



Comply

Contents

Welcome to Puppet Comply®.....	6
Comply terminology.....	6
Comply overview.....	7
Supported CIS Benchmarks.....	8
CIS-CAT Pro Assessor.....	8
CIS-CAT Pro Assessor history.....	8
End-of-life announcement: Comply (SCM) version 2.x.....	27
Release notes.....	28
Comply release notes.....	28
Comply known issues.....	46
Beginner’s guide to Comply.....	47
Puppet Application Manager.....	49
Welcome to Puppet Application Manager (PAM).....	50
Release notes.....	52
PAM release notes.....	52
Known issues.....	65
Architecture overview.....	65
PAM system requirements.....	69
Component versions in PAM releases.....	82
Install PAM.....	83
Install Puppet applications using PAM on a customer-supported Kubernetes cluster.....	84
PAM HA online installation.....	86
PAM HA offline installation.....	90
PAM standalone online installation.....	93
PAM standalone offline installation.....	96
Automate PAM and Puppet application online installations.....	98
Automate PAM and Puppet application offline installations.....	100
Uninstall PAM.....	104
Working with Puppet applications.....	104
Install applications via the PAM UI.....	105
Update a license for online installations.....	106
Update a license for offline installations.....	106
Upgrade an automated online application installation.....	106
Upgrade an automated offline application installation.....	107
Maintenance and tuning.....	108
Upgrading PAM on a Puppet-supported cluster.....	109
Upgrading PAM on a customer-supported cluster.....	114
Backing up PAM using snapshots.....	115
Migrating PAM data to a new system.....	118
Disaster recovery with PAM.....	124
Troubleshooting PAM.....	125

Installing.....	130
System requirements.....	130
Hosting the CIS-CAT Pro Assessor bundle internally.....	131
Set up Comply.....	132
Configure Comply (online environment).....	133
Configure Comply in an offline environment.....	133
Configure Comply TLS certificates.....	134
Configure Comply for a custom NGINX ingress (online environment).....	135
Configure Comply for a custom NGINX ingress (offline environment).....	136
Configure Comply TLS certificates for a custom NGINX ingress.....	137
Install the Comply module.....	139
Classify nodes.....	140
Specify your own Java binary.....	140
Deploy Comply.....	140
Add PE credentials.....	141
Configure inventory refresh interval.....	142
Configure data retention policy.....	142
Uninstall Comply and remove PAM.....	142
Uninstall Comply without removing PAM.....	142
Remove the CIS-CAT Pro Assessor from a node.....	143
Upgrading.....	143
Using the Comply API.....	146
Manage personal access tokens.....	146
Create and manage personal access tokens as a user.....	147
Manage personal access tokens as an admin.....	147
Authenticate public APIs.....	147
REST API.....	148
REST API tutorial.....	148
Extract Compliance results using the Comply API.....	151
Export Comply data using the Comply API.....	151
Synchronize inventory with Puppet Enterprise using the Comply API.....	151
Managing access for Comply users.....	152
Guidelines for running Comply at scale.....	154
Desired compliance.....	156
Custom profiles.....	158
Exceptions.....	159
CIS scans.....	161

Run an ad hoc scan.....	161
Scheduled scans.....	162
View details about a scan schedule.....	162
Pause and resume a scan schedule.....	163
Edit a scan schedule.....	163
Delete a scan schedule.....	163
Add or remove nodes in scheduled scans.....	163
Create a one-off scan schedule.....	164
Create a repeating scan schedule.....	164
CIS scan reports.....	165
CIS scan report details.....	166
Scan results.....	166
Enforce CIS benchmarks.....	169
Troubleshooting.....	169
Introducing the Compliance Enforcement Modules.....	173
CEM for Linux.....	174
Release notes.....	174
Known issues and limitations.....	187
Getting started.....	190
Basic concepts.....	190
Next steps.....	192
Installing CEM.....	192
Prepare to install the module.....	193
Install and evaluate the module in a test environment.....	194
Install the module in the production environment.....	195
Uninstall the module.....	195
Upgrading CEM.....	196
Prepare to upgrade the module.....	196
Upgrade the module.....	197
Configuring CEM.....	198
Overview of configuration options.....	198
How to configure the module: Examples and guidelines.....	201
Auditing and querying issues identified during scans.....	209
Reference: Benchmarks and controls.....	210
Control updates introduced for Red Hat Enterprise Linux 8 STIG, Version 1, Release 11.....	210
Control updates introduced for Red Hat Enterprise Linux 7 STIG, Version 3, Release 12.....	211
CEM for Windows.....	212
Release notes.....	212
Known issues and limitations.....	219
Getting started.....	220
Basic concepts.....	221
Next steps.....	222
Installing CEM.....	222
Prepare to install the module.....	222

Install and evaluate the module in a test environment.....	223
Install the module in the production environment.....	224
Uninstall the module.....	225
Upgrading CEM.....	225
Prepare to upgrade the module.....	225
Upgrade the module.....	227
Configuring CEM.....	227
Overview of configuration options.....	228
How to configure the module: Examples and guidelines.....	229
Reference: Benchmarks and controls.....	233
Control updates introduced for CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0.....	233
Control updates introduced for CIS Microsoft Windows Server 2019 Benchmark v2.0.0.....	234
Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v2.0.0.....	235
Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v1.4.0.....	236

Copyright and trademark notices.....242

Welcome to Puppet Comply®

Puppet Comply is a tool that assesses the infrastructure you manage with Puppet Enterprise® against CIS Benchmarks — the best practices from the Center for Internet Security (CIS) for securely configuring systems.

Using Comply, you can:

- Run scans to check the compliance of your infrastructure against CIS Benchmarks on up to 100,000 nodes.
- Set your desired compliance — a default benchmark and profile that you want your scans to be measured against.
- Customize profiles to specify which rules you want visible in scan reports.
- Identify the cause and source of compliance failures, and determine what configuration changes must be made to which systems.

Comply uses Puppet Enterprise (PE) to retrieve node and fact information. After you install Comply, you must configure it to integrate with PE.

If this is your first time using Comply, try out our [Beginner's guide to Comply](#).

Important: Before you use the product and its documentation, review the [Copyright and trademark notices](#) on page 242.

Puppet Comply docs links	Other useful places
<p>Learn the basics:</p> <ul style="list-style-type: none"> Comply overview Comply terminology Beginner's guide to Comply Release notes <p>Install and configure Comply:</p> <ul style="list-style-type: none"> System requirements Install Puppet Application Manager Set up Comply <p>Run and manage CIS scans:</p> <ul style="list-style-type: none"> Run scans and review reports Set desired compliance Create a custom profile on page 158 View scan results 	<p>Comply videos:</p> <ul style="list-style-type: none"> Comply introduction and demo <p>Related Puppet products:</p> <ul style="list-style-type: none"> Puppet Enterprise Puppet Forge <p>Get help:</p> <ul style="list-style-type: none"> Troubleshooting Support portal

Comply terminology

Learn the key terms that are associated with Puppet Comply.

CIS Benchmarks

Developed by the Center for Internet Security (CIS), *CIS Benchmarks* are internationally recognized standards and best practices for securely configuring systems. For more information, see [CIS Benchmarks](#).

CIS assessor

Comply integrates with the *CIS assessor* (CIS-CAT PRO), the scanner tool that assesses CIS benchmarks. As part of the Comply configuration process, Puppet Enterprise (PE) installs the CIS assessor on your target nodes. For more information on the assessor, see [CIS-CAT Pro](#).

Profiles

CIS Benchmarks include different levels of security settings, called *profiles*. The *Level 1* profiles are the base recommendation for every system, and the *Level 2* profiles are intended for environments requiring greater security. Comply can scan for either profile.

Rules

Each profile contains multiple *rules* that define specific elements of system configuration.

Custom profiles

A *custom profile* is a benchmark profile that you customize to fit your organization's internally defined standards, by specifying which rules you want visible in scan reports. Once you create a custom profile, it appears as an option in Comply when selecting a benchmark and profile.

Desired compliance

Desired compliance is the benchmark and profile that you assign to a node. It becomes the default scan for that node.

For a full list of Puppet® terminology, see the [Puppet Glossary](#).

Comply overview

Welcome to Puppet Comply!

This overview is intended for new users of Comply. We go over what Comply is, how it works, and show a demo of the 1.0.0 release. Before you begin, we recommend familiarizing yourself with our [terminology](#).

What is Comply and how does it work?

Comply is a tool that expands the compliance capabilities of Puppet Enterprise (PE), by integrating with the *CIS assessor* to scan your infrastructure against the latest *CIS Benchmarks*. Comply connects to your PE environment and gathers information about your PE managed nodes, including operating system facts and classification node groups. It uses this information to suggest appropriate scans.

You can choose to run ad-hoc scans or *desired compliance* scans — a default CIS benchmark and profile scan that you assign to a node. Comply can automate desired compliance for you based on the information it gathers about your nodes from PE, or you can manually choose your desired compliance from a list of benchmarks and profiles. You can also create *custom profiles* to fit internally defined standards, by specifying which rules you want visible in scan reports. Most of the time, you only need to set your desired compliance once.

The scans are run as a *task* in PE. Scan results populate in the Comply **Compliance dashboard**, where you can see the number of nodes scanned and their compliance breakdown. In each node listed, there is a further breakdown of rule information which tells you why that rule is important, and steps you can take to fix the rule if it is failing the scans.

To see Comply in action, watch the demo below, or go through the steps yourself in our [beginner's guide](#).

For a full list of features, see the [release notes](#).

Supported CIS Benchmarks

Comply supports the following CIS operating system benchmarks.

Operating system	Supported versions
Alma Linux	8, 9
Amazon Linux	2, 2 STIG, 2023
Azure Compute Microsoft Windows Server	2019, 2022
Debian Linux	11 (v1.0.0), 11 STIG, 12
macOS	x86_64 processor: 12, 13, 14, 15 ARM processor: Not supported
Oracle Linux	8, 9
Red Hat Enterprise Linux (RHEL)	8, 8 STIG, 9, 9 STIG
Rocky Linux	8, 9
SUSE Linux Enterprise Server (SLES)	12, 15
Ubuntu	20.04 LTS, 20.04 LTS STIG, 22.04 LTS, 22.04 LTS STIG, 24.04 LTS, 24.04 LTS STIG
Windows	2016, 2016 STIG, 2019, 2019 STIG, 2019 (for stand-alone systems)**, 2022, 2022 STIG, 2022 (for stand-alone systems)**, 10 (for stand-alone systems)**, 10 Enterprise, 11 (for stand-alone systems)**, 11 Enterprise

**** Microsoft Windows users:** Starting with Comply 2.8.0, you can apply the CIS Microsoft Windows 10 and 11 Stand-alone Benchmarks to stand-alone systems, which are not connected to a domain.

CIS-CAT Pro Assessor

The CIS Benchmarks are set of best practices to help organizations securely configure IT systems, software, networks, and cloud infrastructure. The CIS-CAT Pro Assessor by the [Center for Internet Security](#) allows you to scan systems against these CIS benchmarks and CIS controls and report on their levels of compliance.

