



Troubleshoot

Ansible Automation Platform 2.6



May 12, 2026

Contents

1 Troubleshoot	4
Access troubleshooting information in the UI	4
Collect configuration and diagnostic information	4
Troubleshoot Ansible Automation Platform on OpenShift Container Platform by using the must-gather command	5
Troubleshoot Ansible Automation Platform on VM-based installations by generating an sos report	6
Troubleshoot your containerized deployment	6
Gather logs from your containerized deployment	6
Diagnose the problem	8
Troubleshoot your containerized installation	12
Troubleshoot your containerized configuration	14
Understand the architecture of your containerized deployment	15
Troubleshoot your Operator-based deployment of Ansible Automation Platform	24
Understand the automation controller operator logs	25
View events in the OpenShift Container Platform	25
View operator logs	25
Configure log verbosity	26
Inspect a OpenShift Container Platform resource	28
Core Ansible Automation Platform resources	28
Standard Kubernetes resources	29
Collect diagnostic data	30
Troubleshoot your RPM-based deployment of Ansible Automation Platform	31
Gather Ansible Automation Platform logs	31
Troubleshoot automation controller	32
Unable to login to automation controller through HTTP	33
Unable to run a job	33
Playbooks do not show up in the Job Template list	33

Troubleshoot

Playbook stays in pending	34
Prevent installation failures when reusing an external database.....	34
View private EC2 VPC instances in the automation controller inventory.....	34
Troubleshoot Red Hat Edge Manager	35
View a device's effective target configuration.....	35
Generate a device log bundle	36
Red Hat product documentation legal notices	37
GNU GENERAL PUBLIC LICENSE	38
Apache license	49

1 Troubleshoot

Access troubleshooting information in the UI

You can use the **Troubleshooting** page to enable or disable certain flags that aid in debugging issues within Ansible Automation Platform.

Procedure

1. From the navigation panel, select **Settings > Automation Execution > Troubleshooting**.
2. The **Troubleshooting** page is displayed.
3. Click **Edit**.
4. You can select the following options:
 - **Enable or Disable tmp dir cleanup**: Select this to enable or disable the cleanup of tmp directories generated during execution of a job after job execution completes.
 - **Debug Web Requests**: Select this to enable or disable web request profiling for debugging slow web requests.
 - **Release Receptor Work**: Select this to turn on or off the deletion of job pods after they complete or fail. This can be helpful in debugging why a job failed.
 - **Keep receptor work on error**: Select this to prevent receptor work from being released when an error is detected.
5. Click **Save** to save your changes.

Collect configuration and diagnostic information

To start troubleshooting Ansible Automation Platform, use the `must-gather` command on OpenShift Container Platform or the `sos` utility on a VM-based installation to collect configuration and diagnostic information. You can attach the output of these utilities to your support case.

Troubleshoot Ansible Automation Platform on OpenShift Container Platform by using the `must-gather` command

The `oc adm must-gather` command line interface (CLI) command collects information from your Ansible Automation Platform installation deployed on OpenShift Container Platform. It gathers information that is often needed for debugging issues, including resource definitions and service logs.

Before you begin

- The OpenShift CLI (`oc`) is installed.

Running the `oc adm must-gather` CLI command creates a new directory containing the collected data that you can use to troubleshoot or attach to your support case.

If your OpenShift environment does not have access to `registry.redhat.io` and you cannot run the `must-gather` command, then run the `oc adm inspect` command instead.

Procedure

1. Log in to your cluster:

```
oc login <openshift_url>
```

2. Run one of the following commands based on your level of access in the cluster:

- Run `must-gather` across the entire cluster:

```
oc adm must-gather --image=registry.redhat.io/ansible-automation-platform-26/aap-must-gather-rhel9 --dest-dir <dest_dir>
```

- `--image` specifies the image that gathers data
- `--dest-dir` specifies the directory for the output

- Run `must-gather` for a specific namespace in the cluster:

```
oc adm must-gather --image=registry.redhat.io/ansible-automation-platform-26/aap-must-gather-rhel9 --dest-dir <dest_dir> -- /usr/bin/ns-gather <namespace>
```

- `-- /usr/bin/ns-gather` limits the `must-gather` data collection to a specified namespace

- To attach the `must-gather` archive to your support case, create a compressed file from the `must-gather` directory created before and attach it to your support case.
 - For example, on a computer that uses a Linux operating system, run the following command, replacing `<must-gather-local.5421342344627712289/>` with the `must-gather` directory name:

```
$ tar cvaf must-gather.tar.gz <must-gather-local.5421342344627712289/>
```

Related information

[Installing the OpenShift CLI](#)

[ocm adm inspect](#)

Troubleshoot Ansible Automation Platform on VM-based installations by generating an sos report

The `sos` utility collects configuration, diagnostic, and troubleshooting data from your Ansible Automation Platform on a VM-based installation.

For more information about installing and using the `sos` utility, see [Generating an sos report for technical support](#).

Troubleshoot your containerized deployment

Use this information to troubleshoot your containerized Ansible Automation Platform installation.

Gather logs from your containerized deployment

Use the `sos` utility to collect configuration and diagnostic data for Red Hat Technical Support. An `sos` report is the standard starting point for troubleshooting.

You can collect an `sos` report for each host in your containerized Ansible Automation Platform deployment by running the `log_gathering` playbook with the appropriate parameters.

Procedure

- Go to the Ansible Automation Platform installation directory.

