

Installing Flow Enterprise Server 2022.1.1

Before you begin your Flow Enterprise Server 2022.1.1 installation, carefully read the [system requirements](https://help.pluralsight.com/help/2022-1-1-system-requirements) (<https://help.pluralsight.com/help/2022-1-1-system-requirements>).

This article covers installing Flow Enterprise for the first time. Read more about [upgrading Flow](https://help.pluralsight.com/help/upgrading-to-flow-enterprise-server-2022-1-1) (<https://help.pluralsight.com/help/upgrading-to-flow-enterprise-server-2022-1-1>) if you want to upgrade to 2022.1.1 from a previous version.

The Flow application is based on two components: Replicated's KOTS platform and the Pluralsight Flow application stack which runs on top of it. During installation, KOTS is installed first, then Flow is added as an application on this platform.

In this article

[Preinstallation instructions](#)

[Install Flow Enterprise](#)

[Configure the KOTS app](#)

Preinstallation instructions

- Read and follow the [Flow Enterprise Server 2022.1.1 system requirements](https://help.pluralsight.com/help/2022-1-1-system-requirements) (<https://help.pluralsight.com/help/2022-1-1-system-requirements>).
- If you're using an external database, have your database administrator complete the necessary steps to [prepare the database](https://help.pluralsight.com/help/manual-installation-preparing-the-database) (<https://help.pluralsight.com/help/manual-installation-preparing-the-database>).
- Download the latest copy of the `flow-enterprise-tools` package. Version 2.1.x is required. This is provided by your Pluralsight support representative. Install this package on each node of the Flow cluster.
- For the host server, copy `flow-enterprise-tools-[-airgap]-.tar.gz` to the home directory of the user account used for the installation on the host server.
- Extract the tools file using `tar xvf flow-enterprise-tools-[-airgap]-.tar.gz`.

Note: You can run any tool from the bin directory by running `cd /path/to/flow-enterprise-tools/bin ./[tool name]`. Install the tools package with the `install-enterprise-tools.sh` script `cd /path/to/flow-enterprise-tools ./install-enterprise-tools.sh`. The script will ask where to install the components. The default is `/usr/local/share/flow-enterprise-tools`.

- Install the `flow-enterprise-tools` package using the script `install-enterprise-tools.sh`. This allows installation of the scripts and binaries in the system where users can access them from standard `PATH` settings. You can also run the scripts of this package in-place.

Important: For the root user, `/usr/local/bin` must be in the `PATH` environment variable. The root user

should be set up as a Flow user after the installation of Flow. This is critical if your OS is hardened.

- Create the following directories. These can all be bundled in a single volume
app directory (default: `/opt/flow`)
`mkdir -p [app_directory]`
`chmod 755 [app_directory]`
Once the app directory is created, change the ownership of the directory to the following:
`chown -R 37355:37355 [app_directory]`
- If you're using an embedded database, make a database directory (default `/opt/flow/database`) on the primary host server only:
`mkdir -p [database_directory]`
`chmod 725 [database_directory]`
Once the database directory is created, change the ownership of the directory to the following:
`chown -R 1000:1000 [database_directory]`

If your version of Flow Enterprise Server is airgapped, download the airgap bundle from Replicated. A password is required. If you can't access Replicated, contact Support for assistance.

Tip: Depending on how you install Flow, you need to download a few different packages. `flow-enterprise-tools` contains the tools for installation and maintenance of the Kubernetes framework. Download the airgap version for airgapped installations. The application airgap bundle, only for airgap installations, is downloaded from Replicated. It contains the Flow application files, which are installed into the Kubernetes framework provided by `flow-enterprise-tools`.

[back to top](#)

Install Flow Enterprise

Now you are ready to install Flow Enterprise 2022.1.1. For an online installation, run `sudo flow-tools install`. For an airgapped installation, run `sudo flow-tools install -a`. This only needs to be done on the primary node.

Important: A raw block device for each node is required for all Flow installations.

