis Cyber 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 Acronis Cyber 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 Acronis Cyber Infrastructure disk space via iSCSI.

The figure shows two volumes located on redundant storage provided by Acronis Cyber Infrastructure. The volumes are attached as LUNs to a group of two targets running on Acronis Cyber 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.

5.1.1 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 Networks and Traffic Types* (page 2).
- 2. Create a target group on chosen nodes. See Creating Target Groups (page 149).
- 3. Create volumes and attach them to the target group as LUNs. Typically you do this while creating the target group. However, you can also do this later as described in *Managing Volumes* (page 156).
- 4. Optionally, enable CHAP and ACL authorization for the target group: create CHAP accounts and assign them to the target group, populate group's access control list. Typically, you do this while creating the target group. However, you can also do this later as described in *Restricting Access to Target Groups* (page 162).
- 5. Connect initiators to targets using standard tools of your operating system or product (consult the Storage User's Guide). To view target IQNs, click the target group name.

5.1.1.1 Managing Legacy iSCSI Targets

After the upgrade to Acronis Cyber Infrastructure 2.5, you can run older iSCSI targets created on version 2.4 alongside new targets. For each older target, a target group is automatically created and iSCSI LUNs are moved to iSCSI volumes. Such a target group have a name in the format group:<target_name>. For example, the iSCSI target with the IQN iqn.2014-06.com.vstorage:target1 will be placed in the target group named group:target1.

In the admin panel, you can manage older targets only by deleting them and detaching volumes from them. The full functionality is available using the vstorage-iscsi utility and described in the Acronis Storage 2.4 Administrator's Command Line Guide.

As older iSCSI targets do not support the ALUA mode, their LUNs are not highly available. To enable high availablity for them, detach a volume from an older target group and attach it to a newly created one as described in *Detaching Volumes* (page 161) and *Attaching Volumes to Target Groups* (page 158).

5.1.2 Managing Target Groups

This section explains how to create and manage groups of iSCSI targets.

5.1.2.1 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, open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS** and click **Create target group**. A wizard will open where you need to do the following:

1. On Name and type, enter a target group name and select a type: iSCSI or Fibre Channel.

Cre	eate target group			>
•	Name and type	Name and type		
0		Name	Type	
0		Target group 1	iSCSI	~

2. On **Nodes**, select nodes to add to the target group. On these nodes, iSCSI targets will run. You can only choose nodes with network interfaces that are assigned the **iSCSI** traffic type. It is recommended to have at least two nodes in the target group to achieve high availability. If you plan to use multiple iSCSI initiators, you should have as many nodes in the target group.

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.

If node network interfaces are not configured, click the cogwheel icon, select networks as required, and click **Apply**.

Cre	eate target group					×					
•	Name and type	Nod	es t nodes where iSCSI tar	gets will run. You o	an only choose nodes	with network					
•	Nodes	interf least	nterfaces that are assigned the "ISCSI public" traffic type. It is recommended to select at east two nodes to achieve high availability. If you plan to use multiple ISCSI initiators, select								
0		Searc	h	Q							
Ĭ	Volumes		Name 🦊	Node sta	IP address	Network state					
°			node001.vsto	Healthy	10.37.130.249	Configured					
0			node002.vsto	Healthy	10.37.130.27	🕑 Configured					
			node004.vsto	Healthy	10.37.130.44	🕑 Configured					
						Back Next					

3. On **Targets**, select iSCSI interfaces to add to the target group. You can choose from a list of network interfaces that are assigned the **iSCSI** traffic type. If you plan to use multiple iSCSI initiators, you should select as many interfaces per node. One interface can be added to multiple target groups, although it may reduce performance.

Cr	eate target group	×
	Name and type Nodes	Targets On this step, you need to select ISCSI interfaces to add to the target group. You can choose from a list of network interfaces that are assigned the "ISCSI public" traffic type. It is recommended to select at least two interfaces on different nodes for high availability. If you plan to use multiple ISCSI initiators, select as many interfaces per node. One interface can
•	Targets	be added to multiple target groups, although it may reduce performance.
0		iqn.2014-06.com.vstorage: target1
0		✓ eth0-10.94.18.147
0		
		node001.vstoragedomain. iqn.2014-06.com.vstorage: target2
		Back Next

Croate target group

4. On **Volumes**, select volumes to attach to target group LUNs. You can choose from a list of volumes that are not attached to any target groups. If no volumes are available, you can create them on this step so they are attached to the target group automatically or later and attach them manually.

Cr	eate target group						×		
	Name and type Nodes	Volu On thi from a	Volumes On this step, you need to select volumes to attach to target group LUNs. You can choose from a list of volumes that are not attached to any target groups.						
•	Targets	Searcl	h	Q			+ Create		
	Volumes		Name 🤳	ID 🕇	Policy	Size	LUN ID		
Ĩ	• volumes		🖨 tg1-vol1	b5c21e1	Tier 0, Fa	8 KiB of 1 GiB	0		
0			tg1-vol2	f9a9ada	Tier 0, Fa	0 bytes of 1 Gi	1		
0			🖨 tg1-vol3	f1963fa0	Tier 0, Fa	0 bytes of 1 Gi	2		
						Back	Next		

 On Access control, configure access to the target group. It is recommended to use CHAP or ACL in untrusted public networks. Without access control, any connections to the target group are allowed.
 For more information, see *Restricting Access to Target Groups* (page 162).

Create target group		×
 Name and type Nodes 	Access control On this step, you can configure access to the target group. It is recommended to use CH or ACL in untrusted public networks.	IAP
 Targets 	ACL CHAP Populate the access control list with iSCSI initiator IQNs that will be allowed to communi with the target group.	icate
Volumes	CHAP user (optional)	
Access control	user1	
o Summary	+ Create user	
	Search Q +	Add
	IQN ↑ Initiator name LUNs	
	iqn.1991-05.com.mic initiator1 0,1,2	
	Back Ne	xt

6. On **Summary**, review the target group details. You can go back to change them if necessary. Click **Create**.

The created target group will appear on the **TARGET GROUPS** tab. Its targets will start automatically.

5.1.2.2 Adding Targets

To add a target to a target group, do the following:

 Open STORAGE SERVICES > Block storage > TARGET GROUPS, click the name of the desired target group to open it.

NODE	S TARGET G	ROUPS VOLUMES	CHAP USERS		
Searc	:h	۹		+ Create target (group
	Na 🦊	Target states	LUNs 🕇	Nodes 🕇	ø
	<mark>⊨⊒</mark> I‰G1		2	3	

- 2. On the **TARGETS** tab, click **Add target**. The **Create target** wizard will open.
- 3. On **Nodes**, select nodes to add to the target group. On these nodes, iSCSI targets will run. You can only choose nodes with network interfaces that are assigned the **iSCSI** traffic type. It is recommended to have at least two nodes in the target group to achieve high availability. If you plan to use multiple iSCSI initiators, you should have as many nodes in the target group.

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.

If node network interfaces are not configured, click the cogwheel icon, select networks as required, and click **Apply**.

•	Nodes	Node Select that a	Nodes Select nodes where iSCSI targets will run. You can only choose nodes with network interfaces that are assigned the "iSCSI" traffic type. It is recommended to select at least two nodes to				
0		achieve high availability. If you plan to use multiple iSCSI initiators, select as many nodes.					
		•	Name 🦊	Node sta	IP address	Network state	
			node001.vsto	Healthy	10.37.130.249	📀 Configured	٥
			node002.vsto	Healthy	10.37.130.27	📀 Configured	¢
			node004.vsto	Healthy	10.37.130.44	📀 Configured	٥
						Nex	rt

Create target

4. On **Targets**, select iSCSI interfaces to add to the target group. You can choose from a list of network interfaces that are assigned the **iSCSI** traffic type. If you plan to use multiple iSCSI initiators, you should select as many interfaces per node. One interface can be added to multiple target groups, although it may reduce performance.

X

Create target	×
Nodes	Targets On this step, you need to select iSCSI interfaces to add to the target group. You can
 Targets 	choose from a list of network interfaces that are assigned the "iSCSI public" traffic type. It is recommended to select at least two interfaces on different nodes for high availability. If
o Summary	you plan to use multiple ISCSI initiators, select as many interfaces per node. One interface can be added to multiple target groups, although it may reduce performance.
	iqn.2014-06.com.vstorage: target6 eth0-10.94.18.146
	Back Next

5. On **Summary**, review the target details. You can go back to change them if necessary. Click **Next**.

The created target will appear on the Targets tab.

5.1.2.3 Starting and Stopping Targets

To start or stop all targets in a target group, open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS**, click the ellipsis icon of the desired target group, and click **Start targets** or **Stop targets**, respectively.

NODE	S TARGET GROU	JPS VOLUMES	CHAP USERS		
Search Q + Create target group					
	Name 🦊	Target states	LUNs 🕇	Nodes ↑	ø
	TG1		2	3	
				+	Add LUNs
				ш ш	Delete

5.1.2.4 Deleting Targets

To delete a target from a target group, do the following:

1. Open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS**, click the name of the desired target group to open it.

NOD	ES TARGET G	ROUPS VOLUME	S CHAP USERS	;	
Sea	rch	Q		+ Create target (group
	Na ↓	Target states	LUNs 🕇	Nodes ↑	ø
	<mark>≣I</mark> ∰G1		2	3	

2. On the **Targets** tab, click the ellipsis button of the desired target then click **Delete**.

Target	ts LUNs			
Searc	ch	Q	+ Add	target
	Name 🦊	IQN 1	Node 🕇	¢
	🔒 target1	iqn.2014-06.com.vst	node001.vstoraged	
	🔒 target2	iqn.2014-06.com.vst	node002.vstoraged	•••
	🔒 target3	iqn.2014-06.com.vst	node004.vsto 👖 Delete	Ф

3. Click **Delete** in the confirmation window. Check the **Force** box to delete a target with active connections.

If you delete a target on the Active/Optimized path (indicated in LUN details), said path will switch to another target.

5.1.2.5 Deleting Target Groups

To delete a target group, open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS**, click the ellipsis icon of the desired target group, and click **Delete**.

NODE	S TARGET GRO	UPS VOLUMES	CHAP USERS		
Search Q + Create target group				et group	
	Name ↓	Target states	LUNs ↑	Nodes ↑	¢
	TG1		2	3	
				+ Add OO Stop	UNs targets
				m Delet	e 🖑

Click **Delete** in the confirmation window. Check the **Force** box to delete a target group with active connections.

5.1.3 Managing Volumes

This section describes how to create and manage volumes to be exported via iSCSI.

5.1.3.1 Creating Volumes

While it is convenient to create desired volumes while creating a target group, you can also do this at any time afterwards:

- 1. Open **STORAGE SERVICES** > **Block storage** > **VOLUMES** and click **Create volume**. A wizard will open.
- 2. On **Name and size**, enter a volume name and specify a size in gigabytes. Note that volumes can be extended later but not shrunk.

Create volume	×
 Name and size 	Name and size Enter a name and size for the volume. Note that volumes can be expanded later but not shrunk.
Storage policy Summary	Name tg1-vol4 Size (GB) Min. 1 GB, 1 Max. 2000 GB
	Next

3. On **Storage policy**, select a redundancy mode, a storage tier, and a failure domain. To benefit from high availability, select a mode other than **No redundancy** and failure domain other than **Disk**.

Create volume		×
Name and size	Storage policy	failure domain. To benefit from high
Storage policy	availability, select a mode other than "No redund	dancy" and failure domain other than "Disk".
• Summary	Tier Tier 0 ~ Type Erasure coding • Replication	Failure domain Host
	O No redundancy	
	2 replicas 100% overhead	
	• 3 replicas 200% overhead	
		Back Next

4. On **Summary**, review the volume details. You can go back to change them if necessary. Click **Create**.

5.1.3.2 Attaching Volumes to Target Groups

To add a volume as a LUN to a target group, do the following:

1. Open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS**, click the ellipsis icon of the desired target group, and click **Add LUNs**.

NODE	S TARGET GROUP	PS VOLUMES	CHAP USERS		
Searc	ch	۹		+ Cre	ate target group
	Name 🦊	Target states	LUNs ↑	Nodes 🕇	¢
	TG1		2	3	
				+	
				0	0 Stop targets
				Ľ	🗓 Delete

2. In the **Attach** window that opens, select volumes to attach to the target group (create them if needed) and click **Attach**.

Atte	ach				×		
Selec selec	Select volumes to attach to target group LUNs. You can choose from a list of volumes that are not attached to any target groups. LUN IDs will selected automatically, but you can change them later in target group LUN settings.						
Sear	:h	Q			+ Create		
	Name 🦊	ID 🕇	Policy	Size	LUN ID		
	🖨 tg1-vol3	f1963fa0-56e5-419	Tier 0, Failure dom	0 bytes of 1 GiB	2		
					Attach		

Alternatively, you can do the same on the **VOLUMES** tab:

1. Click the ellipsis icon of the desired volume then click **Attach**.

NODE	S TARGET GROUP	PS VOLUMES CHAI	P USERS		
Searc	ch	٩		+ Create v	olume
	Name 🦊	Target group ↑	LUN ID 🕇	Size	¢
	🖨 tg1-vol1	TG1	0	8 KiB of 1 GiB	
	➡ tg1-vol2	TG1	1	0 bytes of 1 GiB	
	tg1-vol3	N/A	N/A	0 bytes of 1 GiB	
				→ Attach	6
					2

2. In the **Attach** window that opens, select a target group and click **Attach**.

Attach	×
Select a target group to attach the volume " tg1-vol3 " to	
Select group TG1	~
	Attach

5.1.3.3 Setting LUN Limits

To set a read/write limit for a volume attached to a target group as a LUN, do the following:

 Open STORAGE SERVICES > Block storage > TARGET GROUPS, click the name of the desired target group to open it, and switch to LUNS.

NODE	S TARGET G	ROUPS VOLUN	MES CHAP USER	S	
Sear	ch	Q		+ Create target	group
	Na 🦊	Target states	LUNs 🕇	Nodes ↑	¢
	压G1		2	3	

2. Click the desired LUN to open its details, then click the **Limits** pencil icon.

I Im Delete × Size 0 bytes of 1 GiB Limits Read: UNLIMITED Write: UNLIMITED 6	Targets	LUNS			
Size 0 bytes of 1 GiB Limits Read: UNLIMITED	1			<u>च</u> Delete	×
Limits Read: UNLIMITED C	Size		0 bytes of 1 GiB		
	Limits		Read: UNLIMITED Write: UNLIMITED		6

3. In the **Set LUN limit** window that opens, enter limit values and click **Save**.

Set LUN limit	×
IOPS Unlimited Set limit Read limit (IOPS) 100 Write limit (IOPS) 100)
Throughput O Unlimited O Set limit	
Read limit (MB/s) 10 Write limit (MB/s 10	5)
	Cancel Save

Set limits will be shown in LUN details.

5.1.3.4 Detaching Volumes

To detach a volume from a target group, do the following:

1. Open **STORAGE SERVICES** > **Block storage** > **TARGET GROUPS**, click the name of the desired target group to open it, and switch to **LUNs**.

NODI	ES TARGET G	ROUPS VOLUM	ES CHAP USERS	;	
Sear	rch	Q		+ Create target g	group
	Na ↓	Target states	LUNs 🕇	Nodes 🕇	ø
	IIII KG1		2	3	

2. Click the ellipsis button of the desired LUN then click **Detach**.

Target	s LUNs			
Searc	h	Q	÷	Attach
	LUN ID 👃	Volume ↑	Size	¢
	0	tg1-vol1	8 KiB of 1 GiB	
	1	tg1-vol2	0 bytes of 1 GiB	
			←I Detacl	h 🖑

Alternatively, you can open **STORAGE SERVICES** > **Block storage** > **VOLUMES**, click the ellipsis icon of the desired volume, and click **Detach**.

