

Flow Enterprise Server 2022.2.1 system requirements

These system requirements apply to Flow Enterprise Server version 2022.2.1

In this article

[System requirements](#)

[Operating system](#)

[Database requirements](#)

[Network requirements](#)

System requirements

Flow uses Kubernetes for orchestration. The Flow stack must run in its own instance.

Cluster setup

Flow strongly recommends using three or more nodes to ensure proper [disk quorum](https://rook.io/docs/rook/v1.5/ceph-mon-health.html#monitor-quorum) and system stability. Having three or more nodes provides protection from node failure, disk failure, and volume corruption.

If you have fewer than three nodes, Flow shows a warning message on the KOTS admin console, but you can still proceed.

The minimum requirements for each node are listed below:

Computing resource requirements

- CPU: Minimum of 16 CPU cores per node. 32 cores recommended.
- Memory: Minimum of 32 GB per node. 64 GB recommended.

Storage requirements

Unless otherwise noted, the storage requirement applies to all nodes in the cluster, both primary and worker.

- Installation cache (`/tmp`): 15 GB
 - This is only used during installation.
- Raw block storage: 50 GB
 - This is required for both new installations and upgrades.
 - This storage is expected to expand in the future as new technology is added to the Flow stack.
- Infrastructure Directories: Allocate 115 GB for infrastructure required by Flow. This includes a 30% buffer for future growth. Specific directories are subject to change, but are currently:
 - Containerd (`/var/lib/containerd`): ~60 GB
 - Container logs (`/var/log/pods`): ~10 GB

- Kubernetes (`/var/lib/kubelet`): ~10 GB
- Add-on packages (`/var/lib/kurl`): ~10 GB
- Application directories
 - Replicated (`/opt/replicated`): 5 GB
 - Application directory (default `/opt/flow`)

Note: The application directory is configurable via settings. `[app_directory]` is used as a placeholder.

- Repository cache (default `[app_directory]/repository_cache`):
 - Worker node(s): 100 GB or the sum of the total sizes of your code repositories, multiplied by 2.5. Choose whichever is greater.
- Application logs (default `[app_directory]/logs`): 10 GB
- Embedded database (optional; default `[app_directory]/database`): 2 GB

Additional storage requirement details

Solid state drive (SSD) type storage with high IOPS for the volumes used for Flow Enterprise Server is recommended. Having separate disks for volumes like `/var/lib/containerd` and `/var/lib/kubelet` increases performance during report generation and data processing.

[back to top](#)

Operating system

Flow Enterprise Server can be run on compatible operating systems, including:

- Ubuntu 18.04, 20.04
- RHEL 8.1, 8.2, 8.3, 8.4, 8.5
 - For RHEL-based distributions, the file system and storage drive must support the `overlay2` storage driver for Containerd services.

Operating systems with hardened kernel images containing PAM (privilege access manager agents) or SELINUX policies may interfere with Flow installation and normal operation of the software stack. Those agents and policies may need to be disabled for you to install Flow or for it to function normally. Support for such hardened operating systems is provided on a "best effort" basis.

[back to top](#)

Database requirements

Flow Enterprise Server requires a PostgreSQL database server with solid state drives (SSDs). The size of your

database server will vary greatly depending on the size of your repositories, the amount of activity on your repositories, the number of active contributors you have, and the total number of concurrent users on Flow reports.

We recommend a minimum of:

- PostgreSQL version 12 with minor version 12.1 or greater
- 8 CPU cores or more
- 32 GB of RAM or more
- Disk space of approximately 25% of the space calculated for the Repository cache (default `[app_directory]/repository_cache`), as defined above

Installations with hundreds of thousands of repositories have required very large servers with 64 cores and 512 GB RAM.

The database performance greatly impacts the performance of the overall Flow application. Proactive monitoring of the CPU, memory, and IOPS health are key. Engage your Database team to monitor and tune your database.

Note: It is possible to use an embedded database for small trial installations, but you should use a standalone database server for your production installation. Your installation consultant can provide more information.

