Getting Started With IBM Netcool Operations Insight on IBM Cloud Private

Netcool Operations Insight v1.5

IBM Cloud Private v3.1.0

IBM Training

January 2019 edition

Contents

About this lab guide

This guide describes the installation of Netcool Operations Insight (NOI) v1.5 on IBM Cloud Private (ICP) v3.1.0. This document is intended to be a complete, end-to-end walk-through of all the tasks required to install and configure the following components:

- Docker Community Edition (CE)
- IBM Cloud Private (ICP) v3.1.0
- GlusterFS
- Netcool Operations Insight (NOI) v1.5, part number CNX04EN

This guide also shows you how to prepare your ICP hosts before the installation, and demonstrates all required post-installation tasks.

Example architecture

The example architecture that is referenced throughout this lab guide consists of four hosts, which is the recommended minimum for running NOI 1.5 on ICP. This minimum configuration will not provide production-like performance; however, it will be sufficient for demonstration purposes.

The following diagram shows the host name, IP address and primary role of each host in the example architecture.

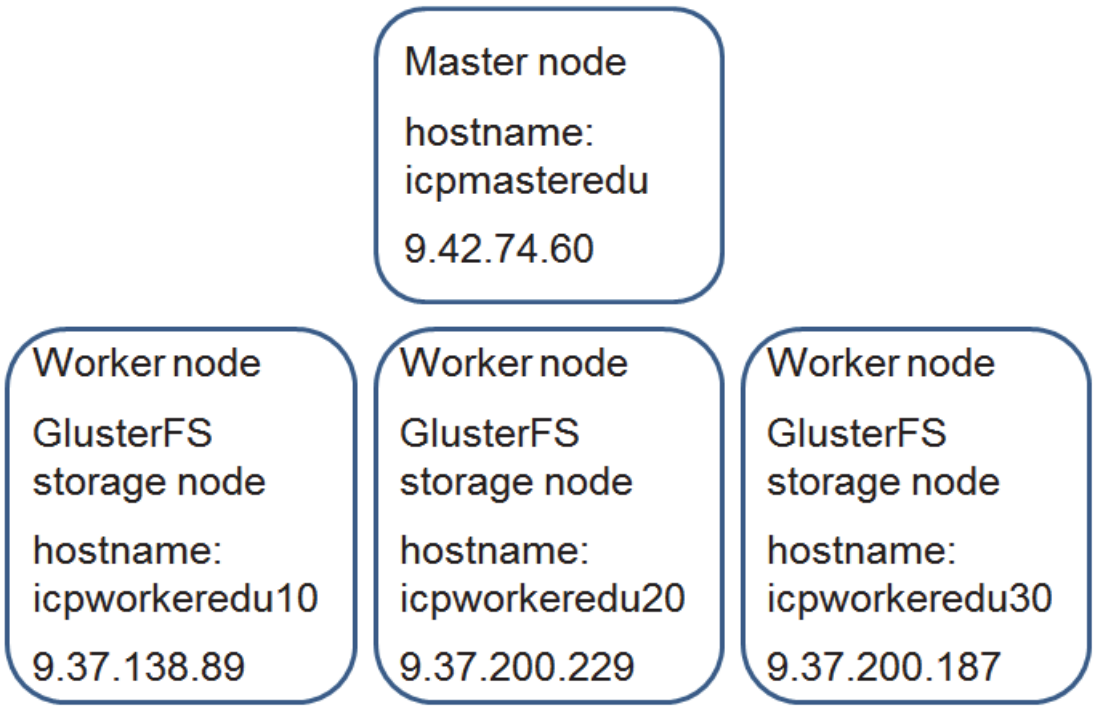

The following table lists details about each host in the example architecture. Throughout this document, you will find references to these host names and IP addresses. Of course, when you are installing NOI on ICP, you must use the details of your own environment, rather than the example values in this document.

Host	Resources	ICP Role	Operating system
icpmasteredu	8 CPUs 32 GB RAM 200 GB disk	• boot node • master node • proxy node • management node	Ubuntu 16.04 LTS
icpworkeredu10	8 CPUs 32 GB RAM 200 GB disk Extra 80 GB unused disk	• worker node • GlusterFS node	Ubuntu 16.04 LTS
icpworkeredu20	8 CPUs 32 GB RAM 200 GB disk Extra 80 GB unused disk	• worker node • GlusterFS node	Ubuntu 16.04 LTS
icpworkeredu30	8 CPUs 32 GB RAM 200 GB disk Extra 80 GB unused disk	• worker node • GlusterFS node	Ubuntu 16.04 LTS

Important: Four hosts, each with 8 CPUs, and 32 GB RAM are the minimum resources needed to install NOI on ICP.

In the example environment, all hosts have Internet access.

Notice the extra 80 GB unused disk on each of the worker nodes. These disks are required to install and configure GlusterFS storage. The storage device that you use for GlusterFS must be a raw disk. It must not be formatted, partitioned, or used for file system storage needs.

This guide assumes all hosts have a fresh installation of Ubuntu 16.04 LTS.

Note: The instructions in this document have been tested on Ubuntu 16.04 LTS. You can use many of the commands and tools on other operating systems. However, some of the commands used in this guide must be adjusted for your own environment, for example `apt-get` vs `yum`.

Preparing the hosts

In this section, you prepare your hosts to install IBM Cloud Private (ICP). These steps assume that your hosts have a fresh installation of Ubuntu 16.04 LTS. Some of these tasks might already be completed in your environment, such as installing curl, FTP, and so on. Review these steps carefully and make sure your environment meets all of the prerequisites in this section.

Important: Perform all of the steps in this section on every host in your environment.

1. As the Ubuntu default user, set the password for the root user and then switch to the root user account.
   a. Run the following command as the default Ubuntu user.

   ```
   sudo -i passwd root
   ```

   b. Enter the password for the default Ubuntu user, then enter the password you want for root twice. In this example, the default Ubuntu user is `ubuntu`.

   ```
   [sudo] password for ubuntu:
   Enter new UNIX password:
   Retype new UNIX password:
   ```

   c. Run the following command to switch to the root user account, then enter the root password that you just set.

   ```
   su - root
   Password:
   ```

Note: All further commands and actions in this section are run as the root user.

2. Install and configure vsftpd. This step is not mandatory, but it will make it easier to transfer files to your hosts.
   a. Run the following command to download the latest package lists and information in the default set of Ubuntu repositories.

   ```
   apt-get update
   ```

   b. Run the following command to install vsftpd.

   ```
   apt-get install vsftpd
   ```

   c. Open the vsftpd.conf file with a text editor. This example uses vi.

   ```
   vi /etc/vsftpd.conf
   ```

d. Find the following line and remove the comment character in front of the line. This change allows you to write files to the host. Save and close the file when you are finished.

```
write_enable=YES
```

e. Run the following commands to restart the vsftpd server, set it to start when the host boots, and verify that it is running.

```
systemctl stop vsftpd.service
systemctl start vsftpd.service
systemctl enable vsftpd.service
systemctl status vsftpd.service
```

3. Run the following command to install curl.

```
apt-get -y install curl
```

4. Edit your hosts file and add all of the hosts in your environment.
   a. Open the /etc/hosts file with a text editor. This example uses vi.

```
vi /etc/hosts
```

   b. Add the IP address, fully-qualified domain name, and an optional alias of all the hosts in your environment. Of course, the hosts in your environment will be different than the following example. Save and close the file when you are finished.

```
9.42.74.60      icpmasteredu.rtp.raleigh.ibm.com icpmasteredu
9.37.138.89     icpworkeredu10.rtp.raleigh.ibm.com icpworkeredu10
9.37.200.229    icpworkeredu20.rtp.raleigh.ibm.com icpworkeredu20
9.37.200.187    icpworkeredu30.rtp.raleigh.ibm.com icpworkeredu30
```

5. Disable the firewall on your hosts.
   a. If you are using Uncomplicated Firewall (ufw), which is the default firewall for Ubuntu, run the following commands to disable ufw and verify that it is inactive.

```
ufw disable
ufw status
```

   b. If you are using firewalld, run the following commands to stop the firewall, disable it from starting when the host boots, and verify that it is inactive.

```
systemctl stop firewalld
systemctl disable firewalld
systemctl status firewalld
```

6. Verify that python is installed. Install it if necessary.
   a. Run the following command to verify that python is installed and to show the version. The supported versions of python are v2.6, v2.7, or v3.5 or higher.

```
python --version
```

   b. Run the following command to install python, if necessary.

```
apt-get install -y python
```

7. Ubuntu 16.04 activates time synchronization by default. Verify that the clocks on each host are synchronized.
   a. Run the following command to display the synchronization status of your host.

   ```
   timedatectl status
   ```

   b. Look for the following output. This output verifies that the host's clock is synchronized.

   ```
   ...
   Network time on: yes
   NTP synchronized: yes
   ...
   ```

8. Run the following command to install Java.

   ```
   apt-get -y install default-jre
   ```

9. Increase the maximum map count kernel parameter.
   a. Run the following command to edit the /etc/sysctl.conf file.

   ```
   vi /etc/sysctl.conf
   ```

   b. Add the following line to the bottom of the file. Save and close the file when you are finished.

   ```
   vm.max_map_count=262144
   ```

   c. Run the following command to reload the sysctl parameters.

   ```
   sysctl -p
   ```

10. Install Docker Community Edition (CE).

Note: ICP v3.1.0 supports Docker CE versions 1.12 to 18.03.1 on Ubuntu 16.04. This guide shows you how to install Docker CE v18.03.1.

   a. Use the following command to download and add a Docker repository key.

   ```
   curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo apt-key add -
   ```

   b. Run the following command to add the Docker repository to your package sources. Run the command on a single line.

   ```
   add-apt-repository "deb [arch=amd64] https://download.docker.com/linux/ubuntu
   $(lsb_release -cs) stable"
   ```

   c. Run the following command to update your system package lists.

   ```
   apt-get update
   ```

   d. Verify that Docker server version 18.03.1-ce is available to install. Run the following command. Look for `18.03.1~ce-0~ubuntu` in the output of the command.

   ```
   apt-cache policy docker-ce
   ```

e. Run the following command to install Docker CE v18.03.1.

```
apt-get install -y docker-ce=18.03.1~ce-0~ubuntu
```

f. Run the following commands to start Docker and verify that it is running. Look for the message: `active (running).`

```
service docker start

service docker status
```

Press Ctrl+C to exit the status output.

g. Run the following command to verify that Docker starts when the host boots. Look for the plus symbol `[+]` next to docker.

```
service --status-all

...
[ + ]  docker
...
```

11. Install socat.

a. Download the socat installation archive from the following URL. In this example, the socat version is 2.0.0-b9, and the file name is socat-2.0.0-b9.tar.gz.

http://www.dest-unreach.org/socat/download/

b. Change to the directory where you downloaded socat. Decompress the installation archive.

```
tar -zxvf socat-2.0.0-b9.tar.gz
```

c. Change to the socat-2.0.0-b9/ sub-directory.

```
cd socat-2.0.0-b9/
```

d. Open the nestlex.c file with a text editor. This example uses vi.

```
vi nestlex.c
```

e. Add the line: `#include "stddef.h"` at the top of the #include section. Save and close the file when you are finished.

```
/* a function for lexical scanning of nested character patterns */
#include "stddef.h"
#include "config.h"
#include "mytypes.h"
```

f. Run the confgure script within the socat-2.0.0-b9/ sub-directory.

```
./configure
```

g. Run the following command to build the socat installation binary files. You can ignore all warning messages.

```
make
```

h. Run the following command to install socat.

```
make install
```

i. Verify that socat is installed with the following command.

```
socat -V
```

Preparing the GlusterFS nodes

In this example environment, the three worker nodes are also storage nodes for GlusterFS. The following instructions include the steps required to prepare the storage nodes. You do not install the GlusterFS cluster in this section; this section only shows you how to prepare the hosts and their disks. You install and configure the GlusterFS cluster later in this guide, when you deploy ICP.

