

VOTIRO

Votiro Cloud V9.9

Installation Guide

January 2025

Copyright Notice

The material herein is proprietary to Votiro CyberSec Ltd. This document is for informative purposes. Any unauthorized reproduction, use or disclosure of any part of this document is strictly prohibited.

Votiro CyberSec's name and logos are trademarks of Votiro CyberSec Ltd., its subsidiaries or affiliates. All other company or product names are the trademarks of their respective holders.

www.votiro.com

Contents

1 Installing Votiro Cloud	4
1.1 Deployment Specifications	4
1.1.1 Item retention	5
1.2 Prerequisites and Considerations	5
1.2.1 Ports	5
1.2.2 Virtual Appliance Communication Settings	6
1.2.3 Syncing with an NTP Server	7
1.2.4 Using an External Storage Server	7
1.2.5 Load Balancing	8
1.2.6 Votiro Registry in Azure	8
1.3 Deploying an OVF	8
1.4 Deploying Votiro Cloud	11
1.4.1 General Guidelines	11

1 Installing Votiro Cloud

To install Votiro Cloud quickly into your organization we will create a cluster of virtual machines (VM). Each VM requires dedicated resources, see [Deployment Specifications below](#).

To install Votiro Cloud and login to start using the Management Dashboard, follow these four steps:

- Deploy an OVF
- Configure the Network Environment
- Deploy Votiro Cloud
- Login to the Management Dashboard

IMPORTANT!

You may need to determine in advance of your installation the following:

- Unique IP addresses: see [Deployment Specifications](#);
- Hostname for FQDN (use lower case alphanumeric characters).

1.1 Deployment Specifications

The following deployment specifications are for the installation of Votiro Cloud with 3 and 5 node clusters. Scale specifications as you increase the number of nodes in your cluster.

The expected maximum performance for this configuration after a fresh install is 35,000 files (emails) per hour for the 3 node cluster and 90,000 files (emails) per hour for the 5 node cluster.

Table 1 Deployment Specifications

Votiro Cloud	3 Node Cluster	5 Node Cluster	Azure Dv4 series - Instance D8 v4
CPU Cores	8 per node	8 per node	8v CPUs
RAM	24 GB per node	32 GB per node	32 GiB
Drive Capacity	500 GB per node, SSD	500 GB per node, SSD	Attached SSD
Remote Storage Support	File Storage Network. For example, SAN, NAS	File Storage Network. For example, SAN, NAS	
Hypervisor Support	VMWare ESXi 6, ESXi 7.0, ESXi 8.0, Amazon Web Services, Microsoft Azure	VMWare ESXi 6, ESXi 7.0, ESXi 8.0, Amazon Web Services, Microsoft Azure	

Votiro Cloud	3 Node Cluster	5 Node Cluster	Azure Dv4 series - Instance D8 v4
Network Adapters	5 - 1 per node + 1 for the LB VIP + 1 for the internal k8 LB	7 - 1 per node + 1 for the LB VIP + 1 for the internal k8 LB	

1.1.1 Item retention

By default, our product supports item retention for 90 days. Items include all of the events records (email/files events that arrived at our product).

Note:
If a customer wants to change the default value of 90 days, they must contact Votiro support to change it.

1.2 Prerequisites and Considerations

There are both prerequisites and a number of topics for you to consider when implementing Votiro Cloud into your environment. See sections for more details:

- [Ports](#)
- [Virtual Appliance Communication Settings](#)
- [Syncing with an NTP Server](#)
- [Using an External Storage Server](#)
- [Load Balancing](#)
- [Votiro Registry in Azure](#)

1.2.1 Ports

Network connectivity requirements enabling secure outbound and inbound communications with Votiro Cloud are detailed in the tables below.

Table 2 Outbound Firewall Rules

Outbound	Source	Destination	Port Number	Transport Protocol
Releasing Files	ovf_network	Exchange / Edge	25	tcp
Active Directory	ovf_network	Domain Controller <ul style="list-style-type: none"> ■ LDAP ■ LDAPS 	<ul style="list-style-type: none"> ■ 389 ■ 636 	<ul style="list-style-type: none"> ■ tcp ■ tcp
SIEM	ovf_network	SIEM Server	514	udp

Table 3 Inbound Firewall Rules

Inbound	Source	Destination	Port Number	Transport Protocol
SSH, SCP	Any	ovf_network	22	tcp
Processing Request	API Client	ovf_network	443	tcp
Monitoring Grafana	Grafana	ovf_network		
Monitoring Prometheus	Prometheus			