NODE	S TARGET GR	OUPS VOLUMES	CHAP USERS	
Searc	:h	Q		+ Create volume
	Name 🦊	Target gr ↑	LUN ID 🕇	Size 🗘
	⊟ tg1	Target group 1	0	8 KiB of 1 GiB
	⊟ tg1	N/A	N/A	0 bytes 🔶 Detach 네
	⊟ tg1	N/A	N/A	0 bytes 🔟 Delete

5.1.3.5 Deleting Volumes

To delete a volume that is not attached to a target group, open **STORAGE SERVICES** > **Block storage** > **VOLUMES**, click the ellipsis icon of the desired volume, and click **Delete**.

NODE	S TARGET GROU	JPS VOLUMES C	HAP USERS	
Searc	:h	۹		+ Create volume
	Name 🦊	Target group 🕇	LUN ID 🕇	Size 🔅
	🔒 tg1-vol1	TG1	0	8 KiB of 1 →I Attach
	₿ tg1-vol2	TG1	1	0 bytes of 🔟 Delete 👆
	🔒 tg1-vol3	N/A	N/A	0 bytes of 1 GiB

5.1.4 Restricting Access to Target Groups

You can restrict access to entire target groups (and all volumes attached to them) by way of ACL-based authorization as well as password-based authentication (CHAP).

5.1.4.1 Managing Access Control Lists

An access control list (ACL) 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 add an initiator to a target group's ACL, do the following:

- Open STORAGE SERVICES > Block storage > TARGET GROUPS and click the desired target group in the list (anywhere except group's name).
- 2. In group details that open, click **Access control** and then click the pencil icon.

NODES	TARGET GROUPS	VOLUMES	CHAP USERS		
Target g	group 1			🔟 Delete	×
	Overview	Proj	perties	Access control	
CHAI	^o authentication				ı
CHAP	user	user1			
ACL					€
IQN	Ť	Initiator na	ime	LUNs	

3. In the Access control window that opens, check the ACL box and click Add.

Access control	×
ACL CHAP Populate the access control list be allowed to communicate wi	t with iSCSI initiator IQNs that will th the target group.
+ Add	
	Cancel Save

4. In the window that opens, specify initiator's IQN, enter an alias, select LUNs that it will be able to access. Click **Add**. The initiator will appear in the ACL.

Add ACL	×
^{IQN} iqn.1991-05.com.microsoft:example.com	
Initiator name initiator1	
LUNs O	~
Cancel	Add

5. Having populated the ACL with initiators, click **Save**.

To edit or delete initiators in the ACL, click the pencil icon in target group details. In the **Access control** window that opens, click the pencil icon of the desired initiator then click **Edit** or **Delete**. Having changed the ACL, click **Save**.

5.1.4.2 Managing CHAP Users

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 restrict access to a target group to a specific CHAP user, do the following:

- Open STORAGE SERVICES > Block storage > TARGET GROUPS and click the desired target group in the list (anywhere except group's name).
- 2. In group details that open, click **Access control** and then click the pencil icon.

NODES	TARGET GROUPS	VOLUMES	CHAP USERS		
Target g	group 1			🔟 Delete	×
	Overview	Pro	perties	Access control	
CHAF	^o authentication				0
CHAP	user	user1			
ACL					Ъ
IQN	1	Initiator na	me	LUNs	

3. In the **Access control** window that opens, check the **CHAP** box and click **Create user**.



4. In the **Create CHAP user** window that opens, enter a user name and a password (12 to 16 characters long). Click **Create**.

Create CHAP user		\times
Name		
user1		
Password		
	Cancel	e

5. Back on the **Access control** screen, select the desired CHAP user and click **Save**.

Access control		×
ACL CHAP Select CHAP	i iSCSI initiators will conne	ect to the target group.
CHAP user (optional) user1	~	
		Cancel Save

To change the password of a CHAP user, open **STORAGE SERVICES** > **Block storage** > **CHAP USERS**, click a user to open details, and click the pencil icon. In the **Edit CHAP user** window that opens, specify a new password and click **Apply**.

Edit CHAP user	×
Name user1 Password	
	Cancel Apply

To delete a CHAP user that is not added to any ACLs, open **STORAGE SERVICES** > **Block storage** > **CHAP USERS**, click the ellipsis icon of the user, and click **Delete**.

NODE	S TARGET GROUP	S VOLUMES	CHAP USERS	
Searc	:h (٦	+ Create user	
	Name 🤳	Та	arget group	
	🙎 user1	1		
	🙎 user2	0]
			🛅 Delete 🕁	

5.1.5 Monitoring Target Groups

After creating a target group, you can monitor it on the **Overview** tab. The charts show the read and write I/O activity and latency across all LUNs attached to the target group.

Target group 1

 \times

🗓 Delete



5.1.5.1 Advanced iSCSI Monitoring via Grafana

For advanced monitoring of target groups, go to the **MONITORING** > **Dashboard** screen and click **Grafana dashboard**. A separate browser tab will open with preconfigured Grafana dashboards. Two of them are dedicated to the iSCSI service. To see a detailed description for each chart, click **i** in chart's top left corner.

On the **iSCSI overview** dashboard, note the following charts:

• **iSCSI availability**. The chart shows target availability. Time periods when the targets have not been available will be highlighted in red. In this case, check

/var/log/vstorage/iscsi/vstorage-target-monitor.log on the nodes with the failed service and report
a problem.

• **Latency**. The chart shows the time spent on read and write I/O operations across all iSCSI LUNs. It should average a few dozens of milliseconds with peak values below 1s.



The **iSCSI details** dashboard is intended for troubleshooting by the technical support team. To monitor a particular target group, target, session, or LUN, select it from a drop-down list above.



5.2 Exporting Storage via S3

Acronis Cyber Infrastructure allows you to export cluster disk space to customers in the form of an S3-like object-based storage.

Acronis Cyber Infrastructure is implemented as an Amazon S3-like API, which is one of the most common object storage APIs. End users can work with Acronis Cyber Infrastructure as they work with Amazon S3. You can use the usual applications for S3 and continue working with it after the data migration from Amazon S3 to Acronis Cyber Infrastructure.

Object storage is a storage architecture that enables managing data as objects (like in a key-value storage) as opposed to files in file systems or blocks in a block storage. Except for the data, each object has metadata that describes it as well as a unique identifier that allows finding the object in the storage. Object storage is optimized for storing billions of objects, in particular for application storage, static web content hosting, online storage services, big data, and backups. All of these uses are enabled by object storage thanks to a combination of very high scalability and data availability and consistency.

Compared to other types of storage, the key difference of object storage is that parts of an object cannot be modified, so if the object changes a new version of it is spawned instead. This approach is extremely important for maintaining data availability and consistency. First of all, changing an object as a whole eliminates the issue of conflicts. That is, the object with the latest timestamp is considered to be the current version and that is it. As a result, objects are always consistent, i.e. their state is relevant and appropriate.

Another feature of object storage is eventual consistency. Eventual consistency does not guarantee that reads are to return the new state after the write has been completed. Readers can observe the old state for an undefined period of time until the write is propagated to all the replicas (copies). This is very important for storage availability as geographically distant data centers may not be able to perform data update synchronously (e.g., due to network issues) and the update itself may also be slow as awaiting acknowledges from all the data replicas over long distances can take hundreds of milliseconds. So eventual consistency helps hide communication latencies on writes at the cost of the probable old state observed by readers. However, many use cases can easily tolerate it.

5.2.1 S3 Storage Infrastructure Overview

The object storage infrastructure consists of the following entities: object servers (OS), name servers (NS), S3 gateways (GW), and the block-level backend.

These entities run as services on the Acronis Cyber Infrastructure nodes. Each service should be deployed on multiple Acronis Cyber Infrastructure nodes for high availability.


- An object server stores actual object data received from S3 gateway. The data is packed into special containers to achieve high performance. The containers are redundant, you can specify the redundancy mode while configuring object storage. An object server also stores its own data in block storage with built-in high availability.
- A name server stores object metadata received from S3 gateway. Metadata includes object name, size, ACL (access control list), location, owner, and such. Name server (NS) also stores its own data in block storage with built-in high availability.
- An S3 gateway is a data proxy between object storage services and end users. It receives and handles Amazon S3 protocol requests and S3 user authentication and ACL checks. The S3 gateway uses the NGINX web server for external connections and has no data of its own (i.e. is stateless).
- The block-level backend is block storage with high availability of services and data. Since all object storage services run on hosts, no virtual environments (and hence licenses) are required for object storage.

5.2.2 Planning the S3 Cluster

Before creating an S3 cluster, do the following:

- 1. Define which nodes of the storage cluster will run the S3 storage access point services. It is recommended to have all nodes available in Acronis Cyber Infrastructure run these services.
- 2. Configure the network so that the following is achieved:
 - All components of the S3 cluster communicate with each other via the S3 private network. All nodes of an S3 cluster must be connected to the S3 private network. Acronis Cyber Infrastructure internal network can be used for this purpose.
 - The nodes running S3 gateways must have access to the public network.
 - The public network for the S3 gateways must be balanced by an external DNS load balancer.

For more details on network configuration, refer to the Installation Guide.

 All components of the S3 cluster should run on multiple nodes for high-availability. Name server and object server components in the S3 cluster are automatically balanced and migrated between S3 nodes.
 S3 gateways are not automatically migrated; their high availability is based on DNS records. You need to maintain the DNS records manually when adding or removing S3 gateways.

5.2.3 Sample S3 Storage

This section shows a sample object storage deployed on top of a storage cluster of five nodes that run various services. The final setup is shown on the figure below.



5.2.4 Creating the S3 Cluster

To set up object storage services on a cluster node, do the following:

- On the INFRASTRUCTURE > Networks screen, make sure that the OSTOR private and S3 public traffic types are added to your networks.
- 2. In the left menu, click **STORAGE SERVICES** > **S3**.
- 3. Select one or more nodes and click **Create S3 cluster** in the right menu. To create a highly available S3 cluster, select at least three nodes. It is also recommended to enable HA for the management node prior to creating the S3 cluster. See *Enabling High Availability* (page 237) for more details.
- 4. Make sure the correct network interface is selected in the drop-down list.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

node001	¢
Object Storage private	S3
eth1 - 10.37.130.250 🗸	eth0 - 10.94.17.81
node002	\$
Object Storage private	S3
eth1 - 10.37.130.28 💙	eth0 - 10.94.18.146
🚟 node003	ö
Object Storage private	S3
eth1 - 10.37.130.45 💙	eth0 - 10.94.18.147
PROCEED	

× Create S3 cluster

Click **Proceed**.

On the Volume Parameters tab, select the desired tier, failure domain, and data redundancy mode.
 For more information, refer to Understanding Storage Tiers, Understanding Failure Domains, and
 Understanding Data Redundancy.

< Volume parar	neters		
Tier:			
Tier 0			~
Data redundancy:		Failure domain:	
C Erasure coding		Host	~
Replication			
No redundancy			
2 replicas	100% overhe	ead	
3 replicas	200% overhe	ead	
PROCEED			

For replication, you can change the redundancy scheme later on the **S3** > **OVERVIEW** > **Settings** panel. For erasure coding, changing redundancy scheme is disabled, 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.

Click **Proceed**.

6. Specify the external (publicly resolvable) DNS name for the S3 endpoint that will be used by the end users to access the object storage. For example, s3.example.com. Click **Proceed**.

Important: Configure your DNS server according to the example suggested in the admin panel.

7. From the drop-down list, select an S3 endpoint protocol: HTTP, HTTPS or both.

< Protocols	
S3 endpoint protocols:	
HTTPS	v
Endpoint URL: https://s3.example.com/ <i>bucke</i>	etname/objectname
Generate self-signed certifie	cate
SSL certificate: Upload None	
DONE	CONFIGURE NOTARY

It is recommended to use only HTTPS for production deployments.

If you have selected HTTPS, do one of the following:

• Check **Generate self-signed certificate** to get a self-signed certificate for HTTPS evaluation purposes.

Take note of the following:

- S3 geo-replication requires a certificate from a trusted authority. It does not work with self-signed certificates.
- To access the data in the S3 cluster via a browser, add the self-signed certificate to browser's exceptions.
- Acquire a key and a trusted wildcard SSL certificate for endpoint's bottom-level domain. For example, the endpoint s3.storage.example.com would need a wildcard certificate for *.s3.storage.example.com with the subject alternative name s3.storage.example.com.

If you acquired an SSL certificate from an intermediate certificate authority (CA), you should have an end-user certificate along with a CA bundle that contains the root and intermediate certificates. To be able to use these certificates, you need to merge them into a chain first. A certificate chain includes the end-user certificate, the certificates of intermediate CAs, and the certificate of a trusted root CA. In this case, an SSL certificate can only be trusted if every certificate in the chain is properly issued and valid. For example, if you have an end-user certificate, two intermediate CA certificates, and a root CA certificate, create a new certificate file and add all certificates to it in the following order:



Upload the prepared certificate, and, depending on its type, do one of the following:

- specify the passphrase (PKCS#12 files);
- upload the SSL key.

You can change the redundancy mode later on the **S3** > **OVERVIEW** > **Protocol settings** panel. Click **Proceed**.

- 8. If required, click **Configure Acronis Notary** and specify **Notary DNS name** and **Notary user key**.
- 9. Click **Done** to create an S3 cluster.

After the S3 cluster is created, open the **S3 Overview** screen to view cluster status, hostname, used disk capacity, the number of users, I/O activity, and the state of S3 services.

To check if the S3 cluster is successfully deployed and can be accessed by users, visit https://<S3_DNS_name> or http://<S3_DNS_name> in your browser. You should receive the following XML response:



To start using the S3 storage, you will also need to create at least one S3 user.

5.2.5 Managing S3 Users

The concept of S3 user is one of the base concepts of object storage along with those of object and bucket (container for storing objects). The Amazon S3 protocol uses a permission model based on access control lists (ACLs) where each bucket and each object is assigned an ACL that lists all users with access to the given resource and the type of this access (read, write, read ACL, write ACL). The list of users includes the entity owner assigned to every object and bucket at creation. The entity owner has extra rights compared to other users. For example, the bucket owner is the only one who can delete that bucket.

User model and access policies implemented in Acronis Cyber Infrastructure comply with the Amazon S3 user model and access policies.

User management scenarios in Acronis Cyber Infrastructure are largely based on the Amazon Web Services user management and include the following operations: create, query, and delete users as well as generate and revoke user access key pairs.

5.2.5.1 Adding S3 Users

To add an S3 user, do the following:

- 1. On the **STORAGE SERVICES** > **S3** > **USERS** screen, click **ADD USER**.
- 2. Specify a valid email address as login for the user and click **ADD**.



5.2.5.2 Managing S3 Access Key Pairs

Each S3 user has one or two key pairs (access key and secret key) for accessing the S3 cloud. You can think of the access key as login and the secret key as password. (For more information about S3 key pairs, refer to the Amazon documentation.) The access keys are generated and stored locally in the storage cluster on S3

name servers. Each user can have up to two key pairs. It is recommended to periodically revoke old and generate new access key pairs.

To view, add, or revoke the S3 access key pairs for an S3 user, do the following:

× Kevs

1. Select a user in the list and click **Keys**.

-
S3 access key ID
9110226866ee63caWB1D
S3 secure access key
M9QQgiwZ5lTY3A52bQvU9lrDGVowiOiFb4nciP5W
🕮 <u>Revoke</u>
0
Generate access key

- 2. The existing keys will be shown on the **Keys** panel.
 - To revoke a key, click **Revoke**.
 - To add a new key, click **Generate access key**.

