

Adding a new node to the Flow Enterprise cluster

If needed, add additional nodes to your cluster. This is usually done during the initial installation of Flow, or during an upgrade. If you need additional CPU or memory after installation, you can add additional nodes then.

Important: Beginning with Flow Enterprise Server 2021.1.1 and continuing with Flow Enterprise Server 2021.3.1, new installations require a raw block device for the primary node and all worker nodes, in addition to existing [storage requirements](https://help.pluralsight.com/help/flow-enterprise-server-2021-3-1-1-system-requirements) (<https://help.pluralsight.com/help/flow-enterprise-server-2021-3-1-1-system-requirements>). If you're upgrading to 2021.1.1 or 2021.3.1 and an exported configuration is detected during installation, you can proceed without the raw block device requirement. Raw block devices will be a requirement for future upgrades.

Note: Some of the commands below need to run as `sudo` or non-`sudo`, depending on the flags you're using.

In this article

[Adding additional device filters for the worker node](#)

[Joining a node with a raw block device](#)

[Joining a node without a raw block device](#)

Adding additional device filters for the worker node

All nodes in the cluster should use the same operating system, Kernel version, and storage types for consistency. It is possible that the raw device name varies from primary to worker nodes. If so, update the Ceph cluster on the primary node to add the additional device filters for the worker nodes. This must happen before running the join command on the worker node.

Note: Filters must be compatible with go-lang regular expressions.

In this example, the primary node is running Red Hat 8 with a device filter of `'^xvd[b-f]'` but the worker node being added is running Red Hat 7.9 with a raw device of type `'^nvme[1-9]n1'`. To make sure Ceph uses the worker node storage correctly, update the primary node with the command `./flow-tools join -f '^nvme[1-9]n1'`.

Every platform and operating system may have variations of device names. Learn more about device names for [instances running on Amazon Web services](#)

(https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/device_naming.html).

Below is an example output from the command.

```
[root@master bin]$ ./flow-tools join -f '^nvme[1-9]n1'

[INFO] Current ceph device filters are : xvd

[INFO] Updated device filter for the cluster will be : xvd|^nvme[1-9]n1

Warning: Updating Ceph filters for the cluster. Continue (Yy|Nn)? y

[INFO] Updating new filter to xvd|^nvme[1-9]n1

cephcluster.ceph.rook.io/rook-ceph configured

[INFO] Ceph Filter updated to rook-ceph CRD : xvd|^nvme[1-9]n1

[INFO] Restarting Ceph Operator

[INFO] Ceph Operator Pod = rook-ceph-operator-9c7456b5b-8gdrz

pod "rook-ceph-operator-9c7456b5b-8gdrz" deleted
```

[back to top](#)

Joining a node with a raw block device

Note: Join tokens and credentials expire after 24 hours. Generate a new token by using the join command again as required.

On the primary node, run `sudo ./flow-tools join -g`.

Copy the join command output. You may need to add additional flags to this command to reflect your proxy setup or other configurations. Use the `--help` option to see all available flags.

```
[root@host.domain bin]$ sudo ./flow-tools join -g

[INFO] Existing kots/containerd based Flow install detected

[INFO] Existing ceph disk found : /dev/xvdb

Kubernetes connection credentials for worker node. Expires in 24 hours

Kubernetes Connection String : kubernetes-master-address=192.168.1.40:6443 kubeadm-
token=xthyba.hc61xn6c3gj959jo kubeadm-token-ca-
hash=sha256:6edcbe6899406c173730e4e6e508da94220669a7c422f9bd9f0cfa84a8eaed5b kubernetes-
```

```
version=1.19.7 docker-registry-ip=10.96.0.49
```

You may add additional command line options to the flow-tools join command.

Run `./flow-tools join --help` for all available flags and options like `[-a|-f|-k|-n|--proxy]` etc.