2. Run the `log_gathering` playbook. This playbook connects to each host in the inventory file, installs the `sos` tool, and generates the `sos` report.

```
$ ansible-playbook -i <path_to_inventory_file>
ansible.containerized_installer.log_gathering
```

3. Optional: To define additional parameters, specify them with the `-e` option. For example:

```
$ ansible-playbook -i <path_to_inventory_file>
ansible.containerized_installer.log_gathering -e
'target_sos_directory=<path_to_files>' -e 'case_number=00000000' -e
'clean=true' -e 'upload=true' -s
```

- a. You can use the `-s` option to step through each task in the playbook and confirm its execution. This is optional but can be helpful for debugging.
- b. The following is a list of the parameters you can use with the `log_gathering` playbook:

Parameter reference

Parameter name	Description	Default
<code>target_sos_directory</code>	Used to change the default location for the <code>sos</code> report files.	<code>/tmp</code> directory of the current server.
<code>case_number</code>	Specifies the support case number if relevant to the log gathering.	
<code>clean</code>	Obfuscates sensitive data that might be present on the <code>sos</code> report.	<code>false</code>
<code>upload</code>	Automatically uploads the <code>sos</code> report data to Red Hat.	<code>false</code>

4. Gather the `sos` report files described in the playbook output and share them with the support engineer or directly upload the `sos` report to Red Hat using the `upload=true` additional parameter.

Related information

[What is an sos report and how to create one in Red Hat Enterprise Linux?](#)

Diagnose the problem

For general container-based troubleshooting, you can inspect the container logs for any running service to help troubleshoot underlying issues.

Identifying the running containers

To get a list of the running container names run the following command:

```
$ podman ps --all --format "{{.Names}}"
```

Container details

Component group	Container name	Purpose
Automation controller	automation-controller-rsyslog	Handles centralized logging for automation controller.
Automation controller	automation-controller-task	Manages and runs tasks related to automation controller, such as running playbooks and interacting with inventories.
Automation controller	automation-controller-web	A web server that provides a REST API for automation controller. This is accessed and routed through platform gateway for user interaction.
Event-Driven Ansible	automation-eda-api	Exposes the API for Event-Driven Ansible, allowing external systems to trigger and manage event-driven automations.
Event-Driven Ansible	automation-eda-daphne	A web server for Event-Driven Ansible, handling WebSocket connections and serving static files.
Event-Driven Ansible	automation-eda-web	A web server that provides a REST API for Event-Driven Ansible. This is accessed and routed through platform gateway for user interaction.
Event-Driven Ansible	automation-eda-worker- <number>	These containers run the automation rules and playbooks based on incoming events.
Event-Driven Ansible	automation-eda-activation-worker- <number>	These containers manage the activation of automation rules, ensuring they run when specific conditions are met.

Component group	Container name	Purpose
Event-Driven Ansible	automation-eda-scheduler	Responsible for scheduling and managing recurring tasks and rule activations.
Platform gateway	automation-gateway-proxy	Acts as a reverse proxy, routing incoming requests to the appropriate Ansible Automation Platform services.
Platform gateway	automation-gateway	Responsible for authentication, authorization, and overall request handling for the platform, all of which is exposed through a REST API and served by a web server.
Automation hub	automation-hub-api	Provides the API for automation hub, enabling interaction with collection content, user management, and other automation hub functionality.
Automation hub	automation-hub-content	Manages and serves Ansible Content Collections, roles, and modules stored in automation hub.
Automation hub	automation-hub-web	A web server that provides a REST API for automation hub. This is accessed and routed through platform gateway for user interaction.
Automation hub	automation-hub-worker- <code><number></code>	These containers handle background tasks for automation hub, such as content synchronization, indexing, and validation.
Performance Co-Pilot	pcp	If Performance Co-Pilot Monitoring is enabled, this container is used for system performance monitoring and data collection.
PostgreSQL	postgresql	Hosts the PostgreSQL database for Ansible Automation Platform.
Receptor	receptor	Facilitates secure and reliable communication within Ansible Automation Platform.
Redis	redis- <code><suffix></code>	Responsible for caching, real-time analytics and fast data retrieval.

Inspecting the logs

Containerized Ansible Automation Platform uses `journald` for Podman logging. To inspect any running container logs, run the `journalctl` command:

```
$ journalctl CONTAINER_NAME=<container_name>
```

Example command with output:

```
$ journalctl CONTAINER_NAME=automation-gateway-proxy

Oct 08 01:40:12 aap.example.org automation-gateway-proxy[1919]: [2024-10-08
00:40:12.885][2][info][upstream] [external/envoy/source/common/upstream/cds_ap>

Oct 08 01:40:12 aap.example.org automation-gateway-proxy[1919]: [2024-10-08
00:40:12.885][2][info][upstream] [external/envoy/source/common/upstream/cds_ap>

Oct 08 01:40:19 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T00:40:16.753Z] "GET /up HTTP/1.1" 200 - 0 1138 10 0 "192.0.2.1" "python-
>
```

To view the logs of a running container in real-time, run the `podman logs -f` command:

```
$ podman logs -f <container_name>
```

Controlling container operations

You can control operations for a container by running the `systemctl` command:

```
$ systemctl --user status <container_name>
```

Example command with output:

```

$ systemctl --user status automation-gateway-proxy
● automation-gateway-proxy.service - Podman automation-gateway-proxy.service
   Loaded: loaded (/home/user/.config/systemd/user/automation-gateway-proxy.service; enabled; preset: disabled)
   Active: active (running) since Mon 2024-10-07 12:39:23 BST; 23h ago
     Docs: man:podman-generate-systemd(1)
   Process: 780 ExecStart=/usr/bin/podman start automation-gateway-proxy (code=exited, status=0/SUCCESS)
   Main PID: 1919 (conmon)
      Tasks: 1 (limit: 48430)
     Memory: 852.0K
         CPU: 2.996s
   CGroup: /user.slice/user-1000.slice/user@1000.service/app.slice/automation-gateway-proxy.service
           └─1919 /usr/bin/conmon --api-version 1 -c
             2dc3c7b2cecd73010bad1e0aaa806015065f92556ed3591c9d2084d7ee209c7a -u
             2dc3c7b2cecd73010bad1e0aaa80>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:02.926Z] "GET /api/galaxy/_ui/v1/settings/ HTTP/1.1" 200 - 0 654
58 47 ">

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:03.387Z] "GET /api/controller/v2/config/ HTTP/1.1" 200 - 0 4018 58
44 "1>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:03.370Z] "GET /api/galaxy/v3/plugin/ansible/search/collection-
versions/?>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:03.405Z] "GET /api/controller/v2/organizations/?
role_level=notification_>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:04.366Z] "GET /api/galaxy/_ui/v1/me/ HTTP/1.1" 200 - 0 1368 79 40
"192.1>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:04.360Z] "GET /api/controller/v2/workflow_approvals/?
page_size=200&statu>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:04.379Z] "GET /api/controller/v2/job_templates/7/ HTTP/1.1" 200 -
0 1356>

Oct 08 11:44:10 aap.example.org automation-gateway-proxy[1919]:
[2024-10-08T10:44:04.378Z] "GET /api/galaxy/_ui/v1/feature-flags/ HTTP/1.1" 200 - 0
207 81>

Oct 08 11:44:13 aap.example.org automation-gateway-proxy[1919]: [2024-10-08
10:44:13.856][2][info][upstream] [external/envoy/source/common/upstream/cds_ap>

Oct 08 11:44:13 aap.example.org automation-gateway-proxy[1919]: [2024-10-08
10:44:13.856][2][info][upstream] [external/envoy/source/common/upstream/cds_ap