To access a bucket, a user will need the following information:

- admin panel IP address,
- DNS name of the S3 cluster specified during configuration,
- S3 access key ID,
- S3 secret access key,
- SSL certificate if the HTTPS protocol was chosen during configuration.

The certificate file can be found in the /etc/nginx/ssl/ directory on any node hosting the S3 gateway service.

To automatically log in to S3 with user credentials using the generated keys, select a user and click **Browse**.

To **Browse** using an SSL certificate, make sure it is valid or, in case of a self-signed one, add it to browser's exceptions.

5.2.6 Managing S3 Buckets

All objects in Amazon S3-like storage are stored in containers called "buckets". Buckets are addressed by names that are unique in the given object storage, so an S3 user of that object storage cannot create a bucket that has the same name as a different bucket in the same object storage. Buckets are used to:

- group and isolate objects from those in other buckets,
- provide ACL management mechanisms for objects in them,
- set per-bucket access policies, for example, versioning in the bucket.

In the current version of Acronis Cyber Infrastructure, you can enable and disable Acronis Notary for object storage buckets and monitor the space used by them on the **STORAGE SERVICES** > **S3** > **Buckets** screen. You cannot create and manage object storage buckets from Acronis Cyber Infrastructure admin panel. However, you can do it via the Acronis Cyber Infrastructure user panel or by using a third-party application. For example, the applications listed below allow you to perform the following actions:

- CyberDuck: create and manage buckets and their contents.
- MountainDuck: mount object storage as a disk drive and manage buckets and their contents.
- Backup Exec: store backups in the object storage.

5.2.6.1 Listing S3 Bucket Contents

You can list bucket contents with a web browser. To do this, visit the URL that consists of the external DNS name for the S3 endpoint that you specified when creating the S3 cluster and the bucket name. For example, mys3storage.example.com/mybucket or mybucket.mys3storage.example.com (depending on DNS configuration).

You can also copy the link to bucket contents by right-clicking it in CyberDuck, and then selecting Copy URL.

5.2.6.2 Managing Acronis Notary in S3 Buckets

Acronis Cyber Infrastructure offers integration with the Acronis Notary service to leverage blockchain notarization and ensure the immutability of data saved in object storage clusters. To use Acronis Notary in user buckets, you need to set it up in the S3 cluster and enable it for said buckets.

To set up Acronis Notary, do the following:

- 1. Get the DNS name and the user key for the notary service from your sales contact.
- 2. On the **STORAGE SERVICES** > **S3** screen, click **Notary settings**.
- 3. On the **Notary Settings** screen, specify the DNS name and user key in the respective fields and click **Done**.

× Notary settings
Notary DNS name
http://ask-sales-for-notary-provider.acronis.com
Notary user key

DONE

To enable or disable blockchain notarization for a bucket, select a bucket on the **STORAGE SERVICES** > **S3** > **Buckets** screen and click **Enable Notary** or **Disable Notary**, respectively.

Notarization is disabled for new buckets by default.

Once you enable notarization for a bucket, certificates are created automatically only for the newly uploaded files. The previously uploaded files are left unnotarized. Once a file was notarized, it will remain notarized even if you disable notarization later.

5.2.7 Best Practices for Using S3 in Acronis Cyber Infrastructure

This section offers recommendations on how to best use the S3 feature of Acronis Cyber Infrastructure.

5.2.7.1 S3 Bucket and Key Naming Policies

It is recommended to use bucket names that comply with DNS naming conventions:

- can be from 3 to 63 characters long
- must start and end with a lowercase letter or number
- can contain lowercase letters, numbers, periods (.), hyphens (-), and underscores (_)
- can be a series of valid name parts (described previously) separated by periods

An object key can be a string of any UTF-8 encoded characters up to 1024 bytes long.

5.2.7.2 Improving Performance of PUT Operations

Object storage supports uploading objects as large as 5 GB per single PUT request (5 TB via multipart upload). Upload performance can be improved by splitting large objects into pieces and uploading them concurrently (thus dividing the load between multiple OS services) with multipart upload API.

It is recommended to use multipart uploads for objects larger than 5 MB.

5.2.8 Replicating S3 Data Between Datacenters

Acronis Cyber Infrastructure can store replicas of S3 cluster data and keep them up-to-date in multiple geographically distributed datacenters with S3 clusters based on Acronis Cyber Infrastructure. Geo-replication reduces the response time for local S3 users accessing the data in a remote S3 cluster or remote S3 users accessing the data in a local S3 cluster as they do not need to have an Internet connection.

Geo-replication schedules the update of the replicas as soon as any data is modified. Geo-replication performance depends on the speed of Internet connection, the redundancy mode, and cluster performance.

If you have multiple datacenters with enough free space, it is recommended to set up geo-replication between S3 clusters residing in these datacenters.

Important: Each cluster must have its own SSL certificate signed by a global certificate authority.

To set up geo-replication between S3 clusters, exchange tokens between datacenters as follows:

 In the admin panel of a remote datacenter, open the STORAGE SERVICES > S3 > GEO-REPLICATION screen.



s3.example.com

- 2. In the section of the home S3 cluster, click **TOKEN** and, on the **Get token** panel, copy the token.
- In the admin panel of the local datacenter, open the STORAGE SERVICES > S3 > GEO-REPLICATION screen and click ADD DATACENTER.

× Add datacenter

To replicate data from another datacenter, insert the token obtained from its management panel.

Token

eyJ1c2VyX3NlY3JldF9rZXkiOiAiOVdLRjB3ZkJlQU85Y2JqTWJVRUdGU0p Fb0Z1dGNMemhJSGF2ZGg2SCIsICJ1aWQiOiAiNWIxNWQyMmM3MD Q4Yjg5YyIsICJyZWFkYWJsZV9uYW1lIjogInN0b3IxIiwgInVzZXJfa2V5X2Ik IjogImRjNGI0MTg2ZjU0MzU4MTZTRUhCIiwgInVybCl6ICJodHRwczovL 3MzLmV4YW1wbGUuY29tOjQ0MyIsICJpc19zZWxmIjogdHJ1ZX0=



- 4. Enter the copied token and click **Done**.
- 5. Configure the remote S3 cluster the same way.

5.2.9 Monitoring S3 Access Points

The S3 monitoring screen enables you to inspect the availability of each S3 component as well as the performance of NS and OS services (which are highly available).

If you see that some of the NS or OS services are offline, it means that the S3 access point does not function properly. In this case, please contact the techical support or consult the Administrator's Command Line Guide. S3 gateways are not highly available, but DNS load balancing should be enough to avoid downtime if the gateway fails.

The performance charts represent the number of operations that the OS/NS services are performing.

5.2.10 Releasing Nodes from S3 Clusters

Before releasing a node, make sure that the cluster has enough nodes running name and object servers as well as gateways left.

Warning: When the last node in the S3 cluster is removed, the cluster is destroyed, and all the data is deleted.

To release a node from an S3 cluster, do the following:

- 1. On the **STORAGE SERVICES** > **S3** > **Nodes** screen, check the box of the node to release.
- 2. Click Release.

5.2.11 Supported Amazon S3 Features

Besides basic Amazon S3 operations like GET, PUT, COPY, DELETE, the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol supports the following features:

- multipart upload
- access control lists (ACLs)
- versioning
- signed URLs

- object locking
- geo-replication
- blockchain notarization via Acronis Notary

All the supported Amazon S3 operations, headers, and authentication schemes are listed further.

5.2.11.1 Supported Amazon S3 REST Operations

The following Amazon S3 REST operations are currently supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

Supported service operations: GET Service.

Supported bucket operations:

- DELETE/HEAD/PUT Bucket
- GET Bucket (List Objects)
- GET/PUT Bucket acl
- GET Bucket location (returns US East)
- GET Bucket Object versions
- GET/PUT Bucket versioning
- List Multipart Uploads

Supported object operations:

- DELETE/GET/HEAD/POST/PUT Object
- Delete Multiple Objects
- PUT Object Copy
- GET/PUT Object acl
- Delete Multiple Objects
- Abort Multipart Upload
- Complete Multipart Upload
- Initiate Multipart Upload

- List Parts
- Upload Part

Note: For more information on Amazon S3 REST operations, see Amazon S3 REST API documentation.

5.2.11.2 Supported Amazon Request Headers

The following Amazon S3 REST request headers are currently supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

- Authorization
- Content-Length
- Content-Type
- Content-MD5
- Date
- Host
- x-amz-content-sha256
- x-amz-date
- x-amz-security-token
- x-amz-object-lock-retain-until-date
- x-amz-object-lock-mode
- x-amz-object-lock-legal-hold
- x-amz-bypass-governance-retention
- x-amz-bucket-object-lock-enabled

The following Amazon S3 REST request headers are ignored:

- Expect
- x-amz-security-token

Note: For more information on Amazon S3 REST request headers, see the Amazon S3 REST API documentation.

5.2.11.3 Supported Amazon Response Headers

The following Amazon S3 REST response headers are currently supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

- Content-Length
- Content-Type
- Connection
- Date
- ETag
- x-amz-delete-marker
- x-amz-request-id
- x-amz-version-id
- x-amz-object-lock-retain-until-date
- x-amz-object-lock-mode
- x-amz-object-lock-legal-hold

The following Amazon S3 REST response headers are not used:

- Server
- x-amz-id-2

Note: For more information on Amazon S3 REST response headers, see the Amazon S3 REST API documentation.

5.2.11.4 Supported Amazon Error Response Headers

The following Amazon S3 REST error response headers are currently supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

- Code
- Error
- Message
- RequestId
- Resource

The following Amazon S3 REST error response headers are not supported:

- RequestId (not used)
- Resource

Note: For more information on Amazon S3 REST response headers, see the Amazon S3 REST API documentation.

5.2.11.5 Supported Authentication Schemes

The following authentication schemes are supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

- Signature Version 2
- Signature Version 4

The following authentication methods are supported by the Acronis Cyber Infrastructure implementation of the Amazon S3 protocol:

- Using the authorization header
 - Transferring payload in a single chunk
- Using query parameters
- Browser-based uploads using POST

The following authentication method is not supported:

• Transferring payload in multiple chunks

5.3 Exporting Storage via NFS

Acronis Cyber Infrastructure allows you to organize nodes into a highly available NFS cluster in which you can create NFS shares. In Acronis Cyber Infrastructure terms, an NFS share is an access point for a volume and as such it can be assigned an IP address or DNS name. The volume, in turn, can be assigned the usual properties: redundancy type, tier, and failure domain. In each share you can create multiple NFS exports which are actual exported directories for user data. Each export has, among other properties, a path that, combined with share's IP address, uniquely identifies the export on the network and allows you to mount it using standard commands.

On the technical side, NFS volumes are based on object storage. Aside from offering high availability and scalability, object storage eliminates the limit on the amount of files and the size of data you can keep in the NFS cluster. Each share is perfect for keeping billions of files of any size. However, such scalability implies IO overhead that is wasted on file size changes and rewrites. For this reason, an NFS cluster makes a perfect cold and warm file storage but is not recommended for hot and high performance, often rewritten data (like running virtual machines). Integration of Acronis Cyber Infrastructure with solutions from VMware, for example, is best done via iSCSI to achieve better performance.

Note: Acronis Cyber Infrastructure only supports NFS version 4 and newer, including pNFS.

5.3.1 Setting Up an NFS Cluster

Since NFS is based on object storage, creating an NFS cluster is similar to creating an S3 one. Do the following:

- On the INFRASTRUCTURE > Networks screen, make sure that the OSTOR private and NFS traffic types are added to your networks.
- 2. In the left menu, click **STORAGE SERVICES** > **NFS**.
- 3. Select one or more nodes and click **Create NFS cluster** in the right menu.
- 4. Make sure the correct network interface is selected in the drop-down list.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

5. Click CREATE.

After the NFS cluster has been created, you can proceed to creating NFS shares.

5.3.2 Creating NFS Shares

To create an NFS share, do the following:

- 1. On the **STORAGE SERVICES** > **NFS** > **SHARES** screen, click **ADD NFS SHARE**.
- 2. On the **Add NFS Share** panel, specify a unique name and an IP address, which must be unused and, if authentication is enabled, domain resolvable. In addition, this IP address should be within the network subnet of the node's interface. Click **PROCEED**.
- 3. In **Share size**, specify the size of the share in gigabytes. For users accessing exports, this value will be the filesystem size.
- 4. Select the desired tier, failure domain, and data redundancy type. For more details on these volume properties, see the Installation Guide.

You will be able to change the redundancy mode later.

5. Click DONE.

After the share has been created, you can proceed to creating NFS exports.

5.3.3 Creating NFS Exports

The process of creating NFS exports includes the following steps:

- 1. Creating a root export that will contain user exports.
- 2. Mounting the root export.
- 3. Creating user exports in the mounted root export.

5.3.3.1 Creating the Root Export

To create a root NFS export, do the following:

- 1. On the **STORAGE SERVICES** > **NFS** > **SHARES** screen, click the number in the **Exports** column in the row of the desired share. This will open the share screen.
- 2. On the share screen, click **ADD EXPORT**, specify root as the export name and / as path and select the read and write access mode.

Important: Do not use other names or paths for the root export.

imes Add Export	
Name	
Path]
/ Access	
 Read and write Read 	
DONE	

This will create a directory with a default path that designates export location inside the share. This path is automatically generated based on the share name and used (alongside the share's IP address) to mount the export.

Important: Do not give the users access to the root export.

The root export will be shown in the export list.

After creating the root export, mount it as described in the Storage User's Guide.

Warning: Do not mount NFS shares on cluster nodes. It may lead to node freeze.

5.3.3.2 Creating User Exports

After creating and mounting the root export, you can proceed to creating user NFS exports. To do this:

- 1. In the mounted root export, create a subdirectory for a user export, e.g., export1.
- 2. On the share screen, click **ADD EXPORT**, enter a user export name, specify /export1 as path, and select the access mode.
- 3. Click Done.

The user export will appear in the export list.

5.3.4 Setting Up User Authentication and Authorization

Acronis Cyber Infrastructure allows you to authenticate users for access to specific NFS shares via Kerberos and authorize them to access specific NFS exports inside these shares via LDAP.

5.3.4.1 Authenticating NFS Share Users with Kerberos

To enable user authentication in an NFS share, do the following:

- 1. Assign a forward and reverse resolvable FQDN (fully qualified domain name) to share's IP address.
- 2. On the **SETTINGS** > **Security** > **KERBEROS** tab, specify the following Kerberos information:
 - 2.1. In **Realm**, your DNS name in uppercase letters.
 - 2.2. In **KDC service**, the DNS name or IP address of the host running the realm's KDC (key distribution center) service.
 - 2.3. In **KDC administration service**, the DNS name or IP address of the host running the realm's KDC administration service.

Usually, the KDC and its administration service run on the same host.

- 3. On the Kerberos server, perform these steps:
 - 3.1. Log in as administrator to the Kerberos database administration program.
 - 3.2. Add a principal for the share with the command addprinc -randkey nfs/<share_FQDN>@<realm>. For example:

addprinc -randkey nfs/share1.example.com@example.com

3.3. Generate a keytab (key table) for the principal and save it to a directory you can upload from. For example:

ktadd -k /tmp/krb5.keytab nfs/share1.example.com@example.com

- 4. On the **STORAGE SERVICES** > **NFS** > **SHARE** tab, select a share and click **Authentication**.
- 5. Upload the keytab file and click **SAVE**.

