

Deploy an ExtraHop recordstore on Linux KVM

Published: 2024-07-22

In this guide, you will learn how to deploy an ExtraHop virtual recordstore on a Linux kernel-based virtual machine (KVM) and to join multiple recordstore to create a cluster. You should be familiar with basic KVM administration before proceeding.

- Important:** If you want to deploy more than one ExtraHop virtual sensor, create the new instance with the original deployment package or clone an existing instance that has never been started.

System requirements

Your environment must meet the following requirements to deploy a virtual recordstore:

- Important:** ExtraHop tests virtual clusters on local storage for optimal performance. ExtraHop strongly recommends deploying virtual clusters on continuously available, low latency storage, such as a local disk, direct-attached storage (DAS), network-attached storage (NAS), or storage area network (SAN).

- A KVM hypervisor environment capable of hosting the virtual recordstore. The virtual recordstore is available in the following configurations:

EXA Manager Node	EXA-XS	EXA-S	EXA-M	EXA-L
4 CPUs	4 CPUs	8 CPUs	16 CPUs	32 CPUs
8 GB RAM	8 GB RAM	16 GB RAM	32 GB RAM	64 GB RAM
4 GB boot disk	4 GB boot disk	4 GB boot disk	4 GB boot disk	4 GB boot disk
12 GB	250 GB or smaller datastore disk	500 GB or smaller datastore disk	1 TB or smaller datastore disk	2 TB or smaller datastore disk

The hypervisor CPU should provide Streaming SIMD Extensions 4.2 (SSE4.2) and POPCNT instruction support.

- Note:** The EXA manager node is preconfigured with a 12 GB datastore disk. You must manually configure a second virtual disk to the other EXA configurations to store record data.

Consult with your ExtraHop sales representative or Technical Support to determine the datastore disk size that is best for your needs.

- Note:** For KVM deployments, virtio-scsi interface is recommended for the boot and datastore disks.

- A virtual recordstore license key.
- The following TCP ports must be open:
 - TCP port 443: Enables browser access to Administration settings. Requests sent to port 80 are automatically redirected to HTTPS port 443.
 - TCP port 9443: Enables recordstore nodes to communicate with other nodes in the same cluster.

Package contents

The installation package for KVM systems is a tar.gz file that contains the following items:

EXA-5100v-<x>.xml

The domain XML configuration file

EXA-5100v-<x>.xml.md5

The domain XML checksum file

extrahop-boot.qcow2

The boot disk

extrahop-boot.qcow2.md5

The boot disk checksum file

Deploy the virtual recordstore

To deploy the virtual recordstore, complete the following procedures:

- [Determine the best virtual bridge configuration for your network](#)
- [Edit the domain XML configuration file and create your virtual appliance](#)
- [Create the datastore disk](#)
- [Start the VM](#)
- [Configure the Explore appliance](#)

Determine the best bridge configuration

Identify the bridge through which you will access the management interface of your recordstore.

1. Make sure the management bridge is accessible to the virtual recordstore and to all users who must access the management interface.
2. If you need to access the management interface from an external computer, configure a physical interface on the management bridge.

Edit the domain XML configuration file

After you identify the management bridge, edit the configuration file, and create the virtual recordstore.

1. Contact [ExtraHop Support](#) to obtain and download the Explore KVM package.
2. Extract the tar.gz file that contains the installation package.
3. Copy the extrahop-boot.qcow2 file to your KVM system.
4. Open the domain XML configuration file in a text editor and edit the following values:
 - a) Change the VM name to a name for your ExtraHop virtual recordstore.

For example:

```
<name>ExtraHop-EXA-S</name>
```

- b) Change the source file path ([PATH_TO_STORAGE]) to the location where you stored the virtual disk file in step 3.

```
<source file='/images/extrahop-boot.qcow2' />
```

- c) Change the source bridge for the management network (`ovsbr0`) to match the name of your management bridge.

```
<interface type='bridge'>
  <source bridge='ovsbr0' />
  <model type='virtio' />
  <alias name='net0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03'
    function='0x0' />
</interface>
```

- d) Optional: If your virtual bridge is configured through Open vSwitch virtual switch software, add the following virtualport type setting to the interface (after the source bridge setting):

```
<virtualport type='openvswitch'>
</virtualport>
```

5. Save the XML file.

Create the datastore disk

Create the datastore disk so that the allotted space is large enough to store the type of records you want to store for the amount of lookback desired.

To create the datastore disk, run the following command:

```
qemu-img create -f qcow2 <path to storage location> <size>
```

Where `<size>` is the size of the disk in gigabytes. This example creates a qcow2 image with a maximum size of 2 TB:

```
qemu-img create -f qcow2 /home/extrahop/extrahop-data.qcow2 2000G
```

Create the recordstore

Create the virtual recordstore with your revised domain XML configuration file by running the following command:

```
virsh define <EXA-5100v-<x>.xml>
```

Where `<EXA-5100v-<x>.xml>` is the name of your domain XML configuration file.

Start the VM

1. Start the VM by running the following command:

```
virsh start <vm_name>
```

Where `<vm_name>` is the name of your ExtraHop virtual recordstore you configured in step 4 of the [Edit the domain XML file](#) section.

2. Log in to the KVM console and view the IP address for your new ExtraHop virtual recordstore by running the following command:

```
virsh console <vm_name>
```

(Optional) Configure a static IP address

By default, the ExtraHop system is configured with DHCP enabled. If your network does not support DHCP, you must configure a static address manually.