```

Getting container information about the execution plane

To get container information about automation controller, Event-Driven Ansible, and `execution_nodes` nodes, prefix any Podman commands with either:

```
CONTAINER_HOST=unix://run/user/<user_id>/podman/podman.sock
```

or

```
CONTAINERS_STORAGE_CONF=<user_home_directory>/aap/containers/storage.conf
```

Example with output:

```
$ CONTAINER_HOST=unix://run/user/1000/podman/podman.sock podman images
```

REPOSITORY	IMAGE ID	CREATED	SIZE	TAG
registry.redhat.io/ansible-automation-platform-25/ee-supported-rhel8	59d1bc680a7c	6 days ago	2.24 GB	latest
registry.redhat.io/ansible-automation-platform-25/ee-minimal-rhel8	a64b9fc48094	6 days ago	338 MB	latest

Troubleshoot your containerized installation

Use this information to troubleshoot your containerized installation of Ansible Automation Platform.

The installation takes a long time, or has errors, what should I check?

1. Ensure your system meets the minimum requirements as outlined in [System requirements](#). Factors such as improper storage choices and high latency when distributing across many hosts will all have an impact on installation time.
2. Review the installation log file which is located by default at `./aap_install.log`. You can change the log file location within the `ansible.cfg` file in the installation directory.
3. Enable task profiling callbacks on an ad hoc basis to give an overview of where the installation program spends the most time. To do this, use the local `ansible.cfg` file. Add a callback line under the `[defaults]` section, for example:

```
$ cat ansible.cfg
[defaults]
callbacks_enabled = ansible.posix.profile_tasks
```

Automation controller returns an error of 413

This error occurs when `manifest.zip` license files that are larger than the `controller_nginx_client_max_body_size` setting. If this error occurs, update the inventory file to include the following variable:

```
controller_nginx_client_max_body_size=5m
```

The default setting of `5m` should prevent this issue, but you can increase the value as needed.

When attempting to install containerized Ansible Automation Platform in Amazon Web Services you receive output that there is no space left on device

```
TASK [ansible.containerized_installer.automationcontroller : Create the receptor
container] *****

fatal: [ec2-13-48-25-168.eu-north-1.compute.amazonaws.com]: FAILED! => {"changed":
false, "msg": "Can't create container receptor", "stderr": "Error: creating
container storage: creating an ID-mapped copy of layer
\"98955f43cc908bd50ff43585fec2c7dd9445eaf05eec1e3144f93ffc00ed4ba\": error during
chown: storage-chown-by-maps: lchown usr/local/lib/python3.9/site-packages/azure/
mgmt/network/v2019_11_01/operations/__pycache__/_
_available_service_aliases_operations.cpython-39.pyc: no space left on device: exit
status 1\n", "stderr_lines": ["Error: creating container storage: creating an ID-
mapped copy of layer
\"98955f43cc908bd50ff43585fec2c7dd9445eaf05eec1e3144f93ffc00ed4ba\": error during
chown: storage-chown-by-maps: lchown usr/local/lib/python3.9/site-packages/azure/
mgmt/network/v2019_11_01/operations/__pycache__/_
_available_service_aliases_operations.cpython-39.pyc: no space left on device: exit
status 1"], "stdout": "", "stdout_lines": []}
```

If you are installing a `/home` filesystem into a default Amazon Web Services marketplace RHEL instance, it might be too small since `/home` is part of the root `/` filesystem. To resolve this issue you must make more space available. For more information about the system requirements, see [System requirements](#).

"Install container tools" task fails due to unavailable packages

This error can be seen in the installation process output as the following:

```
TASK [ansible.containerized_installer.common : Install container tools]
*****
*****

fatal: [192.0.2.1]: FAILED! => {"changed": false, "failures": ["No package crun
available.", "No package podman available.", "No package slirp4netns available.",
"No package fuse-overlayfs available."], "msg": "Failed to install some of the
specified packages", "rc": 1, "results": []}

fatal: [192.0.2.2]: FAILED! => {"changed": false, "failures": ["No package crun
available.", "No package podman available.", "No package slirp4netns available.",
"No package fuse-overlayfs available."], "msg": "Failed to install some of the
specified packages", "rc": 1, "results": []}

fatal: [192.0.2.3]: FAILED! => {"changed": false, "failures": ["No package crun
available.", "No package podman available.", "No package slirp4netns available.",
"No package fuse-overlayfs available."], "msg": "Failed to install some of the
specified packages", "rc": 1, "results": []}

fatal: [192.0.2.4]: FAILED! => {"changed": false, "failures": ["No package crun
available.", "No package podman available.", "No package slirp4netns available.",
"No package fuse-overlayfs available."], "msg": "Failed to install some of the
specified packages", "rc": 1, "results": []}

fatal: [192.0.2.5]: FAILED! => {"changed": false, "failures": ["No package crun
available.", "No package podman available.", "No package slirp4netns available.",
"No package fuse-overlayfs available."], "msg": "Failed to install some of the
specified packages", "rc": 1, "results": []}
```

To fix this error, run the following command on the target hosts:

```
sudo subscription-manager register
```

Troubleshoot your containerized configuration

Use this information to troubleshoot your containerized Ansible Automation Platform configuration.