Additional Port Connectivity Requirements when Connecting to External Storage

When there’s a firewall between the cluster to the external NFS-based storage or the connection is somehow restricted on the customer’s end, the following ports should be opened/allowed when trying to connect to the external storage:

- Port 111 TCP\UDP – PortMapper (mandatory).
- Port 2049 TCP\UDP – NFS service (mandatory).
- Port 635 TCP\UDP – Mount daemon (mandatory only when working with NetApp).
- Port 4045 TCP\UDP – NFS lock manager (mandatory only when working with NetApp).
- Port 4046 TCP\UDP – NFS status (mandatory only when working with NetApp).
- Port 4049 TCP\UDP – NFS quota daemon (mandatory only when working with NetApp).

1.2.2 Virtual Appliance Communication Settings

Internal Communication Settings

For internal communications between nodes of each machine inside the VLAN, the following settings are required:

Port number	Protocol	Description
22	TCP	During init, node 1 will communicate with node 2 and 3 and will update keys, username, etc.
25	TCP	Required when enabling "Release" function from management console (email integration)
389	TCP (LDAP)	LDAP - Active Directory integration
636	TCP (LDAPS)	LDAPS - Secure Active Directory integration
2379-2380	TCP	etcd server-client API (used by kube-apiserver, etcd)
6443	TCP	Kubernetes API server
10250-10252	TCP	Kubelet API

Port number	Protocol	Description
10255	TCP	Worker node read-only Kubelet API
24007-24008	TCP	GlusterFS (daemon+management) (note it is 24007-24008)
49152-49154	TCP	GlusterFS (for each brick in a volume)
123	UDP	Require to enable Network Time Protocol (NTP) (See Syncing with an NTP Server below)
514	UDP	On-prem Syslog integration
8472	UDP	Flannel overlay network (K8s requirement)
51820	UDP	
51821	UDP	
5001	TCP	

External Communication Settings

For external communications, the following settings are required:

- 22/tcp
- 443/tcp

1.2.3 Syncing with an NTP Server

When using an NTP server, as a pre-requisite you must sync with it using port **123/udp**.

1.2.4 Using an External Storage Server

In addition to the virtual appliance machines' internal storage, you can use an external storage server. Votiro Cloud can be configured to communicate with your storage server, using a mount from the external storage to the virtual appliance machines.

When external storage is configured it is used as the main storage area. Storage will contain a set of original and processed files.

The mount created results in the true storage type, such as SAN and NAS, being transparent, leading to Votiro Cloud supporting all External Storage types.

For instructions on how to configure External Storage, contact Votiro's Support team.

Note

* The internal storage requirement remains at 300 GB per node. It is available for use should the external storage server link fail. Stored files are transferred from the VM to the external storage server when it becomes available.

* Read / Write permissions should be granted to user **1000** for the relevant path.

* Cluster IPs should be added under Policy-Export rules.

1.2.5 Load Balancing

Votiro Cloud automatically supports load balancing using a basic internal load balancer.

Note: An external hardware-based load balancer is required in your production environment to balance between the nodes of your VM.

WARNING!

Our product supports high-availability when a node fails. The system will continue to sanitize but at reduced performance. There could be a minimal downtime of one minute. We recommend recovering the failing node as soon as possible to restore the system to maximum sanitization performance.

1.2.6 Votiro Registry in Azure

This consideration is relevant when your Votiro Cloud installation includes an online environment.

To enable secure communication with your Votiro appliance, the proxy server ACL must include permission for the Votiro registry in the Azure URL.


1.3 Deploying an OVF

In this step you will create the virtual machines. You will require a virtual machine for each node in your cluster.


1. Deploy **OVFs**, using these specifications:

- ◆ 8 CPU
- ◆ 24 GB Memory
- ◆ 500 GB Storage


- Name each **node** uniquely using your corporate naming conventions.