Important: Each share and client (user that mounts the export) must have their own principal and keytab.

5.3.4.2 Authorizing NFS Export Users with LDAP

By configuring access to a user directory via LDAP, you can control which users can access which NFS exports. You will need a directory of user accounts with desired NFS access parameters.

To configure access to an LDAP server, do the following:

- 1. On the **SETTINGS** > **Security** > **LDAP** tab, specify the following information:
 - Address, the IP address of the LDAP server;
 - Base DN, the distinguished name of the search starting point;
- 2. Click Save.

5.4 Connecting Acronis Backup Software to Storage Backends via Backup Gateway

The Backup Gateway storage access point (also called "gateway") is intended for service providers who use Acronis Backup Cloud and/or Acronis Backup Advanced and want to organize an on-premises storage for their clients' backed-up data.

Backup Gateway enables a service provider to easily configure storage for the proprietary deduplication-friendly data format used by Acronis.

Backup Gateway supports the following storage backends:

- storage clusters with software redundancy by means of erasure coding
- NFS shares
- public clouds, including a number of S3 solutions as well as Microsoft Azure, OpenStack Swift, and Google Cloud Platform

While your choice should depend on scenario and requirements, it is recommended to keep Acronis backup data in the local storage cluster. In this case, you can have the best performance due to WAN optimizations and data locality. Keeping backups in an NFS share or a public cloud implies the unavoidable data transfer and other overhead, which reduces overall performance.

Take note of the following:

- When configuring Backup Gateway, you will need to provide the credentials of your administrator account in the Acronis backup software.
- In cases when not local but external storage (e.g., NFS) is used with Backup Gateway, redundancy has to be provided by the said external storage. Backup Gateway does not provide data redundancy or perform data deduplication itself.
- To be able to register Backup Gateway in Acronis Backup Cloud, two-factor authentication (2FA) should be disabled for your partner account.

5.4.1 Understanding the Infrastructure

The Backup Gateway storage access point runs as services on the Acronis Cyber Infrastructure nodes. It is recommended to deploy it on two or more nodes for high availability.



5.4.2 Connecting to the Local Storage Cluster via Backup Gateway

Before you proceed, make sure that the destination storage has enough space for both existing and new backups.

To set up Backup Gateway, do the following:

- On the INFRASTRUCTURE > Networks screen, make sure that the ABGW private and ABGW public traffic types are added to the networks you intend to use.
- 2. In the left menu, click **STORAGE SERVICES** > **Backup storage**.
- 3. Select the node(s) to run the gateway services on and click **Create gateway** in the right menu.

4. Select **This Acronis Cyber Infrastructure cluster** as storage type.

.

5. Make sure the correct network interface is selected in the drop-down list. Click **NEXT**.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

< Configure netwo	rk	
node001		٠
ABGW private	ABGW public	
eth1 - 10.37.130.250 🗸	eth0 - 10.94.17.81	
ABGW private	ABGW public	٥
eth1 - 10.37.130.28	etnu - 10.94.18.146	
	ВАСК	NEXT

 On the Volume Parameters tab, select the desired tier, failure domain, and data redundancy mode.
 For more information, refer to Understanding Storage Tiers, Understanding Failure Domains, and Understanding Data Redundancy.

< Volume parameters

Tier:		
Tier 0		~
Data redundancy:	Failure domain:	
• Erasure coding	Host	~
Encoding 1+0	0% overhead	
Encoding 1+1	100% overhead	
Encoding 1+2	200% overhead	
	BACK	хт

Redundancy by replication is not supported for Backup Gateway. For erasure coding, changing redundancy scheme is disabled, 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.

Click **NEXT**.

7. On the DNS Configuration tab, specify the external DNS name for this gateway, e.g, backupgateway.example.com. Make sure that each node running the gateway service has a port open for outgoing Internet connections and incoming connections from your Acronis backup software. Backup agents will use this address and port to upload the backup data.

Important: Configure your DNS server according to the example suggested in the admin panel.

Important: Each time you change nodes in the Backup Gateway cluster, adjust the DNS settings accordingly.

< DNS configuration
DNS name
backup.example.com
This may require changing the DNS server configuration, which may look as follows:
 \$TTL 1h (a) IN SOA ns1.myhoster.com. root.backup.example.com. (2018120313 ; serial 1h ; refresh 30m ; retry 7d ; expiration 1h) ; minimum
BACK NEXT

Click Next.

8. On the **Register in backup software** pane, specify the following information for your Acronis product:

Important: Make sure that two-factor authentication (2FA) is disabled for your partner account. You can also disable it for a specific user within a 2FA-enabled tenant as described in Acronis Cyber Cloud documentation and specify the user credentials.

- In Address, specify the address of the Acronis Backup Cloud management portal (e.g., https://cloud.acronis.com/) or the hostname/IP address and port of the Acronis Backup Advanced management server (e.g., http://192.168.1.2:9877).
- In **Account**, specify the credentials of a partner account in the cloud or of an organization administrator on the local management server.
- 9. Finally, click **DONE**.

5.4.3 Connecting to External NFS Shares via Backup Gateway

Take note of these limitations:

- Acronis Cyber Infrastructure does not provide data redundancy on top of NFS volumes. Depending on the implementation, NFS shares may use their own hardware or software redundancy.
- In the current version of Acronis Cyber Infrastructure, only one cluster node may store backups on an NFS volume.

Before you proceed, make sure that:

- 1. The NFS share has enough space for backups.
- 2. Each NFS export is used by only one gateway. In particular, do not configure two Acronis Cyber Infrastructure installations to use the same NFS export for backup storage.

To set up Backup Gateway, do the following:

- On the INFRASTRUCTURE > Networks screen, make sure that the ABGW private and ABGW public traffic types are added to the networks you intend to use.
- 2. In the left menu, click **STORAGE SERVICES** > **Backup storage**.
- 3. Select the node(s) to run the gateway services on and click **Create gateway** in the right menu.
- 4. Select **Network File System** as storage type.
- 5. Make sure the correct network interface is selected in the drop-down list. Click **NEXT**.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

< Configure netv	vork
anode001	0
ABGW private	ABGW public
eth1 - 10.37.130.109	✓ eth0 - 10.94.16.12

6. On the **Volume Parameters** tab, specify the hostname or IP address of the NFS share as well as the export name. Click **NEXT**.

Volume parameters

NFS hostname or IP				
nfs.example.com				
Export name				
/path/to/volume				
 NFS3 (no clustering) 				
O NFS4				
	٢	ВАСК	NE	хт

7. On the DNS Configuration tab, specify the external DNS name for this gateway, e.g, backupgateway.example.com. Make sure that each node running the gateway service has a port open for outgoing Internet connections and incoming connections from your Acronis backup software. Backup agents will use this address and port to upload the backup data.

Important: Configure your DNS server according to the example suggested in the admin panel.

Important: Each time you change nodes in the Backup Gateway cluster, adjust the DNS settings accordingly.

< DNS configuration

ba	ckup.example.com
This	s may require changing the DNS server configuration, which may
ool	< as follows:
\$T	TL 1h
@	IN SOA ns1.myhoster.com. root.backup.example.com. (
	2018120313 ; serial
	30m ; retry
	7d ; expiration
	1h) : minimum

Click Next.

8. On the **Register in backup software** pane, specify the following information for your Acronis product:

Important: Make sure that two-factor authentication (2FA) is disabled for your partner account. You can also disable it for a specific user within a 2FA-enabled tenant as described in Acronis Cyber Cloud documentation and specify the user credentials.

- In Address, specify the address of the Acronis Backup Cloud management portal (e.g., https://cloud.acronis.com/) or the hostname/IP address and port of the Acronis Backup Advanced management server (e.g., http://192.168.1.2:9877).
- In **Account**, specify the credentials of a partner account in the cloud or of an organization administrator on the local management server.
- 9. Finally, click **DONE**.

5.4.4 Connecting to Public Cloud Storage via Backup Gateway

With Backup Gateway, you can have Acronis Backup Cloud or Acronis Backup Advanced store backups in a number of public clouds and on-premises object storage solutions:

- Amazon S3
- IBM Cloud
- Alibaba Cloud
- IIJ
- Cleversafe
- Cloudian
- Microsoft Azure
- Swift object storage
- Softlayer (Swift)
- Google Cloud Platform
- Wasabi
- Other solutions using S3

However, compared to the local storage cluster, storing backup data in a public cloud increases the latency of all I/O requests to backups and reduces performance. For this reason, it is recommended to use the local storage cluster as storage backend.

Backups are cold data with a specific access pattern: the data is not accessed frequently but is expected to be available immediately when accessed. For this use case, it is cost-efficient to choose storage classes intended for long-term storage with infrequently accessed data. The recommended storage classes include the following:

- Infrequent Access for Amazon S3
- Cool Blob Storage for Microsoft Azure
- Nearline and Coldline storage for Google Cloud Platform

Archive storage classes like Amazon S3 Glacier, Azure Archive Blob, or Google Archive cannot be used for backup because they do not provide instant access to data. High access latency (several hours) makes it

technically impossible to browse archives, restore data fast, and create incremental backups. Even though the archive storage is usually very cost-efficient, keep in mind that there are a number of different cost factors. In fact, the total cost of public cloud storage consists of payments for storing data, operations, traffic, data retrieval, early deletion and so on. For example, an archive storage service can charge six months' storage payment for just one data recall operation. If the storage data is expected to be accessed more frequently, the added costs increase significantly the total cost of data storage. In order to avoid the low data retrieval rate and to cut expenses, we recommend using Acronis Cyber Cloud for storing backup data.

5.4.4.1 Important Requirements and Restrictions

- When working with public clouds, Backup Gateway uses the local storage as the staging area as well as to keep service information. It means that the data to be uploaded to a public cloud is first stored locally and only then sent to the destination. For this reason, it is vital that the local storage is persistent and redundant so the data does not get lost. There are multiple ways to ensure the persistence and redundancy of local storage. You can deploy Backup Gateway on multiple cluster nodes and select a good redundancy mode. If Acronis Cyber Infrastructure with the gateway is deployed on a single physical node, you can make the local storage redundant by replicating it among local disks. If Acronis Cyber Infrastructure with the gateway is deployed in a virtual machine, make sure it is made redundant by the virtualization solution it runs on.
- Make sure the local storage cluster has plenty of logical space for staging. For example, if you perform backup daily, provide enough space for at least 1.5 days' worth of backups. If the daily backup total is 2TB, provide at least 3TB of logical space. The required raw storage will vary depending on the encoding mode: 9TB (3TB per node) in the 1+2 mode, 5TB (1TB per node) in the 3+2 mode, etc.
- If you are to store backups in an Amazon S3 cloud, keep in mind that Backup Gateway may sometimes block access to such backups due to the eventual consistency of Amazon S3. It means that Amazon S3 may occasionally return stale data as it needs time to render the most recent version of the data accessible. Backup Gateway detects such delays and protects backup integrity by blocking access until the cloud updates.
- Use a separate object container for each Backup Gateway cluster.

5.4.4.2 Setting Up Backup Gateway

Before you proceed, make sure that the destination storage has enough space for both existing and new backups.

To set up Backup Gateway, do the following:

- On the INFRASTRUCTURE > Networks screen, make sure that the ABGW private and ABGW public traffic types are added to the networks you intend to use.
- 2. In the left menu, click **STORAGE SERVICES** > **Backup storage**.
- 3. Select the node(s) to run the gateway services on and click **Create gateway** in the right menu.
- 4. Select **Public Cloud** as storage type.
- 5. Make sure the correct network interface is selected in the drop-down list. Click **NEXT**.

Configure network

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

node001	¢
ABGW private	ABGW public
eth1 - 10.37.130.250 🗸	eth0 - 10.94.17.81
node002	۵
ABGW private	ABGW public
eth1 - 10.37.130.28 🗸	eth0 - 10.94.18.146
	BACK

- 6. On the **Public cloud parameters** pane, do the following:
 - 6.1. Select a public cloud provider. If your provider is S3 compatible but not in the list, try AuthV2 compatible (S3) or AuthV4 compatible (S3).
 - 6.2. Depending on the provider, specify **Region**, **Authentication** (keystone) URL, or **Endpoint** URL.

- 6.3. In case of Swift object storage, specify the authentication protocol version and attributes required by it.
- 6.4. Specify user credentials. In case of Google Cloud, select a JSON file with keys to upload.
- 6.5. Specify the folder (bucket, container) to store backups in. The folder must be writeable.

Use a separate object container for each Backup Gateway cluster.

Click **NEXT**.

7. On the **Register in backup software** pane, specify the following information for your Acronis product:

Important: Make sure that two-factor authentication (2FA) is disabled for your partner account. You can also disable it for a specific user within a 2FA-enabled tenant as described in Acronis Cyber Cloud documentation and specify the user credentials.

- In Address, specify the address of the Acronis Backup Cloud management portal (e.g., https://cloud.acronis.com/) or the hostname/IP address and port of the Acronis Backup Advanced management server (e.g., http://192.168.1.2:9877).
- In **Account**, specify the credentials of a partner account in the cloud or of an organization administrator on the local management server.
- 8. Finally, click **DONE**.

5.4.5 Updating certificate for Backup Gateway

When you register a Backup Gateway in Acronis Backup Cloud or Acronis Backup Advanced, they exchange certificates that are valid for three years. One and a half months before expiration, you will be alerted about the expiring certificate in the admin panel. To update the certificate, you need to connect to your backup software and renew the certificate. Do the following:

- 1. On the **STORAGE SERVICES** > **Backup storage** screen, click **Update certificate**.
- 2. On the **Connect to backup software** pane, specify the following information for your Acronis product:

Important: Make sure that two-factor authentication (2FA) is disabled for your partner account. You can also disable it for a specific user within a 2FA-enabled tenant as described in Acronis Cyber Cloud documentation and specify the user credentials.

- In Address, specify the address of the Acronis Backup Cloud management portal (e.g., https://cloud.acronis.com/) or the hostname/IP address and port of the Acronis Backup Advanced management server (e.g., http://192.168.1.2:9877).
- In **Account**, specify the credentials of a partner account in the cloud or of an organization administrator on the local management server.

X Connect to backup software
Connect to backup software where this storage is registered.
Address
https://cloud.acronis.com
Enter the URL of the cloud management portal or <name address:port="" ip=""> of the local management server.</name>
Account
admin
Enter the credentials of a partner account in the cloud or of an organization administrator on the local management server.
BACK

3. Click **NEXT**.

4. On all nodes included into the Backup Gateway cluster, restart the service:
systemctl restart vstorage-abgw

5.4.6 Re-registering Backup Gateway in a New Acronis Backup Advanced

To switch a configured Backup Gateway to a different Acronis Backup Advanced instance, re-register the gateway with that instance. To do this:

- 1. On the **STORAGE SERVICES** > **Backup storage** screen, click **Re-register**.
- 2. On the **Re-registration in Acronis Backup** tab, specify the following:
 - In Address, specify the hostname/IP address of the target management server and the port 9877 (e.g., http://192.168.1.2:9877). Note that the address must be provided using the HTTP protocol, not HTTPS.
 - In **Account**, specify the credentials of the management server administrator account.
- 3. Click DONE.

5.4.7 Migrating Backups from Older Acronis Solutions

By means of Backup Gateway, you can migrate backups from Acronis Storage 1.5 and Acronis Storage Gateway 1.6 and 1.7 to a storage backend of your choice: the local storage cluster, external NFS, or public cloud.

Migration to NFS backends is not available, however, if multiple nodes are selected as Backup Gateway.

