


Deploy the ExtraHop recordstore with VMware

Published: 2024-07-22


In this guide, you will learn how to deploy a virtual ExtraHop recordstore with the vSphere client running on a Windows machine and to join multiple recordstores to create a recordstore cluster. You should be familiar with administering VMware ESX and ESXi environments before proceeding.

The virtual recordstore is distributed as an OVA package that includes a preconfigured virtual machine (VM) with a 64-bit, Linux-based operating system (OS) that is optimized to work with VMware ESX and ESXi version 6.5 and later.

 **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 ExtraHop 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).
- An existing installation of VMware ESX or ESXi server version 6.5 or later capable of hosting the virtual recordstore. The virtual recordstore is available in the following configurations:

Recordstore Manager-Only Node	5100v Extra-Small	5100v Small	5100v Medium	5100v Large
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 recordstore manager-only node is preconfigured with a 12 GB datastore disk. You must manually configure a second virtual disk to the other recordstore 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.

- A vSphere client
- A virtual recordstore license key.
- The following TCP ports must be open:
 - TCP ports 80 and 443: Enables you to administer the recordstore. Requests sent to port 80 are automatically redirected to HTTPS port 443.
 - TCP port 9443: Enables recordstore nodes to communicate with other recordstore nodes in the same cluster.

Deploy a virtual ExtraHop recordstore

Before you begin

If you have not already done so, download the virtual ExtraHop recordstore OVA file for VMware from the [ExtraHop Customer Portal](#).



Note: If you must migrate the virtual machine (VM) to a different host after deployment, shut down the virtual recordstore first and then migrate with a tool such as VMware VMotion. Live migration is not supported.

1. Start the VMware vSphere client and connect to your ESX server.
2. From the **File** menu, select **Deploy OVF Template**.
3. Follow the Virtual Machine wizard prompts to deploy the OVF template:
For most deployments, the default settings are sufficient.
 - a) Browse to the location of the downloaded OVA file, select the file, and then click **Next**.
 - b) Review and verify the OVF template details and click **Next**.
 - c) Type the VM name and location. Give the VM a unique and specific name for the ESX Inventory and then click **Next**.
 - d) For Disk Format, select **Thick Provision Lazy Zeroed** and then click **Next**.
 - e) Map the OVF-configured network interface labels with the correct ESX-configured interface labels and then click **Next**.
 - a) Verify the configuration, do not select the **Power on after deployment** checkbox, and then click **Finish** to complete the deployment.

When the deployment is complete, you can see the unique name you assigned to the recordstore VM instance in the inventory tree for the ESX server to which it was deployed.
4. Click the new recordstore VM instance in the directory tree.
5. From the Actions drop-down list, select **Edit Settings...** to configure the disk where the recordstore data is stored.
6. From the New device drop-down list, select **New Hard Disk**, confirm that **Thick Provision Lazy Zeroed** is selected for Disk Provisioning, and then click **Add**.
7. In the New Hard disk field, type the size of your virtual storage disk and then click **OK**.
8. From the Actions drop-down list, select **Power On**.
9. From the Actions drop-down list, select **Open Console**.
10. Log in with the `shell` user account and the password `default`.
11. Run the `show ipaddr` command to display the IP address of the virtual recordstore.
12. Exit the console window.

Configure a static IP address through the CLI

The ExtraHop system is configured by default with DHCP enabled. If your network does not support DHCP, no IP address is acquired, and you must configure a static address manually.

You can manually configure a static IP address for the ExtraHop system from the CLI.



Important: We strongly recommend [configuring a unique hostname](#). If the system IP address changes, the ExtraHop console can re-establish connection easily to the system by hostname.

1. Access the CLI through an SSH connection, by connecting a USB keyboard and SVGA monitor to the physical ExtraHop appliance, or through an RS-232 serial (null modem) cable and a terminal emulator program. Set the terminal emulator to 115200 baud with 8 data bits, no parity, 1 stop bit (8N1), and hardware flow control disabled.

2. At the login prompt, type `shell` and then press ENTER.
3. At the password prompt, type `default`, and then press ENTER.
4. 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 interface configuration mode:

```
interface
```

- e) 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 interface configuration mode:

```
exit
```

- g) Save the running configuration file:

```
running-config save
```

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

Configure the recordstore

After you obtain the IP address for the ExtraHop 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 ExtraHop recordstore](#)
- Review the [Recordstore 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 ExtraHop recordstores in a cluster.

When you create a recordstore cluster, be sure to deploy all nodes, including manager nodes, in the same location or data center. For more information on supported recordstore cluster configurations, see [Recordstore cluster guidelines](#).




Important: If you are creating a recordstore cluster with six to nine nodes, you must configure the cluster with at least three manager-only nodes. For more information, see [Deploying manager-only 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.



10. In the Status and Diagnostics section, click **Explore Cluster Status**.

Wait for the Status field to change to Green before adding the next node.


11. Repeat steps 5 through 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.

12. 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.



13. 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 ExtraHop 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 recordstore 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.