Sometimes the post install for seeding my Ansible Automation Platform content errors out

This could manifest itself as output similar to this:

```
TASK [infra.controller_configuration.projects : Configure Controller Projects | Wait
for finish the projects creation] *****
```

```
Friday 29 September 2023 11:02:32 +0100 (0:00:00.443) 0:00:53.521 *****
```

```
FAILED - RETRYING: [daap1.lan]: Configure Controller Projects | Wait for finish the
projects creation (1 retries left).
```

```
failed: [daap1.lan] (item={'failed': 0, 'started': 1, 'finished': 0,
'ansible_job_id': '536962174348.33944', 'results_file': '/home/aap/.ansible_async/
536962174348.33944', 'changed': False, '__controller_project_item': {'name': 'AAP
Config-As-Code Examples', 'organization': 'Default', 'scm_branch': 'main',
'scm_clean': 'no', 'scm_delete_on_update': 'no', 'scm_type': 'git',
'scm_update_on_launch': 'no', 'scm_url': 'https://github.com/user/repo.git'},
'ansible_loop_var': '__controller_project_item'}) =>
{"__projects_job_async_results_item": {"__controller_project_item": {"name": "AAP
Config-As-Code Examples", "organization": "Default", "scm_branch": "main",
"scm_clean": "no", "scm_delete_on_update": "no", "scm_type": "git",
"scm_update_on_launch": "no", "scm_url": "https://github.com/user/repo.git"},
"ansible_job_id": "536962174348.33944", "ansible_loop_var":
 "__controller_project_item", "changed": false, "failed": 0, "finished": 0,
"results_file": "/home/aap/.ansible_async/536962174348.33944", "started": 1},
"ansible_job_id": "536962174348.33944", "ansible_loop_var":
 "__projects_job_async_results_item", "attempts": 30, "changed": false, "finished":
0, "results_file": "/home/aap/.ansible_async/536962174348.33944", "started": 1,
"stderr": "", "stderr_lines": [], "stdout": "", "stdout_lines": []}
```

The `infra.controller_configuration.dispatch` role uses an asynchronous loop with 30 retries to apply each configuration type, and the default delay between retries is 1 second. If the configuration is large, this might not be enough time to apply everything before the last retry occurs.

Increase the retry delay by setting the `controller_configuration_async_delay` variable to 2 seconds for example. You can set this variable in the `[all:vars]` section of the installation program inventory file.

Re-run the installation program to ensure everything works as expected.

Understand the architecture of your containerized deployment

Use this information to understand the architecture for your containerized Ansible Automation Platform deployment.

Can you give details of the architecture for the Ansible Automation Platform containerized design?

We use as much of the underlying Red Hat Enterprise Linux technology as possible. Podman is used for the container runtime and management of services.

Use `podman ps` to list the running containers on the system.

Use `podman images` to display information about locally stored images.

Containerized Ansible Automation Platform runs as rootless containers for enhanced security by default. This means you can install containerized Ansible Automation Platform by using any local unprivileged user account. Privilege escalation is only needed for certain root level tasks, and by default is not needed to use root directly.

The installation program adds the following files to the filesystem where you run the installation program on the underlying Red Hat Enterprise Linux host:

```
$ tree -L 1
.
├── aap_install.log
├── ansible.cfg
├── collections
├── galaxy.yml
├── inventory
├── LICENSE
├── meta
├── playbooks
├── plugins
├── README.md
├── requirements.yml
└── roles
```

The installation root directory includes other containerized services that make use of Podman volumes.

Here are some examples for further reference:

The `containers` directory includes some of the Podman specifics used and installed for the execution plane:

```

containers/
├── podman
├── storage
│   ├── defaultNetworkBackend
│   ├── libpod
│   ├── networks
│   ├── overlay
│   ├── overlay-containers
│   ├── overlay-images
│   ├── overlay-layers
│   ├── storage.lock
│   └── usersns.lock
└── storage.conf

```

The `controller` directory has some of the installed configuration and runtime data points:

```

controller/
├── data
│   ├── job_execution
│   ├── projects
│   └── rsyslog
├── etc
│   ├── conf.d
│   ├── launch_awx_task.sh
│   ├── settings.py
│   ├── tower.cert
│   └── tower.key
├── nginx
│   └── etc
├── rsyslog
│   └── run
└── supervisor
    └── run

```

The `receptor` directory has the automation mesh configuration:

```
receptor/  
├─ etc  
│ └─ receptor.conf  
└─ run  
    ├─ receptor.sock  
    └─ receptor.sock.lock
```

After installation, you will also find other files in the local user's `/home` directory such as the `.cache` directory:

```
.cache/  
├─ containers  
│ └─ short-name-aliases.conf.lock  
└─ rhsm  
    └─ rhsm.log
```

As services are run using rootless Podman by default, you can use other services such as running `systemd` as non-privileged users. Under `systemd` you can see some of the component service controls available:

The `.config` directory:

```

.config/
├─ cni
│  └─ net.d
│     └─ cni.lock
├─ containers
│  ├─ auth.json
│  └─ containers.conf
└─ systemd
   └─ user
      ├─ automation-controller-rsyslog.service
      ├─ automation-controller-task.service
      ├─ automation-controller-web.service
      ├─ default.target.wants
      ├─ podman.service.d
      ├─ postgresql.service
      ├─ receptor.service
      ├─ redis.service
      └─ sockets.target.wants

```

This is specific to Podman and conforms to the Open Container Initiative (OCI) specifications. When you run Podman as the root user `/var/lib/containers` is used by default. For standard users the hierarchy under `$HOME/.local` is used.