Important: Before you proceed, make sure that the destination storage has enough space for both existing and new backups.

The migration procedure can be described as follows:

- 1. Root credentials for SSH access to the chosen source storage are provided to Backup Gateway.
- 2. Backup Gateway sets up a proxy on the source storage that starts redirecting requests incoming from Acronis Backup Agents from the source storage to Backup Gateway.

 Backup Gateway starts relocating backups to the chosen storage backend. The data that remains to be migrated is shown in the Migration Backlog section on the Backup Gateway Overview screen. When the backlog empties, all data has been migrated.

After the migration has started, the data of new and incremental backups is stored on the destination storage. Backups from the source storage are pulled in the background. The entire process is transparent to backup agents, which continue working uninterrupted.

- 4. To be able to dispose of the source storage after migration completes, requests from Acronis Backup Agents are directed straight to Backup Gateway, bypassing the proxy on the source storage. Steps that you need to take depend on how the source storage is registered in Acronis Backup Cloud: under the IP address or DNS name.
 - If the source storage is already registered under the DNS name, you need to change the IP address behind it to those of the Backup Gateway nodes.
 - If the source storage is registered under the IP address, it is strongly recommended to re-register Backup Gateway in Acronis Backup Cloud under a DNS name that resolves into the IP addresses of Backup Gateway nodes. Using a DNS name will provide a smoother transition and you will not need to reconfigure Acronis Backup Cloud even if you change nodes in the Backup Gateway (you will still need to adjust the IP addresses behind the DNS name accordingly).

Alternatively, if you do not want to use a DNS name, you need to wait for the migration to complete, shut down both the source and destination machines, and reconfigure your network so that the public interface of the destination machine gets the IP address of the source machine.

The concrete steps that you need to perform in the admin panel to initiate backup migration are described in the next subsections.

5.4.7.1 Migrating Backups from Acronis Storage 1.5

- Update all Acronis Storage 1.5 nodes to version 1.5.65665 or newer as earlier versions are not eligible for migration. To do this, log in to the Acronis Storage web console, proceed to SETTINGS > Software Update, upload the latest ISO image, and click Update.
- Log in to the new storage cluster and on the STORAGE SERVICES > Backup storage > NODES screen, select one or more nodes and click Migrate.
- 3. Select Acronis Storage 1.5 and click NEXT.
- 4. Specify the DNS name of the source storage registered in Acronis Backup Cloud and click **NEXT**.

< Enter source storage	DNS (2/9)	
Specify the DNS name of the source stor Backup Cloud. DNS name	age registered in Acronis	
source.example.com		
	BACK	

- 5. Provide the credentials for the cloud management portal of the Acronis Backup Cloud installation that the source storage is registered in and click **NEXT**.
- 6. Enable SSH access on all FES nodes of Acronis Storage 1.5 as instructed and click **NEXT**.
- 7. Map the public IP addresses of FES nodes accessible via SSH to their private IP addresses and click **NEXT**. This step is required to access FES nodes via SSH through tunnels.

Set up IP mappin	g for FES nodes (5/9)
Listed below are public IP address storage. For each FES node, specit connections.	ses of the FES nodes in the source fy its private IP address open for SSH
Public IP address (FES)	Private IP address (SSH)
10.28.74.3	10.28.74.1:2001
10.28.74.9	10.28.74.1:2002
	BACK NEXT

- 8. Choose a storage type to create a gateway to one of the destinations:
 - local storage cluster

- external NFS
- public cloud
- 9. Make sure the correct network interface is selected in the drop-down list. Click **NEXT**.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

< Configure netwo	rk	
node001		¢
ABGW private	ABGW public	
eth1 - 10.37.130.250 🗸	eth0 - 10.94.17.81	
ABGW private	ABGW public eth0 - 10 94 18 146	٥
etni - 10.37.130.28	ВАСК	NEXT

- 10. Configure the destination storage backend:
 - For a storage cluster, select the desired tier, failure domain, and redundancy mode.
 - For NFS, specify a hostname or IP address, an export name and path, and choose the NFS version.

< Volume parameters

NFS hostname or IP		
nfs.example.com		
Export name		
/path/to/volume		
 NFS3 (no clustering) 		
NF54		
	ВАСК	NEXT

• For a public cloud, select a public cloud provider, specify credentials, and the name of the folder (bucket, container).

Use a separate object container for each Backup Gateway cluster.

< Public cloud parameters	
Select the object storage time	
Select the object storage type	
Amazon S3	~
Region	
us-east-1	~
Access key ID	
Secret Access key	
Bucket	
bocket	
acronis-us-west-gateway-files	

Click **NEXT**.

11. Review the source and destination storages and click **PROCEED**.

12. On the next panel, follow the instructions to point the source storage DNS name to the IP addresses of your new storage cluster. Having updated the DNS configuration, wait for 24 hours for all backup agents to cache the new IP addresses. Until this happens, the **START MIGRATION** button will be disabled. After all backup agents have been rerouted to the new cluster, the button will become enabled and you can click it to start migration.

Reconfigure DNS

Before migration can start, all traffic between backup agents and source storage must be rerouted via a TCP proxy that has been set up in this cluster. For this, you will need to reconfigure your DNS server as suggested below to map source storage's DNS name **source.example.com** to this storage cluster's IP address(es). After that, all backup agents must cache the new IP address(es), which may take about a day.

```
Suggested DNS configuration
                                                                                 Copy to clipboard
 $TTL 1h
 0
    IN SOA ns1.myhoster.com. source.example.com (
    2018042013 ; serial
    1h ; refresh
    30m ; retry
    7d ; expiration
    1h) ; minimum
 ; primary name server
 NS ns1.myhoster.com.
; secondary name server
 NS ns2.myhoster.com.
 A 10.248.64.99
X Cancel migration and reset settings
                                                                              START MIGRATION
```

Depending on data size, migration may take as long as several days.

5.4.7.2 Migrating Backups from Acronis Storage Gateway 1.6 and 1.7 (NFS)

- 1. Disable the firewall or explicitly open TCP port 44446 on the source Acronis Storage Gateway.
 - To disable the firewall, run

systemctl stop firewalld

- To open TCP port 44446 in the firewall, do the following:
- 1.1. Find out the zone where port 44445 is open:

firewall-cmd --list-all-zones | grep active
mix_eth0 (active)

1.2. Add the required port to the same zone:

#	firewall-cmd	zone=mix_eth0	permanent	add-port=44446/tcp
#	firewall-cmd	reload		

- In the admin panel of the backup gateway node, proceed to STORAGE SERVICES > Backup storage > NODES, select the node(s) to run the gateway services on, and click Migrate.
- 3. Select the source storage version and click **NEXT**.
- 4. Specify the connection details for the source storage and click **NEXT**.

<	Connect to source (2/7)
Spec Cloue	ify the address of the source storage (as registered in Backup d) and the root password to that machine.
Host	name or IP address
Pass	word
Make conn	e sure the SSH service is running and port 22 is open for incoming ections.

5. Provide the credentials for the cloud management portal of the Acronis Backup Cloud installation that the source storage is registered in and click **NEXT**.

6. If the source storage is registered in Acronis Backup Cloud under an IP address, you will see the DNS configuration screen. On it, click **RE-REGISTER WITH DNS** and specify the source storage DNS name (recommended, see above). Or, if you want to keep using the IP address, click **PROCEED WITH IP**.

If you specified a DNS name, configure your DNS server according to the suggested example.

Important: Each time you change nodes in the Backup Gateway cluster, adjust the DNS settings accordingly.

- 7. Choose a storage type to create a gateway to one of the destinations:
 - local storage cluster
 - external NFS
 - public cloud
- 8. Make sure the correct network interface is selected in the drop-down list. Click **NEXT**.

If necessary, click the cogwheel icon and configure node's network interfaces on the **Network Configuration** screen.

Configure network		
node001	¢	
ABGW private	ABGW public	
eth1 - 10.37.130.250 🗸	eth0 - 10.94.17.81	
Image: node002ABGW privateeth1 - 10.37.130.28	ABGW public eth0 - 10.94.18.146	
	BACK	

- 9. Configure the destination storage backend:
 - For a storage cluster, select the desired tier, failure domain, and redundancy mode.
 - For NFS, specify a hostname or IP address, an export name and path, and choose the NFS version.

< Volume parameters

NFS hostname or IP		
nfs.example.com		
Export name		
/path/to/volume		
 NFS3 (no clustering) 		
O NF54		
	ВАСК	NEXT

• For a public cloud, select a public cloud provider, specify credentials, and the name of the folder (bucket, container).

Public cloud parameters

Use a separate object container for each Backup Gateway cluster.

<

Select the object storage type	
Amazon S3	~
Region	
us-east-1	~
Access key ID	
Secret Access key	
Bucket	
acronis-us-west-gateway-files	

Click **NEXT**.

10. Review the source and destination storages and click **START MIGRATION**.

Depending on data size, migration may take as long as several days.

5.4.8 Managing Geo-Replication for Backup Gateway

Acronis Cyber Infrastructure allows you to enable Backup Gateway replication between two geographically distributed datacenters registered in the Cloud Management Panel. It provides backup data protection against the primary datacenter failure. You can enable geo-replication for Backup Gateways that are set up on different storage backends: a local storage cluster, NFS share, or public cloud.

For successful geo-replication, the following requirements must be met:

- Two storage clusters with Backup Gateways are deployed.
- All storage clusters are updated to the latest version.
- All storage clusters are registered in the Cloud Management Panel.
- All storage clusters can reach each other via domain names on TCP port 44445.

5.4.8.1 Enabling Geo-Replication

To set up geo-replication between two storage clusters, primary and secondary, do the following:

1. On the cluster that will be configured as secondary, click the copy icon next to the **DNS name** and **UID** fields to copy its DNS name and UID to clipboard.

OVERVIEW	NODES	GEO-REPLICATION	
Geo-rep	lication		Configure replication
DNS name		slave.example.com	E
UID		e4519719924e23a7ca433e8ad0	a4584e-1560428430 🗈

- 2. On the cluster that will be configured as primary, click **Configure replication** and do the following in the **Configure replication** window:
 - 2.1. Paste the DNS name and UID of the secondary cluster into the corresponding fields.
 - 2.2. Click **Download configuration file** to download the configuration file of the primary cluster to your local server.

2.3. Click **Done**.

Configure replication	×
Select configuration type	
Primary cluster	
DNS name of the secondary cluster slave.example.com	
Secondary cluster UID e4519719924e23a7ca433e8ad0a4584e-1560428430	
Make sure the following prerequisites are met:	
1. Two storage clusters with Backup Gateways are deployed.	
2. All storage clusters are updated to the latest version.	
3. All storage clusters are registered in the Cloud Management Panel.	
4. All storage clusters can reach each other via domain names on TCP port 44445.	
Cancel	one

The primary cluster is now configured and ready to be connected to the secondary one, which needs to be configured next.

- 3. On the secondary cluster, click **Configure replication** and do the following in the **Configure replication** window:
 - 3.1. Select the **Secondary cluster** configuration type.
 - 3.2. Upload the the configuration file of the primary cluster from your local server.

3.3. Click **Done**.

Configure replication	×
Select configuration type Primary cluster Secondary cluster 	
Specify the configuration file obtained from the primary cluster.	
Configuration file downloaded from the primary cluster dc-configs.tar.bz2	Browse
 Make sure the following prerequisites are met: 1. Two storage clusters with Backup Gateways are deployed. 2. All storage clusters are updated to the latest version. 3. All storage clusters are registered in the Cloud Management Panel. 4. All storage clusters can reach each other via domain names on TCP port 44445. 	
Cancel	Done

The secondary cluster is now also configured and ready to be connected to the primary one.

If after configuring the secondary cluster, you need to change the configuration of the primary cluster for some reason, download the new configuration and upload it to the secondary cluster by clicking the upload icon next to the **Configuration file** field. Before doing so, make sure the primary cluster UID has not been changed.

4. Back on the primary cluster, click **Connect** to enable replication between the two datacenters.

Primary cluster ME	🛃 Download configuration file
DNS name	master.example.com
UID	751b6cd668a7e7b511459c87b918e005-1560428660
📰 Secondary cluster	
DNS name	slave.example.com
UID	e4519719924e23a7ca433e8ad0a4584e-1560428430
 Upload the configuration fil 	e to the secondary datacenter. Confirm by clicking "Connect".
	Cancel

5.4.8.2 Performing a Failover

If the primary cluster becomes unavailable, you can perform a manual failover by promoting the secondary cluster to primary. This operation will switch the configuration of the secondary cluster, including its DNS name, to the configuration of the primary one. Failover of the primary cluster can be performed in the following cases:

- The current primary cluster is completely non-operational and isolated from the Internet and any backup agents.
- Backup agents are unable to communicate with the current primary cluster.
- The DNS name of the primary cluster has been reconfigured to its IP addresses.

Warning: Promoting the secondary cluster to primary is an irreversible operation that will invalidate all data on the primary cluster. Use it only in case of emergency.

To perform a failover, click **Promote to primary** on the secondary cluster and then **Failover** in the confirmation window.

📰 Secondary cluster м	🖽 Promote to primary cluster
DNS name	slave.example.com
UID	e4519719924e23a7ca433e8ad0a4584e-1560415462
Configuration file	Uploaded June 13, 2019 9:07 AM 1

If the current primary cluster is still operational, forcibly release all its nodes from Backup Gateway first and then perform a failover.

5.4.8.3 Updating the Geo-replication Configuration

Once a year you need to renew the Backup Gateway certificate. The certificate update changes the cluster configuration, which in turn requires updating the geo-replication configuration. Do the following:

- 1. On the primary cluster, update the certificate as described in *Updating certificate for Backup Gateway* (page 206)
- 2. On the primary cluster, click **Download configuration file** to download its new configuration to your local server.
- 3. On the secondary cluster, click the upload icon next to the **Configuration file** field to upload the new configuration to the secondary cluster.

5.4.8.4 Disabling Geo-replication

To disable geo-replication, click **Disable replication** on the primary cluster. To remove the secondary cluster from the geo-replication configuration, gracefully release all its nodes from Backup Gateway (see *Releasing Nodes from Backup Gateway* (page 226)).

Primary cluster ME	🛃 Download configuration file
DNS name	master.example.com
UID	751b6cd668a7e7b511459c87b918e005-1560428660
📰 Secondary cluster	O Disable replication
DNS name	slave.example.com
UID	e4519719924e23a7ca433e8ad0a4584e-1560428430

5.4.9 Monitoring Backup Gateway

After you create a Backup Gateway, you can monitor it on the **STORAGE SERVICES** > **Backup storage** > **OVERVIEW** screen. The charts show the following information:

- the performance of Backup Gateway services
- the geo-replication speed and backlog (the amount of data waiting to be replicated)
- object storage speed and backlog (the amount of data waiting to be uploaded to public cloud)
- migration speed and backlog (the amount of data waiting to be migrated)
- how many files are left in migration queue

If you migrate backups from Acronis Storage 1.5 or 1.7, migration backlog will be larger than the amount of data on the source storage. The reason is that Acronis Storage versions prior to 2.x use the old backup (FES) protocol that sends more data over network. The difference between source data size and backlog also very much depends on the retention policy utilized by the backup solution. Despite this, the resulting space occupied by migrated data on the destination will be similar to that on the source.

If backlogs do not decrease over time, it means the data cannot be replicated, migrated, or uploaded fast enough. The reason may be insufficient network transfer speed, and you may need to check or upgrade your network.