Puppet Comply uses the CIS-CAT Pro Assessor as its scanner tool to assess compliance with CIS Benchmarks. The assessor is installed as an integral part of the Comply setup process. During the Comply configuration, Puppet Enterprise automatically installs the CIS-CAT Pro Assessor on the target nodes you want scanned for compliance. This enables comprehensive security assessments across various operating systems and platforms. The assessor's regular updates and wide benchmark coverage allow you to perform up-to-date compliance scans based on different security profiles.

The CIS-CAT Pro Assessor uses an embedded Java Runtime Environment (JRE) for its operation. This embedded JRE allows the assessor to run independently without relying on the system's installed Java version, ensuring consistency across different environments where Comply is deployed. For more information, please visit the [CIS-CAT Pro Assessor page](#) from the Center for Internet Security.

CIS-CAT Pro Assessor history

Comply supports the latest and previous version of the CIS-CAT Pro Assessor.

CIS-CAT Pro Assessor v4.55.0

Checksum Linux: ba91327eadba75c8a566c595e2225085fa9e761f3f82636b5d8cf206c5135b24

Checksum Mac: 9f4c904b5d01181653b8519fe733c4c2d580b5c52a3dbf0fdd848501b934ca49

Checksum Windows: 94fd9a38a6eddf57fbfb78fe4205c800a8043bd5233e1bbad4cd4c90d173a75

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS AlmaLinux OS 9 Benchmark v2.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Amazon Linux 2023 Benchmark v1.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Cloud-tailored Benchmark v1.1.0
- CIS Apple macOS 13.0 Ventura Benchmark v3.1.0
- CIS Apple macOS 14.0 Sonoma Benchmark v2.1.0
- CIS Apple macOS 14.0 Sonoma Cloud-tailored Benchmark v1.1.0
- CIS Apple macOS 15.0 Sequoia Benchmark v1.1.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Debian Linux 12 Benchmark v1.1.0
- CIS Microsoft Windows 10 Enterprise Benchmark v4.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v4.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v4.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v4.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Benchmark v4.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v4.0.0
- CIS Microsoft Windows Server 2022 Stand-alone Benchmark v1.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v2.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 9 STIG Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS Rocky Linux 9 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.2.1
- CIS SUSE Linux Enterprise 15 Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS Benchmark v3.0.0
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS STIG Benchmark v1.0.0
- CIS Ubuntu Linux 24.04 LTS Benchmark v1.0.0
- CIS Ubuntu Linux 24.04 LTS STIG Benchmark v1.0.0

CIS-CAT Pro Assessor v4.52.0**Checksum Linux:** 5b0773a5b27759d55d6f01b45051f09f41b8550be91239c1faecda2a49c6629e**Checksum Mac:** d7ebcd072f371798dded1f1448dff4bd628dcc1c4cc6b99d1464f4d1b34aec40**Checksum Windows:** e055ec6183e75181e0822abcec5ce0535c4d37b9e4e9c5276b2a3a6fe38f5aa7**Benchmarks:**

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS AlmaLinux OS 9 Benchmark v2.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Amazon Linux 2023 Benchmark v1.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Cloud-tailored Benchmark v1.1.0
- CIS Apple macOS 13.0 Ventura Benchmark v3.0.0
- CIS Apple macOS 14.0 Sonoma Benchmark v2.0.0
- CIS Apple macOS 14.0 Sonoma Cloud-tailored Benchmark v1.1.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Debian Linux 12 Benchmark v1.1.0
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v4.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Benchmark v3.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v3.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v2.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS Rocky Linux 9 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.2.0
- CIS SUSE Linux Enterprise 15 Benchmark v2.0.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.2.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v3.0.0

- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0
- CIS Ubuntu Linux 24.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.47.0

Checksum Linux: 9f88ff4faa3b9ae03a95c2dfb4d84ada7c8daf07f2622be1149d4447719c383d

Checksum Mac: 29fe0194de02a6d6bb22ed266d13db3d259ee75f2489470c349d6ccfeada68ac

Checksum Windows: 172c46cc2c8b25c8037ba937a64a002c42f69736212f55d2772fc9456d0deed5

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS AlmaLinux OS 9 Benchmark v2.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Amazon Linux 2023 Benchmark v1.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.1.0
- CIS Apple macOS 13.0 Ventura Benchmark v3.0.0
- CIS Apple macOS 14.0 Sonoma Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Debian Linux 12 Benchmark v1.1.0
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Benchmark v3.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v3.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS Rocky Linux 9 Benchmark v2.0.0

- CIS SUSE Linux Enterprise 12 Benchmark v3.2.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.2.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0
- CIS Ubuntu Linux 24.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.43.0

Checksum Linux: b165c4fda3bfa1dccc5d17016486f2207b5842c46aecbb8807b1a47552a5e229

Checksum Mac: 47b844b7f251de1174287cd45d1f5e49c6cf0cbea174556e536a71d238fe3fb7

Checksum Windows: 0574ffa8c274148365d7824835f1ecbd3fa0d208b0283dbb174f287546294946

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.1.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.1.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Debian Linux 12 Benchmark v1.0.1
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v3.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v3.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0

- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.2.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0

CIS-CAT Pro Assessor v4.42.0

Checksum Linux: 17ff1f330a33abc6d50fc7b0ea1485134a0ee150f6cb1c78957bedef6f3382c0

Checksum Mac: 9c6136e324309ce39ecbb89322562e108b76602eccdb16abc319ae0fa70946bb

Checksum Windows: 94d3959a78c826559a4fc5ac6047fbe9a130bb7fe009c03eb004a4d1895856e2

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Debian Linux 12 Benchmark v1.0.1
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v3.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
- CIS Microsoft Windows Server 2022 Benchmark v3.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0

- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.2.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0

CIS-CAT Pro Assessor v4.41.0

Checksum Linux: af3826fb58bb5b2eab8af9ae31984d02d528f783b14254354824da6353366b44

Checksum Mac: 79fd572025e27dbc2512d3d40b647d3d8fe338ddec85c2b455765fa6de155a54

Checksum Windows: 7d9fe127362a1d2b6791aa6f50a6064f5e9f202ff87dac64bdfa1c1b63b6cd4d

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v3.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v3.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v3.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.2.0

- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v2.0.0

CIS-CAT Pro Assessor v4.39.0

Checksum Linux: ec693a8701bc4e5c8ae2b1cba8dad4df53b4d00079dcff90ec85b5b696b9347

Checksum Mac: 7d7a6ea94a704664dda2f45a87558bd5d8ac1275632a03fe510581e446b36deb

Checksum Windows: d7f30296434a0a0f4d942edfd08c4a40308a4a7a9b3c1f8459aa5e95bcbc7f56

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v3.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0

- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.38.0

Checksum Linux: cc42bdf5e07e46b348843cc109d463591db7515b5680c7edcc4925d15b4747ce

Checksum Mac: 2ddf7a327e9714beaa7ae59b79e4180e4c912fd59da5bc7d8be35f2aaad0df5d

Checksum Windows: 6dc6c85da4ad93ef054faa30f27d2737d864e73c2e3a534bb98e7d5e0ddc8c63

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.37.0**Checksum Linux:** 925facfca8484abe8f424380b8c95679dfc6359237e17a0086375dcb49474ed6**Checksum Mac:** 586dbb1f959b6571f294e2de74e72f8a5bb5d86b08bd5aee9da5c0f82bb15804**Checksum Windows:** 87760bb71c37d0cafca094f5176e80eedb4496cd0583eb0efdc99230692f4632**Benchmarks:**

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v3.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v4.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0
- CIS Oracle Linux 7 Benchmark v4.0.0
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.36.0**Checksum Linux:** 31f28068c77bf49dc4538c16b25f1437e3c10bad8bbb205eaed49fa7c05a5d56

Checksum Mac: 4461315f7e9caa081a7dca935af55ff94e1cf5167c16abd4da332016db30c0ac

Checksum Windows: 4ff47d969e4bba6c826035bdb25c54a1c87ca0ed09d4dcd624b325bff8e5ff74

Benchmarks:

- CIS AlmaLinux OS 8 Benchmark v3.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0 ***REMOVED***
- CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0 ***NEW***
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Oracle Linux 8 Benchmark v3.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v2.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

CIS-CAT Pro Assessor v4.34.0

Checksum Linux: 807f5c2868b70928d66d18b5aced9381d63fcff5d1fbbdcc775e756a5d886ded

Checksum Mac: fd1f7ba6997c22b733a619b6a2b2957c7ed18654664d99efd522f344b080bc09

Checksum Windows: 9f09693f73bde1569c8fe4f14ea8659af219b65ba05c880dc50b48b598849746

Benchmarks:

- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.1.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.1.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS Alma Linux OS 8 Benchmark v2.0.0

CIS-CAT Pro Assessor v4.33.0

Checksum Linux: 6ecaa39bbfb2c54ddd7c17a0007b4c16ca8513ce62acf63b1eb34e34b409b44

Checksum Mac: 6ecaa39bbfb2c54ddd7c17a0007b4c16ca8513ce62acf63b1eb34e34b409b44

Checksum Windows: 72ea2030badc8ff742e10f281913537d1d0d40a14246dd717c884e3418951a09

Benchmarks:

- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.1.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.1.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v2.0.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0
- CIS Debian Linux 11 STIG Benchmark v1.0.0

CIS-CAT Pro Assessor v4.32.0

Checksum Linux: 27d08f3b004a4b98654fe63246ba17643a2ef23814fdeed1f2a8d56f9fdaba47

Checksum Mac: 27d08f3b004a4b98654fe63246ba17643a2ef23814fdeed1f2a8d56f9fdaba47

Checksum Windows: ce3b06d528345dc5aa59defa79ad81c096d24af7bfc6f3cda215b82376dee065

Benchmarks:

- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0

- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.1.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.1.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

CIS-CAT Pro Assessor v4.30.0

Checksum Linux: 21b9fd770c0b04aa4a39bade909f5abdd63f9f88bd4a27525b0753ef11f64069

Checksum Mac: 21b9fd770c0b04aa4a39bade909f5abdd63f9f88bd4a27525b0753ef11f64069

Checksum Windows: 763055fd2bb47f4b86473e65171d0f909d72f8e33e3cb4cf160718cc4944dad2

Benchmarks:

- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v2.0.0
- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0

- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Microsoft Windows Server 2019 Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.1
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Microsoft Windows Server 2022 Benchmark v2.0.0
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Stand-alone Benchmark v2.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows Server 2016 Benchmark v2.0.0
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

CIS-CAT Pro Assessor v4.28.0

Checksum Linux: 1b393f6efcb0908a6b62d0a2db522dc0875f4df3089d338cae48948410f2585f

Checksum Mac: 1b393f6efcb0908a6b62d0a2db522dc0875f4df3089d338cae48948410f2585f

Checksum Windows: 99e950263e4f0aa0e2beafcbdb2cf1650056e26cfed7aee43489756b326dd958