The `.local` directory:

```

.local/
└─ share
   └─ containers
      ├─ cache
      ├─ podman
      └─ storage

```

As an example `.local/storage/volumes` contains what the output from `podman volume ls` provides:

```
$ podman volume ls

DRIVER      VOLUME NAME
local       d73d3fe63a957bee04b4853fd38c39bf37c321d14fdab9ee3c9df03645135788
local       postgresql
local       redis_data
local       redis_etc
local       redis_run
```

The execution plane is isolated from the control plane main services to ensure it does not affect the main services.

Control plane services run with the standard Podman configuration and can be found in:

```
~/.local/share/containers/storage.
```

Execution plane services (automation controller, Event-Driven Ansible and execution nodes) use a dedicated configuration found in `~/aap/containers/storage.conf`. This separation prevents execution plane containers from affecting the control plane services.

You can view the execution plane configuration with one of the following commands:

```
CONTAINERS_STORAGE_CONF=~/aap/containers/storage.conf podman <subcommand>
```

```
CONTAINER_HOST=unix:///run/user/<user uid>/podman/podman.sock podman <subcommand>
```

How can I see host resource utilization statistics?

Run the following command to display host resource utilization statistics:

```
$ podman container stats -a
```

Example output based on a Dell sold and offered containerized Ansible Automation Platform solution (DAAP) install that utilizes ~1.8 GB RAM:

ID	NET IO	NAME BLOCK IO	PIDS	CPU % CPU TIME	MEM USAGE / LIMIT AVG CPU %	MEM %
0d5d8eb93c18	0B / 0B	automation-controller-web 0B / 0B	16	0.23% 20.885142s	959.1MB / 3.761GB 1.19%	25.50%
3429d559836d	0B / 0B	automation-controller-rsyslog 0B / 0B	6	0.07% 4.099565s	144.5MB / 3.761GB 0.23%	3.84%
448d0bae0942	0B / 0B	automation-controller-task 0B / 0B	33	1.51% 34.285272s	633.1MB / 3.761GB 1.93%	16.83%
7f140e65b57e	0B / 0B	receptor 0B / 0B	7	0.01% 1.010613s	5.923MB / 3.761GB 0.06%	0.16%
c1458367ca9c	0B / 0B	redis 0B / 0B	5	0.48% 9.074042s	10.52MB / 3.761GB 0.47%	0.28%
ef712cc2dc89	0B / 0B	postgresql 0B / 0B	21	0.09% 15.571059s	21.88MB / 3.761GB 0.80%	0.58%

How much storage is used and where?

The container volume storage is under the local user at `$HOME/.local/share/containers/storage/volumes`.

1. To view the details of each volume, run the following command:

```
$ podman volume ls
```

2. Run the following command to display detailed information about a specific volume:

```
$ podman volume inspect <volume_name>
```

For example:

```
$ podman volume inspect postgresql
```

Example output:

```
[
  {
    "Name": "postgresql",
    "Driver": "local",
    "Mountpoint": "/home/aap/.local/share/containers/storage/volumes/postgresql/
_data",
    "CreatedAt": "2024-01-08T23:39:24.983964686Z",
    "Labels": {},
    "Scope": "local",
    "Options": {},
    "MountCount": 0,
    "NeedsCopyUp": true
  }
]
```

Several files created by the installation program are located in `$HOME/aap/` and bind-mounted into various running containers.

1. To view the mounts associated with a container run the following command:

```
$ podman ps --format "{{.ID}}\t{{.Command}}\t{{.Names}}"
```

Example output:

```

89e779b81b83    run-postgresql    postgresql
4c33cc77ef7d    run-redis         redis
3d8a028d892d    /usr/bin/receptor...    receptor
09821701645c    /usr/bin/launch_a...    automation-controller-rsyslog
a2ddb5cac71b    /usr/bin/launch_a...    automation-controller-task
fa0029a3b003    /usr/bin/launch_a...    automation-controller-web
20f192534691    gunicorn --bind 1...    automation-eda-api
f49804c7e6cb    daphne -b 127.0.0...    automation-eda-daphne
d340b9c1cb74    /bin/sh -c nginx ...    automation-eda-web
111f47de5205    aap-eda-manage rq...    automation-eda-worker-1
171fcb1785af    aap-eda-manage rq...    automation-eda-worker-2
049d10555b51    aap-eda-manage rq...    automation-eda-activation-worker-1
7a78a41a8425    aap-eda-manage rq...    automation-eda-activation-worker-2
da9afa8ef5e2    aap-eda-manage sc...    automation-eda-scheduler
8a2958be9baf    gunicorn --name p...    automation-hub-api
0a8b57581749    gunicorn --name p...    automation-hub-content
68005b987498    nginx -g daemon o...    automation-hub-web
cb07af77f89f    pulpcore-worker    automation-hub-worker-1
a3ba05136446    pulpcore-worker    automation-hub-worker-2

```

2. Run the following command:

```
$ podman inspect <container_name> | jq -r .[].Mounts[].Source
```

Example output:

```
/home/aap/.local/share/containers/storage/volumes/receptor_run/_data
/home/aap/.local/share/containers/storage/volumes/redis_run/_data
/home/aap/aap/controller/data/rsyslog
/home/aap/aap/controller/etc/tower.key
/home/aap/aap/controller/etc/conf.d/callback_receiver_workers.py
/home/aap/aap/controller/data/job_execution
/home/aap/aap/controller/nginx/etc/controller.conf
/home/aap/aap/controller/etc/conf.d/subscription_usage_model.py
/home/aap/aap/controller/etc/conf.d/cluster_host_id.py
/home/aap/aap/controller/etc/conf.d/insights.py
/home/aap/aap/controller/rsyslog/run
/home/aap/aap/controller/data/projects
/home/aap/aap/controller/etc/settings.py
/home/aap/aap/receptor/etc/receptor.conf
/home/aap/aap/controller/etc/conf.d/execution_environments.py
/home/aap/aap/tls/extracted
/home/aap/aap/controller/supervisor/run
/home/aap/aap/controller/etc/uwsgi.ini
/home/aap/aap/controller/etc/conf.d/container_groups.py
/home/aap/aap/controller/etc/launch_awx_task.sh
/home/aap/aap/controller/etc/tower.cert
```

3. If the `jq` RPM is not installed, install it by running the following command:

```
$ sudo dnf -y install jq
```

Troubleshoot your Operator-based deployment of Ansible Automation Platform

Use the commands and procedures in this section to diagnose and resolve common issues with your Ansible Automation Platform deployment on OpenShift Container Platform. Topics include viewing logs, inspecting resources, and collecting diagnostic data for Red Hat Support.