Acronis Backup Gateway



5.4.9.1 Advanced Backup Gateway Monitoring via Grafana

For advanced monitoring of the Backup Gateway cluster, go to the **MONITORING** > **Dashboard** screen and click **Grafana dashboard**. A separate browser tab will open with preconfigured Grafana dashboards, two of which are dedicated to Acronis Backup Gateway. To see a detailed description for each chart, click the **i** icon on its left corner.

On the **Acronis Backup Gateway** dashboard, you need to pay attention to the following charts:

• **Availability**. Any time period during which the gateways have not been available will be highlighted in red. In this case, you will need to look into logs on the nodes with the failed service and report a problem. To see the Backup Gateway log, use the following command:

zstdcat /var/log/vstorage/abgw.log.zst

- **Migration/Replication throughput**. The migration chart should be displayed during migration or if the cluster serves as master in a geo-replication configuration. The replication chart should mirror the ingress bandwidth chart.
- **Migration/replication backlog**. The migration chart should decrease over time. The replication chart should be near zero, high values are indicative of network issues.
- **Rate limiting/ingress throttling**. If the chart is not empty, it means the underlying storage lacks free space and the Backup Gateway is throttling user requests to slow down the data flow. Add more storage space to the cluster to solve the issue. For more information, see https://kb.acronis.com/content/62823.
- **New client connections**. A high rate of failed connections due to SSL certificate verification problems on the chart means that clients uploaded an invalid certificate chain.
- **IO watchdog timeouts**. If the chart is not empty, it means the underlying storage is not healthy and cannot deliver the required performance.



To see the charts for a particular client request, file, and I/O operation, choose them from the drop-down menus above. A high rate of failed requests or operations and high latencies on these charts indicate that

the Backup Gateway experiences issues that need to be reported. For example, you can check charts for the "Append" request:

- The **Append rate** chart displays the backup data flow from Backup agents to the storage in operations per second (one operation equals one big block of backup data; blocks can be of various size).
- The **Append latency** chart shows the time spent on processing requests and should average several tens of milliseconds with peak values below one second.



The **Acronis Backup Gateway Details** dashboard is intended for low-level troubleshooting by the support team. To monitor a particular node, client request, file, and I/O operation, choose them from the drop-down menus above. On the dashboard, you can make sure the **Event loop inactivity** chart is empty. Otherwise, the Backup Gateway is not healthy on this node and the issue needs to be reported.



5.4.10 Releasing Nodes from Backup Gateway

Backup Gateway is meant to provide access to one specific storage backend. If you need to switch the backend, e.g., from a public cloud to a local storage cluster or one public cloud bucket to another, you need to delete the Backup Gateway by releasing all its nodes and create a new one.

To release one or more nodes from the Backup Gateway cluster, select them on the **STORAGE SERVICES** > **Backup storage** > **NODES** screen and click **Release**. The Backup Gateway cluster will remain operational until there is at least one node in it.

When the Backup Gateway is deleted, it is also unregistered from your Acronis backup software, which loses access to the storage backend.

Do the following to release the last node in the gateway:

- 1. On the **STORAGE SERVICES** > **Backup storage** > **NODES** screen, select the node and click **Release**.
- 2. On the **Unregister from backup software** panel, choose one of the following:
 - **Graceful release** (recommended, see note below). Releases the node, deletes the Backup Gateway and unregisters it from your Acronis backup software.
 - **Force release**. Releases the node, deletes the Backup Gateway but does not unregister it from your Acronis backup software.

Important: Choose this option only if you are sure that the gateway has already been unregistered from your Acronis backup software. Otherwise, you will need to register a new gateway in your Acronis backup software and for that you will need to delete and recreate not just the Backup Gateway but also the entire storage cluster.

× Unregister from backup software
 Graceful release () Forced release ()
Unregister this storage from backup software.
Administrator account
admin
Enter the credentials of a partner account in the cloud or of an organization administrator on the local management server.
NEXT

3. Specify the credentials of your administrator account in your Acronis backup software and click **NEXT**. In case the release is forced, simply click **NEXT**.

CHAPTER 6

Managing General Settings

This chapter describes how you can configure Acronis Cyber Infrastructure settings.

6.1 Managing Licenses

Acronis Cyber Infrastructure comes with a trial license that allows you to evaluate its features. The trial license has no expiration date but limits storage capacity to 1TB.

Acronis Cyber Infrastructure supports the following licensing models for production environments:

- License key. Implementing the provisioning model, keys are time-limited (subscription) or perpetual and grant a certain storage capacity. If a commercial license is already installed, a key augments its expiration date or storage limit.
- Services provider license agreement (SPLA). SPLA implements the pay-as-you-go model: it grants unlimited storage capacity and customers are charged for the actual usage of these resources. With SPLA, Acronis Cyber Infrastructure automatically sends reports to Acronis Cyber Cloud once every four hours. If no reports have been received for two weeks, the license expires. For reports to reach destination, the cluster must be able to access the Acronis datacenter that has been used to enable SPLA. Make sure that TCP port 443 is open.

Note: SPLA license is valid for Cloud Partners. If SPLA is enabled, you can connect Backup Gateway only to Acronis Backup Cloud and not to Acronis Backup 12.5 or Acronis Backup Advanced 12.5. To connect Backup Gateway to these products, you will need to use license keys. Furthermore, Acronis Backup Gateway usage is not counted in SPLA in Acronis Cyber Infrastructure. SPLA only counts

universal usage that is not related to backup. Backup usage is shown in the Acronis Backup Cloud section of Acronis Cyber Cloud.

You can switch the licensing model at any time:

- Switching from a license key to SPLA terminates the key even if it has not yet expired. Terminated keys cannot be used anymore.
- Switching from SPLA to a license key changes the licensing model to subscription or perpetual. After doing so, ask your service provider to terminate your SPLA by either disabling the Storage application for your account or deleting the account.

Important: If a license expires, all write operations to the storage cluster stop until a valid license is installed.

6.1.1 Installing License Keys

To install a license key, do the following:

- 1. If you are switching from SPLA, ask your service provider to terminate the agreement by either disabling the Acronis Cyber Infrastructure application for your account or deleting the account.
- 2. On the **SETTINGS** > **Licenses** screen, click **Upgrade** and **Register key**.

ter product k	еу
XXXXXXX-XXX XXXXXXX-XXX	xxxxx-xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx

3. In the **Register license key** window, paste the license key and click **REGISTER**.

41.52 MB of 1 TB 🗦 5 TB space used
ACTIVATE CANCEL

4. Back on the **Licenses** screen, click **Activate** if you are activating from a trial or choose one of the following:

41.52 MB of 5 TB space used
Choose an activation option:
O Upgrade
 Prolong
ACTIVATE CANCEL

- **Upgrade**, to add storage capacity to the active license.
- **Prolong**, to prolong the soon-to-be-expired license.

And click **Activate**.

The expiration date or storage capacity will change according to what the key grants.

6.1.2 Installing SPLA Licenses

To install a SPLA license, do the following:

- 1. On the **SETTINGS** > **Licenses** screen, click **Upgrade** and **Use SPLA**.
- 2. In the **Use SPLA** window, select a region from the drop-down list and click **Activate**. You will be redirected to a login page of Acronis Cyber Cloud.

Note: For more information on datacenters, see https://kb.acronis.com/servicesbydc.

- 3. Log in to Acronis Cyber Cloud.
- 4. In the **Register cluster** window, accept the license agreement.
- 5. In the registration confirmation window, click **Done**.

The registered cluster will show up in Acronis Cyber Cloud. You will be able to monitor its resource usage and download reports.

6.2 Managing Add-On Services

You can deploy the following additional services for the compute cluster:

• The Kubernetes service allows you to deploy scalable and production-ready Kubernetes clusters with pre-integrated persistent storage. For more details, see *Managing Kubernetes Clusters* (page 143).

Important: To be able to deploy and work with Kubernetes clusters, make the following services accessible:

- the etcd discovery service at https://discovery.etcd.io—from all management nodes and the public network with the VM public traffic type
- the public Docker Hub repository at https://registry-1.docker.io—from the public network with the
 VM public traffic type
- the compute API—from the public network with the VM public traffic type
- the Kubernetes API at the public or floating IP address of the Kubernetes load balancer or master VM on port 6443—from all management nodes

If the **Compute API** traffic type is added to a private network that is inaccessible directly from the network with the **VM public** traffic type but exposed to public networks via NAT and available publicly via the DNS name, you need to set the DNS name for the compute API as described in Setting a DNS Name for the Compute API.

Note: Installing Kubernetes automatically installs the load balancer service as well.

- The load balancer service enables workload scaling and improves application availability and security. To learn how to manage load balancers, see *Managing Load Balancers* (page 129).
- The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. For more information, refer to Using Metering for Compute Resources.

Note: The metering service will only take into account compute objects created after it has been enabled.

On the **SETTINGS** > **Add-on services** screen, click **Install** for the service you want to deploy. One service can be installed at a time.

Note: The load balancer service is installed along with the Kubernetes service as a dependency.

🕸 Kubernetes service	Install
The Kubernetes service allows you to deploy scalable and production-ready Kubernetes clusters with pre-integrated persistent stor	rage.
品 Load balancer service	Install
The load balancer service enables workload scaling and improves application availability and security.	
Billing metering service	Install
The billing metering service collects, stores, and provides usage metrics for resources consumed by end users in their projects. The can be accessed via the Gnocchi API.	meters

In the current version of Acronis Cyber Infrastructure, installed services cannot be removed.

6.3 Managing Tier Encryption

Acronis Cyber Infrastructure can encrypt data stored on disks with the AES-256 standard, so if a disk gets lost or stolen the data will be safe. Acronis Cyber Infrastructure stores disk encryption keys in cluster's metadata (MDS).

Encryption can be enabled or disabled only for the newly created chunk services (CS). Once tier encryption is enabled, you can decrypt disks (CSs) by manually releasing them from encrypted tiers. Correspondingly, simply enabling encryption on the disk's tier will not encrypt its data (CS). To encrypt a disk, you must assign it to an encrypted tier.

Take note of the following:

- 1. Acronis Cyber Infrastructure does not encrypt data transmitted over the internal network.
- 2. Enabled encryption slightly decreases performance.



To enable tier encryption, do the following:

- 1. Navigate to **SETTINGS** > **Advanced settings** > **DISK**.
- 2. Enable Enable AES-256 encryption for data stored on disks.
- 3. Select tiers that you want to encrypt.
- 4. Click Save.

6.4 Enabling RDMA

Warning: This feature is experimental and not recommended for production until version 4.0.

Acronis Cyber Infrastructure supports remote direct memory access (RDMA) over Converged Ethernet (RoCE), Internet Wide-area RDMA Protocol (iWARP), or InfiniBand (IB) for the storage backend network. The RDMA technology allows servers in this network to exchange data in main memory without involving their processors, cache or operating systems, thus freeing up resources and improving throughput and performance.

Your RDMA network infrastructure must be ready before you install Acronis Cyber Infrastructure.

Important: In the current version of Acronis Cyber Infrastructure, you can only enable (or disable) RDMA before creating the storage cluster.

By default, RDMA is disabled. Before enabling it, make sure that each network adapter connected to a network with the **Storage** traffic type supports RDMA.

To enable or disable RDMA, use the switcher on the **SETTINGS** > **Advanced settings** > **RDMA** tab. Changing this option may temporarily affect cluster availability.

Encryption SNMP Notifications RDMA

Acronis Cyber Infrastructure supports RDMA over Converged Ethernet (RoCE) or InfiniBand for the storage backend network. RDMA reduces I/O latency and improves cluster performance. Before enabling this feature, make sure that each network adapter that handles the "Storage" traffic type supports RDMA. This option can only be changed before creating the storage cluster.

Enable RDMA

6.4.1 Configuring InfiniBand Devices

Note: As the admin panel only shows IP states and does not show InfiniBand connection states, it may report plugged in but yet unconfigured IB devices as **UNPLUGGED**. The status will change to **OK** once you assign an IP address to such a device.

If you have an InfiniBand infrastructure, do the following before enabling RDMA:

- Assign the traffic type Storage to an empty network (without any other traffic types) on the INFRASTRUCTURE > Networks screen. Create a new network if needed by clicking Edit > Create network.
- 2. Configure each IB device on registered nodes:
 - 2.1. Open **INFRASTRUCTURE** > **Nodes** > <node> > **Network**, and select the device.
 - 2.2. Click **Configure**. In the pane that opens:
 - Assign an IP address (Manually will be selected by default).
 - Specify the gateway.
 - Check the **Connected mode** box.
 - Set MTU to 65520.

× Configure	
 Automatically (DHCP) Automatically (DHCP address or Manually 	ly)
10.0.0.1/32	
Gateway	
10.248.64.1	
Connected mode	
MTU	
65520	
05520	

Click **DONE**.

2.3. Click **Assign network**. In the pane that opens, select the network with just the **Storage** traffic type and click **DONE**.

6.4.2 Configuring RoCE and iWARP Devices

If you have a RoCE or iWARP infrastructure, do the following for each network device in it before enabling RDMA:

- 1. Open **INFRASTRUCTURE** > **Nodes** > <node> > **Network**, and select the device.
- 2. Click **Configure**. In the pane that opens, assign an IP address if the device does not have one yet and specify the gateway. Click **DONE**.
- Click Assign network. In the pane that opens, select the network with the Storage traffic type (and possibly other traffic types like Internal management, OSTOR private, and ABGW private) and click DONE.

6.5 Enabling High Availability

High availability keeps Acronis Cyber Infrastructure services operational even if the node they are located on fails. In such cases, services from a failed node are relocated to healthy nodes according to the Raft consensus algorithm. High availability is ensured by:

- Metadata redundancy. For a storage cluster to function, not all but just the majority of MDS servers must be up. By setting up multiple MDS servers in the cluster you will make sure that if an MDS server fails, other MDS servers will continue controlling the cluster.
- Data redundancy. Copies of each piece of data are stored across different storage nodes to ensure that the data is available even if some of the storage nodes are inaccessible.
- Monitoring of node health.

To achieve the complete high availability of the storage cluster and its services, we recommend that you do the following:

- 1. deploy three or more metadata servers,
- 2. enable management node HA, and
- 3. enable HA for the specific service.

Note: The required number of metadata servers is deployed automatically on recommended hardware configurations; Management node HA must be enabled manually as described in the next subsection; High availability for services is enabled by adding the minimum required number of nodes to that service's cluster.

On top of highly available metadata services and enabled management node HA, Acronis Cyber Infrastructure provides additional high availability for the following services:

- Admin panel. If the management node fails or becomes unreachable over the network, an admin panel instance on another node takes over the panel's service so it remains accessible at the same dedicated IP address. The relocation of the service can take several minutes. Admin panel HA is enabled manually along with management node HA (see *Enabling Management Node High Availability* (page 238)).
- Virtual machines. If a compute node fails or becomes unreachable over the network, virtual machines hosted on it are evacuated to other healthy compute nodes based on their free resources. The compute cluster can survive the failure of only one node. By default, high availability for virtual

machines is enabled automatically after creating the compute cluster and can be disabled manually, if required (see the *Configuring Virtual Machine High Availability* (page 98)).