Benchmarks:

- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Microsoft Windows Server 2019 Benchmark v1.3.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0

- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Debian Linux 10 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows Server 2016 Benchmark v1.4.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v1.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v1.1.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Microsoft Windows 11 Enterprise Benchmark v2.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows 10 Stand-alone Benchmark v1.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

CIS-CAT Pro Assessor v4.27.0

Checksum (Linux): 3686201a61b6b23b8fe473ba04d6c1235619906bff90943b1abf474a1339b1e6

Checksum (macOS): 3686201a61b6b23b8fe473ba04d6c1235619906bff90943b1abf474a1339b1e6

Checksum (Windows): b8b07a7a98b4c886ce7eccc22a81a6f0c2627cfce9ef6b67899cee179df0cea

Benchmarks:

- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Azure Compute Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Microsoft Windows Server 2019 Benchmark v1.3.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0

- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Debian Linux 10 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows Server 2016 Benchmark v1.4.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v1.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Microsoft Windows 11 Enterprise Benchmark v1.0.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v1.1.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows 10 Stand-alone Benchmark v1.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v1.12.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

Java version: The CIS-CAT Pro Assessor uses embedded JRE. The JRE version details for this release are as follows:

- ```
./java -version
openjdk version "1.8.0_352"
```

OpenJDK Runtime Environment (Temurin)(build 1.8.0\_352-b08)

OpenJDK 64-Bit Server VM (Temurin)(build 25.352-b08, mixed mode)

### **CIS-CAT Pro Assessor v4.25.0**

**Checksum:** ab0bc45d6a7f1c9a297f6c9ecd82c067d9d80c666990309123ba985b4bbaf155

#### **Benchmarks:**

- CIS Oracle Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Microsoft Windows Server 2019 Benchmark v1.3.0

- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 9 Benchmark v1.0.0
- CIS Ubuntu Linux 16.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows Server 2016 Benchmark v1.4.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows 11 Stand-alone Benchmark v1.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Microsoft Windows 11 Enterprise Benchmark v1.0.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v1.1.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v2.0.0
- CIS Apple macOS 11.0 Big Sur Benchmark v3.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows 10 Stand-alone Benchmark v1.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v1.12.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

### **CIS-CAT Pro Assessor v4.23.0**

**Checksum:** e8c1ae0519c89c420c27622d4497c12e772d82c3ea2c6952abd9185acafe532f

#### **Benchmarks:**

- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Microsoft Windows Server 2019 Benchmark v1.3.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Amazon Linux Benchmark v2.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v1.1.0

- CIS Ubuntu Linux 16.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows Server 2016 Benchmark v1.4.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0
- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 10.15 Catalina Benchmark v2.1.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Microsoft Windows 11 Enterprise Benchmark v1.0.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v1.1.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows 10 Stand-alone Benchmark v1.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v1.12.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Apple macOS 11.0 Big Sur Benchmark v2.1.0
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

### **CIS-CAT Pro Assessor v4.22.0**

**Checksum:** 9331f00f1b481fdd161512f258dcc0002ca5893d5a7c807dfc0379751071408

#### **Benchmarks:**

- CIS Ubuntu Linux 22.04 LTS Benchmark v1.0.0
- CIS Microsoft Windows Server 2019 Benchmark v1.3.0
- CIS Red Hat Enterprise Linux 8 STIG Benchmark v1.0.0
- CIS Amazon Linux 2 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 R2 Benchmark v2.6.0
- CIS Amazon Linux Benchmark v2.0.0
- CIS Apple macOS 12.0 Monterey Benchmark v1.1.0
- CIS Ubuntu Linux 16.04 LTS Benchmark v2.0.0
- CIS Debian Linux 10 Benchmark v1.0.0
- CIS SUSE Linux Enterprise 12 Benchmark v3.1.0
- CIS Microsoft Windows Server 2016 Benchmark v1.4.0
- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 8 Benchmark v2.0.0

- CIS Debian Linux 11 Benchmark v1.0.0
- CIS Amazon Linux 2 STIG Benchmark v2.0.0
- CIS Apple macOS 10.15 Catalina Benchmark v2.1.0
- CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0
- CIS Microsoft Windows Server 2012 (non-R2) Benchmark v2.4.0
- CIS Microsoft Windows Server 2016 STIG Benchmark v1.2.0
- CIS Microsoft Windows 11 Enterprise Benchmark v1.0.0
- CIS Ubuntu Linux 20.04 LTS Benchmark v1.1.0
- CIS Microsoft Windows Server 2019 STIG Benchmark v1.1.0
- CIS CentOS Linux 7 Benchmark v3.1.2
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Microsoft Windows Server 2022 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
- CIS Azure Compute Microsoft Windows Server 2019 Benchmark v1.0.0
- CIS Rocky Linux 8 Benchmark v1.0.0
- CIS Red Hat Enterprise Linux 7 Benchmark v3.1.1
- CIS Microsoft Windows 10 Stand-alone Benchmark v1.0.1
- CIS Ubuntu Linux 20.04 LTS STIG Benchmark v1.0.0
- CIS Microsoft Windows 10 Enterprise Benchmark v1.12.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS SUSE Linux Enterprise 15 Benchmark v1.1.1
- CIS Ubuntu Linux 18.04 LTS Benchmark v2.1.0
- CIS Oracle Linux 7 Benchmark v3.1.1
- CIS Apple macOS 11.0 Big Sur Benchmark v2.1.0
- CIS Debian Linux 9 Benchmark v1.0.1
- CIS Alma Linux OS 8 Benchmark v2.0.0

## End-of-life announcement: Comply (SCM) version 2.x

---

Announcing the end-of-life (EOL) for Comply version 2.x.

Version 2.x of Comply reaches end-of-life (EOL) on February 5, 2026. Going forward, development will focus on the latest version, SCM 3.x, which will continue to receive full support, updates, and enhancements.

### Next steps

- **[Migrate to SCM 3.x](#)**

If your organization is currently using Comply (SCM) version 2.x, we recommend migrating to SCM 3.x as soon as possible to avoid disruption and ensure continued support. Migrating to SCM 3.x reduces dependency on external applications, improving efficiency and reliability, while retaining or enhancing all core features.

- **Active maintenance period**

Until February 5, 2026, we'll continue to provide quarterly updates for critical bugs and security fixes for Comply (SCM) version 2.x, as well as migration support for customers transitioning to SCM 3.x.

- **Extended maintenance option**

Customers who require more time to complete their migration may opt into an additional six months of extended support for Comply (SCM) 2.x, available as a premium offering. Please contact your account team for more information.

- **End of support**

After the extended support period ends, no further updates, fixes or support will be provided for Comply (SCM) 2.x. We strongly encourage planning your migration in advance to avoid disruption.

## We're here to help

Our Professional Services, Support, and Engineering teams are available to assist with migration planning, technical questions, and any challenges you may encounter during this transition.

# Release notes

---

Learn about the new features, enhancements, and resolved issues for the Puppet Comply 2.x release series.

## Comply release notes

---

These are the new features, enhancements, and resolved issues for the Puppet Comply 2.x release series.

### Comply 2.25.0

Released 11 July 2025.

New in this release:

- **Secret management for Podman.** Secrets no longer stored in plain text files under `/etc/puppetlabs/comply/` on the install target.
- **Specify your own Java binary.** Comply 2.25.0 adds the ability to [specify your own Java binary](#).
- **CIS-CAT Pro Assessor v4.55.0.** Comply 2.25.0 contains the CIS-CAT Pro Assessor v4.55.0.
- Benchmarks added in this release:
  - Apple macOS 15.0 Sequoia Benchmark v1.1.0
  - Microsoft Windows Server 2022 Stand-alone Benchmark v1.0.0
  - Red Hat Enterprise Linux 9 STIG Benchmark v1.0.0
  - Ubuntu Linux 22.04 LTS STIG Benchmark v1.0.0
  - Ubuntu Linux 24.04 LTS STIG Benchmark v1.0.0
- Benchmarks updated in this release:
  - Apple macOS 13.0 Ventura Benchmark v3.1.0
  - Apple macOS 14.0 Sonoma Benchmark v2.1.0
  - Microsoft Windows 10 Enterprise Benchmark v4.0.0
  - Microsoft Windows 10 Stand-alone Benchmark v4.0.0
  - Microsoft Windows 11 Stand-alone Benchmark v4.0.0
  - Microsoft Windows Server 2019 Benchmark v4.0.0
  - Microsoft Windows Server 2022 Benchmark v4.0.0
  - SUSE Linux Enterprise 12 Benchmark v3.2.1
- Benchmarks removed in this release:
  - Apple macOS 11.0 Big Sur Benchmark v4.0.0
  - CentOS Linux 7 Benchmark v4.0.0
  - Debian Linux 10 Benchmark v2.0.0
  - Oracle Linux 7 Benchmark v4.0.0
  - Red Hat Enterprise Linux 7 Benchmark v4.0.0
  - Red Hat Enterprise Linux 7 STIG Benchmark v2.0.0
  - Ubuntu Linux 18.04 LTS Benchmark v2.2.0

Resolved in this release:

- **Installation could stall when trying to bring up containers.** Fixed an issue where the installation of could stall when IPV6 was improperly enabled. SCM now properly disables IPV6 when an IPV6 network is not detected to ensure IPV6 installs are more robust.

- **PostgreSQL container can be unexpectedly stopped.** Fixed an issue where on a sufficiently slow startup, it was possible that the PostgreSQL container could prematurely time out before the startup sequence completes.
- **After upgrading to Puppet Agent 8.13.1, agent runs now return errors.** Fixed an issue where Puppet Agent would return errors after upgrading Puppet Agent to 8.13.1. Comply now properly parses facts as UTF8 as required within Puppet 8.

Security fixes in this release:

- **CVE-2025-23419.** Updated NGINX to 1.28.0 to address this vulnerability.
- **CVE-2025-3501.** Updated Keycloak to 26.2.5 to address this vulnerability.

## Comply 2.24.0

Released 30 April 2025.

New in this release:

- **CIS-CAT Pro Assessor v4.52.0.** Comply 2.24.0 contains the CIS-CAT Pro Assessor v4.52.0.
- Benchmarks updated in this release:
  - Apple macOS 12.0 Monterey v4.0.0
  - Apple macOS 12.0 Monterey Cloud-tailored v1.1.0
  - Apple macOS 14.0 Sonoma Cloud-tailored v1.1.0
  - Microsoft Windows 11 Enterprise v4.0.0
  - Microsoft Windows Server 2019 STIG v3.0.0
  - Microsoft Windows Server 2022 STIG v2.0.0
  - Red Hat 8 STIG v2.0.0
  - SUSE Linux Enterprise 15 v2.0.1
  - Ubuntu Linux 20.04 LTS v3.0.0
- Benchmarks removed in this release:
  - CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
  - CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
- Java update:
  - The JRE included in the assessor bundle is updated to Amazon Corretto v8.442.06.1.