Important: Each GlusterFS node must have a disk dedicated for GlusterFS storage. The storage device that you use for GlusterFS must be a raw disk. It must not be formatted, partitioned, or used for file system storage.

Run all of the commands and actions in this section as the root user.

Perform all of the steps in this section on your **worker nodes only**. These steps are not necessary on the master node.

1. Remove all data, partitions, and logical volumes from the disk that you want to use for GlusterFS storage.

   In this example, the target disk is named /dev/sdb. This example disk contains data, a partition, a logical volume, and is mounted to /data. The next steps show you how to completely clean the disk. The disk that you use for your GlusterFS node might not need to be prepared in exactly this way.

   a. Identify the disk you want to use for GlusterFS. Run the `lsblk` command to view all of your storage devices, partitions, and logical volumes. In this example, the target device is `/dev/sdb`. The example output shows the target device has a partition (`sdb1`) and a logical volume (`vg_data-lv_data`). The logical volume in this example is mounted to `/data`.

```
lsblk

NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
fd0                             2:0    1     4K  0 disk
sda                             8:0    0   200G  0 disk
├─sda1                          8:1    0   487M  0 part /boot
├─sda2                          8:2    0     1K  0 part
├─sda3                          8:3    0   120G  0 part
│ └─UB16--4--64SVR--vg-root   252:0    0 195.5G  0 lvm  /
└─sda5                          8:5    0  79.5G  0 part
  ├─UB16--4--64SVR--vg-root   252:0    0 195.5G  0 lvm  /
  └─UB16--4--64SVR--vg-swap_1 252:1    0     4G  0 lvm  [SWAP]
sdb                             8:16   0    80G  0 disk
└─sdb1                          8:17   0    80G  0 part
  └─vg_data-lv_data           252:2    0    80G  0 lvm  /data
sr0                            11:0    1  1024M  0 rom
```

Note: In this example, the target disk has a partition, a logical volume, and a mount point. The next steps show you how to remove these objects. These steps might not be necessary in your environment, depending on your storage configuration.

b. If your target disk is mounted, run a command like the following example to unmount the disk. In this example, the disk is mounted to `/data`.

```
umount /data
```

c. In the preceding example, the `/dev/sdb` device has a partition named `sdb1` and a logical volume named `vg_data-lv_data` that must be removed. Run the following command to find the volume group that corresponds to the unwanted logical volume. In this example the unwanted volume group is named `vg_data`, which corresponds to the logical volume named `vg_data-lv_data`.

```
vgdisplay

--- Volume group ---
  VG Name               vg_data
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  2
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                1
  Open LV               0
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               80.00 GiB
  PE Size               4.00 MiB
  Total PE              20479
  Alloc PE / Size       20479 / 80.00 GiB
  Free  PE / Size       0 / 0
  VG UUID               xpqxnb-o53v-cXUU-YYwy-Ch85-FH3F-nn9Q1i
```

d. Run the following command to remove the volume group. Enter y when you are prompted. In this example, the unwanted volume group is named `vg_data`. You can ignore any errors about failure to write.

```
vgremove vg_data
```

e. Run the following command to verify that the unwanted volume group is no longer present.

```
vgdisplay
```

f. Run the following command to remove all signatures, including partition information, from the device. In this example, the device is `/dev/sdb`.

```
wipefs --all --force /dev/sdb
```

g. Run the following command to overwrite all partitions, data, and boot records on your target disk with zeros. In this example, the device is `/dev/sdb`. You can ignore any errors about no space left on the device. This command runs for several minutes, depending on the size of your target disk and your system resources.

```
dd if=/dev/zero of=/dev/sdb bs=10M
```

h. Run the `lsblk` command again to verify that there are no partitions or logical volumes associated with your target disk. In this example, the device is `/dev/sdb`.

```
lsblk
```

```
NAME                          MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
fd0                             2:0    1     4K  0 disk
sda                             8:0    0   200G  0 disk
├─sda1                          8:1    0   487M  0 part
├─sda3                          8:3    0   120G  0 part
│ └─UB16--4--64SVR--vg-root   252:0    0 195.5G  0 lvm  /
└─sda5                          8:5    0  79.5G  0 part
  ├─UB16--4--64SVR--vg-root   252:0    0 195.5G  0 lvm  /
  └─UB16--4--64SVR--vg-swap_1 252:1    0     4G  0 lvm  [SWAP]
sdb                             8:16   0    80G  0 disk
sr0                            11:0    1  1024M  0 rom
```

2. If your host automatically mounts the target disk, remove the corresponding entry in the /etc/fstab file.

a. Open the /etc/fstab file with a text editor. This example uses vi.

```
vi /etc/fstab
```

b. Find the line that corresponds to the unwanted disk, logical volume, or mount point. In this example, the logical volume named `vg_data-lv_data` is mounted to `/data`. Remove the line, or exclude it with the comment character.

```
# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point>   <type>  <options>       <dump>  <pass>
/dev/mapper/UB16--4--64SVR--vg-root /           ext4    errors=remount-ro 0
1
# /boot was on /dev/sda1 during installation
UUID=90342336-b5ce-426d-a1cc-cb1f29cf8c41 /boot          ext2    defaults
0       2
/dev/mapper/UB16--4--64SVR--vg-swap_1 none          swap    sw              0
0
/dev/fd0        /media/floppy0  auto    rw,user,noauto,exec,utf8 0       0
#/dev/mapper/vg_data-lv_data  /data   ext4   defaults 0 0
```

c. Save and close the file when you are finished.

3. Use the following command to reboot your storage nodes.

```
init 6
```

4. After your worker nodes have restarted, run the following command to verify that no file system is mounted to your target disk. The output of this command should not contain any reference to your target device or its former mount point.

```
cat /proc/mounts
```

5. Run the following command to find the symlink and link path of your target disk device. In this example, the target device is `/dev/sdb`, and the symlink that points to `../../sdb` is `pci-0000:03:00.0-scsi-0:0:1:0`. The link path is `/dev/disk/by-path`.

   ```
   ls -altr /dev/disk/*
   ```

   The symlink of your device can be in any of these link paths:

   - /dev/disk/by-path
   - /dev/disk/by-id
   - /dev/disk/by-uuid
   - /dev/disk/by-label

   /dev/disk/by-path:

   ```
   total 0
   lrwxrwxrwx 1 root root   9 Oct 16 13:50 pci-0000:00:07.1-ata-2 -> ../../sr0
   lrwxrwxrwx 1 root root   9 Oct 16 13:50 pci-0000:03:00.0-scsi-0:0:0:0 ->
   ../../sda
   lrwxrwxrwx 1 root root   9 Oct 16 13:50 pci-0000:03:00.0-scsi-0:0:1:0 ->
   ../../sdb
   lrwxrwxrwx 1 root root  10 Oct 16 13:50 pci-0000:03:00.0-scsi-0:0:0:0-part1 ->
   ../../sda1

   /dev/disk/by-uuid:
   total 0
   lrwxrwxrwx 1 root root  10 Oct 16 13:50 c6ad8ec9-dcdf-446a-9ef3-bea012e43b80 ->
   ../../dm-0
   lrwxrwxrwx 1 root root  10 Oct 16 13:50 90342336-b5ce-426d-a1cc-cb1f29cf8c41 ->
   ../../sda1

   /dev/disk/by-id:
   total 0
   lrwxrwxrwx 1 root root   9 Oct 16 13:50
   ata-VMware_Virtual_IDE_CDROM_Drive_10000000000000000001 -> ../../sr0
   lrwxrwxrwx 1 root root  10 Oct 16 13:50 dm-name-UB16--4--64SVR--vg-root ->
   ../../dm-0
   lrwxrwxrwx 1 root root  10 Oct 16 13:50
   lvm-pv-uuid-K0gVeN-ZbhN-uMGc-Npqo-Dkfx-xJce-ZF1Mrm -> ../../sda5
   ```

Important: Make a note of the link path and symlink for the target disk. Record the link path and symlink for each of your storage nodes. You use them in the next section when you deploy ICP.

6. If your symlink contains the colon character (:), your GlusterFS deployment might fail. Also, some operating systems do not generate a symlink. If your environment has a colon character (:) in the symlink, or no symlink for your device, follow the next steps.
   a. Run the following command to query for information about your target device. In this example, the target device is `/dev/sdb`.

```
udevadm info --root --name=/dev/sdb
```

   b. In the output of the preceding command, gather the following three values:
      - DEVTYPE
      - SUBSYSTEM
      - DEVPATH

   In the example environment, the output of the preceding command is:

```
P:
/devices/pci0000:00/0000:00:15.0/0000:03:00.0/host2/target2:0:1/2:0:1:0/block/s
db
N: sdb
S: disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
E: DEVLINKS=/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
E: DEVNAME=/dev/sdb
E:
DEVPATH=/devices/pci0000:00/0000:00:15.0/0000:03:00.0/host2/target2:0:1/2:0:1:0
/block/sdb
E: DEVTYPE=disk
E: ID_BUS=scsi
E: ID_MODEL=Virtual_disk
E: ID_MODEL_ENC=Virtual\x20disk\x20\x20\x20\x20
E: ID_PATH=pci-0000:03:00.0-scsi-0:0:1:0
E: ID_PATH_TAG=pci-0000_03_00_0-scsi-0_0_1_0
E: ID_REVISION=1.0
E: ID_SCSI=1
E: ID_TYPE=disk
E: ID_VENDOR=VMware
E: ID_VENDOR_ENC=VMware\x20\x20
E: MAJOR=8
E: MINOR=16
E: SUBSYSTEM=block
E: TAGS=:systemd:
E: USEC_INITIALIZED=1812347
```

   c. Create a custom udev rules file. In this example, the rules files is named **10-custom-icp.rules**.

```
vi /lib/udev/rules.d/10-custom-icp.rules
```

d. Add a single line to your rules file similar to the following example. Of course, replace the DEVTYPE, SUBSYSTEM, and DEVPATH values with actual values from your environment. Save and close the file when you are finished.

```
ENV{DEVTYPE}=="<your_devtype>", ENV{SUBSYSTEM}=="<your_subsystem>",
ENV{DEVPATH}=="<your_devpath>" SYMLINK+="disk/gluster-disk-1"
```

For example:

```
ENV{DEVTYPE}=="disk", ENV{SUBSYSTEM}=="block",
ENV{DEVPATH}=="/devices/pci0000:00/0000:00:15.0/0000:03:00.0/host2/target2:0:1/
2:0:1:0/block/sdb" SYMLINK+="disk/gluster-disk-1"
```

e. Run the following two commands to reload the udev rules.

```
udevadm control --reload-rules
udevadm trigger --type=devices --action=change
```

f. Run the following command to find the new symlink of your target disk device. In this example, the target device is `/dev/sdb`, and the new symlink that points to `../sdb` is `/dev/disk/gluster-disk-1`.

```
ls -altr /dev/disk/*

lrwxrwxrwx 1 root root   6 Oct 16 15:25 /dev/disk/gluster-disk-1 -> ../sdb

/dev/disk/by-path:
total 0
drwxr-xr-x 2 root root 180 Oct 16 13:50 .
lrwxrwxrwx 1 root root  10 Oct 16 15:25 pci-0000:03:00.0-scsi-0:0:0:0-part5 ->
../../sda5
lrwxrwxrwx 1 root root  10 Oct 16 15:25 pci-0000:03:00.0-scsi-0:0:0:0-part3 ->
../../sda3
...
```

g. Make a note of the symlink you found with the preceding command. Record the symlink for each of your storage nodes. You use them in the next section when you deploy ICP.