The device filter is in GoLang regular expression format. Device names vary between distributions. Add multiple filters if the devices are named differently in worker nodes. An example of such a case may look like `sd[b-z]|nvme[1-3]|xvd[b-c]`.

```
admin-user@primary-node:~$ sudo ./flow-tools install
```

[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] systemctl command : OK

[INFO] IPV6 Kernel module: LOADED

[INFO] IPV6 Check : OK

[INFO] IPv4 Forwarding: ENABLED

[INFO] Check IPtable Rules: OK

[INFO] Checking ufw firewall: INACTIVE

[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/containerd: OK

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

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

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

[INFO] Space check in /tmp: 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] bash Umask setting: OK

[INFO] /etc/profile Umask setting: OK

[INFO] Checking PATH for /usr/local/bin: OK

[INFO] Checking distro...

[INFO] No existing ceph raw disks detected

[INFO] Installation type is : NEW

=== Discovered Block Devices ===

/dev/nvme1n1

Above is the list of block devices found during valid device discovery

Please provide pattern to match devices that should be used for K8s volume storage: nvme1n1

[INFO] Validating block storage device filter...

Device match: /dev/nvme1n1

Device size: 100G

Device status: valid

[INFO] Total valid block storage: 100G

[INFO] Block storage: OK

[INFO] Adding patch to use raw ceph block devices for installation

[INFO] Installing KOTS application

[INFO] Saving environment

[INFO] Fetching kurl.sh installation script from: <https://k8s.kurl.sh/version/v2022.01.05-0/flow-enterprise>

[INFO] Fetching join script from: <https://k8s.kurl.sh/version/v2022.01.05-0/flow-enterprise/join.sh>

Downloading package kurl-bin-utils-v2022.01.05-0.tar.gz

SELinux is not installed: no configuration will be applied

The installer will use network interface 'ens5' (with IP address '192.168.1.50')

Downloading package weave-2.6.5.tar.gz

Downloading package rook-1.0.4.tar.gz

Downloading package contour-1.14.1.tar.gz

Downloading package registry-2.7.1.tar.gz

Downloading package prometheus-0.46.0.tar.gz

Downloading package kotsadm-1.47.3.tar.gz

⚙ Running host preflights

[TCP Port Status] Running collector...

[CPU Info] Running collector...

[Amount of Memory] Running collector...

[Block Devices] Running collector...

[Ephemeral Disk Usage /var/lib/kubelet] Running collector...

[Kubernetes API TCP Port Status] Running collector...

[ETCD Client API TCP Port Status] Running collector...

[ETCD Server API TCP Port Status] Running collector...

[ETCD Health Server TCP Port Status] Running collector...

[Kubelet Health Server TCP Port Status] Running collector...

[Kubelet API TCP Port Status] Running collector...

[Kube Controller Manager Health Server TCP Port Status] Running collector...

[Kube Scheduler Health Server TCP Port Status] Running collector...

[Filesystem Latency Two Minute Benchmark] Running collector...

[Weave Network Policy Controller Metrics Server TCP Port Status] Running collector...

[Weave Net Metrics Server TCP Port Status] Running collector...

[Weave Net Control TCP Port Status] Running collector...

[Block Devices] Running collector...

[Node Exporter Metrics Server TCP Port Status] Running collector...

[Can Access Replicated API] Running collector...

[PASS] Number of CPUs: This server has at least 4 CPU cores

[PASS] Amount of Memory: The system has at least 8G of memory

[PASS] Ephemeral Disk Usage /var/lib/kubelet: The disk containing directory /var/lib/kubelet has at least 10Gi of disk space available and is not nearly full

[PASS] Kubernetes API TCP Port Status: Port 6443 is open

[PASS] ETCD Client API TCP Port Status: Port 2379 is open

[PASS] ETCD Server API TCP Port Status: Port 2380 is open

[PASS] ETCD Health Server TCP Port Status: Port 2381 is available

[PASS] Kubelet Health Server TCP Port Status: Port 10248 is available

[PASS] Kubelet API TCP Port Status: Port 10250 is open

[PASS] Kube Controller Manager Health Server TCP Port Status: Port 10257 is available