Resolved in this release:

- **Export service is not functioning as expected.** Fixed an issue affecting the export service, exports are now properly exported and can be downloaded on the **Exported data** page.
- **User session is not invalidated after logging out.** Fixed an issue to correctly invalidate a user session upon logging out.
- **Scan shows success, but no results shown.** Fixed an issue where, during a scan, the console output indicated a success status yet the results were reported as 0%, suggesting no assessments were performed. This was due to an OS mismatch between the node's OS and the benchmark being used. To address this, a puppet fact `ignore_platform_mismatch = true` can be applied to the node group where the node belongs. When set to true this allows the scan to proceed despite platform differences, ensuring results are generated.

Security fixes in this release:

- **CVE-2024-21208, CVE-2024-21210, CVE-2024-21235, CVE-2024-21217.** Updated Amazon Corretto to 8.442.06.1 to address these vulnerabilities.
- **CVE-2024-21538.** Updated `cross-spawn` to 7.0.5 to address this vulnerability.
- **CVE-2024-45337.** Updated `golang.org/x/crypto` to v0.31.0 to address this vulnerability.
- **CVE-2025-30153.** Updated `kin-openapi` to v0.131.0 to address this vulnerability.

## Comply 2.23.0

Released 11 December 2024.

New in this release:

- **Additions to the Comply API.** Improved accessibility and efficiency of Comply functions by adding two new endpoints to the Comply API:
  - **Profiles API.** Added four new endpoints (`/v1/profiles`, `/v1/profiles/{id}`, `/v1/custom-profiles`, `/v1/custom-profiles/{id}`) to retrieve information about benchmark profiles in Comply.
  - **Custom Scan API.** Added a new endpoint (`/v1/custom-scan`) to create ad-hoc custom scans, which can be run with a custom profile or a profile and benchmark.
- **Added a task to remove old versions of the CIS-CAT Pro Assessor.** Added the `remove_assessor` task that allows you to remove old versions of the CIS-CAT Pro Assessor that are no longer in use.
- **CIS-CAT Pro Assessor v4.47.0.** Comply 2.23.0 contains the CIS-CAT Pro Assessor v4.47.0.
- Benchmarks updated in this release:
  - AlmaLinux OS 9 Benchmark v2.0.0
  - Amazon Linux 2023 Benchmark v1.0.0
  - Apple macOS 13.0 Ventura Benchmark v3.0.0
  - Apple macOS 14.0 Sonoma Benchmark v2.0.0
  - Debian Linux 12 Benchmark v1.1.0
  - Microsoft Windows Server 2016 STIG Benchmark v3.0.0
  - Rocky Linux 9 Benchmark v2.0.0
  - SUSE Linux Enterprise 12 Benchmark v3.2.0
  - Ubuntu Linux 24.04 LTS Benchmark v1.0.0

Resolved in this release:

- **Server failing to complete Comply scan.** Fixed an issue where scans would not complete when running against machines running CIS-CAT Pro Assessor versions 4.37.0 and later.

**Note:** For Linux and macOS users, this fix requires the `timeout` package, which is installed on Linux hosts by default. Users on macOS may need to install this package.

- **Remove broken link for Assessor download.** Removed a broken link for the Assessor download from the **Settings** page.
- **Default Windows 2022 benchmark incorrectly set.** Fixed an issue affecting assessor versions 4.39.0, 4.41.0, 4.42.0, and 4.43.0 where the default benchmark for Windows 2022 was incorrectly set to Azure.
- **Comply fails to install when IPv6 is disabled.** Fixed an issue where Comply NGINX components incorrectly depended on IPv6, preventing installation on systems with IPv6 disabled.
- **Restore from backup failing.** Fixed a permission issue causing restore from backup to fail due to issues copying database backups into the container on the target host.
- **Incorrect number of scanned nodes displayed.** Fixed an issue where the **Nodes** tab on the **Scan reports** page could display an incorrect number of nodes: sometimes showing more or fewer nodes than were scanned.

Security fixes in this release:

- **CVE-2013-2028, CVE-2021-23017.** Updated NGINX to 1.27.2 to address these vulnerabilities.

## Comply 2.22.0

Released 16 August 2024.

New in this release:

- **Added application configuration to Comply.** You can now set the `inventory refresh interval` and `data retention policy` from the **Settings** page.
- **Scheduled scans now support node groups.** Node groups can now be added to scheduled scans. When a node is added to a node group, it is automatically included in the next scheduled scan for that node group.
- **CIS-CAT Pro Assessor v4.43.0.** Comply 2.22.0 contains the CIS-CAT Pro Assessor v4.43.0.

- Benchmarks updated in this release:
  - Apple macOS 12.0 Monterey Benchmark v3.1.0
  - Apple macOS 13.0 Ventura Benchmark v2.1.0
  - Microsoft Windows Server 2019 Stand-alone v2.0.0
  - Oracle Linux 9 Benchmark v2.0.0
  - Red Hat Enterprise Linux 9 Benchmark v2.0.0

Security fixes in this release:

- **CVE-2023-2976.** Updated KeyCloak to 25.0.0 to address this vulnerability.

## Comply 2.21.0

Released 27 June 2024.

New in this release:

- **Desired compliance can be set for operating systems.** You can now [set the desired compliance defaults for each operating system](#). Any node added to the operating system is automatically assigned the benchmark and profile you set for that operating system.
- **REST API documentation updated.** Added a [REST API tutorial](#) to the Comply documentation.
- **CIS-CAT Pro Assessor v4.42.0.** Comply 2.21.0 contains the CIS-CAT Pro Assessor v4.42.0.
- Benchmarks updated in this release:
  - Debian Linux 12 Benchmark v1.0.1
  - Microsoft Windows 11 Stand-alone Benchmark v3.0.0
  - Microsoft Windows Server 2019 Benchmark v3.0.1

Resolved in this release:

- **Exceptions disappearing upon upgrade to 2.20.0.** Fixed an issue that caused existing rule exceptions to disappear after upgrading.
- **Search box on exceptions page not accepting input.** Fixed an issue affecting the search bar on the exceptions page.
- **macOS not getting desired benchmark assigned.** Fixed an issue that was causing macOS nodes to be listed as Darwin on the **Inventory** page, which prevented the desired compliance from being set for those nodes.
- **Upgrades from older versions of Comply to 2.20 not working.** Fixed an issue where the scarp container would not start following an upgrade from an older version.

Security fixes in this release:

- **CVE-2024-4068.** Updated braces to address this vulnerability.
- **CVE-2024-2961, CVE-2024-33599, CVE-2024-2700, CVE-2024-1132, CVE-2024-1249, CVE-2024-2419, CVE-2024-3656, GHSA-69fp-7c8p-crjr.** Updated KeyCloak to address these vulnerabilities.
- **CVE-2023-5363.** Updated oauth2-proxy to address this vulnerability.

## Comply 2.20.0

Released 7 May 2024.

New in this release:

- **CIS-CAT Pro Assessor v4.41.0.** Comply 2.20.0 contains the CIS-CAT Pro Assessor v4.41.0.

- Benchmarks updated in this release:
  - Debian Linux 11 Benchmark v2.0.0
  - Microsoft Windows 10 Stand-alone Benchmark v3.0.0
  - Microsoft Windows Server 2016 Benchmark v3.0.0
  - Microsoft Windows Server 2019 Benchmark v3.0.0
  - Microsoft Windows Server 2022 Benchmark v3.0.0
  - Ubuntu Linux 18.04 LTS Benchmark v2.2.0
  - Ubuntu Linux 22.04 LTS Benchmark v2.0.0

Resolved in this release:

- **Unable to reset the desired compliance when a node changes operating systems.** Fixed an issue where you could not change the desired compliance after changing the OS on a node. You can now reset the desired compliance on a node when the OS of the node changes.

Security fixes in this release:

- Resolved security vulnerabilities present in embedded, third-party dependencies of the CIS-CAT Pro Assessor v4.41.0:
  - PostgreSQL updated to v42.7.2.
  - xmlsec updated to v4.0.1.
  - cxf-core updated to v3.5.8.
  - bouncycastle updated to v1.78.

## Comply 2.19.0

Released 14 March 2024.

New in this release:

- **Additions to the Comply API.** Improved accessibility and efficiency of Comply functions by adding two new endpoints to the Comply API:
  - Exports API. You can use the Exports API to create, retrieve, download, and delete exports of data from Comply.
  - Inventory API. You can use the Inventory API to initiate a PE inventory sync.
- **CIS-CAT Pro Assessor v4.39.0.** Comply 2.19.0 contains the CIS-CAT Pro Assessor v4.39.0.
- Benchmarks updated in this release:
  - Microsoft Windows 10 Enterprise v3.0.0
  - Microsoft Windows 11 Enterprise v3.0.0

Security fixes in this release:

- Resolved security vulnerabilities present in embedded, third-party dependencies of the CIS-CAT Pro Assessor v4.39.0:
  - commons-compress updated to v1.26.0.
- Resolved security vulnerability CVE-2023-26159, present in the dependency follow-redirects-1.15.2, by upgrading to follow-redirects-1.15.4.

## Comply 2.18.2

Released 22 February 2024.

New in this release:

- **CIS-CAT Pro Assessor v4.38.0.** Comply 2.18.2 contains the CIS-CAT Pro Assessor v4.38.0.

Resolved in this release:

- **Exception page when viewing a custom profile as comply-viewer.** Previously, the console would display an unknown error when viewing a custom profile with the viewer role. This has been fixed.

Security fixes in this release:

- Resolved security vulnerabilities present in embedded, third party dependencies of the CIS-CAT Pro Assessor v4.37.0:
  - unzip.exe updated to v6
  - ion-java updated to v1.11.1. CVE-2024-21634

## Comply 2.18.1

Released 18 January 2024.

New in this release:

- **Public API specifications.** Public API information is available at `https://<COMPLY-HOSTNAME>>/openapi.json``, where COMPLY-HOSTNAME is your Comply server.
- **CIS-CAT Pro Assessor v4.37.0.** Comply 2.18.1 contains the CIS-CAT Pro Assessor v4.37.0.
- Benchmarks updated in this release:
  - CIS Amazon Linux 2 Benchmark v3.0.0
  - CIS Microsoft Windows Server 2019 STIG Benchmark v2.0.0
  - CIS CentOS Linux 7 Benchmark v4.0.0
  - CIS Oracle Linux 7 Benchmark v4.0.0
  - CIS Red Hat Enterprise Linux 7 Benchmark v4.0.0
- Benchmarks added in this release:
  - CIS Microsoft Windows Server 2019 Stand-alone v1.0.0
  - CIS Microsoft Windows Server 2022 STIG Benchmark v1.0.0

Resolved in this release:

- **GraphQL pod health check fails and does not recover.** Added a probe to the GraphQL pod that restarts the pod if it is not live.