7. Load the dm_thin_pool kernel module and make it persist after a system restart.

a. Run the following command to configure the dm_thin_pool kernel module.

```
modprobe dm_thin_pool
```

b. Run the following command to make the module persist after a system restart.

```
echo dm_thin_pool | sudo tee -a /etc/modules
```

c. Run the following command and verify that the dm_thin_pool kernel module is listed.

```
cat /etc/modules
```

8. Install GlusterFS Client.

Note: ICP v3.1.0 supports GlusterFS version 4.0.2. The next steps show you how to install GlusterFS Client v4.0.2.

a. Run the following command to install the Ubuntu package: software-properties-common.

```
apt-get install software-properties-common
```

b. Run the following command to add the GlusterFS 4.0 repository to your package sources. Press **Enter** to confirm.

```
add-apt-repository ppa:gluster/glusterfs-4.0
```

c. Run the following command to update your system package lists.

```
apt-get update
```

d. Verify that GlusterFS Client v4.0.2 is available to install. Run the following command. Look for `4.0.2-ubuntu1~xenial1` in the output of the command.

```
apt-cache policy glusterfs-client
```

e. Run the following command to install GlusterFS Client v4.0.2.

```
apt-get install -y glusterfs-client
```

f. Use the following command to verify that GlusterFS Client was successfully installed.

```
glusterfs --version
```

Installing IBM Cloud Private

In this section, you customize your IBM Cloud Private (ICP) environment by editing configuration files. You then deploy IBM Cloud Private.

Important: Run all of the commands and actions in this section as the root user.

Perform all of the steps in this section on your master node only. These steps are not necessary on the worker nodes.

1. Generate an SSH key pair on your master node and share that key with your worker nodes.
   a. Run the following command to generate an SSH key.

```
ssh-keygen -b 4096 -f ~/.ssh/id_rsa -N ""
```

   b. Use the following command to add the public key to the authorized keys file.

```
cat ~/.ssh/id_rsa.pub | sudo tee -a ~/.ssh/authorized_keys
```

   c. Copy the public key to each worker node in your environment. Of course, the hosts and IP addresses in your environment will be different than the following examples. Enter **yes** and enter the root password of the remote host when you are prompted.

```
ssh-copy-id -i ~/.ssh/id_rsa.pub root@9.37.138.89

ssh-copy-id -i ~/.ssh/id_rsa.pub root@9.37.200.229

ssh-copy-id -i ~/.ssh/id_rsa.pub root@9.37.200.187
```

   d. From your master note, connect to each of your hosts as the root user with SSH, including the master node. Verify that you are not prompted for a password. Log out of each host after you verify it. Of course, the hosts and IP addresses in your environment will be different than the following examples.

```
ssh root@9.42.74.60
exit

ssh root@9.37.138.89
exit

ssh root@9.37.200.229
exit

ssh root@9.37.200.187
exit
```

2. Run the following command to pull the IBM Cloud Private-CE installer image from Docker Hub.

```
docker pull ibmcom/icp-inception:3.1.0
```

3. Run the following commands to create a directory for ICP, then change to the new directory. In this example, ICP will be installed into the **/opt/ibm-cloud-private-ce-3.1.0** directory.

```
mkdir /opt/ibm-cloud-private-ce-3.1.0

cd /opt/ibm-cloud-private-ce-3.1.0
```

4. Run the following command on a single line to extract the ICP configuration files. This action creates a sub-directory named **cluster**.

```
docker run -e LICENSE=accept -v "$(pwd)":/data ibmcom/icp-inception:3.1.0 cp -r cluster /data
```

5. Change to the sub-directory named **cluster**.

```
cd /opt/ibm-cloud-private-ce-3.1.0/cluster
```

6. Run the following command to copy your SSH private key file to the sub-directory named **cluster**.

```
cp ~/.ssh/id_rsa ssh_key
```

7. Customize your **cluster/hosts** file.
   a. Run the following commands to back up the original hosts file.

```
cd /opt/ibm-cloud-private-ce-3.1.0/cluster

cp hosts hosts.BAK
```

   b. Edit the **cluster/hosts** file and change it to look like the following example. Replace the example IP addresses with the IP addresses in your environment. Keep the file open when you are finished.

   Notice the following changes in this example:

   - The comment character has been removed from the **[management]** section, because we want to install management services.
   - The master node, the proxy node, and management node are all the same host.
   - Each of the three worker nodes are listed in the **[worker]** section.
   - The **[va]** section is excluded by comment character.

```
vi /opt/ibm-cloud-private-ce-3.1.0/cluster/hosts

[master]
9.42.74.60

[worker]
9.37.138.89
9.37.200.229
9.37.200.187

[proxy]
9.42.74.60

[management]
9.42.74.60

#[va]
#5.5.5.5
```

c. Add lines like the following example to the bottom of your **cluster/hosts** file. These lines define a custom host group named **glusterfs**. This host group identifies your worker nodes as GlusterFS storage nodes. Replace the example IP addresses with the IP addresses of the worker nodes in your environment. Save and close the file when you are finished.

```
[hostgroup-glusterfs]
9.37.138.89
9.37.200.229
9.37.200.187
```

When you are finished, your **cluster/hosts** file should look like the following example:

```
[master]
9.42.74.60

[worker]
9.37.138.89
9.37.200.229
9.37.200.187

[proxy]
9.42.74.60

[management]
9.42.74.60

#[va]
#5.5.5.5

[hostgroup-glusterfs]
9.37.138.89
9.37.200.229
9.37.200.187
```

8. Customize your **cluster/config.yaml** file.
   a. Run the following commands to back up the original config.yaml file.

```
cd /opt/ibm-cloud-private-ce-3.1.0/cluster

cp config.yaml config.yaml.BAK
```

   b. Open the config.yaml file with a text editor.

Important: The indentations in the config.yaml file are important. If the lines in your config.yaml are not indented correctly, your ICP installation might fail. If you copy text from this document and paste it to your config.yaml file, be sure to insert the proper indentations. Use spaces to indent lines in your config.yaml file, not tabs.

c. Find the **management_services** section. Edit your management_services section to look like the following example. Keep the file open when you are finished.

NOTE: The services that have been disabled in this example are not needed for NOI. If you are deploying other applications into your ICP environment, you might need to enable some of these services.

```
management_services:
  istio: disabled
  vulnerability-advisor: disabled
  storage-glusterfs: enabled
  storage-minio: disabled
```

d. Find and uncomment the following line. Remove all spaces in front of the line.

```
metrics_max_age: 1
```

e. Find and uncomment the following line. Remove all spaces in front of the line.

```
logs_maxage: 1
```

f. Add the following lines to the bottom of your **cluster/config.yaml** file. These lines allow applications to be deployed from the docker.io/ibmcom and the local cluster registries.

```
image-security-enforcement:
  clusterImagePolicy:
    - name: "docker.io/ibmcom/*"
      policy:
    - name: "mycluster.icp:8500/*"
      policy:
```

g. Add the following line to the bottom of your **cluster/config.yaml** file. This line allows you to use the worker nodes as GlusterFS nodes.

```
no_taint_group: ["hostgroup-glusterfs"]
```

h. Add the following lines to the bottom of your **cluster/config.yaml** file. These lines configure a GlusterFS cluster named **glusterfs**. Add the IP addresses of your worker nodes in the **nodes** section. Add the link path and symlink that you found earlier in this document to the **devices** sections.

```
storage-glusterfs:
  nodes:
    - ip: <YOUR_FIRST_WORKER_IP>
      devices:
        - <link path>/<symlink of your storage device>
    - ip: <YOUR_SECOND_WORKER_IP>
      devices:
        - <link path>/<symlink of your storage device>
    - ip: <YOUR_THIRD_WORKER_IP>
      devices:
        - <link path>/<symlink of your storage device>
  storageClass:
    create: true
    name: glusterfs
    isDefault: true
    volumeType: replicate:3
    reclaimPolicy: Delete
    volumeBindingMode: Immediate
    volumeNamePrefix: icp
    additionalProvisionerParams: {}
    allowVolumeExpansion: true
  gluster:
    resources:
      requests:
        cpu: 500m
        memory: 512Mi
      limits:
        cpu: 1000m
        memory: 1Gi
  heketi:
    backupDbSecret: heketi-db-backup
    authSecret: "heketi-secret"
    resources:
      requests:
        cpu: 500m
        memory: 512Mi
      limits:
        cpu: 1000m
        memory: 1Gi
  prometheus:
    enabled: false
    path: "/metrics"
    port: 8080
  nodeSelector:
    key: hostgroup
    value: glusterfs
  podPriorityClass: "system-cluster-critical"
  tolerations: []
```

i. Save and close the **cluster/config.yaml** file when you are finished.

9. Install IBM Cloud Private (ICP).
   a. Change to the **cluster** sub-directory.

   ```
   cd /opt/ibm-cloud-private-ce-3.1.0/cluster
   ```

   b. Run the following command on a single line to install ICP. The installation process runs for 30 to 60 minutes, depending on your system resources.

   ```
   docker run --net=host -t -e LICENSE=accept -v "$(pwd)":/installer/cluster
   ibmcom/icp-inception:3.1.0 install
   ```

   c. After the installation is finished, you see a message like the following example. Record the Dashboard URL and the default user name and password.

   ```
   The Dashboard URL: https://9.42.74.60:8443, default username/password is
   admin/admin

   Playbook run took 0 days, 0 hours, 28 minutes, 57 seconds
   ```

IBM Cloud Private post-installation tasks

In this section, you install the following tools:

- The Kubernetes command line tool
- The IBM Cloud Private command line interface
- Helm, which is a Kubernetes package manager

Important: Run all of the commands and actions in this section as the root user.

Perform all of the steps in this section on your master node only. These steps are not necessary on the worker nodes.