Default non-airgapped join command with ceph raw disks for this cluster is below, change device filter as needed:

For example: Red Hat, filters could be like `'xvd[b-z]'`, for amazon ami with nvme, `'nvme[1-9]n1'`

```
sudo ./flow-tools join kubernetes-master-address=192.168.1.40:6443 kubeadm-  
token=xthyba.hc61xn6c3gj959jo kubeadm-token-ca-  
hash=sha256:6edcbe6899406c173730e4e6e508da94220669a7c422f9bd9f0cfa84a8eaed5b  
kubernetes-version=1.19.7 docker-registry-ip=10.96.0.49
```

Default air-gapped node join command for this cluster is below:

```
sudo ./flow-tools join -a kubernetes-master-address=192.168.1.40:6443 kubeadm-  
token=xthyba.hc61xn6c3gj959jo kubeadm-token-ca-  
hash=sha256:6edcbe6899406c173730e4e6e508da94220669a7c422f9bd9f0cfa84a8eaed5b  
kubernetes-version=1.19.7 docker-registry-ip=10.96.0.49
```

Next, log in to the worker node(s) you're adding to the cluster. Copy the `flow-enterprise-tools` package of the same version on the primary node to the worker node. Extract the file. Change the directory to the `bin` folder such as `./flow-enterprise-tools/bin`.

Run the join command output you copied above in the location where you installed `flow-enterprise-tools`. Change the raw disk volume name according to your Linux distribution.

```
$ cd ./flow-enterprise-tools/bin
```

```
$ sudo ./flow-tools join kubernetes-master-address=192.168.1.40:6443 kubeadm-  
token=xthyba.hc61xn6c3gj959jo kubeadm-token-ca-  
hash=sha256:6edcbe6899406c173730e4e6e508da94220669a7c422f9bd9f0cfa84a8eaed5b  
kubernetes-version=1.19.7 docker-registry-ip=10.96.0.49
```

```
[INFO] Verifying installation environment...
```

```
[INFO] HTTP command (curl): OK
```

```
[INFO] Archive command (tar): OK
```

```
[INFO] Swarm does not exist: OK
```

```
[INFO] Verifying system requirements...
```

```
[INFO] Checking networking...
```

```
[INFO] sysctl command : OK
```

[INFO] IPV6 Kernel module: LOADED

[INFO] IPV6 Check : OK

[INFO] IPv4 Forwarding: ENABLED

[INFO] Check IPTable Rules: OK

[INFO] Detecting proxy: NOT DETECTED

[INFO] https://replicated.app site check : OK

[INFO] Checking hardware...

[INFO] CPU: OK

[INFO] Memory: OK

[INFO] Space check in /var/lib/docker: OK

[INFO] Space check in /var/lib/kubelet: OK

[INFO] Space check in /opt/replicated: OK

[INFO] Space for Repo cache in /opt/flow: 499 GB

[INFO] Disk Space Check: OK

[INFO] Non SSD Disks: NOT DETECTED

[INFO] Checking filesystem and permissions...

[INFO] Login restrictions check: OK

[INFO] Selinux Status: enabled

[INFO] Selinux Current mode: permissive

...

[kubelet-start] Writing kubelet configuration to file "/var/lib/kubelet/config.yaml"

[kubelet-start] Writing kubelet environment file with flags to file "/var/lib/kubelet/kubeadm-flags.env"

[kubelet-start] Starting the kubelet

[kubelet-start] Waiting for the kubelet to perform the TLS Bootstrap...

This node has joined the cluster:

* Certificate signing request was sent to apiserver and a response was received.

* The Kubelet was informed of the new secure connection details.

Run 'kubectl get nodes' on the control-plane to see this node join the cluster.

⚙ Node joined successfully

Installation

Complete ✓

[INFO] Additional configuration steps:

[INFO] Please run the following command on the primary node to set the proper labels for the Flow app:

```
kubectl label nodes worker- --selector='node-role.kubernetes.io/master'
```

```
kubectl label nodes worker= --selector='!node-role.kubernetes.io/master'
```

On the primary node, confirm the node was successfully joined by running `kubectl get nodes` .

```
$ kubectl get nodes
```

```
[root@master ~]$ kubectl get nodes
```

```
NAME           STATUS    ROLES    AGE   VERSION
master.domain.com Ready    master   6d1h  v1.19.7
worker.domain.com Ready     5d20h   v1.19.7
```