Understand the automation controller operator logs

If an **AutomationController** instance fails to deploy, check the `automation-controller-operator` container logs. These logs contain the installer role output required to debug deployment issues.

View events in the OpenShift Container Platform

You can view events in the OpenShift Container Platform web console to monitor for errors and troubleshoot issues. This helps you quickly diagnose problems by examining the status of custom resources and their related events.

You can debug by first reviewing the status conditions of the Ansible Automation Platform custom resource (CR) and then checking any nested CRs for errors.

Procedure

1. Log in to the OpenShift Container Platform web console.
2. In the navigation menu, select **Home > Events**.
3. Select your project from the project list.
4. To view events for a specific resource, navigate to that resource's page. Many resource pages, such as pods and deployments, have their own **Events** tab.
5. Select a resource to bring you to the **Pod Details** page.

Result

Check the **Conditions** section on the **Pod details** page to confirm no errors are listed in the **Message** column.

View operator logs

The following procedure is an example of how to view the logs for an `automation-controller-operator` pod.

Procedure

1. To find the pod name, run:

```
oc get pods | grep operator
```

2. To view the logs for the pod, run:

```
oc logs <operator-pod-name> -f
```

- a. Alternatively, to view the logs without first getting the pod name, run:

```
oc logs deployments/automation-controller-operator-controller-manager -c
automation-controller-manager -f
```

Configure log verbosity

You can enable task output for debugging on any custom resources (CRs) by setting `no_log` to `false` in the component section of the `AnsibleAutomationPlatform` CR.

The logs then show output for any failed tasks that originally had `no_log` set to `true`. All Ansible Automation Platform components (automation controller, automation hub, and Event-Driven Ansible) support the `no_log` setting.

Procedure

1. Edit the Ansible Automation Platform CR and set the `no_log` field to `false` for the component you want to debug.

For automation controller:

```
apiVersion: aap.ansible.com/v1alpha1
kind: AnsibleAutomationPlatform
metadata:
  name: myaap
spec:
  controller:
    no_log: false
```

For automation hub:

```

apiVersion: aap.ansible.com/v1alpha1
kind: AnsibleAutomationPlatform
metadata:
  name: myaap
spec:
  hub:
    no_log: false

```

For Event-Driven Ansible

```

apiVersion: aap.ansible.com/v1alpha1
kind: AnsibleAutomationPlatform
metadata:
  name: myaap
spec:
  eda:
    no_log: false

```

NOTE:

This might expose sensitive data in the logs. On production clusters, this value must generally be set to `true` unless you are actively debugging an issue.

2. To increase the Ansible Playbook verbosity from the operator, set the verbosity level using an annotation on the Ansible Automation Platform CR:

```

apiVersion: aap.ansible.com/v1alpha1
kind: AnsibleAutomationPlatform
metadata:
  name: myaap
  annotations:
    ansible.sdk.operatorframework.io/verbosity: "4"
spec:
  # ... component configuration ...

```

Inspect a OpenShift Container Platform resource

To inspect a OpenShift Container Platform resource, you must use the `oc` command to get a summary or the full YAML definition of the resource.

Procedure

1. To view a human-readable summary of a resource, run:

```
oc describe -n <namespace> <resource> <resource-name>
```

2. To view the complete YAML definition of a resource, use the `-o yaml` flag:

```
oc get -n <namespace> <resource> <resource-name> -o yaml
```

- For example, to get the YAML for the `automationcontroller` custom resource, run:

```
oc get -n aap automationcontroller aap -o yaml
```

Core Ansible Automation Platform resources

The following table lists and describes the core custom resources (CRs) that the Ansible Automation Platform Operator manages. Understanding these resources will help you with advanced troubleshooting and configuration.

Resource name	Description
<code>ansibleautomationplatform</code>	CR for deploying the entire Ansible Automation Platform.
<code>ansibleautomationplatformbackup</code>	CR for creating backups of the entire Ansible Automation Platform instance.
<code>ansibleautomationplatformrestore</code>	CR for restoring the entire Ansible Automation Platform instance from a backup.

Resource name	Description
<code>automationcontroller</code>	CR defining the desired state of an automation controller instance.
<code>automationcontrollerbackup</code>	CR for creating backups of automation controller data and configuration.
<code>automationcontrollerrestore</code>	CR for restoring the automation controller from a backup.
<code>automationhub</code>	CR for deploying an automation hub (Galaxy) instance.
<code>automationhubbackup</code>	CR for creating backups of automation hub data and configuration.
<code>automationhubrestore</code>	CR for restoring automation hub from a backup.
<code>eda</code>	CR for deploying an Event-Driven Ansible (EDA) instance.
<code>edabackup</code>	CR for creating backups of EDA data and configuration.
<code>edarestore</code>	CR for restoring EDA from a backup.
<code>ansiblelightspeed</code>	CR for deploying an Red Hat Ansible Lightspeed instance.

Standard Kubernetes resources

Standard Kubernetes resources are a core part of the OpenShift Container Platform. The following table describes the standard resources you can inspect to troubleshoot the state and configuration of an application.