Security fixes in this release:

- Resolved security vulnerabilities present in the CIS-CAT Pro Assessor:
  - logback-classic and core updated to 1.2.13
  - jackson-databind updated to 2.16.0
  - bouncy castle(bcprov) libraries updated to 1.74

## Comply 2.18.0

Released 13 December 2023.

New in this release:

- **Comply API**
  - Puppet's open & integrated approach ensures data sharing with enterprise tools (Such as risk management, systems of records etc.), improves productivity and cross-team collaboration leveraging the same data to ensure transparency.
  - The Comply API allows you to automate actions, retrieve Comply data, and share Comply data with other groups and tooling. To use the API, you first must create a personal access token, after which you can access API endpoints.
  - The new Comply API improves productivity and resource management by providing access to existing data and functionality.
  - Added a new API allowing users to automate actions, retrieve Comply data, and share Comply data with other groups and tooling.
  - Users can create personal API access tokens and access endpoints depending on their user permissions.
  - Admins can view and revoke users' access tokens.
  - Added an API endpoint for extracting compliance results from Puppet Comply. Users can extract both summary and raw data results for one, many, or all nodes up to 100,000 nodes.
- **CIS-CAT Pro Assessor v4.36.0.** Comply 2.18.0 contains the CIS-CAT Pro Assessor v4.36.0. Benchmarks updated in this release:
  - CIS AlmaLinux OS 8 Benchmark v3.0.0
  - CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
  - CIS Apple macOS 11.0 Big Sur Benchmark v4.0.0
  - CIS Apple macOS 12.0 Monterey Benchmark v3.0.0
  - CIS Microsoft Windows Server 2012 (non-R2) Benchmark v3.0.0
  - CIS Microsoft Windows Server 2012 R2 Benchmark v3.0.0
  - CIS Microsoft Windows Server 2016 STIG Benchmark v2.0.
  - CIS Oracle Linux 8 Benchmark v3.0.0
  - CIS Red Hat Enterprise Linux 8 Benchmark v3.0.0
  - CIS Rocky Linux 8 Benchmark v2.0.0

Benchmarks added in this release:

- CIS Apple macOS 13.0 Ventura Benchmark v2.0.0

Benchmarks removed in this release:

- CIS Apple macOS 10.15 Catalina Benchmark v3.0.0

Resolved in this release:

- **Node results page shows compliance score without exceptions.** Fixed an issue where the **node results** page showed the compliance score of nodes without exceptions included, rather than the 'adjusted compliance score', which accounts for exceptions.
- **Security fix.** Resolved security vulnerabilities present in embedded, third party dependencies for io.netty.

## Comply 2.17.0

Released 2 November 2023.

New in this release:

- **Add parameter to prepend `scan_cmd` in module 2.17.1**

The Comply module version 2.17.1 allows users to limit CPU usage on individual nodes when running a scan.

To configure this feature, add the `limits` parameter to the `comply` class via PE. The parameter's format is JSON.

Example: `{ "systemd-run" : { "CPUQuota" : 50 } }` limits CPU usage to 50% of one CPU core. Values greater than 100 are allowed if you want to use more than one CPU core.

Example: `{ "nice" : { "increment" : 10 } }` increments the `niceness` value of the process.

`systemd-run` is supported on Linux systems where `/usr/bin/systemd-run` is present, and `nice` is supported on Linux and macOS. Neither is supported on Windows.

**Note:** The Comply module 2.17.1 was released with Comply 2.17.0, and it can be used with Comply versions 2.17.0 and later. Upgrading the Comply module from 2.17.0 to 2.17.1 is optional, and version 2.17.0 is still available.

- **Scalability improvements.** Puppet Comply now runs on a maximum of 100,000 nodes.
- **Embedded JRE for MacOS.** The CIS-CAT Pro Assessor for MacOS now contains embedded Java. It is no longer necessary to install Java when running the Assessor on MacOS.
- **CIS-CAT Pro Assessor v4.34.0.** Comply 2.17.0 contains the CIS-CAT Pro Assessor v4.34.0. Benchmarks updated in this release:

- CIS Ubuntu Linux 20.04 LTS Benchmark v2.0.1

Benchmarks removed in this release:

- CIS Red Hat Enterprise Linux 6 Benchmark v3.0.0
- CIS CentOS Linux 6 Benchmark v3.0.0
- CIS Oracle Linux 6 Benchmark v2.0.0
- CIS Debian Linux 9 Benchmark v1.0.1

Resolved in this release:

- **Security fix.** CVE-2023-3635 has been resolved in the CIS-CAT Pro Assessor v4.34.0.
- **The ``comply-auth`` pod should have liveness and/or startup probes defined.** New startup and liveness probes have been added to the ``comply-auth`` pod. This change fixes an issue encountered on upgrade where the auth pod never went into a ready state and did not restart, as it only had a readiness probe. With the new startup probe, the pod has a set period of time to start up. If it fails to start up within the allotted time it restarts. The liveness probe restarts the pod if the pod appears to be unreachable for 30 seconds.

## Comply 2.16.0

Released 21 September 2023.

New in this release:

- **Scalability improvements.** Puppet Comply now runs on a maximum of 75,000 nodes.
- **CIS-CAT Pro Assessor v4.33.0.** Comply 2.16.0 includes the CIS-CAT Pro Assessor v4.33.0. Benchmarks updated in this release:
  - Ubuntu Linux 20.04 LTS STIG v2.0.0
  - Debian Linux 11 STIG v1.0.0

Resolved in this release:

- **Inventory sync improvements.** Made the following improvements to inventory sync:
  - Fixed an issue where background inventory syncs were not reflected on the **Settings** page.
  - Fixed an issue where the progress bar displayed NaN briefly at the beginning of an inventory sync.
  - Fixed an issue where inventory sync errored at the node groups stage due to the host not existing in Comply inventory.
  - Fixed an issue where the refresh data button was not surfacing.
  - Fixed an issue where inventory sync was locking up for at least 20 minutes.
- **Puppet Application Manager documentation updated.** Fixed import of Puppet Application Manager documentation into Comply documentation.

### Comply 2.15.1

Released 4 September 2023.

Resolved in this release:

- **Data not refreshing: new nodes and reports missing from console.** Resolved an issue where the process that syncs inventory into Puppet Comply would hang if it encountered an error in communication with Puppet Enterprise.
- **Empty node groups showing all nodes.** Puppet Comply now only shows PE node groups that contain Comply managed nodes.

### Comply 2.15.0

Released 10 Aug 2023.

New in this release:

- **Scalability improvements.** Exporting raw scan data now works with up to 50,000 nodes. The raw scan data export is generated as an archive of one or more gzip-compressed CSV files, with each CSV file including the results for up to 1,000 nodes. The nodes are sorted in ascending order of name, and each CSV file is named according to the range of nodes it covers.
- **CIS-CAT Pro Assessor v4.32.0.** Comply 2.15.0 includes the CIS-CAT Pro Assessor v4.32.0. Benchmarks updated in this release:
  - Apple macOS 11 v3.1.0
  - Apple macOS 12 v2.1.0

Resolved in this release:

- **Improvements to the inventory sync with Puppet Enterprise.** Previously, inventory sync made paging requests without ordering, which can lead to the ingest retrieving fewer hosts from Puppet Enterprise than expected. This has been fixed. Also improved efficiency and accuracy in the database.
- **Compliance over time chart missing a day.** Missing days when filters are applied in the Compliance over time chart on the Comply dashboard have been fixed.
- **Active exceptions count on dashboard not matching with exceptions page.** Fixed an issue where the active exceptions count was mismatched between the dashboard and the exceptions page.

### Comply 2.14.0

Released 29 June 2023.

New in this release:

- **Compliance at scale.** Comply now supports up to 50,000 nodes.
- **Identity and access management - RBAC integration.** Comply now integrates with Puppet Enterprise (PE) for role-based access control (RBAC). Using PE you can create new Comply users or import them from LDAP. You can assign users to roles in the PE Console. There are three default roles provided for Comply users: comply-

admin, comply-operator, and comply-viewer. Each role has different permissions and a different view of the Comply console.

**Important:** When upgrading to 2.14.0, you must ensure that your PE account has permissions to View and Create User Roles. For more information, visit <https://www.puppet.com/docs/comply/2.x/configure-comply-with-pe.html>.

- **CIS-CAT Pro Assessor v4.30.0.** Comply 2.14.0 includes the CIS-CAT Pro Assessor v4.30.0. Benchmarks updated in this release:
  - Debian Linux 10 v2.0.0
  - Microsoft Windows 10 Standalone v2.0.0
  - Microsoft Windows 11 Standalone v2.0.0
  - Ubuntu Linux 20.04 LTS v2.0.0
  - Azure Compute Microsoft Windows Server 2019 v1.0.1
  - Microsoft Windows Server 2016 v2.0.0
  - Microsoft Windows Server 2019 v2.0.0
  - Microsoft Windows Server 2022 v2.0.0

Resolved in this release:

- **Vulnerability in gin-v1.9.0.** Resolved security vulnerabilities to address CVE-2023-29401.
- **Overall compliance score over time card is not working.** Previously, the overall compliance score over time card in the Comply dashboard displayed the average compliance score based on scans performed on each day. It now shows the average compliance score based on the latest scan result available on each day.
- **OS and environment filters appear as options for Node Results exports.** Removed options for **Operating system** and **Environment** filters on the **Node Results** page to match actual functionality.

## Comply 2.13.0

Released 4 May 2023.

New in this release:

- **Redesign Comply dashboard.** Redesigned and added new features to the compliance dashboard, including numbers of nodes and exceptions, graphs for better understanding compliance score, and quickly accessible action steps.
- **CIS-CAT Pro Assessor v4.28.0.** Comply 2.13.0 includes the CIS-CAT Pro Assessor v4.28.0. Benchmarks updated in this release:
  - Windows 10 Enterprise v2.0.0
  - Windows 11 Enterprise v2.0.0

Resolved in this release:

- **Vulnerabilities in the oauth2-proxy container.** Updated Comply oauth2-proxy container to address CVEs.
- **Vulnerabilities in the Redis container.** Updated Comply Redis container to address CVEs.
- **Snakeyaml vulnerabilities.** The CIS-CAT Pro Assessor v4.28.0 resolves security vulnerabilities in embedded, third-party dependency snakeyaml. This library has moved to version 2.0.0.
- **ciscat.pp fails to apply if Puppet agent version is prior to 6.24 or 7.9.** The CIS-CAT Pro Assessor can now be downloaded with any 6.x or 7.x version of Puppet agent.
- **Performance fixes.** Improved performance and scalability.

## Comply 2.12.0

Released 23 March 2023.