[PASS] Kube Scheduler Health Server TCP Port Status: Port 10259 is available

[PASS] Filesystem Performance: Write latency is ok (p99 target < 10ms, actual: 3.520561ms)

[PASS] NTP Status: System clock is synchronized

[PASS] NTP Status: Timezone is set to UTC

[PASS] Weave Network Policy Controller Metrics Server TCP Port Status: Port 6781 is available

[PASS] Weave Net Metrics Server TCP Port Status: Port 6782 is available

[PASS] Weave Net Control TCP Port Status: Port 6783 is open

[PASS] Block Devices: One available block device

[PASS] Node Exporter Metrics Server TCP Port Status: Port 9100 is available

[PASS] Can Access Replicated API: Connected to https://replicated.app

⚙ Host preflights success

Found pod network: 10.32.0.0/20

Found service network: 10.96.0.0/22

⚙ Addon containerd 1.4.4

⚙ Installing host packages containerd.io

* Applying /etc/sysctl.d/10-console-messages.conf ...

kernel.printk = 4 4 1 7

* Applying /etc/sysctl.d/10-ipv6-privacy.conf ...

net.ipv6.conf.all.use_tempaddr = 2

net.ipv6.conf.default.use_tempaddr = 2

* Applying /etc/sysctl.d/10-kernel-hardening.conf ...

kernel.kptr_restrict = 1

* Applying /etc/sysctl.d/10-link-restrictions.conf ...

fs.protected_hardlinks = 1

fs.protected_symlinks = 1

* Applying /etc/sysctl.d/10-lxd-inotify.conf ...

fs.inotify.max_user_instances = 1024

* Applying /etc/sysctl.d/10-magic-sysrq.conf ...

kernel.sysrq = 176

* Applying /etc/sysctl.d/10-network-security.conf ...

net.ipv4.conf.default.rp_filter = 1

net.ipv4.conf.all.rp_filter = 1

```
net.ipv4.tcp_syncookies = 1
```

```
* Applying /etc/sysctl.d/10-ptrace.conf ...
```

```
kernel.yama.ptrace_scope = 1
```

```
* Applying /etc/sysctl.d/10-zero-page.conf ...
```

```
vm.mmap_min_addr = 65536
```

```
* Applying /usr/lib/sysctl.d/50-default.conf ...
```

```
net.ipv4.conf.all.promote_secondaries = 1
```

```
net.core.default_qdisc = fq_codel
```

```
* Applying /etc/sysctl.d/99-cloudimg-ipv6.conf ...
```

```
net.ipv6.conf.all.use_tempaddr = 0
```

```
net.ipv6.conf.default.use_tempaddr = 0
```

```
* Applying /etc/sysctl.d/99-sysctl.conf ...
```

```
net.ipv4.conf.all.forwarding = 1
```

```
* Applying /etc/sysctl.conf ...
```

```
net.ipv4.conf.all.forwarding = 1
```

```
⚙ Install kubelet, kubectl and cni host packages
```

```
kubelet command missing - will install host components
```

```
Downloading package kubernetes-1.19.13.tar.gz
```

```
...
```

```
label "velero.io/exclude-from-backup" not found.
```

```
persistentvolumeclaim/kotsadm-postgres-kotsadm-postgres-0 labeled
```

```
/var/lib/kurl/addons/kotsadm/1.47.3/assets /var/lib/kurl ~/flow-enterprise-tools/bin
```

```
/var/lib/kurl ~/flow-enterprise-tools/bin
```

```
configmap/kurl-current-config patched
```


configmap/kurl-config created

No resources found

Installation

Complete ✓

Kotsadm: http://192.168.1.50:8800

Login with password (will not be shown again): xjIPU1FcQ

The UIs of Prometheus, Grafana and Alertmanager have been exposed on NodePorts 30900, 30902 and 30903 respectively.