Restart the rook-ceph operator on the primary node using `./flow-tools join -r`.

```
[root@master.domain.com bin]$ ./flow-tools join -r
```

```
[INFO] Deleting Ceph prepare pod jobs
```

```
pod "rook-ceph-osd-prepare-master.domain-opr7qlf" deleted
```

```
[INFO] Restarting Ceph Operator
```

```
[INFO] Ceph Operator Pod = rook-ceph-operator-9c7456b5b-f7hct
```

```
pod "rook-ceph-operator-9c7456b5b-f7hct" deleted
```

```
Checking OSD prepare jobs to complete.....OSD prepare completed
```

```
[INFO] Waiting on OSD to be configured on new node..
```

```
done
```

Confirm a new [OSD storage node](https://docs.ceph.com/en/latest/rados/configuration/storage-devices/) was added to the Ceph cluster by running `kubectl -n rook-ceph get pods` on the primary node.

```
[root@master ~]$ kubectl -n rook-ceph get pods
```

NAME	READY	STATUS	RESTARTS	AGE
rook-ceph-agent-b4nd6	1/1	Running	0	25m
rook-ceph-agent-n4zcn	1/1	Running	0	43h
rook-ceph-mgr-a-b594f79bf-wqfzr	1/1	Running	0	43h
rook-ceph-mon-a-848b69b66d-249dc	1/1	Running	0	43h
rook-ceph-operator-9c7456b5b-v2xqr	1/1	Running	0	7m35s
rook-ceph-osd-0-66d47fb645-qsk52	1/1	Running	0	42h
rook-ceph-osd-1-54cdc7fccf-cr8tb	1/1	Running	0	23m
rook-ceph-osd-prepare-master.domain-opzdtlv	0/2	Completed	0	6m39s
rook-ceph-osd-prepare-worker.domain-skpgj	0/2	Completed	0	6m37s
rook-ceph-rgw-rook-ceph-store-a-867496989f-jmdmw	1/1	Running	0	42h
rook-discover-7ntfs	1/1	Running	0	43h
rook-discover-k78vd	1/1	Running	0	25m

Label the node for the worker using the command at the end of the latest join command output.

```
$ kubectl label node worker.domain.com label worker=
```

```
$ kubectl label node master.domain.com label worker-
```

Check to make sure the nodes are labeled appropriately by running `kubectl get nodes --show-labels`.

```
[root@master ~]$ kubectl get nodes --show-labels
```

NAME	STATUS	ROLES	AGE	VERSION	LABELS
master.domain.com	Ready	master	6d1h	v1.19.7	beta.kubernetes.io/arch=amd64, beta.kubernetes.io/os=linux, gui=, kubernetes.io/arch=amd64, kubernetes.io/hostname=master.domain.com, kubernetes.io/os=linux, kurl.sh/cluster=true, node-role.kubernetes.io/master=

worker.domain.com Ready 5d20h v1.19.7

beta.kubernetes.io/arch=amd64, beta.kubernetes.io/os=linux, kubernetes.io/arch=amd64,
kubernetes.io/hostname=worker.domain.com, kubernetes.io/os=linux, kurl.sh/cluster=true,worker=

[back to top](#)

Joining a node without a raw block device

Note: You can only add nodes without raw block devices if you're upgrading and an existing configuration export is detected during installation.

The entire Flow cluster should follow the same storage specifications across all nodes. Do not mix nodes with raw block devices with ones without raw block devices.

If all nodes are using filesystem disk volumes, still use the node join command to join a node to the cluster. The join command would be similar to the one above, but with an `--override-check` modification . Such a modification looks like:

```
$ sudo ./flow-tools join --override-check kubernetes-master-address=192.168.1.40:6443  
kubeadm-token=xthyba.hc61xn6c3gj959jo kubeadm-token-ca-  
hash=sha256:6edcbe6899406c173730e4e6e508da94220669a7c422f9bd9f0cfa84a8eaed5b  
kubernetes-version=1.19.7 docker-registry-ip=10.96.0.49
```

From there, follow the same steps as for joining a node with a raw block device, Your next step will be labeling nodes.

[back to top](#)

If you need help, please email support@pluralsight.com () for 24/7 assistance.