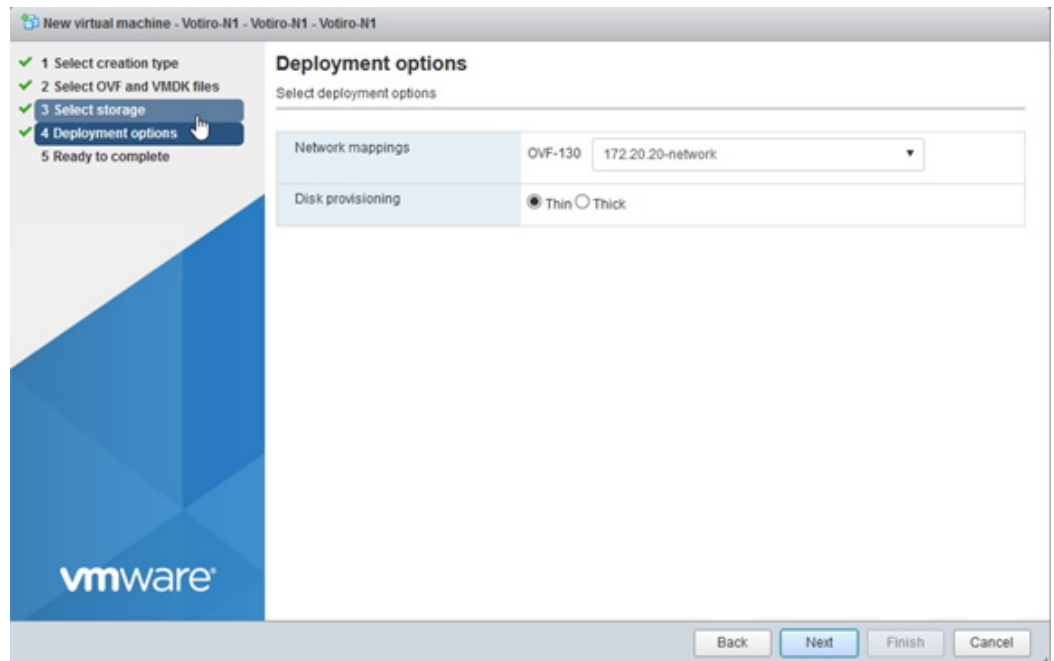
- Select the **OVF** and **VMDK** files during deployment.




- Select your preferred storage location. It is recommended you use **SSD storage**.



- Select the network you would like to deploy the appliances on. You may select **Thin** or **Thick** provisioning. 500GB of storage is required for each appliance.



6. To complete the deployment, click **Finish**.



There are now three or five virtual machines (VM).

1.4 Deploying Votiro Cloud

1.4.1 General Guidelines

To run a successful installation, you need to relate parameters in **inventory.yaml**. The file is located at **/root/ansible-initcluster/inventory.yaml**. That's the only file that needs to be updated prior to running the installation:

```

1 ---|
2 all:
3   children:
4     k3s_cluster:
5       children:
6         server:
7           hosts:
8             1.1.1.1:
9             2.2.2.2:
10            3.3.3.3:
11          # agent:
12          # hosts:
13            # 192.16.35.12:
14            # 192.16.35.13:
15
16 vars:
17   approve_votiro_eula: no # read Votiro eula at: https://votiro.com/eula/ and set to yes to install.
18   controlplane_vip_address: 4.4.4.4
19   paralus_web_vip: 5.5.5.5 # false for external load balancer, or set to a specific ip.
20   votiro_cluster_fqdn: paralus-app.va.votiro.com # fqdn of the paralus application
21   safe_browsing_enabled: false # Online / offline mode for safebrowsing
22   time_zone: Etc/UTC # list of time zones: https://en.wikipedia.org/wiki/List_of_tz_database_time_zones
23   ntp_servers: "pool.ntp.org time.google.com" # list of ntp servers separated by space
24   tenant_id: "" # for migration from older cluster with existing tenant
25   system_id: "" # for migration from older cluster with existing system id
26   # Leave empty to use cluster internal storage.
27   # Both volumes can have same nfs server and path.
28   # example value(can use hostname): 10.10.11.11:/nfs_share_path
29   blob_nfs: ""
30   file_cache_nfs: ""

```

■ 3 Key Configurations before running the script **install-paralus-playbook.yaml**

- ◆ Define Hostname
- ◆ Network Configurations: Use the file named **00-installer-config.yaml** (located under **/etc/netplan**)
- ◆ Define Configuration Parameters: Use the file named **inventory.yaml** (located under **/root/ansible-initcluster**)

Note: You must configure the NTP (Network Time Protocol) server to be valid and accessible at all times to avoid major issues in our product.

Recommendation for external storage: We recommended working with an external NFS (Network File System) for handling larger volumes or to enhance performance (e.g., to achieve a high sanitization rate). External storage can provide greater flexibility and can accommodate larger file sizes beyond the temporary internal storage limit, ensuring that your operations continue smoothly without interruption. Working with internal storage can cause system errors if temporary internal storage exceeds the maximum. In that case, the system will return an error and sanitization will be enabled only after the temporary storage returns to normal operation.