To access Grafana use the generated user:password of admin:gpk1ANE93 .

To access the cluster with kubectl, reload your shell:

```
bash -l
```

```
[INFO] Loading environment
```

```
[INFO] Existing ceph disk found : /dev/nvme1n1
```

```
Kubernetes connection credentials for worker node. Expires in 24 hours
```

```
Kubernetes Connection String : kubernetes-master-address=192.168.1.50:6443 kubeadm-  
token=ebg6jp.jjmx71a8jcf83b kubeadm-token-ca-  
hash=sha256:ed129e637d7b09d970bcbb024cbb948751872bcfc031d8d3a959347f488b9a15 kubernetes-  
version=1.19.13 docker-registry-ip=10.96.0.236
```

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.

Node join command for this cluster is below:

```
sudo ./flow-tools join --channel stable kubernetes-master-address=192.168.1.50:6443 kubeadm-  
token=ebg6jp.jjmx71a8jcf83b kubeadm-token-ca-  
hash=sha256:ed129e637d7b09d970bcbb024cbb948751872bcfc031d8d3a959347f488b9a15 kubernetes-  
version=1.19.13 docker-registry-ip=10.96.0.236
```

node/primary-node.mydomain.com labeled

[INFO] Primary node has been labelled with

gui=true

worker=true

If adding an additional node, please run the following,

after adding a worker node:

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

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

[]

- Reset the admin console password for default

Enter a new password to be used for the Admin Console: ••••••••

- The admin console password has been reset

[INFO] Setting up kubectl command for current user

[INFO] Processing home directory: /home/admin-user

[INFO] Setting up kube-config for user: admin-user

On a single-node cluster, the installation labels the primary node appropriately by default. However, in a multi-node installation, label the worker nodes as directed from the output of the installation command as seen above.

Reset the console password by running `kubectl kots reset-password -n default` if desired.

Next, install Flow Enterprise on each of the worker nodes using the join command output at the end of the installation of the primary node. Read more about [joining a node to the cluster](https://help.pluralsight.com/help/adding-a-new-node) (<https://help.pluralsight.com/help/adding-a-new-node>).

[back to top](#)

Configure the KOTS app

Configure TLS

1. Open your Chrome browser and go to the URL provided at the end of the installation on the primary node. It will look like `http://:8800`.

Note: You will receive an error saying your connection is not private. Flow uses a self-signed SSL/TLS Certificate to secure the communication between your local machine and the admin console during setup. This is secure.


Click **Advanced**, then click **Proceed** to continue to the admin console.

2. On the next screen, enter the following information from the [Before you begin](https://help.pluralsight.com/help/before-you-begin) (<https://help.pluralsight.com/help/before-you-begin>) steps:
 1. Enter your hostname
 2. Upload your key and certificate
3. Click **Upload & Continue**

Log in to the KOTS admin console

Note: The browser should redirect to the main KOTS admin console, where there is a prompt for login. If not, navigate to `https://[hostname]:8800`.

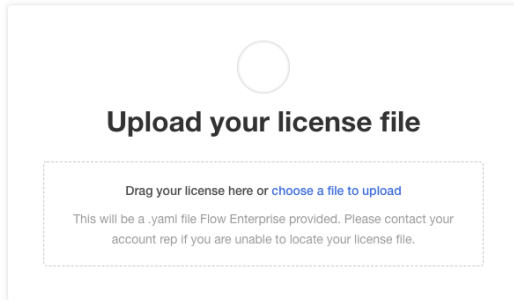
1. Enter the password provided during installation at the end of the installation script output



Log in to Flow Enterprise

Enter the password to access the Flow Enterprise admin console.

2. Upload the license file provided by your Flow representative.



3. On the following screen, there are two options for installing Flow. If your server doesn't have internet access, choose the airgapped environment option and upload the pre-downloaded airgapped binary zip file. Otherwise, click the download from the internet link to continue.