1. Install and configure kubectl, which is the Kubernetes command line tool.
   a. Run the following command on a single line to install kubectl.

   ```
   docker run -e LICENSE=accept --net=host -v /usr/local/bin:/data
   ibmcom/icp-inception-amd64:3.1.0 cp /usr/local/bin/kubectl /data
   ```

   b. Open a browser and go to the ICP Dashboard URL. You obtained the URL at the end of the ICP installation process. In this example, the URL is https://9.42.74.60:8443, where the IP address is the master node.

   c. If you are prompted with a warning about an insecure or non-private connection click **advanced**, add a security exception to your browser, and proceed to the web site.

d. Log in with the user name and password you obtained at the end of the ICP installation process. The default user name is **admin**. The default password is **admin**.

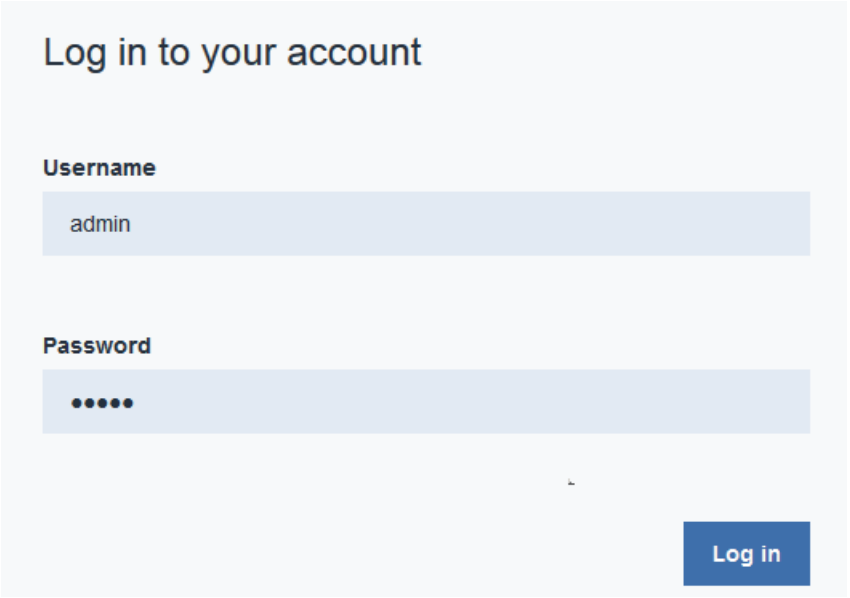

e. At the top right of the page, click the head-and-shoulders icon, then click **Configure client**.

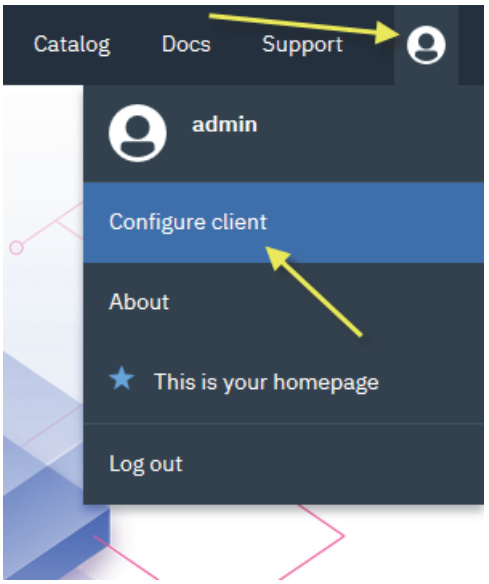

f. Copy, paste, and the five commands in the Configure client window, one at a time, on your master node.

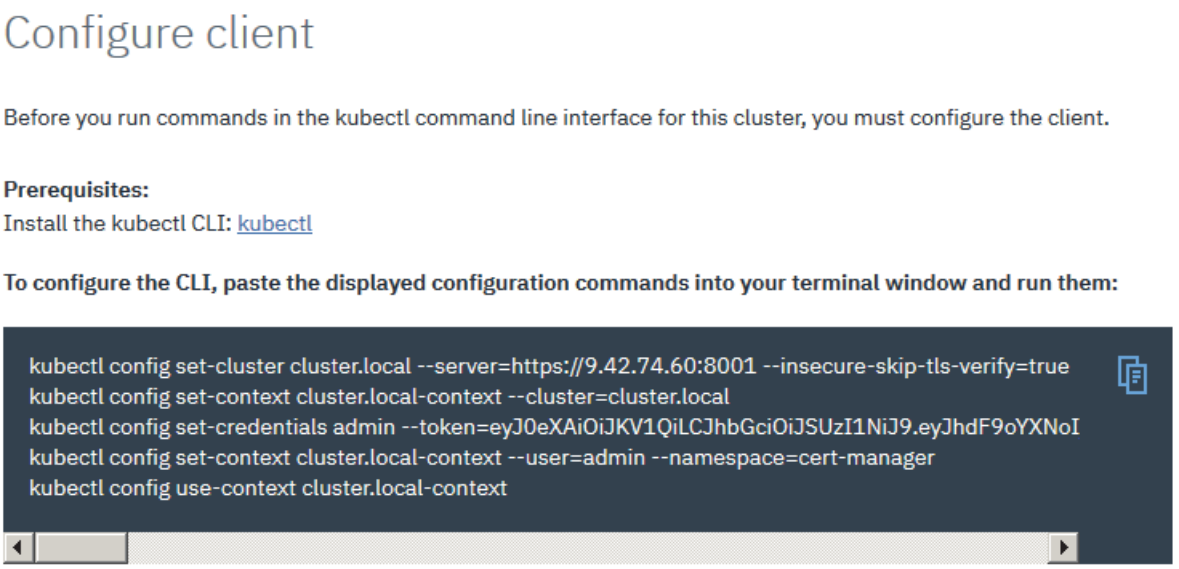

Note: The configuration provided by these commands expires in 12 hours. To continue to use the CLI, you must log in and reconfigure kubectl every 12 hours.

g. Run the following command to verify that kubectl is installed.

```
kubectl version
```

2. Install the IBM Cloud Private command line interface.
   a. Change to the directory where you want to download the IBM Cloud Private command line tool, for example, your Downloads sub-directory.

```
cd ~/Downloads/
```

   b. Return to the ICP Dashboard URL.

c. Click **Menu > Command Line Tools > Cloud Private CLI**.

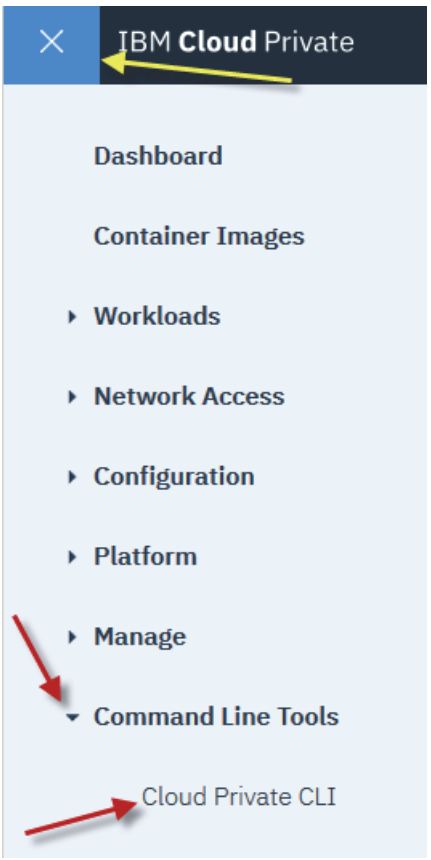

d. Expand **Install IBM Cloud Private CLI and plug-ins**.

e. Choose your operating system. This example uses **Linux (64-bit)**.

f. Copy, paste, and run the command on your master node.

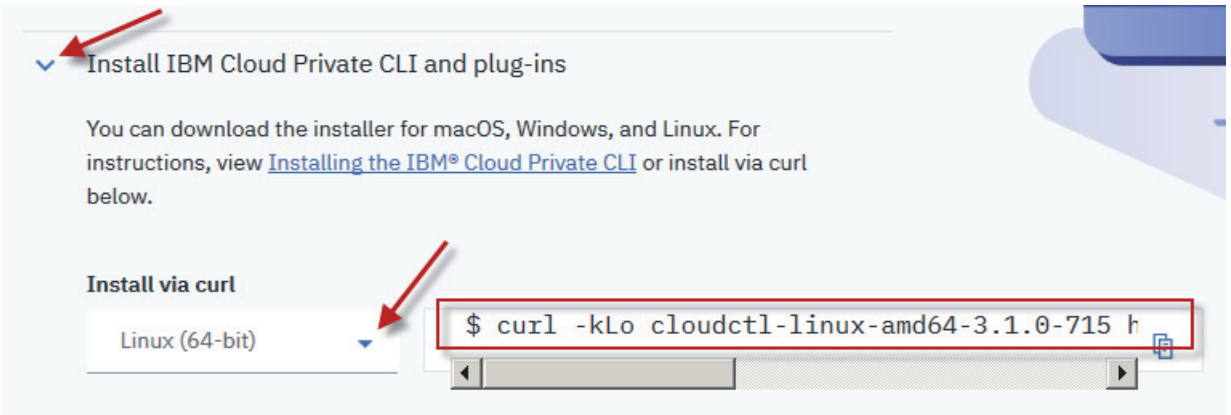

g. Run the following command to relax permissions on the IBM Cloud Private command line tool and make it executable.

```
chmod 755 cloudctl-linux-amd64-3.1.0-715
```

h. Run the following command to move the IBM Cloud Private command line tool to the **/usr/local/bin/** directory and rename it **cloudctl**.

```
mv cloudctl-linux-amd64-3.1.0-715 /usr/local/bin/cloudctl
```

i. Run the following command to verify that the IBM Cloud Private command line tool is installed.

```
cloudctl version
```

j. Use the following command to log in to your cluster. Replace the IP address in the example with the IP address of your master node.

```
cloudctl login -a https://9.42.74.60:8443 --skip-ssl-validation
```

k. Log in with the user name and password you obtained at the end of the ICP installation process. The default user name is **admin**. The default password is **admin**.

```
Username> admin

Password> admin
```

l. Enter **1** to select your account named **mycluster**.

```
Select an account:
1. mycluster Account (id-mycluster-account)
Enter a number> 1
```

m. Enter the number of your **default** namespace. In this example, the default namespace is 2.

```
Select a namespace:
1. cert-manager
2. default
3. istio-system
4. kube-public
5. kube-system
6. platform
7. services
Enter a number> 2
```

3. Install Helm, which is a Kubernetes package manager.
   a. Change to the directory where you want to download the Helm installer, for example, your Downloads sub-directory.

```
cd ~/Downloads/
```

   b. Return to the ICP Dashboard URL.

c. Click **Menu > Command Line Tools > Cloud Private CLI**.