New in this release:

- **Scan wizard changes.** Added filters and removed irrelevant node options when running a scan.
- **Navigation changes.** Made Comply navigation clearer and more streamlined.
- **Exceptions upgrade changes.** Added handling for exceptions during upgrades. Exceptions are upgraded if their benchmark is upgraded. Exceptions that are no longer functional after the upgrade are removed. You can see the status of your exceptions following an upgrade in the **Activity feed**.
- **Custom profiles export.** You can now export one, many, or all of your custom profiles in order to easily gather custom profile details.
- **Scalability improvements.** Comply now supports up to 25000 nodes.
- **CIS-CAT Pro Assessor v4.27.0.** Comply 2.12.0 includes the CIS-CAT Pro Assessor v4.27.0. With this new version, the assessor runs using an embedded JRE, removing the requirement to have a locally installed JRE.

Resolved in this release:

- **Filtering an empty PE Group within Comply displays all nodes.** Filtering an empty PE group now returns 0 nodes instead of all nodes.
- **Broken links in Comply navigation.** Fixed broken links.

## Comply 2.11.0

Released 26 January 2023.

New in this release:

- **Scan wizard redesign.** Improvements to the scan wizard including:
  - For both ad hoc and scheduled scans, you can scan on multiple nodes across different environments. Results for all scanned environments are available in a single report.
  - You can only start scans from the **Scans** page or the **Node detail** page.
- **CIS-CAT Pro Assessor v4.25.0.** Comply 2.11.0 includes the CIS-CAT Pro Assessor v4.25.0. For more information, visit: [CIS-CAT Pro Assessor history](#) on page 8.
- Various enhancements to improve scan reliability and performance with larger node counts.

Resolved in this release:

- **Node groups not imported from Puppet Enterprise (PE) when nodes are not pinned to group.** Previously, node groups were only imported for nodes explicitly pinned to the node group. Comply now also imports groups in which rules match nodes to groups based on facts.
- **Exceptions that are both resolved and expired disappear from the exceptions page.** Exceptions that are resolved before the expiry time is reached now stay in the “expired” tab of the exceptions page.
- **Reports export null data when custom profile filter applied.** Exported reports are no longer empty when a custom profile has been selected on the "Profiles" quick filter.
- **SCE scripts run for too long if they find non-compliant files when packaging and building the Comply module.** SCE scripts now quit immediately upon finding a single non-compliant file.
- **System curl commands fail to download the CIS-CAT Pro Assessor.** Previously, you could not download the CIS-CAT Pro Assessor using your system curl command. This has been fixed.

Security notice:

- Only the latest version of the CIS-CAT Pro Assessor has the latest security fixes. Customers on previous versions of the CIS-CAT Pro Assessor might be vulnerable to security issues. CIS-CAT Pro Assessor v4.25.0 resolves security vulnerabilities present in embedded, third-party dependencies in CIS-CAT Pro Assessor v4.23.0, which was shipped in Comply 2.10.0. For details, see [CIS-CAT Pro Assessor and Dashboard December 2022 Vulnerability Updates](#).

## Comply 2.10.0

Released 1 December 2022

New in this release:

- **CIS-CAT Pro Assessor v4.23.0.** Comply 2.10.0 includes the CIS-CAT Pro Assessor v4.23.0.
- Security notice:
  - The CIS-CAT Pro Assessor v4.23.0 resolves a security vulnerability present in the embedded, third party dependency for the `jackson-databind` mapping functionality. This library has moved to `jackson-databind-2.13.4.jar`.
- **Export scan results.** You can now export the last scan results for all nodes, a subset of nodes, or a single node. All exported data is collected in a single `.csv` file. To export scan results, use the **Export CSV** button on the Node Results pane on the Compliance Dashboard. To view and download previous reports, use the **Generated Reports** button in the Comply navigation pane.
- **Resolve exceptions.** You can now resolve exceptions that are no longer needed. Details about resolved exceptions remain visible in Puppet Comply for reporting purposes. You can resolve an exception for all nodes or for a subset of nodes.
- **Exception details.** You can now view and edit the details of your exceptions.
- **Using old versions of the CIS-CAT Pro Assessor.** You can now upgrade to the latest version of Comply without updating the CIS-CAT Pro Assessor. As of this release the supported versions of the CIS-CAT Pro Assessor are 4.22.0 and 4.23.0. In future releases, the current version and the two previous versions will be supported. All nodes still must run the same version of the CIS-CAT Pro Assessor.
- Security notice:
  - Only the latest version of the CIS-CAT Pro Assessor has the latest security fixes. Customers on older versions of the CIS-CAT Pro Assessor may be vulnerable to security issues.
- **Node group filtering.** Anywhere all nodes are listed, node groups filtering now supports nodes that have been pinned to the node group in PE. Node groups are based on PE classification groups.

Resolved in this release:

- **Exceptions remain active after they are no longer applicable.** Exceptions are now removed if their custom profile is deleted or edited to remove the relevant rule.
- **A deleted exception cannot immediately be re-created.** Previously, if you created an exception for a specified rule and node and then deleted the exception, you could not immediately re-create the exception for the specified rule and node. This has been fixed.

## Comply 2.9.0

Released 20 October 2022

New in this release:

- **CIS-CAT Pro Assessor v4.22.0.** Comply 2.9.0 includes the CIS-CAT Pro Assessor v4.22.0 and the following associated benchmarks:
  - Debian Linux 11, v1.0.0
  - Azure Compute Microsoft Windows Server 2019, v1.0.0
- Security notice:
  - The CIS-CAT Pro Assessor v4.22.0 resolves security vulnerabilities present in embedded, third party dependencies. For details, see [CIS-CAT Pro Dashboard and Assessor September 2022 Vulnerability Updates](#).
- **Create temporary exceptions to rules.** With Comply 2.9.0, you can create a temporary exception to a CIS Benchmark rule and apply that exception to a node, a group of nodes, or all nodes. During the period when the exception is active, the rule's compliance score is excluded from the overall compliance score for the selected nodes. Exceptions are useful in many situations. For example, if you plan to install a software patch on several nodes, but the patch requires additional testing, you can specify a temporary exception for the affected nodes while testing continues. During the next scan, the exception is applied, and the compliance score reflects the exception. When testing is completed, you can apply the software patch to the nodes, and the exception expires automatically on your specified date.
- **View and delete exceptions.** You can go to the new **Exceptions** page to view and delete exceptions.

Resolved in this release:

- **Scans fail to complete processing.** In some cases, when scans were run manually, the scans would remain in the started state and would fail to generate a final report.

## Comply 2.8.0

Released 8 September 2022

New in this release:

- **CIS-CAT Pro Assessor v4.21.0.** Comply 2.8.0 includes the CIS-CAT Pro Assessor v4.21.0 and the following associated benchmarks:
  - Microsoft Windows 11.
  - Microsoft Windows 10 (stand-alone). (A stand-alone system is not connected to a domain and cannot be managed by using Active Directory.)
  - Ubuntu 22.04.
- **Specify a refresh interval to obtain the latest inventory updates from Puppet Enterprise (PE).** By default, the Comply inventory is refreshed every 24 hours with the latest node and fact information from Puppet Enterprise. With Comply 2.8.0, you can customize the refresh interval to meet your organization's requirements.

Resolved in this release:

- **Consistency of scan compliance scores.** To help ensure consistency of compliance scores throughout the Comply user interface, the **Node detail** page and the **Rule detail** page are updated. The donut charts and the accompanying legends now exclude non-scoring statuses. A non-scoring status means that a CIS recommendation is not applicable or cannot be automatically validated. With this change, the charts on the **Node detail** page and **Rule detail** page now provide a more realistic view of compliance.
- **Accurate status for profiles.** The **Profile** column on the **Scan Report** page now reflects the correct status of profiles. Previously, if you hovered over the **Profile** column, you might have seen an invalid message that the profile was deleted.
- **Scheduled scans not running after Comply upgrade.** After upgrading Comply, scheduled scans that were created before the upgrade might not run. After upgrading to Comply v2.8.0, these scans should run as configured.

Security notice:

- This release includes a security update that helps to prevent command injection in the Comply module.

## Comply 2.7.0

Released 27 July 2022

New in this release:

- **CIS-CAT Pro Assessor v4.19.0.** Comply 2.7.0 includes the CIS-CAT Pro Assessor v4.19.0.
- **Learn how to run Comply at scale.** You can scan up to 5000 nodes in a single batch to check the compliance of your infrastructure against Center for Internet Security (CIS) Benchmarks. The documentation is updated to help you configure and run scans at scale. See [Guidelines for running Comply at scale](#) on page 154.
- **Delete a custom profile.** In previous releases, you could create a custom profile based on a CIS Benchmark. In this release, you can also delete one or more custom profiles.

Resolved in this release:

- **Warning messages during preflight checks.** An issue that caused invalid warning messages to be displayed during preflight checks is resolved in this release. The invalid message, `No matching files`, is no longer displayed.

## Comply 2.6.0

Released 16 June 2022

New in this release:

- CIS-CAT Pro Assessor v4.18.0.** Comply 2.6.0 includes the CIS-CAT Pro Assessor v4.18.0 and the following associated benchmarks:

  - Alma Linux 8 v2.0.0
  - Microsoft Windows Server 2016 v1.4.0
  - Microsoft Windows Server 2016 STIG v1.2.0
  - Microsoft Windows Server 2012 v2.4.0
  - Microsoft Windows Server 2012 R2 v2.6.0
  - Possible errors due to renamed benchmarks:** In addition to version changes, CIS renamed two benchmarks in this release. AlmaLinux was renamed to Alma Linux and Microsoft Windows Server 2016 RTM (Release\_1607) was renamed to Microsoft Windows Server 2016. If you are using a benchmark that was renamed, you might see an error message indicating that the benchmark is no longer supported. If your nodes use custom profiles that are based on renamed benchmarks, you must manually update the nodes because they will not be automatically updated during the Comply upgrade process.
- Edit a scheduled scan.** You can edit a scheduled scan to modify the type of scan, the frequency, and the start and end dates.
- Delete a scheduled scan.** You can delete a scheduled scan to permanently remove it.
- Take advantage of enhanced usability for scan reports.** From a scan report, you can navigate to the **Node detail** page, where the **Scan status** pane now includes a legend showing the total number of rules that were run on the node and detailed results. You can hover over the results to see percentages in the donut chart. Similarly, on the **Rule detail** page, the **Scan status** pane now shows the total number of scanned nodes and detailed results. You can hover over the results to see percentages in the donut chart. The **Rule detail** page includes a new **Environment** column so that you can determine the environment (for example, test or production) in which the scan took place. The **Node detail** page includes a new **Last passed on** column, which shows the date and time of the last successful scan for each rule.

Security notice:

- Vulnerability in the 3.14.2-alpine image. The release updates the alpine image to 3.15.4.

## Comply 2.5.1

Released 31 May 2022

Resolved in this release:

- Potential deployment issue for users of Comply 2.4.0 and 2.5.0.** This issue can affect users who install Comply in a Google Kubernetes Engine (GKE) environment and potentially other environments. If you are unable to start Comply after installation, you might be experiencing this issue. To diagnose the issue, review the log for the `comply-scarpy` pod. If the issue is occurring, the pod will be in an `Init:CrashLoopBackOff` state during the attempt to start Comply. Review of the pod will show that the `comply-scarpy-init` container was terminated with an out-of-memory error (`OOMKilled`). To resolve the issue, install Comply 2.5.1. If you do not detect the issue, it is not necessary to install Comply 2.5.1.