- iSCSI service. If the active path to volumes exported via iSCSI fails (e.g., a storage node with active iSCSI targets fails or becomes unreachable over the network), the active path is rerouted via targets located on healthy nodes. Volumes exported via iSCSI remain accessible as long as there is at least one path to them.
- S3 service. If an S3 node fails or becomes unreachable over the network, name server and object server components hosted on it are automatically balanced and migrated between other S3 nodes. S3 gateways are not automatically migrated; their high availability is based on DNS records. You need to maintain the DNS records manually when adding or removing S3 gateways. High availability for S3 service is enabled automatically after enabling management node HA and creating an S3 cluster from three or more nodes. An S3 cluster of three nodes may lose one node and remain operational.
- Backup gateway service. If a backup gateway node fails or becomes unreachable over the network, other nodes in the backup gateway cluster continue to provide access to the chosen storage backend. Backup gateways are not automatically migrated; their high availability is based on DNS records. You need to maintain the DNS records manually when adding or removing backup gateways. High availability for backup gateway is enabled automatically after creating a backup gateway cluster from two or more nodes. Access to the storage backend remains until at least one node in the backup gateway cluster is healthy.
- NFS shares. If a storage node fails or becomes unreachable over the network, NFS volumes located on it are migrated between other NFS nodes. High availability for NFS volumes on a storage node is enabled automatically after creating an NFS cluster.

Also take note of the following:

- 1. Creating the compute cluster prevents (and replaces) the use of the management node backup and restore feature.
- 2. If nodes to be added to the compute cluster have different CPU models, consult Setting Virtual Machines CPU Model.

6.5.1 Enabling Management Node High Availability

To make your infrastructure more resilient and redundant, you can create a high availability configuration of three nodes.

Management node HA and compute cluster are tightly coupled, so changing nodes in one usually affects the other. Take note of the following:

- 1. Each node in the HA configuration must meet the requirements to the management node listed in the Installation Guide. If the compute cluster is to be created, its hardware requirements must be added as well.
- 2. If the HA configuration has been created before the compute cluster, all nodes in it will be added to the compute cluster.
- If the compute cluster has been created before HA configuration, only nodes in the compute cluster can be added to the HA configuration. For this reason, to add a node to HA configuration, add it to the compute cluster first.
- 4. If both the HA configuration and compute cluster include the same three nodes, single nodes cannot be removed from the compute cluster. In such a case, the compute cluster can be destroyed completely, but the HA configuration will remain. This is also true vice versa, the HA configuration can be deleted, but the compute cluster will continue working.

Note: The compute cluster must have at least three nodes to allow self-service users to enable high availability for Kubernetes master nodes.

To enable high availability for the management node and admin panel, do the following:

- 1. Make sure that each node is connected to a network with the **Admin panel** and **Internal management** traffic types.
- 2. On the **SETTINGS** > **Management node** screen, open the **MANAGEMENT HIGH AVAILABILITY** tab.



3. Select three nodes and click **Create HA**. The management node is automatically selected.

4. On Configure network, check that correct network interfaces are selected on each node. Otherwise, click the cogwheel icon for a node and assign networks with the **Internal management** and **Admin** panel traffic types to its network interfaces. Click PROCEED.

× Configure netwo	rk
node001	0
Management	Admin panel
eth1 - 10.37.130.250	br-eth0 - 10.94.17.81
📑 node002	٥
Management	Admin panel
eth1 - 10.37.130.28	br-eth0 - 10.94.18.146
node003	0
Management	Admin panel
eth1 - 10.37.130.45	br-eth0 - 10.94.18.147
PROCEED	

5. Next, on **Configure network**, provide one or more unique static IP addresses for the highly available admin panel, compute API endpoint, and interservice messaging. Click **DONE**.

<	Configure network
Assi	gn unique dedicated virtual IP addresses to these services:
	 Admin panel (public access to this web UI) Compute API (public access to compute APIs) Internal management (private interservice messaging)
In a migi serv	high availability event, virtual IP addresses will automatically rate to a healthy node in the high availability cluster to keep rices accessible.
Virtu Com	ual IP address for npute API, Admin panel
10	.94.129.79
T Virtu Inte	The IP address must belong to the network Public (10.94.0.0/16) ual IP address for rnal management
10	.37.130.100
9 T	The IP address must belong to the network Private (10.37.130.0/24)
	DONE

Once the high availability of the management node is enabled, you can log in to the admin panel at the specified static IP address (on the same port 8888).

As management node HA must include exactly three nodes at all times, removing a node from the HA configuration is not possible without adding another one at the same time. For example, to remove a failed node from the HA configuration, you can replace it with a healthy one as follows:

 On the SETTINGS > Management node > MANAGEMENT HIGH AVAILABILITY tab, select one or two nodes that you wish to remove from the HA configuration and one or two available nodes that will be added into the HA configuration instead and click **Replace**.

MANAGEMENT HIGH	AVAILABILITY	SSL ACCESS	BACKUP			
INCLUDED IN HA CON	FIGURATION	3				Replace
node001		node002		node003	~	
		88				
AVAILABLE NODES	1			Hide av	vailable nodes \land	
node004	~					

 On Configure network, check that correct network interfaces are selected on each node to be added.
 Otherwise, click the cogwheel icon for a node and assign networks with the Internal management and Admin panel traffic types to its network interfaces. Click PROCEED.

To remove nodes from the HA setup, click **Destroy HA**.

6.6 Accessing the Admin Panel via SSL

When configuring various Acronis Cyber Infrastructure features, you may need to enter sensitive information like credentials for user and e-mail accounts, S3 services, and such. The system uses a pre-generated self-signed certificate by default, and you may want to upload one issued by a trusted CA instead.

Note the following before you proceed:

- If you want to upload an SSL certificate before creating the HA cluster, you will need one issued for admin panel's current IP address. If you later create the HA cluster, the admin panel will move to the chosen virtual IP address, and you will need another SSL certificate issued for that address.
- 2. When you create or destroy the HA cluster, the current certificate is overwritten by a self-signed one generated by the system. You will need to re-upload your certificate and key from a trusted CA after completing either operation.
- 3. If you acquired an SSL certificate from an intermediate certificate authority (CA), you should have an end-user certificate along with a CA bundle that contains the root and intermediate certificates. To be able to use these certificates, you need to merge them into a chain first. A certificate chain includes the

end-user certificate, the certificates of intermediate CAs, and the certificate of a trusted root CA. In this case, an SSL certificate can only be trusted if every certificate in the chain is properly issued and valid.

For example, if you have an end-user certificate, two intermediate CA certificates, and a root CA certificate, create a new certificate file and add all certificates to it in the following order:

```
# End-user certificate issued by the intermediate CA 1
----BEGIN CERTIFICATE-----
MIICiDCCAg2gAwIBAgIQNfwmXNmET8k9Jj1X<...>
----END CERTIFICATE-----
# Intermediate CA 1 certificate issued by the intermediate CA 2
  ---BEGIN CERTIFICATE----
MIIEIDCCAwigAwIBAgIQNE7VVyDV7exJ90N9<...>
----END CERTIFICATE-----
# Intermediate CA 2 certificate issued by the root CA
----BEGIN CERTIFICATE----
MIIC8jCCAdqgAwIBAgICZngwDQYJKoZIhvcN<...>
----END CERTIFICATE-----
# Root CA certificate
----BEGIN CERTIFICATE-----
MIIDODCCAiCgAwIBAgIGIAYFFnACMA0GCSqG<...>
   --END CERTIFICATE----
```

To upload an SSL certificate, do the following:

- 1. On the **SETTINGS > Management node > SSL ACCESS** tab, click **Upload**.
- 2. Upload an SSL certificate issued for admin panel's current IP address.
- 3. Upload the private key. This option shows after uploading a valid certificate.
- 4. Click SAVE.

The uploaded certificate will be added to the configuration of the web server that hosts the admin panel and you will be able to access it over HTTPS.

You can also generate a new self-signed certificate instead of the one used by default. However, it will not be trusted and you will have to manually accept it in your browser.

6.7 Backing Up and Restoring Management Database

Node information, statistics, and configurations are stored in a database on the management node (the one with the admin panel). Backups of this database are created automatically via a daily cron job that starts at
3:00 a.m. If the management node is not highly available, restoring such a backup recovers the node in case of failure or database corruption. Backup files are stored in the /mnt/vstorage/webcp/backup/ directory. The retention policy for management node backups is as follows:

- All backups created within the last day are kept.
- From backups created within the last 7 days, the newest for each day is kept.
- From backups created within the last 7-14 days, the oldest one is kept.
- From backups created within the last 14-45 days, the oldest for each week is kept.
- Backups older than 45 days are deleted.

Important: Database backups cannot be restored if management node high availability is enabled.

To back up the database manually, open the **SETTINGS** > **Management node** > **BACKUP** screen and click **BACKUP NOW**.

MAI	NAGEMENT NODE HA CONFIGURATION	SS	L ACCESS	BACKUP
The data <u>How</u>	system stores node information, statistics base that is backed up automatically. <u>/ to recover?</u>	, and	configuration	n in a
ວ 1	Last backup Dec 03, 2018, 3:00 AM Backup location	0	Backup per Daily	iod
,	/mnt/vstorage/webcp/backup/			

Once backup is completed, the **Last backup** date will be refreshed.

Warning: Do not rename the backup file! Otherwise you will not be able to restore the management database from it.

6.7.1 Restoring Management Database from Backup

You can restore a management node database from backup on the same management node or any node in the storage cluster.

Warning: The vstorage-ui-backend service must be running on a single node in the storage cluster. If the management node database is restored on a different node, the previous management node must be re-deployed.

Run the following script on the node where the MN database will be restored:

```
# /usr/libexec/vstorage-ui-backend/bin/restore-management-node.sh \
-x <public_net_iface> -i <private_net_iface> \
-f /mnt/vstorage/webcp/backup/<backup_file>
```

Where <public_net_iface> and <private_net_iface> are interfaces assigned the public and private networks, respectively. If required, you can specify the same network interface in both parameters. If the -f option is omitted, the management node database will be restored from the latest backup.

To access the admin panel, use the public IP address of the node with the restored MN database.

6.7.1.1 Restoring with Existing Compute Cluster

If you have the compute cluster deployed, the management node database must be restored only on one of the compute nodes.

When restoring the MN database with the compute cluster, note the following limitations:

- If you create compute objects after the backup, they will be lost.
- If you modify or delete compute objects after the backup, they will be restored as follows:
 - Compute objects used as configurations (flavors, storage policies, virtual networks, SSH keys) will be fully restored.

• All other compute objects (VMs, volumes, images, etc.) will be partially restored. They will be shown in the admin panel but unusable. You will only be able to remove them from the admin panel.

To restore the MN database on a compute node, run the restoration script with the -n option:

/usr/libexec/vstorage-ui-backend/bin/restore-management-node.sh \
-x <public_net_iface> -i <private_net_iface> -n

where

- The <public_net_iface> and <private_net_iface> are interfaces assigned the public and private networks, respectively.
- The -f option is omitted to restore the MN database from the latest backup.
- The -n option denotes that the compute cluster will be reconfigured to use another management node.

If you restore the MN database on the same node, omit the -n option.

After restoration, virtual machines that resided on the failed management node will be unmanageable in the admin panel and can only be deleted. However, you can rescue them using the vinfra tool as follows:

- VMs with disabled high availability will appear in the "Active" state: evacuate them using the vinfra service compute server evacuate command;
- VMs with enabled high availability will appear in the "Rebuild" state: first reset their state with the vinfra service compute server reset-state command and then evacuate them running vinfra service compute server evacuate.

6.8 Managing Domains, Users, and Projects

Acronis Cyber Infrastructure uses the administrative hierarchy of domains and projects with Role-Based Access Control (RBAC) to manage virtual objects of the compute cluster, such as virtual machines, volumes, private networks, and other. A domain is an isolated container of projects and users with assigned roles. Each project and user can only belong to one domain. A project is an isolated container of virtual objects with defined limits for virtual resources, such as vCPU, RAM, storage and floating IP addresses, and assigned users. A role is global and defines all possible tasks the user may perform at the level of the entire cluster, a specific domain, or project:

- within the cluster, you can perform system administration tasks;
- within a domain, you can create and manage user accounts and assign them to projects;

• within a project, you can create and manage virtual objects.

Such an implementation provides an administrative environment with own users and virtual objects and ensures their isolation from other users and virtual objects.

6.8.1 Managing Domains

During the primary node deployment, the unique **Default** domain is created along with the default user account and project. Only within this domain you can create system administrators with access to the admin panel. The default domain cannot be deleted.

To create a new domain, do the following:

- 1. On the SETTINGS > Projects and users screen, click Create domain.
- 2. In the **Create domain** window, specify the domain name and, optionally, description.

Name			
domain1			
Description (option	al)		
A custom domain	n		

3. Click Create.

Enabling and disabling a domain means allowing and prohibiting access to it, respectively, in the self-service panel.

To edit, disable/enable, or delete a domain, click the ellipsis button next to it and select the desired action. A domain cannot be deleted if it has projects.

6.8.2 Managing Domain Users

A user can be assigned one of the following roles:

- A system administrator has access to the admin panel and can perform system administration tasks depending on assigned permissions. It is the only role that enables creating projects and defining quotas for them. Additionally, a system administrator with domain permissions can manage virtual objects in all projects within the **Default** domain as well as project and user assignment in the self-service panel.
- A domain administrator can manage virtual objects in all projects within the assigned domain as well as project and user assignment in the self-service panel. A domain administrator can only be assigned to one domain.
- A project member acts as a project administrator in a specific domain in the self-service panel. A project member can be assigned to different projects and can manage virtual objects in them.

Within the **Default** domain, the default administrator account is created with the unique **Superuser** permission. The user name for this account is admin and the password is specified during the primary node deployment. This account cannot be deleted or disabled and its permissions cannot be changed. Other than that, admin does not differ from a user assigned the **System administrator** role.

To view and edit existing users of a domain or create new ones, click the desired domain and go the **DOMAIN USERS** tab. Creating a user account differs slightly depending on the user role and is described in sections below.

To edit the user credentials or permissions, click the ellipsis button next to the user and then click **Edit**. Any system administrator can also change their password by clicking the user icon in the top right corner of the admin panel and then clicking **Change password**.

Enabling and disabling a user account means allowing and prohibiting user login, respectively.

To enable/disable or remove a user, click the corresponding ellipsis button and select the desired action.

6.8.2.1 Creating System Administrators

Note: System administrators can be created only within the **Default** domain.

To create a system administrator, do as follows:

- 1. On the **SETTINGS** > **Projects and users** screen, click the **Default** domain.
- 2. Go to the **DOMAIN USERS** tab and click **Create user**.
- 3. In the **Create user** window, specify the user name, password, and, if required, a user e-mail address and description. The user name must be unique within a domain.
- 4. Select the **System administrator** role from the **Role** drop-down menu.
- 5. Choose permissions to be granted to the user account from the **System permission set** section:
 - Full (System administrator): has all permissions and can perform all management operations, including projects creation and other users management;
 - Compute: can create and manage the compute cluster;
 - ISCSI: can create and manage iSCSI targets, LUNs, and CHAP users;
 - S3: can create and manage the S3 cluster;
 - ABGW: can create and manage the Backup Gateway cluster;
 - NFS: can create and manage NFS shares and exports;
 - Cluster: can create the storage cluster, join nodes to it, and manage (assign and release) disks;
 - Network: can modify networks and traffic types;
 - Update: can install updates;
 - SSH: can add and remove SSH keys for cluster nodes access;
 - None (Viewer): can monitor cluster performance and parameters but cannot change any settings.
- 6. Optionally, enable the **Domain permissions set** to be able to manage virtual objects in all projects within the **Default** domain and other users in the self-service panel.
- 7. Click Create.

Create user	×
Login user1	Email (optional) user@example.com
Password	
Description (optional)	
Role System administrator	d services.
System permission set	Full
Can perform all management operations.	
Domain permission set	Full
Cannot manage the domain in the self-ser	vice panel.
	Cancel Create

6.8.2.2 Creating Domain Administrators

To create a domain administrator, do as follows:

- 1. On the **SETTINGS** > **Projects and users** screen, click a domain for which the administrator will be created.
- 2. Go to the **DOMAIN USERS** tab and click **Create user**.