Resource name	Description
<code>pod</code>	Smallest deployable unit containing one or more containers running the application workloads.
<code>deployment</code>	Manages pod configuration and scaling.

Resource name	Description
<code>pvc</code>	A PersistentVolumeClaim (PVC) is a request for storage resources, used for persistent data storage.
<code>service</code>	Exposes pods as network services with stable IP addresses and DNS names within the cluster.
<code>ingress</code>	Manages external HTTP and HTTPS access to services within the cluster.
<code>route</code>	An OpenShift-specific resource for exposing services externally (similar to an ingress).
<code>secrets</code>	Stores sensitive data like passwords, tokens, and certificates.
<code>serviceaccount</code>	Provides identity for processes running in pods to access permissions to other Kubernetes resources.

Collect diagnostic data

Use the `oc adm must-gather` command to collect comprehensive diagnostic data about your cluster and the Ansible Automation Platform components. This data is essential when contacting Red Hat Support.

Procedure

1. To start the `must-gather` tool, run:

```
oc adm must-gather --image=registry.redhat.io/ansible-automation-platform-25/aap-must-gather-rhel8
```

NOTE:

For version 2.6, the base image name changes to `registry.redhat.io/ansible-automation-platform-26/aap-must-gather-rhel9`.

2. View the collected data, use the `omc` tool to query the `must-gather` tarball as if it were a live cluster.

```
omc use <path-to-must-gather>
omc get pods
```

Related information

[How to collect diagnostics data from Ansible Automation Platform running on OpenShift?](#)

[How to generate a sos report within nodes without SSH in OpenShift Container Platform 4](#)

Troubleshoot your RPM-based deployment of Ansible Automation Platform

Resolve common installation issues and errors that can occur when installing RPM-based Ansible Automation Platform. Learn how to generate diagnostic logs to identify problems.

Gather Ansible Automation Platform logs

Run the setup script with the `-s` flag to collect diagnostic data from all nodes. This generates an `sos` report used by Red Hat Technical Support to resolve service requests.

Procedure

1. Access the installation program folder with the inventory file and run the installation program setup script the following command:

```
$ ./setup.sh -s
```

With this command, you can connect to each node present in the inventory, install the `sos` tool, and generate new logs.

NOTE:

If you are running the setup as a non-root user with sudo privileges, you can use the following command:

```
$ ANSIBLE_BECOME_METHOD='sudo'
ANSIBLE_BECOME=True ./setup.sh -s
```

2. *Optional:* If required, change the location of the `sos` report files.

The `sos` report files are copied to the `/tmp` folder for the current server. To change the location, specify the new location by using the following command:

```
$ ./setup.sh -e 'target_sos_directory=/path/to/files' -s
```

Where `target_sos_directory=/path/to/files` is used to specify the destination directory where the `sos` report will be saved. In this case, the `sos` report is stored in the directory `/path/to/files`.

3. Gather the files described on the playbook output and share with the support engineer or directly upload the `sos` report to Red Hat.

To create an `sos` report with additional information or directly upload the data to Red Hat, use the following command:

```
$ ./setup.sh -e 'case_number=0000000' -e 'clean=true' -e 'upload=true' -s
```

Parameter Reference Table

Parameter	Description	Default value
<code>case_number</code>	Specifies the support case number that you want.	-
<code>clean</code>	Obfuscates sensitive data that might be present on the <code>sos</code> report.	<code>false</code>
<code>upload</code>	Automatically uploads the <code>sos</code> report data to Red Hat.	<code>false</code>

Related information

[KCS article:What is an sos report and how to create one in Red Hat Enterprise Linux?](#)

Troubleshoot automation controller

Useful troubleshooting information for automation controller.

Unable to login to automation controller through HTTP

Access to automation controller is intentionally restricted through a secure protocol (HTTPS).

In cases where your configuration is set up to run an automation controller node behind a load balancer or proxy as "HTTP only", and you only want to access it without SSL/TLS (for troubleshooting, for example), you must add the following settings in the `custom.py` file located at `/etc/tower/conf.d` of your automation controller instance:

```
SESSION_COOKIE_SECURE = False
CSRF_COOKIE_SECURE = False
```

If you change these settings to `false` it enables automation controller to manage cookies and login sessions when using the HTTP protocol. You must do this on every node of a cluster installation.

To apply the changes, run:

```
automation-controller-service restart
```

Unable to run a job

If you are unable to run a job from a playbook, review the playbook YAML file. When importing a playbook, either manually or by a source control mechanism, remember that the host definition is controlled by automation controller and should be set to `hosts:all`.

Playbooks do not show up in the Job Template list

If your playbooks are not showing up in the **Job Template** list, check the following:

- Ensure that the playbook is valid YML and can be parsed by Ansible.
- Ensure that the permissions and ownership of the project path (`/var/lib/awx/projects`) is set up so that the "awx" system user can view the files. Run the following command to change the ownership:

```
chown awx -R /var/lib/awx/projects/
```

Playbook stays in pending

If you are attempting to run a playbook job and it stays in the `Pending` state indefinitely, try the following actions:

- Ensure that all supervisor services are running through `supervisorctl status`.
- Ensure that the `/var/` partition has more than 1 GB of space available. Jobs do not complete with insufficient space on the `/var/` partition.
- Run `automation-controller-service restart` on the automation controller server.

If you continue to have issues, run `sosreport` as root on the automation controller server, then file a [support request](#) with the result.

Prevent installation failures when reusing an external database

When reusing an external database for clustered installations, you must manually clear the database before performing subsequent installations.

Instances have been reported where reusing the external database during subsequent installation of nodes causes installation failures.

Example

You perform a clustered installation. Then, you need to do this again and perform a second clustered installation reusing the same external database, only this subsequent installation failed.

When setting up an external database that has been used in a prior installation, you must manually clear the database used for the clustered node before any additional installations can succeed.

View private EC2 VPC instances in the automation controller inventory

By default, automation controller only shows instances in a VPC that have an Elastic IP (EIP) associated with them.