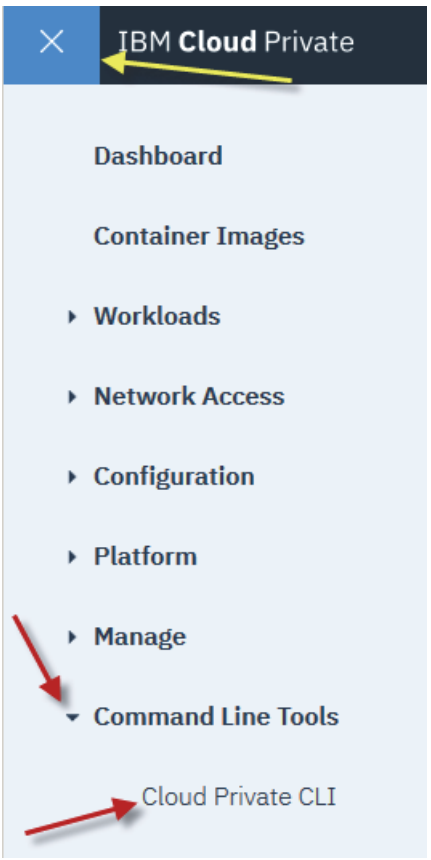

d. Expand **Install Helm**.

e. Choose your operating system. This example uses **Linux (64-bit)**.

f. Copy, paste, and run the command on your master node.

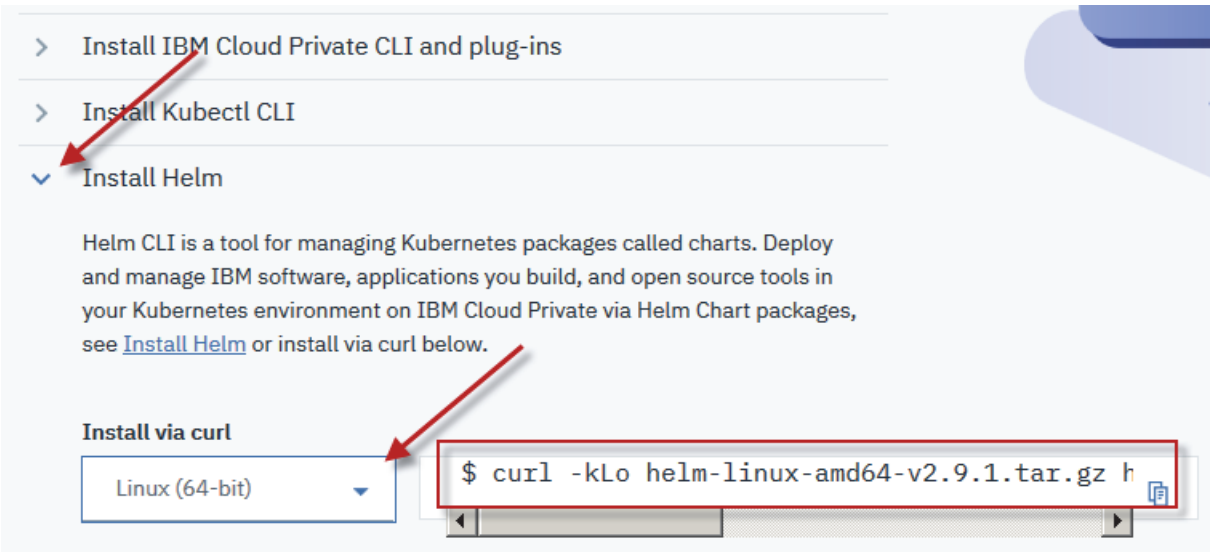

g. Run the following command to set the HELM_HOME environment variable.

```
export HELM_HOME=~/.helm
```

h. Run the following command to initialize the Helm command line interface.

```
helm init --client-only
```

i. Run the following command to verify that Helm is installed.

```
helm version --tls
```

Installing IBM Netcool Operations Insight

Important: Run all of the commands and actions in this section as the root user.

Perform all of the steps in this section on your master node only. These steps are not necessary on the worker nodes.

1. The Netcool Operations Insight installer expects the default storage class to be GlusterFS. Check your default storage class and set it to GlusterFS if necessary.

   a. Run the following command to verify that GlusterFS is set as the default storage class.

   ```
   kubectl get storageclass
   ```

   ```
   NAME                        PROVISIONER                     AGE
   glusterfs (default)         kubernetes.io/glusterfs         22h
   image-manager-storage       kubernetes.io/no-provisioner    22h
   kafka-storage               kubernetes.io/no-provisioner    22h
   logging-storage-datanode    kubernetes.io/no-provisioner    22h
   mariadb-storage             kubernetes.io/no-provisioner    22h
   minio-storage               kubernetes.io/no-provisioner    22h
   mongodb-storage             kubernetes.io/no-provisioner    22h
   zookeeper-storage           kubernetes.io/no-provisioner    22h
   ```

   b. If GlusterFS is not set as the default storage class, run a command like the following example. Run the comand on a single line. In this example, the name of the storage class is **glusterfs**.

   ```
   kubectl patch storageclass glusterfs -p '{"metadata":
   {"annotations":{"storageclass.kubernetes.io/is-default-class":"true"}}}'
   ```

Hint: Before you installed ICP, you set GlusterFS as the default storage class in your config.yaml file.

2. Download the IBM Netcool Operations Insight installation archive file to your master node. This example uses the October 9, 2018 release (part number CNX04EN - IBM Netcool Operations Insight 1.5 Operations Management for IBM Cloud Private English). You must obtain this software yourself from the Software Seller's download site, IBM PartnerWorld, Passport Advantage, or wherever you usually download IBM software.
3. If you have not configured kubectl in the last 12 hours, configure it again.
   a. Open a browser and go to the ICP Dashboard URL.
   b. Log in with the user name and password you obtained at the end of the ICP installation process. The default user name is **admin**. The default password is **admin**.
   c. At the top right of the page, click the head-and-shoulders icon, then click **Configure client**.
   d. Copy, paste, and run the five commands in the Configure client window to your master node.
4. Log in to your cluster.
   a. Use the following command to log in to your cluster. Replace the IP address in the example with the IP address of your master node.

   ```
   cloudctl login -a https://9.42.74.60:8443 --skip-ssl-validation
   ```

   b. Log in with the user name and password you obtained at the end of the ICP installation process. The default user name is **admin**. The default password is **admin**.
   c. Enter the number of your account named **mycluster**.
   d. Enter the number of your **default** namespace.
5. Run the following command to view the current namespace target. Confirm that the target is **default**.

   ```
   cloudctl target

   Namespace:   default
   ```

6. Run the following command to configure the kubectl tool to use your cluster. By default, your cluster is named mycluster.

   ```
   kubectl config set-context mycluster
   ```

7. Use the following command to log in to your local docker repository. This repository was created when you installed ICP. The default user name is **admin**. The default password is **admin**.

   ```
   docker login mycluster.icp:8500
   ```

8. Run the following two commands, one at a time, to unset the proxy connection.

```
unset http_proxy
unset HTTP_PROXY
```

9. Run the following command to load the IBM Netcool Operations Insight installation archive into ICP. In this example, the installation archive file was downloaded to the /home/ibmadmin/Downloads/ directory. This command runs for 30 to 40 minutes, depending on your system resources.

```
cloudctl catalog load-archive --archive
/home/ibmadmin/Downloads/NOI_V1.5_OM_FOR_ICP.tar.gz
```

10. Run the following command create a password for the Netcool/OMNIbus root user. In this example, the password is **netcool**.

```
kubectl create secret generic omni-secret
--from-literal=OMNIBUS_ROOT_PASSWORD=netcool
```

11. Install IBM Netcool Operations Insight.
    a. Open a browser and go to the ICP Dashboard URL, if it is not already open.
    b. Click **Catalog** at the top right of the page.

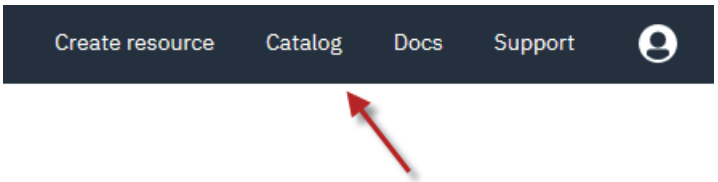

    c. Enter **ibm-netcool** as the filter.
    d. Click the **ibm-netcool-prod** chart.

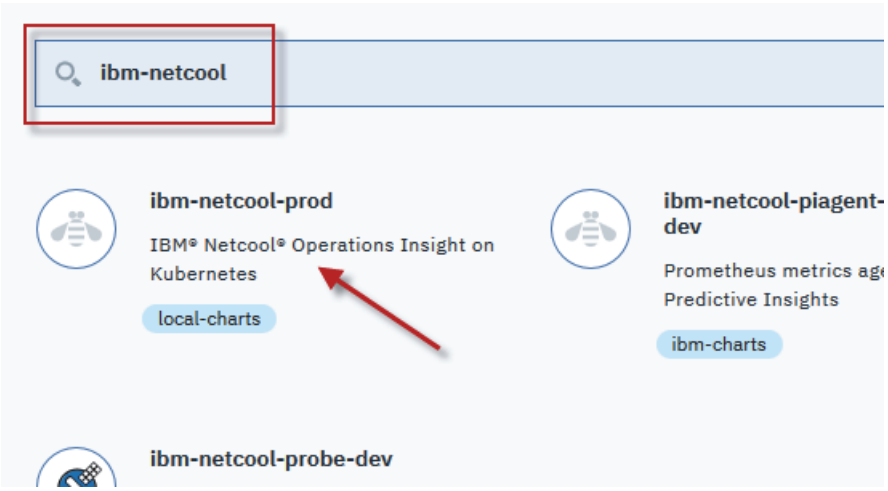

    e. Click the **Configure** button at the bottom right of the page.

f. Enter a name for your NOI environment in the **Helm release name** field. After the installation is finished, the name you enter here becomes part of the URLs that you use to access the NOI user interfaces.

g. Choose **default** as the **Target namespace**.

h. Read and accept the license agreement.

i. Expand **All Parameters**.

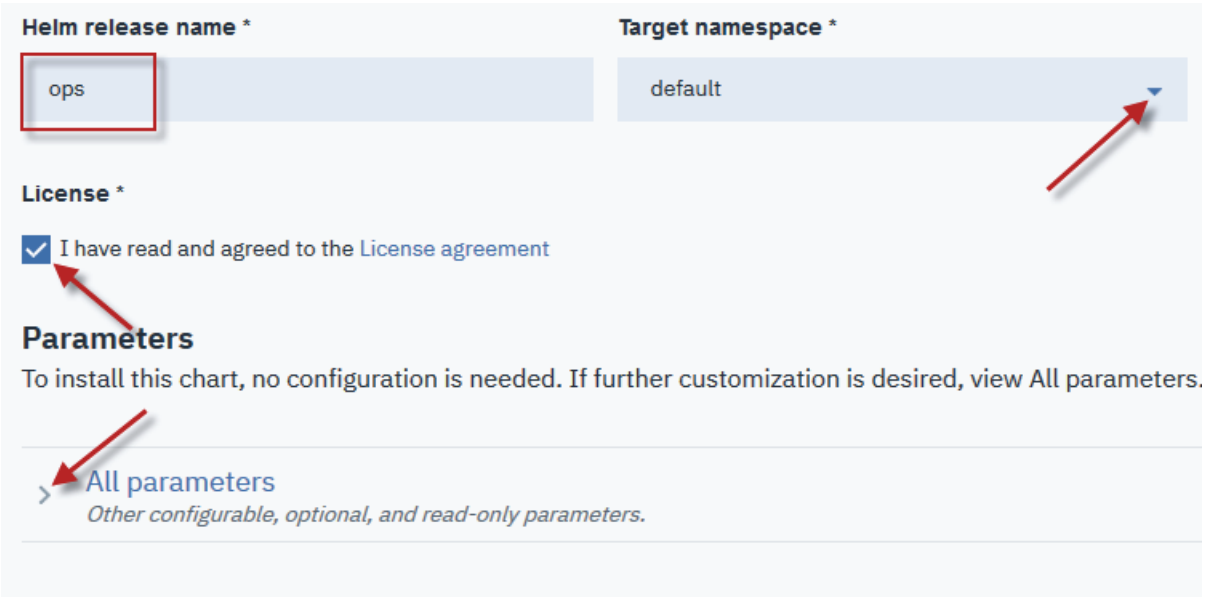

j. Accept the license terms.

k. Enter the fully-qualified domain name of your master node.

l. Deselect **Enable sub-chart resource requests**.