1. Log in to the KVM host.
2. Run the following command to connect to the ExtraHop system through the virtual serial console:

```
virsh console <vm_name>
```

Where `<vm_name>` is the name of your virtual machine.

3. Press ENTER twice to get to the system login prompt.

```
ExtraHop Discover Appliance Version 7.8.2.2116
IP: 192.0.2.81
exampleium login:
```

4. At the login prompt, type `shell`, and then press ENTER.
5. At the password prompt, type `default`, and then press ENTER.
6. To configure the static IP address, run the following commands:

- a) Enable privileged commands:

```
enable
```

- b) At the password prompt, type `default`, and then press ENTER.
- c) Enter configuration mode:

```
configure
```

- d) Enter the interface configuration mode:

```
interface
```

- e) Run the `ip` command and specify the IP address and DNS settings in the following format: `ip ipaddr <ip_address> <netmask> <gateway> <dns_server>`

For example:

```
ip ipaddr 10.10.2.14 255.255.0.0 10.10.1.253 10.10.1.254
```

- f) Leave the interface configuration section:

```
exit
```

- g) Save the running config file:

```
running_config save
```

- h) Type `y` and then press ENTER.

Configure the recordstore

After you obtain the IP address for the recordstore, log in through `https://<explore_ip_address>/admin` and complete the following recommended procedures.



Note: The default login username is `setup` and the password is `default`.

- [Register your ExtraHop system](#) 

- [Connect the console and sensors to ExtraHop recordstores](#)
- [Send record data to the Explore appliance](#)
- Review the [Explore Post-deployment Checklist](#) and configure additional recordstore settings.

Create a recordstore cluster

For the best performance, data redundancy, and stability, you must configure at least three recordstores in a cluster.

- ⓘ **Important:** If you are creating a recordstore cluster with six or more nodes, you must configure the cluster with manager nodes. For manager node instructions, see [Deploying manager nodes](#).

In the following example, the recordstores have the following IP addresses:

- Node 1: 10.20.227.177
- Node 2: 10.20.227.178
- Node 3: 10.20.227.179


You will join nodes 2 and 3 to node 1 to create the recordstore cluster. All three nodes are data nodes. You cannot join a data node to a manager node or join a manager node to a data node to create a cluster.

- ⓘ **Important:** Each node that you join must have the same configuration (physical or virtual) and the same ExtraHop firmware version.


Before you begin

You must have already installed or provisioned the recordstores in your environment before proceeding.


1. Log in to the Administration settings on all three recordstores with the setup user account in three separate browser windows or tabs.
2. Select the browser window of node 1.
3. In the Status and Diagnostics section, click **Fingerprint** and note the fingerprint value. You will later confirm that the fingerprint for node 1 matches when you join the remaining two nodes.
4. Select the browser window of node 2.
5. In the Explore Cluster Settings section, click **Join Cluster**.
6. In the Host field, type the hostname or IP address of data node 1 and then click **Continue**.

 **Note:** For cloud-based deployments, be sure to type the IP address listed in the Interfaces table on the Connectivity page.

7. Confirm that the fingerprint on this page matches the fingerprint you noted in step 3.



8. In the Setup Password field, type the password for the node 1 `setup` user account and then click **Join**. When the join is complete, the Explore Cluster Settings section has two new entries: **Cluster Members** and **Cluster Data Management**.
9. Click Cluster Members. You should see node 1 and node 2 in the list.




- In the Status and Diagnostics section, click **Explore Cluster Status**. Wait for the Status field to change to **Green** before adding the next node.
- Repeat steps 5 - 10 to join each additional node to the new cluster.



Note: To avoid creating multiple clusters, always join a new node to an existing cluster and not to another single appliance.

- When you have added all of your recordstores to the cluster, click **Cluster Members** in the Explore Cluster Settings section. You should see all of the joined nodes in the list, similar to the following figure.




- In the Explore Cluster Settings section, click **Cluster Data Management** and make sure that **Replication Level** is set to **1** and **Shard Reallocation** is **ON**.


Next steps

[Connect the console and sensors to ExtraHop recordstores](#) 

Connect the recordstore to a console and all sensors

After you deploy the recordstore, you must establish a connection from the console and all sensors before you can query records.


 **Important:** Connect the sensor to each recordstore node so that the sensor can distribute the workload across the entire recordstore cluster.

 **Note:** If you manage all of your sensors from a console, you only need to perform this procedure from the console.

1. Log in to the Administration settings on the ExtraHop system through `https://<extrahop-hostname-or-IP-address>/admin`.
2. In the ExtraHop Recordstore Settings section, click **Connect Recordstore**.
3. Click **Add New**.
4. In the Node 1 section, type the hostname or IP address of any recordstore in the cluster.
5. For each additional node in the cluster, click **Add New** and enter the individual hostname or IP address for the node.
6. Click **Save**.
7. Confirm that the fingerprint on this page matches the fingerprint of node 1 of the recordstore cluster.
8. In the Explore Setup Password field, type the password for the node 1 `setup` user account and then click **Connect**.
9. When the Explore Cluster settings are saved, click **Done**.

Send record data to the recordstore

After your recordstore is connected to your console and sensors, you must configure the type of records you want to store.

See [Records](#)  for more information about configuration settings, how to generate and store records, and how to create record queries.