## Comply 2.5.0

Released 5 May 2022

New in this release:

- CIS-CAT Pro Assessor v4.16.1.** Comply 2.5.0 includes the CIS-CAT Pro Assessor v4.16.1 and the following associated benchmarks:

  - Microsoft Windows Server 2019 v1.3.0
  - Microsoft Windows Server 2019 STIG v1.1.0
  - Oracle Linux 8

- Rocky Linux 8  
CIS-CAT Pro Assessor v4.16.1 resolves a security issue (<https://nvd.nist.gov/vuln/detail/CVE-2022-21724>) that does not affect current users of Comply.
- The following CIS benchmarks are at end of life and are no longer supported:
  - CentOS Linux 8
  - SUSE Linux Enterprise Server 11
- **View details about a scheduled scan.** You can select a scheduled compliance scan and view its details, including the creation date, last modification date, affected nodes, start and end times, and frequency. You can also view the scan history, including the number of runs, the date and time of the most recent run, and the date and time of the next scheduled run.
- **Pause, resume, or end a scheduled scan.** On the **Scheduled scan details** page, you can pause, resume, or end a scheduled scan.
- **Assign benchmarks and profiles to multiple nodes simultaneously.** On the **Inventory** page, you can select multiple nodes and assign a benchmark, a profile, and, optionally, a custom profile to all. The selected nodes must be running on the same operating system, and the latest version of the CIS-CAT Pro Assessor must be installed on each node.
- **View a report about scan results for a single rule.** The **Scan rule report** lists the nodes on which the rule was run, the results, and the overall compliance score for the rule.
- **View a report about scan results for a single node.** The **Scan node report** lists the rules that were run on the node, the results, and the overall compliance score for the node.

Resolved in this release:

- **Initial deployment issue on Microsoft Windows Server 2016 and Microsoft Windows Server 2019 operating systems.** In previous releases, the initial deployment of the Comply module sometimes failed with the following error message:

```
Provider wget is not functional on this host
```

## Comply 2.4.0

Released 24 March

New in this release:

- **CIS-CAT Pro Assessor v4.15.0.** Comply 2.4.0 includes the latest version of the CIS-CAT assessor and the following supported associated benchmarks:
  - CentOS Linux 8 (final release)
  - Microsoft Windows 10 v1.12.0.
  - Microsoft Windows Server 2022 v1.0.0
  - Red Hat Enterprise Linux 8 v2.0.0
  - SUSE Linux Enterprise 11 v2.1.1 (final release)

**Note:** The Microsoft Windows 10 benchmark has upgraded from 1.11.0 CIS Microsoft Windows 10 Enterprise Release 21H1 to 1.12.0 CIS Microsoft Windows 10 Enterprise. Comply's 1.12.0 CIS Microsoft Windows 10 Enterprise benchmark is based on Microsoft Windows 10 Enterprise Release 21H2 and is intended for all versions of the Windows 10 operating system, including older versions. If any of your nodes use custom profiles based on the 1.11.0 CIS Microsoft Windows 10 Enterprise Release 21H1 benchmark, you need to resolve these manually, as they will not automatically update during the upgrade process.

- **Profile and Custom profile.** You can view and sort two new columns on the **Inventory** page - **Profile** and **Custom profile**. The columns allow you to see if a node has a default profile or custom profile assigned to it.
- **Benchmark column.** The **Desired compliance** column has been renamed to **Benchmark**.

Resolved in this release:

- **Sync license.** Fixed an issue where a user was logged out of Comply after selecting **Sync license** on the License page.

## Comply 2.3.0

Released 10 February 2022

New in this release:

- **Scheduled scans.** You can now schedule one-off and repeating scans, in addition to running manual ad hoc scans, in Comply.

For more information, see [Scheduled scans](#) on page 162.

- **Environment information.** The Scan list page now shows the scan report environment.

- 

**CIS-CAT Pro Assessor v4.14.0.** Comply

2.3.0 includes the latest version of the CIS-CAT assessor and the following supported associated benchmarks:

- SUSE Linux Enterprise 12 v3.1.0
- SUSE Linux Enterprise 15 v1.1.1

This release of the assessor resolves security vulnerability present in embedded, third party dependencies:

- The **OpenDXL Java** Client library, which includes log4j, is now a derivative work of version 0.2.6 which includes log4j 2.17.1.
- The **logback-core** and **logback-classic** libraries have been moved to version 1.2.10.
- Comply now supports Kubernetes 1.19 to 1.24. Kubernetes 1.17 and 1.18 are no longer supported.

Resolved in this release:

- **Rule details.** Fixed a bug where the last reported time stamp on the rule detail page did not recognize the user's local timezone.
- **Compliance profiles.** Corrected an issue where the default compliance profile was incorrectly assigned for Windows Server versions.

## Comply 2.2.2

Released 20 January 2022

New in this release:

- **Debug mode.** You can now choose to run in debug mode to provide easier access to assessor logs.

For more information, see [Run an ad hoc scan](#) on page 161.

- 

**CIS-CAT Pro Assessor v4.13.1.** Comply

2.2.2 includes the latest version of the CIS-CAT assessor and the following supported associated benchmarks:

- AlmaLinux OS 8 v1.0.0
- Amazon Linux 2 STIG v2.0.0
- Apple macOS 11.0 Big Sur v2.0.0
- Microsoft Windows Server 2012 (non-R2) v2.3.0
- Red Hat Enterprise Linux 8 STIG v1.0.0

**CIS-CAT Pro Assessor v4.13.1** resolved security vulnerabilities present

in the following embedded, third party dependency:

- **log4j-core** - This library was updated to version 2.17.0.

## Comply 2.2.1

Released 20 December 2021

New in this release:

**CIS-CAT Pro Assessor v4.13.0.** Comply 2.2.1 includes the latest version of the CIS-CAT assessor and the following supported associated benchmarks:

- Apple macOS 10.15 Catalina v2.0.0
- Red Hat Enterprise Linux 7 STIG v2.0.0

The following benchmark is at end of life and is no longer supported:

- Mac OS 10.14

Security notice:

- **CIS-CAT Pro Assessor v4.13.0** resolved security vulnerabilities present in the following embedded, third party dependencies:
  - **log4j-core** - This library was updated to version 2.15.0.
  - **bcprov-jdk15on** - This library was updated to version 1.69.
- **Component upgrade to address CVEs.** To address various CVEs, this version includes an upgrade of Kubernetes to 1.19.15.

**Important:** Version 2.15.0 of the log4j-core library addresses the potential escalation of privilege vulnerability. We do not believe Comply is vulnerable to any of the additional risks addressed in the 2.16.0 release, but plan to release an update in the near future which includes version 2.17.0 or later.

## Comply 2.2.0

Released 18 November 2021.

New in this release:

- **Scan Reports improvements.** Scan reporting functionality is extended to include the ability to access a list of historical scans and view scan details. For more information, see [CIS scan report details](#) on page 166.
- **Filtering and sorting.** Filtering and sorting functionality has been implemented on all table columns in the Comply UI.

**Note:** Filter drop-downs display all available options for a given parameter. On pages where multiple filtering options are available, selecting one filter option does not affect the options presented by any other filter drop-down.

- **CIS-CAT Pro Assessor v4.11.0.** Comply 2.2.0 includes the latest version of the CIS-CAT assessor and its associated benchmarks:
  - Microsoft Windows Server 2012 R2 v2.5.0
  - Microsoft Windows Server 2016 STIG v1.1.0
  - SUSE Linux 15 v1.1.0
- **Desired compliance.** The Comply UI has been simplified so that users are no longer required to manually accept the profiles applied by Comply based on fact information from PE.
- **Custom Comply port.** You can now specify a custom Comply port in Puppet Application Manager if you do not want to use the default port (30303). For more information, see [System requirements](#) on page 130.

- **Data retention.** The retention period for scan data can now be set on the Puppet Application Manager **Config** tab. For more information see, [Scan results](#) on page 166.

Resolved in this release:

- **Node Deletion.** A fix was added to ensure that nodes deleted in Puppet Enterprise are no longer listed in Comply as available for scanning.
- **License page node count.** Corrected an issue where the number of nodes displayed on the license page was not updated when a node was deleted in Puppet Enterprise.
- **Required installations page.** The required installations page that was part of the assessor install procedure was removed as it was no longer required.
- **Comply-graphql.** Fixed a known issue where the comply-graphql deployment did not become healthy after restoring Comply using Puppet Application Manager.
- **Rule ordering.** Corrected an issue where rules were not always displayed in the correct numerical order.

## Comply 2.1.0

Released 7 October 2021.

New in this release:

- **Scan Reports.** The Comply UI has a new **Scan Reports** page that provides a report on rules passed/failed and node compliance from the most recent CIS scan. For more information, see [CIS scan report details](#) on page 166.
- **CIS-CAT Pro Assessor v4.9.0.** Comply 2.1.0 includes the latest version of the CIS-CAT assessor and its associated benchmark:
  - CentOS Linux 7 v3.1.2
- **Scanner upgrades.** Scanner upgrade in Comply is not forced but optional to allow better management of PE jobs.

**Note:** By default in Comply 2.1.0, assessor upgrade does not happen automatically when you upgrade Comply. Assessor upgrade takes place when you instigate a Puppet Enterprise (PE) Puppet run job after Comply is upgraded. For more information, see [Upgrade from Comply 2.2.2 to 2.3.0](#) on page 144.

Resolved in this release:

- **Desired compliance upgrades.** Fixed an issue where Windows 10 nodes lost their desired compliance after upgrade to Compliance 2.x
- **Upgrade statistics.** Resolved an issue where statistics were overwritten when multiple upgrades take place.
- **Service start up.** Updated Comply so that it now starts when IPv6 is disabled.
- **Preflight failure.** Fixed an issue where preflight checks failed during install when trailing newline returns were present in certificates.
- **Scan wizard.** The Comply scan wizard was updated to correct an issue where the **environment name** field did not revert to the previous saved value if the scan set up was cancelled.

## Comply 2.0.0

Released August 2021.