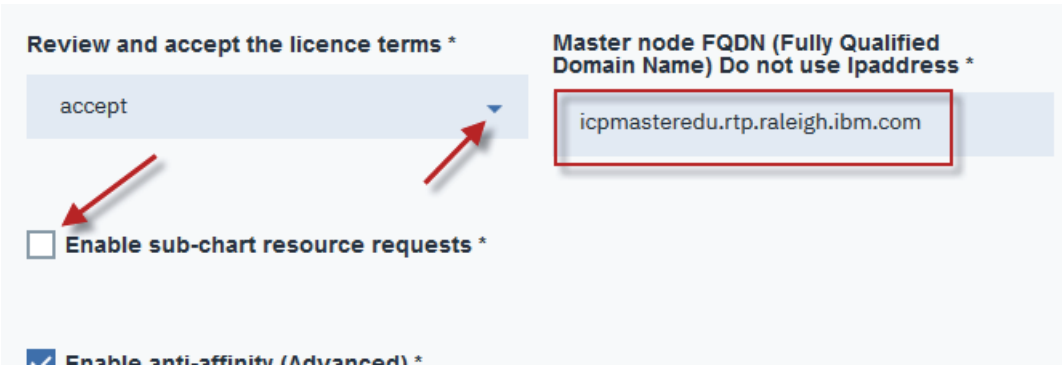

m. Notice that the **LDAP mode** is standalone by default. This means that OpenLDAP will be installed with NOI, and all NOI components will use the dedicated LDAP server. In this example, a standalone LDAP server is used.

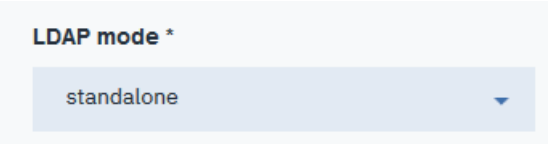

n. Notice the **ASM release name**. If you are planning to install Agile Service Manager (ASM) in the same ICP instance, you must use the same release name that you specify in this field.

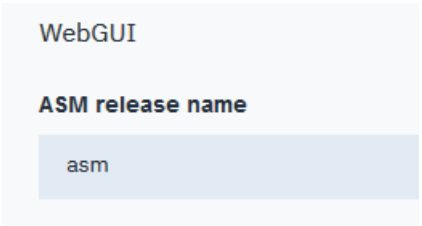

o. Click **Install** at the bottom right of the page.

12. Monitor your installation.
    a. Run the following command to show the status of the NOI pods. The READY column shows the number of pods that will be provisioned on the right, and the number of pods that are finished provisioning on the left.

```
kubectl get pods
```

```
NAME                                READY     STATUS              RESTARTS   AGE
ops-db2ese-8dcbb4dfd-slxv8          0/1       ContainerCreating   0          2m
ops-impactgui-6d68867cc5-thsch      0/1       Init:0/1            0          2m
ops-jdbcgw-9bdc669cc-x4zrd          0/1       Init:0/1            0          2m
ops-ncibackup-5f4c95bfcd-kn6h6      0/1       Init:0/2            0          2m
ops-nciprimary-6965b459d5-chbrf     0/1       Init:0/1            0          2m
ops-ncobackup-776cf468c-mnplb       2/2       Running             0          2m
ops-ncoprimary-686c956546-t5qqv     1/1       Running             0          2m
ops-openldap-7dcd776ccd-n7bgd       1/1       Running             0          2m
ops-scala-bcc96d767-tswvk           0/2       Init:0/1            0          2m
ops-webgui-68f4f89c47-nxfjh         0/1       Init:0/1            0          2m
```

b. After 15 to 20 minutes, run the command again. When all pods are running and ready, the installation has finished.

```
kubectl get pods
```

```
NAME                                READY     STATUS    RESTARTS   AGE
ops-db2ese-8dcbb4dfd-slxv8          1/1       Running   0          15m
ops-impactgui-6d68867cc5-thsch      1/1       Running   0          15m
ops-jdbcgw-9bdc669cc-x4zrd          1/1       Running   0          15m
ops-ncibackup-5f4c95bfcd-kn6h6      1/1       Running   0          15m
ops-nciprimary-6965b459d5-chbrf     1/1       Running   0          15m
ops-ncobackup-776cf468c-mnplb       2/2       Running   5          15m
ops-ncoprimary-686c956546-t5qqv     1/1       Running   1          15m
ops-openldap-7dcd776ccd-n7bgd       1/1       Running   0          15m
ops-scala-bcc96d767-tswvk           2/2       Running   0          15m
ops-webgui-68f4f89c47-nxfjh         1/1       Running   0          15m
```

13. If any of your pods cannot provision correctly, use the `kubectl describe` tool to troubleshoot the pod.

a. Run the following command to show the name of your pods.

```
kubectl get pods
```

```
ops-db2ese-8dcbb4dfd-slxv8          0/1       ContainerCreating   0          2m
ops-impactgui-6d68867cc5-thsch      0/1       Init:0/1            0          2m
ops-jdbcgw-9bdc669cc-x4zrd          0/1       Init:0/1            0          2m
ops-ncibackup-5f4c95bfcd-kn6h6      0/1       Init:0/2            0          2m
ops-nciprimary-6965b459d5-chbrf     0/1       Init:0/1            0          2m
ops-ncobackup-776cf468c-mnplb       2/2       Running             0          2m
ops-ncoprimary-686c956546-t5qqv     1/1       Running             0          2m
ops-openldap-7dcd776ccd-n7bgd       1/1       Running             0          2m
ops-scala-bcc96d767-tswvk           0/2       Init:0/1            0          2m
ops-webgui-68f4f89c47-nxfjh         0/1       Init:0/1            0          2m
```

b. Use a command like the following example to show events from the pod. In this example, the target is the DB2 pod.

```
kubectl describe pod <POD_NAME>
```

```
kubectl describe pod ops-db2ese-8dcbb4dfd-slxv8
```

Events from the pods are at the bottom of the output.

14. OPTIONAL: Look at the logs for your pods.
    a. Run the following command to show the name of your pods.

```
kubectl get pods
```

```
NAME                                READY     STATUS    RESTARTS   AGE
ops-db2ese-8dcbb4dfd-slxv8          1/1       Running   0          23m
ops-impactgui-6d68867cc5-thsch      1/1       Running   0          23m
ops-jdbcgw-9bdc669cc-x4zrd          1/1       Running   0          23m
ops-ncibackup-5f4c95bfcd-kn6h6      1/1       Running   0          23m
ops-nciprimary-6965b459d5-chbrf     1/1       Running   0          23m
ops-ncobackup-776cf468c-mnplb       2/2       Running   5          23m
ops-ncoprimary-686c956546-t5qqv     1/1       Running   1          23m
ops-openldap-7dcd776ccd-n7bgd       1/1       Running   0          23m
ops-scala-bcc96d767-tswvk           2/2       Running   0          23m
ops-webgui-68f4f89c47-nxfjh         1/1       Running   0          23m
```

    b. Use a command like the following example to tail the log of a pod. In this example, the target log is the primary Netcool/Impact log. Press Ctrl+C to stop the tail.

```
kubectl logs -f <POD_NAME>
```

```
kubectl logs -f ops-nciprimary-6965b459d5-chbrf
```

    c. Some pods have more than one container running. For these pods, you must specify the name of the container you want. In this example, an error message is displayed when trying to view the Log Analysis log files.

```
kubectl logs -f ops-scala-bcc96d767-tswvk
```

```
Error from server (BadRequest): a container name must be specified for pod
ops-scala-bcc96d767-tswvk, choose one of: [unity gateway] or one of the init
containers: [wait4db2ese]
```

    To get the logs from the Log Analysis application, use a command like the following example.

```
kubectl logs -f ops-scala-bcc96d767-tswvk unity
```

    To get the logs from the Log Analysis gateway to Netcool/OMNIbus, use a command like the following example.

```
kubectl logs -f ops-scala-bcc96d767-tswvk gateway
```

15. OPTIONAL: Connect to your pods.

    a. Run the following command to show the name of your pods.

```
kubectl get pods
```

```
NAME                                READY     STATUS    RESTARTS   AGE
ops-db2ese-8dcbb4dfd-slxv8          1/1       Running   0          23m
ops-impactgui-6d68867cc5-thsch      1/1       Running   0          23m
ops-jdbcgw-9bdc669cc-x4zrd          1/1       Running   0          23m
ops-ncibackup-5f4c95bfcd-kn6h6      1/1       Running   0          23m
ops-nciprimary-6965b459d5-chbrf     1/1       Running   0          23m
ops-ncobackup-776cf468c-mnplb       2/2       Running   5          23m
ops-ncoprimary-686c956546-t5qqv     1/1       Running   1          23m
ops-openldap-7dcd776ccd-n7bgd       1/1       Running   0          23m
ops-scala-bcc96d767-tswvk           2/2       Running   0          23m
ops-webgui-68f4f89c47-nxfjh         1/1       Running   0          23m
```

    b. Use a command like the following example to connect to a pod. In this example, the target pod is the primary Netcool/OMNIbus server.

```
kubectl exec -ti <POD_NAME> bash
```

```
kubectl exec -ti ops-ncoprimary-686c956546-t5qqv bash
```

    c. Type **exit** to disconnect from the pod.

Important: Most of the changes you can make by connecting to a pod are not persisted to disk. If you make any changes by directly connecting to a pod, those changes will probably be lost if the pod restarts.

IBM Netcool Operations Insight post-installation tasks

In this section, you complete the following tasks:

- Add the NOI user interface endpoints to your hosts file
- Test the NOI user interfaces and basic functions
- Connect to the LDAP administration tool

Important: Run all of the commands in this section as the root user on your master node. You can use any host for the steps in this section that use a browser.