Procedure

1. From the navigation panel, select **Automation Execution > Infrastructure > Inventories**.
2. Select the inventory that has the **Source** set to **Amazon EC2**, and click the **Source** tab.
In the **Source Variables** field, enter:

```
vpc_destination_variable: private_ip_address
```

3. Click **Save** and trigger an update of the group.

Result

Once you complete these steps, you can see your VPC instances.

NOTE:

Automation controller must be running inside the VPC with access to those instances if you want to configure them.

Troubleshoot Red Hat Edge Manager

When working with devices in Red Hat Edge Manager, you might see issues related to configuration, connectivity, or deployment. Troubleshooting these issues requires understanding how device configurations are applied, how to check logs, and how to verify communication between the device and service.

View a device's effective target configuration

The device manifest returned by the `flightctl get device` command still only has references to external configuration and secret objects. Only when the device agent queries the service, the service replaces the references with the actual configuration and secret data.

While this better protects potentially sensitive data, it also makes troubleshooting faulty configurations hard. This is why a user can be authorized to query the effective configuration as rendered by the service to the agent.

- To query the effective configuration, use the following command:

```
flightctl get device/${device_name} --rendered | jq
```

Generate a device log bundle

Use the integrated `flightctl-must-gather` script directly on the device to generate a comprehensive bundle of diagnostic logs. This log bundle, in a standard `.tar` format, provides the necessary data to debug the device agent and assists in efficient troubleshooting and bug reporting.

- Run the following command on the device and include the `.tar` file in the bug report.

NOTE:

This depends on an SSH connection to extract the `.tar` file.

```
sudo flightctl-must-gather
```

Red Hat product documentation legal notices

Copyright © Red Hat

Except as otherwise noted below, the text of and illustrations in this documentation are licensed by Red Hat under the [Creative Commons Attribution–Share Alike 3.0 Unported license](#). If you distribute this document or an adaptation of it, you must provide the URL for the original version. Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Linux® is the registered trademark of Linus Torvalds in the United States and other countries.

XFS is a trademark or registered trademark of Hewlett Packard Enterprise Development LP or its subsidiaries in the United States and other countries.

The OpenStack® Word Mark and OpenStack logo are trademarks or registered trademarks of the Linux Foundation, used under license.

All other trademarks are the property of their respective owners.

GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. link:<https://fsf.org/>. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If

such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS

0. Definitions. "This License" refers to version 3 of the GNU General Public License.

"Copyright" also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

"The Program" refers to any copyrightable work licensed under this License. Each licensee is addressed as "you". "Licensees" and "recipients" may be individuals or organizations.

To "modify" a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a "modified version" of the earlier work or a work "based on" the earlier work.

A "covered work" means either the unmodified Program or a work based on the Program.

To "propagate" a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To "convey" a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays "Appropriate Legal Notices" to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion.

1. Source Code. The "source code" for a work means the preferred form of the work for making modifications to it. "Object code" means any non-source form of a work.

A "Standard Interface" means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The "System Libraries" of an executable work include anything, other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component,

or to implement a Standard Interface for which an implementation is available to the public in source code form. A "Major Component", in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The "Corresponding Source" for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work's System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work.

2. Basic Permissions. All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary.

3. Protecting Users' Legal Rights From Anti-Circumvention Law. No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work's users, your or third parties' legal rights to forbid circumvention of technological measures.

4. Conveying Verbatim Copies. You may convey verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for each copy that you convey, and you may offer support or warranty protection for a fee.

5. Conveying Modified Source Versions. You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

- a) The work must carry prominent notices stating that you modified it, and giving a relevant date.
- b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to "keep intact all notices".
- c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it.
- d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so. A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an "aggregate" if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation's users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate.

6. Conveying Non-Source Forms. You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

- a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange.
- b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge.

- c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b.
- d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements.
- e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects or is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you convey an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying.

7. Additional Terms. “Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

- a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or
- b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or
- c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or
- d) Limiting the use for publicity purposes of names of licensors or authors of the material; or
- e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or
- f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way.

8. Termination. You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10.

9. Acceptance Not Required for Having Copies. You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so.

10. Automatic Licensing of Downstream Recipients. Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An "entity transaction" is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party's predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it.

11. Patents. A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor’s “contributor version”.

A contributor’s “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor’s essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express permission to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient’s use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law.

12. No Surrender of Others' Freedom. If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy both those terms and this License would be to refrain entirely from conveying the Program.

13. Use with the GNU Affero General Public License. Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such.

14. Revised Versions of this License. The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License "or any later version" applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version.

15. Disclaimer of Warranty. THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. Limitation of Liability. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES

SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

17. Interpretation of Sections 15 and 16. If the disclaimer of warranty and limitation of liability provided above cannot be given local legal effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>
```

```
Copyright (C) <year> <name of author>
```

```
This program is free software: you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation, either version 3 of the License, or  
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program. If not, see <https://www.gnu.org/licenses/>.
```

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author>  
This program comes with ABSOLUTELY NO WARRANTY; for details type `show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, your program's commands might be different; for a GUI interface, you would use an "about box".

You should also get your employer (if you work as a programmer) or school, if any, to sign a "copyright disclaimer" for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see link:<https://www.gnu.org/licenses/>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read link:<https://www.gnu.org/licenses/why-not-lgpl.html>.

Apache license

Version 2.0, January 2004

<http://www.apache.org/licenses/>

Terms and Conditions for use, reproduction, and distribution

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, **"control"** means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or **"Your"**) shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, **"submitted"** means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the

Licensors for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as **"Not a Contribution."**

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

1. You must give any other recipients of the Work or Derivative Works a copy of this License; and
2. You must cause any modified files to carry prominent notices stating that You changed the files; and
3. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
4. If the Work includes a **"NOTICE"** text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications,

or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS



Copyright 2026. All rights reserved.

www.redhat.com