Set up Flow configurations

Once the license is validated, the configuration screen appears.

The following configuration changes are recommended:

- Match the Flow URL to the fully qualified domain name registered earlier
- Provide a valid email address to receive email alerts.
- Click **Email settings** and provide the details of the SMTP/Email server port, credentials, protocol, and other details.
- If there is a proxy involved, check **Use an outbound HTTP/HTTPS Proxy** and add in the details.
- Enter the details for your external Postgres database.

Once this screen is configured, scroll to the bottom of the screen and click **Save config**.

Deploy Flow

To deploy Flow, on the next screen, click **Deploy**.

All versions		
2022.1.1 Sequence 0	Online Install	Redeploy
Released January 20, 2022 Release notes	Deployed View preflights	Deployed: January 21, 2022 @ 11:37 am CST

Confirm the pop-up window by clicking **Yes, deploy**.

Redeploy 2022.1.1 (Sequence 0)?



To monitor the deployment process, run `kubectl get pods` on the primary node to see each pod's status. Wait until

all pods are in a running state. Once all pods are running, start Flow.

Then proceed to log in to the Flow Enterprise Application using the original URL for the application.

Additional configurations after installing Flow

- To have additional visibility into logs, turn on the built-in Elasticsearch Kibana stack to view Flow logs. If desired, forward the logs to an external Splunk Enterprise instance.
- To prevent your Flow Enterprise Server cluster from running into a forced eviction due to a lack of disk space, go to Disk Pressure Check Settings and click the checkbox next to Enable Disk Pressure Check and Scale Down.

The screenshot shows the 'Disk Pressure Check Settings' configuration page. On the left is a navigation menu with options like 'Global Worker Settings', 'Project Worker Settings', 'Repository Worker Settings', 'Deletion Worker Settings', 'Misc. Worker Settings', 'Commit Deduplication Settings', 'API Integration Settings', 'API Integration Hacks', 'System Settings', 'APM Settings', 'Disk Pressure Check Settings', 'Beta Features', and 'Analytics Settings'. The 'Disk Pressure Check Settings' section is expanded, showing a list of sub-configuration items: 'Enable Disk Pressure Check and Scale Down', 'Minimum Disk Free Percentage to avoid auto scale down', 'Cronjob schedule to check for disk pressure on cluster', 'Beta Features', and 'Analytics Settings'. The main content area is titled 'Disk Pressure Check Settings' and includes a description: 'This section contains settings that are used to control cluster disk space and automatic scale down.' It features three main settings:

- Enable Disk Pressure Check and Scale Down:** A checkbox that is checked. Description: 'Optional feature to enable automatic scaledown on disk pressure to prevent volume corruption'.
- Minimum Disk Free Percentage to avoid auto scale down:** A required field with a value of '20'. Description: 'Minimum amount of total disk space on cluster available before a shutdown is initiated. Integer value between 10-90.' Default value: 20.
- Cronjob schedule to check for disk pressure on cluster:** A required field with a value of '*/30 * * * *'. Description: 'Schedule to run an utility pod to run the disk pressure scaledown check'. Default value: */30 * * * *.

Below these settings is a 'Beta Features' section with an 'Enable Beta Features' checkbox. At the bottom of the configuration area is a 'Save config' button.

- [Enable monitoring services.](https://help.pluralsight.com/help/flow-monitoring-tools) (https://help.pluralsight.com/help/flow-monitoring-tools)
- [Automate certificate renewal](https://help.pluralsight.com/help/certificate-renewal) (https://help.pluralsight.com/help/certificate-renewal)

Log in to Flow

Navigate to your Flow URL. Enter your organization name and the primary email. Choose a password to log in.

Connect with your Pluralsight Flow representative to talk about next steps for your onboarding experience. For steps to get started with Flow, check out our [Getting Started guide](https://help.pluralsight.com/help/get-your-organization-started-using-pluralsight-flow) (https://help.pluralsight.com/help/get-your-organization-started-using-pluralsight-flow).

[back to top](#)

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