1. Add the NOI user interface endpoints to your hosts file. Complete this step on the host that you will use to browse to the DASH, Netcool/Impact, Log Analysis, and WebSphere user interfaces.
   a. Run the following command to show the host name and IP address of each user interface.

```
kubectl get ingress
```

Note: The host name of each user interface contains the Helm release name that you chose before you installed NOI. In this example, the Helm release name is **ops**.

   b. Open your hosts file with a text editor. The location of your hosts file varies depending on your operating system. Complete this step on the host that you will use to browse to the NOI user interfaces.

   c. Add the IP address and host name of each user interface to your hosts file. Use the format: `IP_ADDRESS  HOSTNAME`, for example:

```
9.42.74.60      impact.ops.icpmasteredu.rtp.raleigh.ibm.com
9.42.74.60      scala.ops.icpmasteredu.rtp.raleigh.ibm.com
9.42.74.60      netcool.ops.icpmasteredu.rtp.raleigh.ibm.com
9.42.74.60      was.ops.icpmasteredu.rtp.raleigh.ibm.com
```

   d. Save and close the file when you are finished.

2. Test the Netcool/Impact user interface.
   a. Browse to the Netcool/Impact user interface. Use a URL similar to the following example, with your Helm release name.

https://impact.<YOUR_RELEASE_NAME>.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

   For example:

https://impact.ops.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

b. If you are prompted with a warning about an insecure or non-private connection click **advanced**, add a security exception to your browser, and proceed to the web site.

c. Log in with the user name **impactadmin** and the password **netcool**.

d. Verify that the Netcool/Impact user interface loads.

e. At the bottom left of the page, click the head-and-shoulders icon then click **Log out**.

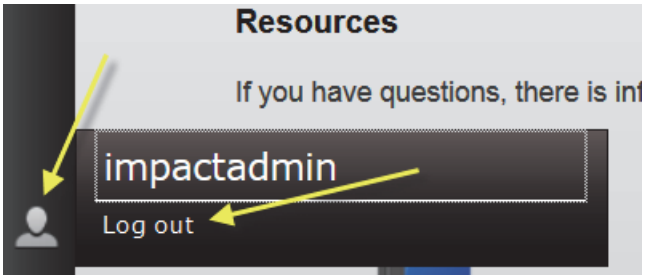

3. Test the Log Analysis user interface.

a. Browse to the Log Analysis user interface. Use a URL similar to the following example, with your Helm release name.

https://scala.<YOUR_RELEASE_NAME>.icpmasteredu.rtp.raleigh.ibm.com/Unity

For example:

https://scala.ops.icpmasteredu.rtp.raleigh.ibm.com/Unity

b. If you are prompted with a warning about an insecure or non-private connection click **advanced**, add a security exception to your browser, and proceed to the web site.

c. Log in with the user name **unityadmin** and the password **unityadmin**.

d. Verify that the Log Analysis user interface loads.

e. At the top right of the page, click **unityadmin** then click **Logout**.

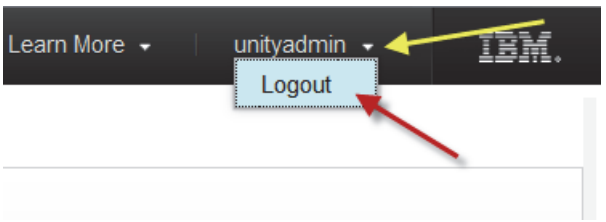

4. Test the WebSphere Administrative Console.

a. Browse to the WebSphere Administrative Console user interface. Use a URL similar to the following example, with your Helm release name.

https://was.<YOUR_RELEASE_NAME>.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

For example:

https://was.ops.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

b. If you are prompted with a warning about an insecure or non-private connection click **advanced**, add a security exception to your browser, and proceed to the web site.

c. Log in with the user name **admin** and the password **netcool**.

d. Verify that the WebSphere Administrative Console loads.

e. At the top right of the page, click **Logout**.

5. Test the DASH user interface.

a. Browse to the DASH user interface. Use a URL similar to the following example, with your Helm release name.

https://netcool.<YOUR_RELEASE_NAME>.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

For example:

https://netcool.ops.icpmasteredu.rtp.raleigh.ibm.com/ibm/console

b. If you are prompted with a warning about an insecure or non-private connection click **advanced**, add a security exception to your browser, and proceed to the web site.

c. Log in with the user name **icpadmin** and the password **netcool**.

d. Verify that the DASH user interface loads.

e. Leave this page open, you use it in the next steps.

6. Test the Event Viewer. This test verifies that WebGUI and the ObjectServer are working.

a. In the DASH user interface, click **Incident > Event Viewer**.

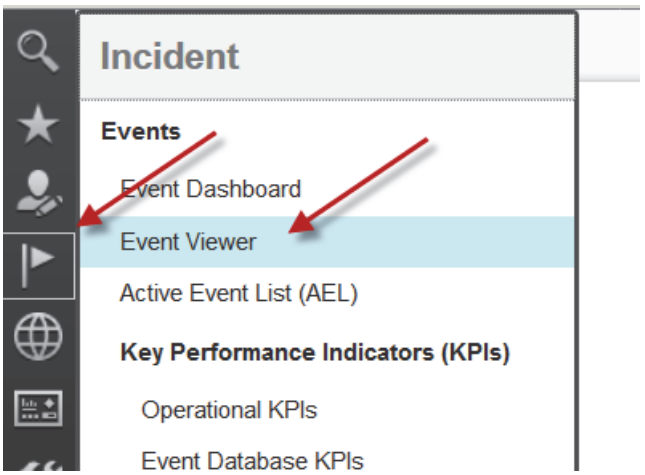

b. Verify that there are events present.

7. Test Log Analysis. This test verifies that Log Analysis and the gateway to Netcool/OMNIbus are working.
   a. Right click any event. Click **Event Search > Search for events by node > 1 day before event**.
   b. Verify that the Log Analysis user interface loads. You might need to allow pop-up windows in your browser. You can ignore any message about no search results.
   c. Enter * as the search string.
   d. Select **Last Day** as the time filter.
   e. Click Search.

   f. Verify that search results are present.
   g. Close the Log Analysis window or tab.
   h. Close the Event Viewer tab.

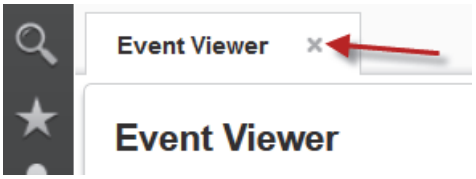

8. Test Netcool/Impact. This test verifies the Netcool/Impact connection to DASH, the ObjectServer, and the historical event database.
   a. Click **Console Integrations > Impact**.

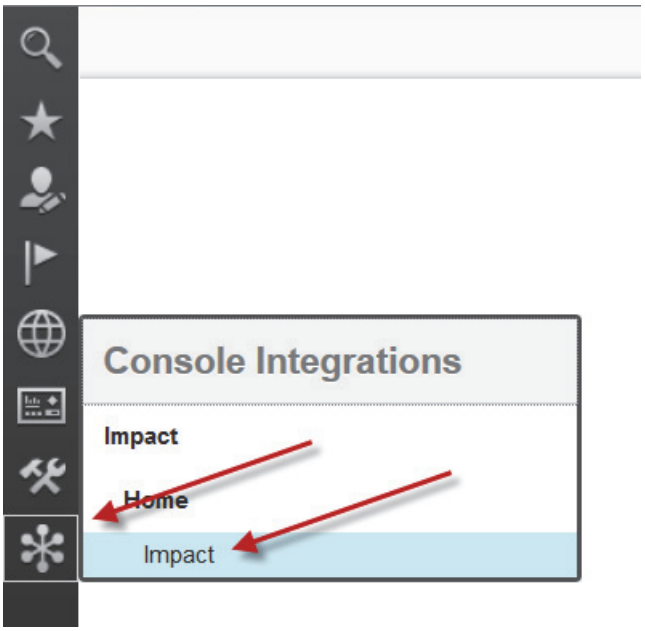

   b. At the top of the page, choose the **Global** project if it is not already selected.

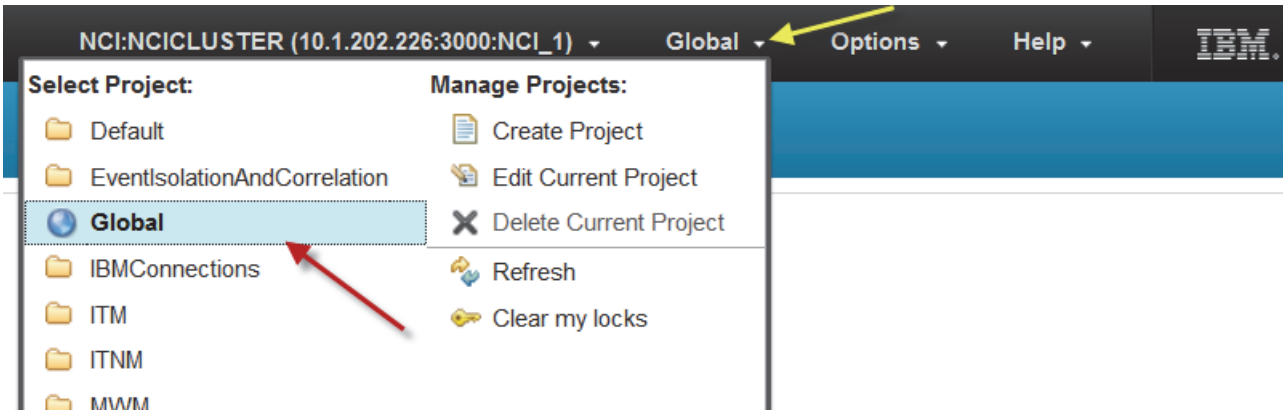

c. Click the **Data Model** tab.

d. Right-click **ObjectServerForNOI**.

e. Click **Test Connection**.

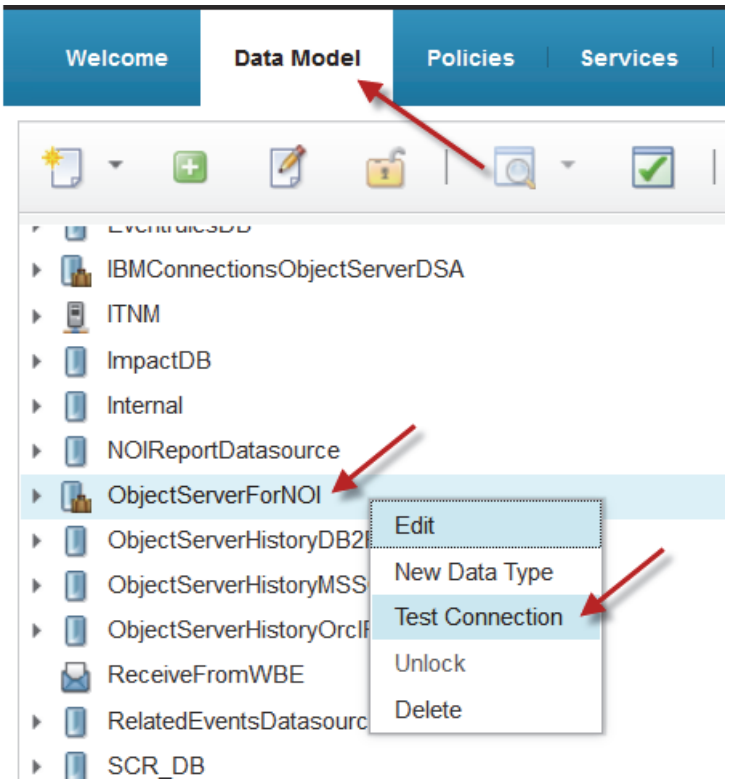

f. Verify that the connection is OK. Click **Close**.

g. Right-click **defaultobjectserver**.

h. Click **Test Connection**.