- Procedure:
 - a. Access the first virtual machine with VMRC.
 - b. Use the credentials that were supplied separately.
 - c. To switch to the root user, type **sudo -i** in the terminal.
 - d. Change hostname command:

```
hostnamectl set-hostname NODE_NAME
```
 - e. `cd /etc/netplan`
 - f. Edit (vi) the file named **00-installer-config.yaml** to edit VM network settings.
 - g. Edit lines 8, 11, and 13 (machine address, gateway, DNS address). You must use the prefix **/** to define the network size. See the example below:

```
1 network:
2   version: 2
3   renderer: networkd
4   ethernets:
5     ens160:
6       optional: true
7       dhcp4: no
8       addresses: [10.130.0.210/23] # machine address
9       routes:
10        - to: 0.0.0.0/0
11          via: 10.130.1.1 # gateway
12       nameservers:
13         addresses: [192.168.11.5] # dns address
```

- h. Save the file.
- i. Apply network configurations:
`netplan apply`
- j. Repeat on the other VMs.

■ **Mandatory configurations for fresh install:**

- a. SSH to the first VM.
 - b. **sudo -i**
 - c. **cd ansible-initcluster**
 - d. **vi inventory.yaml**
- ◆ In the 3 nodes configuration, add the node IPs under the **hosts:** section of the file.
 - ◆ In the 5 nodes configuration, add the first 3 node IPs under the **hosts:** section. This will make them the cluster's master nodes. Uncomment the **agent:** section and add the IPs of the rest of the nodes there. These nodes will be worker nodes.
 - ◆ **approve_votiro_eula:** should be set to **yes**
 - ◆ **controlplane_vip_address:** should receive an unused IP to be used for internal purposes.
 - ◆ **paralus_web_vip:** should either receive an unused IP to be used by the system's load balancer, or leave empty for an external load balancer.
 - ◆ **votiro_cluster_fqdn:** should contain the applicable FQDN for the system

■ **Additional configurations:**

- ◆ **safe_browsing_enabled:** set to true or false for Online / Offline mode for safe browsing.

- ◆ In fresh install mode, the **tenant_id** and **system_id** should be left empty (they are generated automatically).
- ◆ In upgrade mode, fill the **tenant_id** and **system_id** fields with the data from the previous environment.
- ◆ **blob_nfs:** is to be used if a customer wants to save the original and sanitized files in an external storage. Can be left empty for internal blob storage. Example value (can use hostname):

```
10.10.11.11:/nfs_share_path
```

- ◆ **file_cache_nfs:** is used to achieve better performance for the system. Can be left empty for an internal storage usage. Example value (can use hostname):

```
10.10.11.11:/nfs_share_path
```

- e. Save the file.
- f. Take a Snapshot.

■ Run the script

- a. From the **ansible-initcluster** directory, run the next command:
`ansible-playbook install-paralus-playbook.yaml`
- b. Follow the instructions on the screen.
- c. When done, cat the **votiro-setup.log** file to verify successful installation.
- d. Copy the encryption keys and save it in a safe place (see example below).

```
"msg": [
  "Please keep the following encryption keys in a safe place, they cannot be retrieved",
  [
    "salt=D70[REDACTED]",
    "key=7F3[REDACTED]",
    "iv =65[REDACTED]"
  ]
]
}
2024-04-02 07:38:30,959 p=2362 u=root n=ansible | PLAY RECAP *****
*****
2024-04-02 07:38:30,959 p=2362 u=root n=ansible | 10.130.0.210 : ok=74 changed=50 unreachable=0 failed=0
2024-04-02 07:38:30,959 p=2362 u=root n=ansible | 10.130.0.211 : ok=56 changed=42 unreachable=0 failed=0
2024-04-02 07:38:30,959 p=2362 u=root n=ansible | 10.130.0.212 : ok=56 changed=42 unreachable=0 failed=0
2024-04-02 07:38:30,959 p=2362 u=root n=ansible | localhost : ok=72 changed=40 unreachable=0 failed=0
```

- e. Run the health check to verify that the system is running properly.

■ For big file support, do the following:

- ◆ Apply ssh to one of the nodes, and run the **/root/extras/scale-for-large-file/change-memory-limit.sh** script.
- ◆ Edit the **cancellation-service-config** configmap:
`Kubect1 edit cm cancellation-service-config`
and increase the **CancellationTimeout:** field to **01:00:00** (1h)
- ◆ In the management’s policy page, set the large file case to **skip**

Your Votiro Cloud installation has completed successfully.