X

- 3. In the **Create user** window, specify the user name, password, and, if required, a user e-mail address and description. The user name must be unique within a domain.
- 4. Select the **Domain administrator** role from the **Role** drop-down menu.
- 5. Optionally, select the **Image uploading** checkbox. The state of this permission will be inherited by users created by this domain administrator.
- 6. Click **Create**.

Create user

user1	user1@e	xample.com
Password		
	`	
Description (optional)		
Role		
Domain administrator	~	
n create and manage projects	nd services in the assigned	l domain.
Image uploading		

6.8.2.3 Creating Project Members

To create a project member, do as follows:

- 1. On the **SETTINGS** > **Projects and users** screen, click a domain within which the user will be created.
- 2. Go to the **DOMAIN USERS** tab and click **Create user**.
- 3. In the **Create user** window, specify the user name, password, and, if required, a user e-mail address and description. The user name must be unique within a domain.
- 4. Select the **Project member** role from the **Role** drop-down menu.

- 5. Optionally, select the **Image uploading** checkbox. If this option is disabled, this user will not be able to upload images.
- 6. Optionally, click **Assign** and choose a project this user will be assigned to.
- 7. Click **Create**.

Create user			×
Login user1		Email (optional) user1@example.com	
Password	•		
Description (optional)			
Role Project member	~		
Can create and manage services in assign	ned proje	cts.	
Image uploading (i)			
Assign to projects			+ Assign
project1			×
		Cancel	Create

6.8.3 Managing Projects

The **Default** domain has the default **admin** project, which is a bootstrap project for initializing the compute cloud. It cannot be deleted or renamed.

To create a new project, do the following:

1. On the **SETTINGS** > **Projects and users** screen, click a domain within which the project will be created.

- 2. On the **PROJECTS** tab, click **Create project**.
- 3. In the **Create project** window, specify the project name and, optionally, description. The project name must be unique within a domain.
- 4. Optionally, deselect the **Enabled** checkbox to disable the created project.
- 5. Define quotas for virtual resources that will be available inside the project. To specify a certain value for a resource, deselect the **Unlimited** checkbox next to it first.

If you have not yet deployed the compute cluster, you are not able to set project's quotas. Create the compute cluster as described in *Creating the Compute Cluster* (page 69) and return to defining project's quotas as described in *Editing Quotas for Projects* (page 258).

Note: As quotas can exceed the existing virtual resources and virtual resources are not reserved for each project, a system administrator needs to ensure the compute cluster has enough virtual resources for all projects in all domains.

6. Click **Create**.

Creat	e project		×
Name projec	x1	Enabled	
Descri	iption (optional)		
Spec	ify compute quotas		
0	vCPUs	Unlimited	24
000	RAM, GiB	Unlimited	24
8	Storage policy		
	🗹 default, GiB	Unlimited	2048
®	Floating IPs	Unlimited	32
ቆ	Load balancers	Unlimited	8
麥	Kubernetes clusters	Max. 20	8
747		C	ancel

Note: The default storage policy must be shared with projects that will use the Kubernetes-as-a-service feature.

Once the project is created, you can open its panel to view its properties on the **Properties** tab, list its members on the **Members** tab, and monitor its resource consumption on the **Quotas** tab.

project1		×
Edit + Assign members	🕒 Edit quotas 🛛 🛇 Disable	Ū Delete
Properties	Members	Quotas
Details		
Name	project1	
Description	A custom project	
State	Enabled	
Project ID	be9f5fb9e12c4cc4970dcaa	6ec1584ac

Enabling and disabling a project means allowing and prohibiting access to it, respectively, in the self-service panel.

To edit, enable/disable, or delete a project, click the ellipsis button next to it and select the desired action. A project cannot be deleted if it has virtual objects.

6.8.3.1 Assigning Members to Projects

You can manage project members assignment either on the **PROJECTS** tab or **DOMAIN USERS** tab.

To assign a user to a project, do one of the following:

- Within the domain, open the **PROJECTS** tab:
 - 1. Click the project to which you want to assign users.
 - 2. On the project panel, click **Assign members**.
 - 3. In the **Assign members** window, choose one or multiple users to assign to the project. Optionally, click **Create and assign** to create a new project member in a new window. Only user accounts with the **Project member** role are displayed.

4. Click Assign.

Assi	ign membe	ers		×		
Select	Select users to assign as members to the project "project1".					
Sear	ch	Q	+ Creat	te and assign		
2	Login 🛧		Email			
	🙎 user1		user1@example.com			
			Cancel	Assign		

- Within the domain, open the **DOMAIN USERS** tab:
 - 1. Click the user account with the **Project member** role whom you want to assign to the project.
 - 2. On the user panel, click **Assign to project**.
 - 3. On the **Assign user to projects** window, select one or multiple projects and click **Assign**.

Assi	Assign user to projects		×	
Select	projects to assign to the user	"user1".		
Sear	ch	Q		
	Name 🕇	De	escription	
	project1	A	custom project	
			Cancel	Assign

You can monitor user assignment to projects either on the **Members** tab of the project panel or on the **Projects** tab on the user panel.

To unassign a user from a project, do one of the following:

- Within the domain, open the **PROJECTS** tab:
 - 1. Click the project from which you want to unassign users.
 - 2. On the project panel, open the **Members** tab.
 - 3. Click the cross icon next to a user you want to unassign.

project1		>
Edit + Assign members	🕒 Edit quotas 🛛 🛇 Disable	n Delete
Properties	Members (1)	Quotas
Search Q		
Login 🕆	Email	

- Within the domain, open the **DOMAIN USERS** tab:
 - 1. Click the user whom you want to unassign from the project.
 - 2. On the user panel, open the **Projects** tab.
 - 3. Click the cross icon next to the project from which you want to unassign the user.

user1	×
🖉 Edit 🕂 Assign to project 🚫 Disable	T Delete
Properties	Projects (1)
Search Q	
Name 个	Description
project1	A custom project

6.8.3.2 Editing Quotas for Projects

To change resource quotas for a project, do the following:

- 1. Click the project for which you want to edit quotas.
- 2. On the project panel, click **Edit quotas**.
- 3. In the **Edit quotas** window, specify new values for the desired virtual resources.

Note: The default storage policy must be shared with projects that will use the Kubernetes-as-a-service feature.

4. Click **Save** to apply changes.

Х

Edit quotas

Specify compute quotas		
VCPUs	Unlimited	24
📖 RAM, GiB	Unlimited	48
Storage policy		
default, GiB	Unlimited	2048
Ploating IPs	Unlimited	20
Load balancers	Unlimited	10
Kubernetes clusters	Max. 20	10
Placements		
placement1	Unlimited	20
	٢	Cancel Save

6.9 Managing Updates

Acronis Cyber Infrastructure supports non-disruptive rolling updates. Cluster nodes are updated one by one, with the data availability unaffected. During the update, the node enters maintenance mode and its

workloads and VMs are migrated to other nodes. After the update, the node returns to operation. For more information on the maintenance mode, see *Performing Node Maintenance* (page 27).

Take note of the following before you start updating nodes:

- To check for and download updates, the cluster must be healthy and each node in the infrastructure must be able to open outgoing Internet connections. This means, in particular, that cluster DNS must be configured and point to a DNS able to resolve external host names. For more details, see *Adding External DNS Servers* (page 263).
- Unassigned nodes cannot be updated.
- Updates are applied to one storage cluster node at a time.
- If you enable any third-party repositories like EPEL, make sure that packages from the official repository, hci.repo, are never overwritten by packages from the third-party repositories. Disable such repositories when you do not need them anymore. Otherwise, product stability may be at risk.

Note: For details on upgrading to this version, see the Administrator's Guide for version 3.0.

To update the storage cluster from the admin panel, do the following:

 Open the SETTINGS > Updates screen. The date of the last check is displayed in the upper right corner. Click the round arrow to check for new updates. If updates are available for a node, that node's update status changes to Available.



Х

Install

Cancel

- Click **Download** in the upper right corner to get the updates. When the updates are downloaded to a node, its update status changes to **Ready to install**. After the updates have been downloaded for all of the nodes, the button will change to **Install**. Click it to continue.
- 3. In the **Install updates** window, choose how to proceed if a node needs to be rebooted but cannot enter maintenance mode. To avoid downtime, nodes that need to be rebooted are placed in maintenance mode one by one during the update. Nodes that cannot enter maintenance mode are skipped and not updated by default. By selecting the checkbox, you can instead stop the update if some of the nodes cannot enter maintenance mode. Nodes that have already been updated will remain so.

Install updates Updates are ready to install All nodes are eligible for updating. Abort the update if a node cannot enter maintenance ()

While the updates are being installed, you can pause or cancel the process.

4. If you have chosen to abort the update if some of the nodes cannot enter maintenance and it happens, a window will open showing the reasons why the nodes cannot enter maintenance. You will need to decide how to proceed. You can cancel the update, solve the issues, and retry updating without downtime. Or you can force the update on the listed nodes at once. In the latter case, the nodes in question will be rebooted, potentially causing a downtime of services running on them.

After the update is complete, node statuses will change to **Up to date**.

Click Install.

To update the kernel with ReadyKernel, consult Updating Kernel with ReadyKernel.

6.10 Allowing root Access to Cluster Nodes over SSH

In certain situations, you or the technical support team may need root access to cluster nodes via SSH. We recommend using SSH keys as they are generally more secure than passwords. You can generate a key pair on a client from which you will connect to the nodes via SSH. The private key will be stored on the client. Make sure you do not share the private key with anyone for security reasons. The public key will need to be uploaded to Acronis Cyber Infrastructure.

To create and upload a public key, do the following:

1. Obtain an SSH public key from the technical support team, or generate an SSH key pair on a client using the ssh-keygen utility:

ssh-keygen -t rsa

By default, the generated public key is located in /root/.ssh/id_rsa.pub.

- 2. Open the **SETTINGS** > **Security** > **SSH** screen and click **ADD**.
- 3. On the Add public key panel, paste the key and click ADD KEY.

× Add public key



To delete the key after root access is no longer required, select the key and click **Delete**.

6.11 Adding External DNS Servers

Acronis Cyber Infrastructure features a built-in DNS server that enables discovery of all its internal services. For resolving external domain names, you can add DNS servers that already exist in your network infrastructure.

Important: Specify a DNS that belongs to a public network to be able to reach external locations like the updates repository as well as any public networks.

Do the following:

1. On the **SETTINGS** > **Cluster DNS** screen, click **Add**.

Cluster DNS	
Cluster DNS is a flexible and scalable service that provides a D platform and enables discovery of all its internal services. Spec servers that cluster DNS will use to resolve external domain na	NS server for the cify external DNS ames.
DNS servers	+ Add

2. Either specify a static DNS IP address in the **Static** field or select a DHCP-provided DNS IP address from the list. Click **Add** multiple times to specify multiple external DNS servers.

Cluster DNS				
Cluster DNS is a flexible platform and enables d servers that cluster DNS	e and sca iscovery 5 will use	lable service that provides a D of all its internal services. Spe to resolve external domain n	NS server cify exterr ames.	for the nal DNS
DNS servers		+ Add	×	~
10.10.10.10	OR	Obtained via DHCP	~	Ū
Static	OR	10.37.130.2	~	Ū

3. Click the check mark icon to save changes.

6.12 Configuring the Self-Service Panel

The self-service panel is a web-based control panel that allows end users to manage virtual objects, such as virtual machines, volumes, private networks, and other, in isolated administrative environments.

Note: The default system administrator cannot log in to the self-service portal.

To be able to access the self-service panel and manage virtual objects in it, do the following:

- 1. Create the compute cluster as described in *Creating the Compute Cluster* (page 69).
- 2. Create new domains, projects, and user accounts as described in *Managing Domains, Users, and Projects* (page 246).
- 3. Share needed compute images between all projects or grant the permission to upload images to domain administrators and project members.
- 4. Make sure the compute cluster has at least one virtual public network.
- 5. Open TCP port 8800 on the management node as follows:
 - 5.1. On the **INFRASTRUCTURE** > **Networks** screen, click **Edit**.
 - 5.2. Add the **Self-service panel** traffic type to your public network by ticking the corresponding checkbox.
 - 5.3. Click **Save** to apply changes.

You can now access the self-service panel at http://<admin_panel_IP_address>:8800. Use the domain name and user credentials to log in. If high availability for the management node is enabled, log in into the self-service panel using the virtual address for the admin panel:

http://<admin_panel_virtual_IP_address>:8800. You can also use the link in the **Self-service panel URL** field on **SETTINGS** > **Self-service** screen.

Public: 10.94.0.0/16

Self-service panel URL

https://10.94.129.79:8800

To change the virtual IP address of the self-service panel, do the following:

- 1. Make sure high availability for the management node is enabled (see *Enabling Management Node High Availability* (page 238)).
- On the SETTINGS > Self-service screen, click the pencil icon next to the Self-service panel IP address field.
- 3. In the **Edit virtual IP address** window, enter the desired IP address and click **Save**.

Edit virtual IP address	×
You can configure the virtual IP addresses at portal will be accessible.	which the self-service
For the network Public: 10.94.0.0/16 Virtual IP address 10.94.129.70	
	Cancel Save

Note: You cannot change the virtual IP address, if the **Self-service panel** traffic type is assigned along with **Compute API** or **Internal management** to the same network. In this case, you need to destroy the management node HA and re-create it specifying the desired IP address.

On the **SETTINGS** > **Self-service** screen, you can customize the self-service panel appearance as follows:

- In the **Product title** section, click the pencil icon to change the product title displayed on the self-service web-browser tab. In a window that opens, specify a product title and click **Save**.
- In the **Logos** section, click **Upload** to upload logos for the panel header and login screen and select an image file in the PNG, JPG, or SVG format. The image must be 256 x 64 pixels in dimensions and up to 2 MB in size.
- In the **Favicon** section, click the upload icon to upload a favicon for the self-service panel and select an image file in the PNG or ICO format. The image must be 32 x 32 pixels in dimensions and up to 1 MB in size.
- In the **Color scheme** section, click **Change scheme** to choose a color scheme for the self-service panel. In a window that opens, choose the desired color scheme and click **Apply**.

To remove the chosen product title, logos, and favicon from the self-service and reset the theme to default, click **Reset to default**.

Theme 🖸 Reset to default Product title Custom product title 🛛 🔗 Favicon Logos PNG, JPG or SVG format; 256 x 64 pixels; size up to PNG or ICO format; 32 x 32 pixels; size up to 1 MB. 2 MB Favicon Header Left alignment is recommended Custom header logo Color scheme Login screen Centered alignment is recommended Custom login screen logo

6.13 Sending Problem Reports

To send a problem report to the technical support team, do the following:

1. On any screen, click the user icon in the top right corner and select **Report a problem**.

admin	
Report a pro	blem പിന
Language	\mathbf{U}

2. Enter your contact email and a problem description and click **Generate and send**.

Send problem report	×
A problem report contains data about configuration of all cluster services, their usage logs (including debug), details on node hardware (IP/MAC addresses, HWIDs/names, S statuses, etc.), as well as details on software (OS and other data). A report may also co kernel panic and segmentation fault dumps.	statistics, S.M.A.R.T. ontain
After a problem report is generated, contact Acronis Support and mention the report support ticket.	ID in the
Contact email user1@example.com	
Describe the issue you are experiencing (optional)	
Problem description.	
Cancel Generate a	nd send

The problem report will be generated and assigned an ID. Make sure to mention this ID in the support ticket.