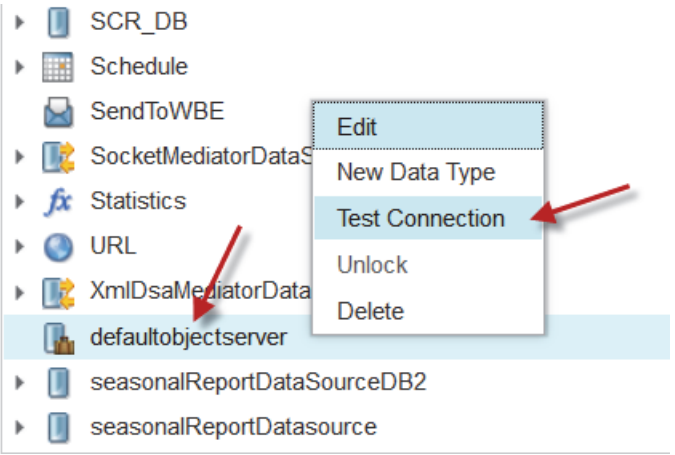

i. Verify that the connection is OK. Click **Close**.

j. Expand **ObjectServerHistoryDB2ForNOI**.

k. Right-click **AlertsHistoryDB2Table**.

l. Click **View Data Items**.

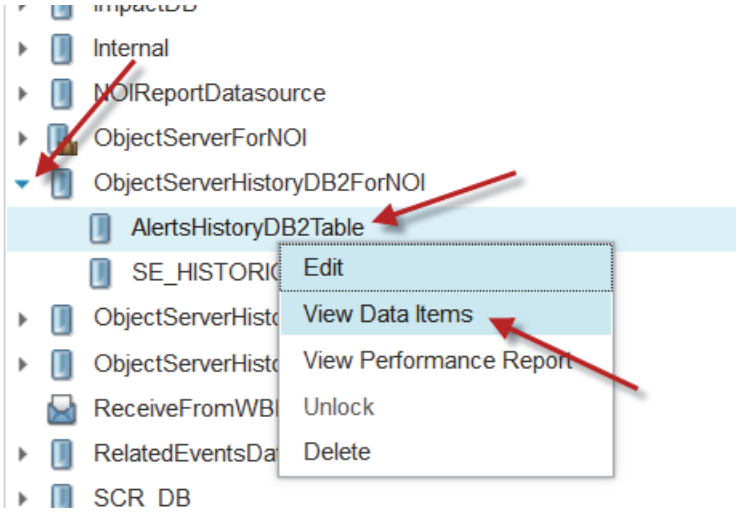

m. Verify that events are present in the table.

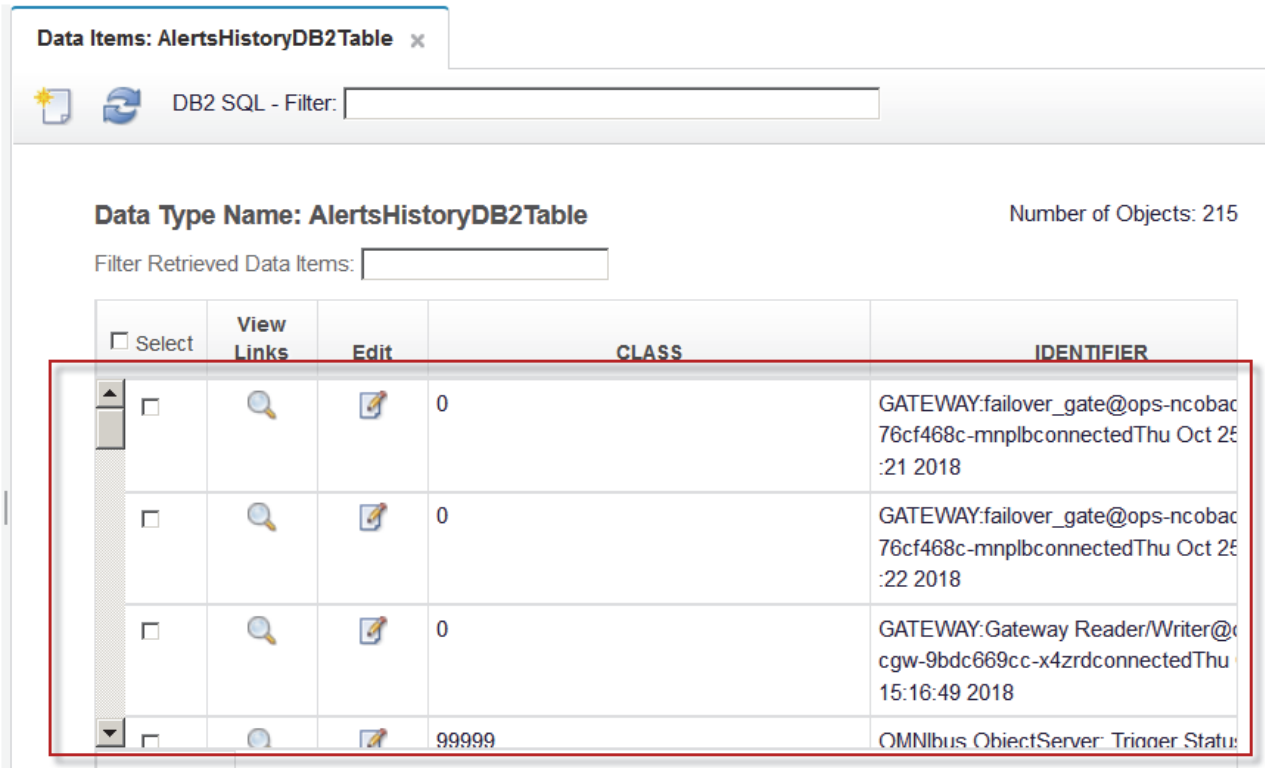

n. At the bottom left of the page, click the head-and-shoulders icon then click **Log out**.

9. Verify that you can log in to the LDAP server as the administrative user.
   a. Return to a terminal where you are connected to your master node.
   b. Run a command similar to the following example to find the LDAP user interface port number. Use your own Helm release name when you run the command.

```
helm status <YOUR_RELEASE_NAME> --tls |egrep -i ldapuisvc
```

      For example:

```
helm status ops --tls |egrep -i ldapuisvc
```

   c. Look for the last port number in the output of the preceding command. In this example, the port number is 32442.

```
ops-ldapuisvc      NodePort   10.0.0.106  <none>       9581:32442/TCP
```

   d. Open a browser and go to a URL similar to the following example. Replace the IP address with the IP address of your master node. Replace the port number with the port number you found in the preceding step.

http://<YOUR_MASTER_NODE_IP>:<YOUR_PORT>/phpldapadmin/

      For example:

http://9.42.74.60:32442/phpldapadmin/

   e. At the top left of the page, click **login**.

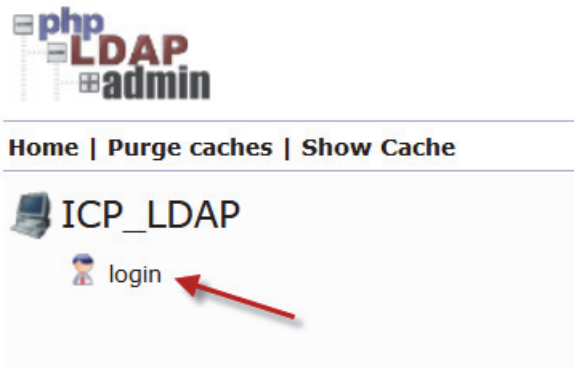

f. Enter **admin** as the password.

g. Click **Authenticate**.

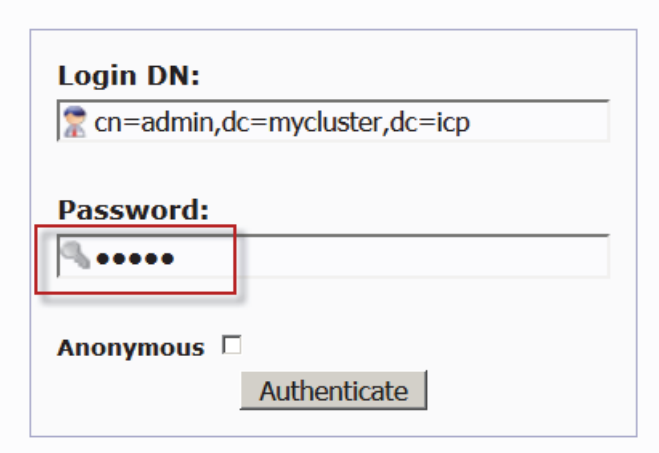

h. Expand **dc=mycluster,dc=icp**.

i. Expand **ou=users**.

j. Verify that users are present. Click logout.

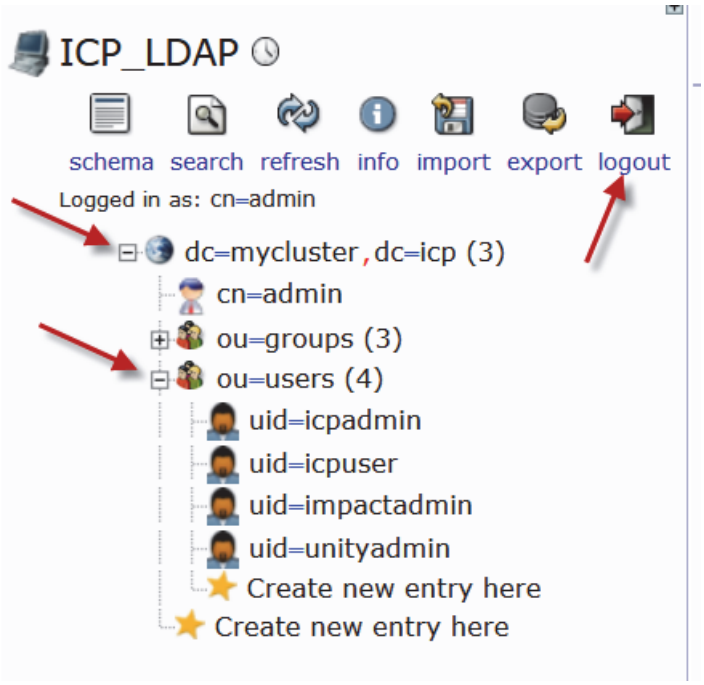

You are now finished installing IBM Cloud Private, GlusterFS, and IBM Netcool Operations Insight. Thank you for your interest in IBM Netcool Operations Insight.

IBM Training