

Heureka High Performance Index Appliance Installation Instructions



Import OVA File

- 1.1 Import the provided .ova into a virtualization environment. Set the CPU and memory settings to the appropriate values for the targeted data and performance.
 - Recommended Settings: 16 core CPU, 32 GiB RAM
 - Minimum Settings: 8 core CPU, 8 GiB RAM
- 1.2 Configure the virtual machine to have both reserved memory and cpu's to minimize performance impact from other virtual machines.
- 1.3 Add an additional disk that will be used to store the index. Do not resize the default hard disk for this purpose, this must be an additional disk that is provisioned to support the index of the endpoint software.
 - This disk should be at least 20% the size the of the targeted data

Configure Appliance

- 2.1. Start up the virtual machine, log in to the virtual machine with the default credentials:
 - Username: heureka
 - Password: heureka
- 2.2 You will be prompted to set a new password for the heureka user, ensure you save this password where it cannot be lost.
- 2.3 Set up networking
 - 2.3.1 Determine what interfaces are installed on your system.

```
nmcli d
```
 - 2.3.2 Make note of the interface that starts with "ens" if it is disconnected, the use the following command to configure it:

```
nmtui
```
 - 2.3.4 Once done, verify networking is working correctly through usual network utilities such as nslookup, ping, and wget. You can use the following command to get the IP address:

```
ip a
```
- 2.4 At this point ssh should be accessible to the machine. It is recommended to ssh in for the rest of the installation.



2.5 The disk added back in step 1.3 can now be added to the appliance and formatted for usage:

2.5.1 Identify the disk to format

```
sudo lsblk
```

2.5.2 Set it to a variable, for example:

```
DEVICE=/dev/sdb
```

2.5.3 Format the disk, for example:

```
sudo mkfs.ext4 -m 0 -F -E \
lazy_itable_init=0,lazy_journal_init=0,discard $DEVICE
```

2.5.5 Determine UUID of disk

```
UUID=$(sudo blkid -s UUID -o value $DEVICE)
```

2.5.6 Backup fstab

```
sudo cp /etc/fstab /etc/fstab.orig
```

2.5.7 Add fstab entry

```
echo "UUID=$UUID /opt/data ext4 discard,defaults,nofail 0 2" \
| sudo tee -a /etc/fstab
```

2.5.8 Mount all entries in the fstab

```
sudo mount -av
```

Endpoint File Share

3.1 Update the appliance operating system to the latest version

```
sudo yum -y update
```

3.3 Mounting a samba share (if necessary)

3.3.1 Put the credentials (if necessary) into `/root/.smbcredentials`:

```
domain=<domain>
username=<username>
password=<password>
```

3.3.2 Set the permissions of the file to restrict access:

```
sudo chmod 600 /root/.smbcredentials
```

3.3.3 Make the directory where the share will be mounted:

```
sudo mkdir /media/<share_name>
```

3.3.4 Determine the correct mount command for your mount point. For example:

```
sudo mount -t cifs //nas/media /media/nas -o \
iocharset=utf8,ro,credentials=/root/.smbcredentials
```

3.3.5 Once you successfully mount your mount point, use the `umount` command to unmount your mount point. For example:

```
sudo umount /media/nas
```

3.3.6 Adding an entry to `/etc/fstab` using the parameters which were successful from step 3.3.4. For example:

```
//nas/media /media/nas cifs iocharset=utf8,ro,credentials=
/root/.smbcredentials 0 0
```

3.3.7 Mount the share:

```
sudo mount -av
```

3.3.8 Check that the mount is working by listing the contents of the mounted directory:

```
ls /media/<share_name>
```

3.3.9 Verify file permissions have read access on the mount:

```
ls -l /media/<share_name>/path/to/a/test/file.txt
```

3.4 Configuring the endpoint software

3.4.1 Copy the `agent.config` and `updater.config` files that were supplied by Heureka into the following location:

```
/var/lib/indexing-service/conf
```

3.4.2 Determine the number of processor cores on the machine using the following command and record the value. It will be used shortly.

```
nproc
```

3.4.3 In the `agent.conf` file, set the `crawlPaths` to include the file share that was added during step 3.3. For example:

```
endpoint.indexer.crawlPaths=/media/nas
```

This value can also be set to a comma delimited list of paths to index. For example:

```
endpoint.indexer.crawlPaths=/media/nas/photos,/media/nas/docs
```



- 3.4.4 In the `agent.conf` file, set the indexing and search threads to match your core count from step 3.4.2 and set the software to use 100% of the processors. For example, if you have 16 cores it should like this:

```
endpoint.indexer.threads=16
endpoint.index.search.threads=16
endpoint.system.cpu.max=1.0
```

- 3.4.5 In the `agent.conf` file, enable shards for greater performance. Shards split the index into smaller pieces which can be searched in parallel:

```
endpoint.index.shards=4
```

- 3.4.6 In the `agent.conf` file, configure the endpoint to use the PostgreSQL database:

```
endpoint.jdbc.url=jdbc:postgresql://localhost:5432/heureka
endpoint.jdbc.password=heureka
endpoint.jdbc.user=heureka
```

- 3.4.7 Determine how much memory in GB is available by running the following command and record the value. It will be used shortly.

```
free -g | grep "Mem:" | awk '{ print $2 }'
```

- 3.4.8 In the `/var/lib/indexing-service/bin/indexing-service` file, change the amount of memory allocated to the endpoint processes. Change the `-Xmx` argument in the `START_COMMAND` variable to be 50% of the physical RAM of the machine. For example, for a machine with 32 GB of RAM change:

```
" -Xmx1024m"\
```

To:

```
" -Xmx16g"\
```

- 3.4.9 In the `/var/lib/indexing-service/bin/indexing-service` file, change the location where the indexing will be written to the mount point mounted in step 2.5.8. Change:

```
" -Duser.dir=${INDEXING_HOME}"\
```

To:

```
" -Duser.dir=/opt/data"\
```

3.5 Start the indexing service

```
sudo systemctl start indexing-service
```

3.6 Tail the logs to ensure the service is working. To stop tailing the logs use CTRL+C.

```
sudo journalctl -u indexing-service --utc -o json-pretty -f \
| jq -r '.MESSAGE'
```

Expanding the index disk

If you ever run out of space on the indexing mount created in step 1.3 it will need to be expanded.

4.1 Shutdown the appliance

```
sudo shutdown now
```

4.2 Expand the disk in the virtualization software

4.3 Start the appliance back up

4.4 Expand the filesystem

4.4.1 Identify the device

```
sudo lsblk
```

4.4.2 Resize the filesystem

```
sudo resize2fs /dev/sdb
```

Editing files

The system comes with both `vi` and `nano` preinstalled. If you want any additional editors, they can be installed using the `yum` command.



Heureka is a technical leader in endpoint search, identify and classification software. Our goal is to bring order to unstructured data by identifying risk while helping you realize the value of unstructured data across all endpoints.

Heureka, Inc
PH. 800.310.0981
info@heureka.com