[back to top](#)

Network requirements

You can implement Flow Enterprise Server without access to the internet. However, we recommend that it be able to connect to the web for access to software packages during the installation and for future updates. You can schedule this to coincide with your regular maintenance. If you require the system to have no external internet access, you can implement and maintain an airgapped installation.

IPv6 must be enabled. This ensures compatibility with the underlying network virtualization layers of the Flow product stack.

For data analysis, your Flow Enterprise system must have access to your Git repositories and ticketing systems.

The following ports should be allowed to these instances:

- HTTP/80 and HTTPS/443: These should be the standard ports for your Git repository and ticket system server data for both Git data and API information.
- SSH/22 (usually): Most Git vendors allow for SSH download of repositories they serve.
- TCP/6443 for KOTS admin API
- TCP/6783 and UDP/6783-6784 for the Weave application service

For access to the system by your users, the following ports should be open to internal users:

- HTTP/80: This must be open for internal health-check pings.
- HTTPS/443: This must be open for users to use the interface.
- HTTPS/8800: This port is used to reach the KOTS admin interface with a web browser. It does not have to be open to general users, but must be available to system administrators.
- SSH/22: System administrators will need access to SSH on the server instance running Flow Enterprise for occasional updates and maintenance.

Additionally, the following ports are used by the application:

- 5432: between the Flow server and the Postgres database
- 25, 587, 465: between the Flow server and email relay
- 53: TCP/UDP: the application must be able to resolve itself against a DNS server

Note: Host file entries or other workarounds will not work.

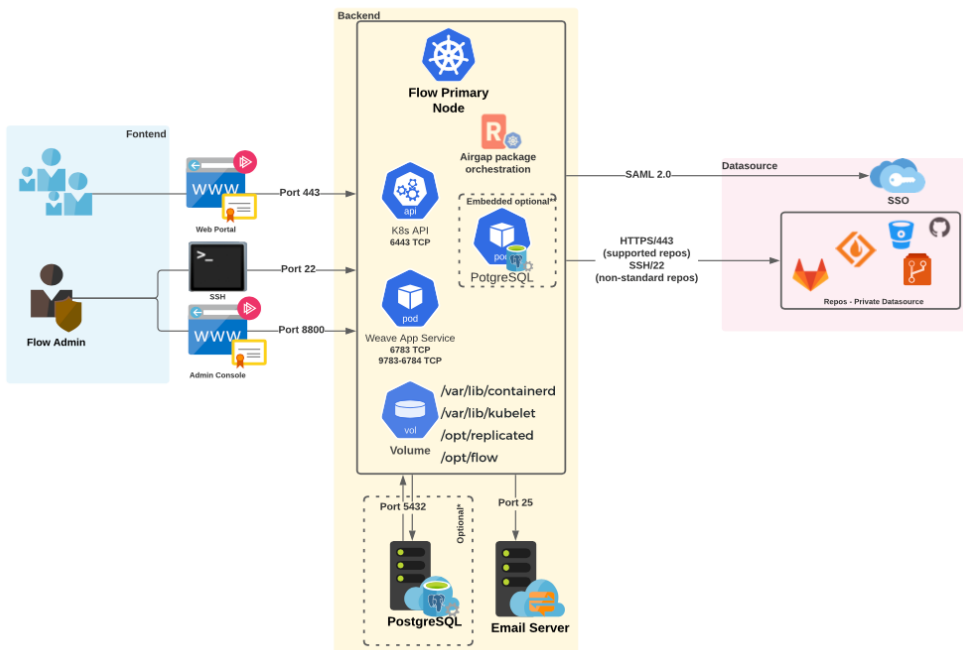
- All ports from the Flow server to itself (usually implicit)

Additional information and limitations

- Only Password Authentication Mode is supported for the KOTS admin console.
- It is possible to install Flow behind an AWS load balancer and a proper security group.

Airgap infrastructure diagram

Flow App - Offline Deployment - airgap



[back to top](#)

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