New in this release:

- **CIS-CAT Pro Assessor v4.8.2.** Comply 2.0.0 includes the latest version of the CIS-CAT assessor and its associated benchmarks:
  - Apple macOS 10.14 v1.4.0
  - Apple macOS 10.15 v1.4.0
  - Apple macOS 11.0 v1.2.0
  - CentOS Linux 7 v3.1.1
  - CentOS Linux 8 v1.0.1
  - Debian Linux 8 v2.0.2
  - Microsoft Windows Server 2019 v1.2.1
  - Microsoft Windows Server 2019 STIG v1.0.1
  - Microsoft Windows 10 20H2 v1.10.1
  - Oracle Linux 7 v3.1.1
  - Oracle Linux 8 v1.0.1
  - Red Hat Linux 7 v3.1.1
  - Red Hat Linux 8 v1.0.1
  - Amazon Linux 2 v2.0.0
  - Microsoft Windows 10 21H1 v1.11.0
  - Microsoft Windows Server 2016 v1.3.0
  - Ubuntu Linux 20.04 LTS STIG v1.0.0
- **Automatic upgrades of the CIS-CAT assessor.** Every time you upgrade your Comply application, the assessor automatically upgrades to the latest version. This update also includes the following changes to how you interact with Comply:
  - You can only run a desired compliance scan against nodes with the latest version of the assessor.
  - You can only run a custom scan against benchmarks with the latest version of the assessor.
  - On the node inventory screen, nodes without the latest assessor are highlighted red to indicate that they need upgrading.
  - You can no longer set a desired compliance benchmark against a node that does not have the latest version of the assessor.
  - When the assessor upgrades, custom profiles are automatically updated to use the new benchmarks and profiles, sending you a notification.
- **Assessor upgrades tab.** The **Assessor upgrades** tab on the **Activity feed** screen provides a summary of assessor upgrades, including the number of nodes that have passed or failed. Note that this only shows the status of your nodes after the upgrade, and does not update again, even if your nodes change to passing.
- **comply module Secure Sockets Layer (SSL).** This includes changes to how you install and upgrade the Comply module.

Resolved in this release:

- **Comply tries to install 7-zip on Windows.** The `comply` module no longer installs 7zip on Windows systems.
- **Windows Server Semi Annual Channel (SAC) builds are assigned the wrong CIS profile.** SAC builds are now assigned the correct Windows 2019 profile.

Security notice:

- **Vulnerability in 12.18.3-alpine image.** The release updates the alpine image to 15.13.0.
- **Vulnerability keycloak:15.0.0.** This release updates keycloak to version 15.0.0.
- **Vulnerability in dependencies.** This release upgrades NodeJS to version 14.17.1 and React to version 17.0.2.

For upgrade instructions, see [Upgrade from Comply 2.2.2 to 2.3.0](#) on page 144.

## Comply known issues

---

These are the known issues for the Puppet Comply 1.x and 2.x releases.

## Scans do not run on Windows Servers

This issue is due to a recently issued Windows Update (April 8, 2025) that is impacting CIS-CAT Pro Assessor.

Please see [this KB article](#) for more information about the issue and how to fix it.

## Possible errors due to renamed benchmarks

The Alma Linux 8 benchmark was renamed by CIS in the CIS-CAT Pro Assessor v4.36.0. As a result the assessor upgrade page reports a warning that 2.0.0 CIS Alma Linux OS 8 is No longer supported.

The new benchmark is 3.0.0 CIS AlmaLinux OS 8, and any nodes previously assigned 2.0.0 CIS Alma Linux OS 8 will have their desired compliance automatically upgraded to the new benchmark.

## Incompatibility with stdlib v9.0 and v9.1

Puppet Comply is not compatible with stdlib versions 9.0 and 9.1.

## Security vulnerability

Comply includes CIS components in its CIS-CAT Pro Assessor. The CIS components of the assessor version 4.34.0, shipped with Comply 2.17.0, contain the security vulnerability CVE-2023-4586. Our investigation into the issue shows that our implementation does not pose risk to Puppet Comply customers. Once the CVE is addressed by CIS, the fix will be included in the next release of Puppet Comply and documented accordingly.

## Security vulnerability

Comply includes CIS components in its CIS-CAT Pro Assessor. The CIS components of the assessor versions 4.30.0, 4.32.0, and 4.33.0, shipped with Comply 2.14.0, 2.15.0, and 2.16.0, contain the security vulnerability CVE-2023-3635. Our investigation into the issue shows that our implementation does not pose risk to Puppet Comply customers, although in certain extreme cases it might cause Comply to stop responding. The CVE has been addressed by CIS, and the fix is included in the CIS-CAT Pro Assessor version 4.34.0, released with Puppet Comply 2.17.0.

## Reports export null data when custom profile filter is applied

Exported reports are empty if a custom profile has been selected on the **Profiles** quick filter.

## Node group filtering does not work for deleted nodes

Deleted nodes do not have node group information available. The **Node Group** quick filters on the **Scan Report** (Nodes tab) and **Rule Detail** pages do not apply to deleted nodes.

## Running scan tasks in Puppet Enterprise (PE)

Comply uses PE tasks to run compliance scans on nodes. Although you can see the scan tasks in PE, we advise **against** running these tasks from PE because this practice can have unforeseen effects on both PE and Comply. Instead, run all CIS scans from Comply. You can view the scan results in both products.

# Beginner's guide to Comply

---

Welcome to the Beginner's guide to Comply! As a new user, you'll need to perform some initial installation and configuration tasks, and then we'll show you how to use the core features of Comply.

You're just a few steps away from enforcing compliant configurations across your infrastructure. Before you begin, we recommend familiarizing yourself with our [terminology](#) and [Comply overview](#) on page 7.

## Step 1: Install and configure Comply

---

Use the main documentation to install and configure Comply. If you already completed these steps, proceed to step 2.

- Install Puppet Application Manager (PAM)
- Set up Comply

### Related concepts

[Install PAM](#) on page 83

You can install Puppet-supported Puppet Application Manager on a single node or in an HA configuration. Both online and offline install packages are available. You can also install it on an existing Kubernetes cluster.

[Set up Comply](#) on page 132

To start using Puppet Comply, you must complete the setup process, using both Puppet Application Manager (PAM) and Puppet Enterprise (PE).

## Step 2: Set desired compliance

---

Desired compliance is the benchmark and profile that you to assign to a particular node. It is what is scanned on that node by default. Most of the time, you only need to set this once for your nodes.

Based on fact information from PE, Comply automatically assigns an appropriate benchmark for each operating system, along with a Level 1 profile, to nodes that have not been set. Accepting this option is the quickest way to get up and running with desired compliance.

Alternatively, you can manually choose your own benchmark and profiles. For more information, see [Manually set desired compliance](#).

## Step 3: Run a CIS scan

---

You are now ready to run a scan.

This topic describes how to run an initial *ad hoc* scan.

1. In Comply, click **Scan reports** and then **Run an ad hoc scan**.
2. In the drop-down menu, select **Desired compliance** or **Custom**.  
If you have not set desired compliance, follow the instructions in [Setting desired compliance](#).
3. If you selected **Custom**, select a benchmark from the **Benchmark** drop-down menu, then select an option from the **Profile** drop-down menu. To use a custom profile for this scan, select the **Use an associated custom profile?** option and choose the relevant option from the **Custom profile** drop-down menu.

- Click **Next** to see the nodes selected for scanning. Use the drop-down menus to filter nodes by operating system, environment, or node group.

To scan only a subset of nodes, deselect any nodes that you want to exclude.

**Debug mode:** By default, assessor logs are set to WARN level. To troubleshoot an issue, you can set the logging level to DEBUG for the scan by clicking **Run in debug mode**. The assessor logs can then be retrieved from the individual node.

On Linux and macOS platforms the assessor log is located at:

```
/opt/puppetlabs/comply/Assessor-CLI/logs/assessor-cli.log
```

On Windows the assessor log is located at:

```
C:/ProgramData/PuppetLabs/comply/Assessor-CLI/logs/assessor-cli.log
```

Note that scanning in debug mode increases the size of the assessor log file significantly.

- Click **Scan**.

You are taken to the **Activity feed**, which lists each scan. Scans are run as a task in PE. Click the scan name to see the scan report, or click the job ID to be taken to PE.

- Optionally, to review the results of your scan, navigate to the **Compliance Dashboard** page.

See [Scan results](#) for a description of the scan data.

Congratulations! You've completed the Beginner's guide to Comply. You're now familiar with the core features and know how to run CIS scans with Comply. To find out how you can enforce and automate CIS benchmarks on your failing nodes, see [Enforce CIS benchmarks](#) on page 169.

#### Related information

[Enforce CIS benchmarks](#) on page 169

Puppet Comply provides visibility into your compliance status, but it cannot fix your failing nodes. Instead, you can use Puppet's Compliance Enforcement Modules (CEM).

[Custom profiles](#) on page 158

A custom profile is a benchmark profile that you customize to fit your organization's internally defined standards. You can base a custom profile on an existing benchmark and profile combination, and then specify which rules to apply.

[Desired compliance](#) on page 156

Set your desired compliance. This is the benchmark and profile that you assign to a particular node and that is scanned on that node by default. Generally, you set compliance only once for your nodes.

## Puppet Application Manager

Before you can begin using Puppet Comply, you must install Puppet Application Manager. Puppet Application Manager is an administrative console that provides tools for managing Comply and other Puppet applications.

**Note:** Puppet Application Manager was previously called the platform admin console in the Comply documentation.

### What does Puppet Application Manager do?

The Puppet Application Manager installation process sets up a managed Kubernetes cluster (or, if you prefer, adds Puppet Application Manager to your existing cluster). Comply runs on this Kubernetes cluster, and Puppet Application Manager manages the cluster for you.

In the Puppet Application Manager UI, you can configure Comply, monitor the cluster's activity, upgrade to the latest version of the software, and back up your installation.

## How do I use Puppet Application Manager to deploy Comply?

Once the cluster is ready, upload your Comply license and provide any needed configuration details about your installation in the Puppet Application Manager UI. You can then deploy the latest version of Comply with one click whenever you're ready.

- [Welcome to Puppet Application Manager \(PAM\)](#) on page 50

Puppet Application Manager is an administrative console where you can install, access, and manage your Puppet applications. It is also where you can go to access upgrades to new Puppet applications releases.

- [Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

- [PAM system requirements](#) on page 69

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

- [Component versions in PAM releases](#) on page 82

These tables show the versions of components included in recent Puppet Application Manager (PAM) releases.

- [Install PAM](#) on page 83

You can install Puppet-supported Puppet Application Manager on a single node or in an HA configuration. Both online and offline install packages are available. You can also install it on an existing Kubernetes cluster.

- [Working with Puppet applications](#) on page 104

You can install and upgrade Puppet applications using the Puppet Application Manager UI.

- [Maintenance and tuning](#) on page 108

Follow these guidelines when you're tuning or performing maintenance on a node running Puppet Application Manager (PAM).

- [Upgrading PAM on a Puppet-supported cluster](#) on page 109

Upgrade Puppet Application Manager (PAM) on a Puppet-supported cluster to take advantage of new features and bug fixes, and to upgrade your cluster to the latest version of Kubernetes when one is available.

- [Upgrading PAM on a customer-supported cluster](#) on page 114

Upgrade Puppet Application Manager (PAM) on your own Kubernetes cluster to take advantage of new features and bug fixes.

- [Backing up PAM using snapshots](#) on page 115

Snapshots are point-in-time backups of your Puppet Application Manager (PAM) deployment, which can be used to roll back to a previous state or restore your installation into a new cluster for disaster recovery.

- [Migrating PAM data to a new system](#) on page 118

By using a snapshot, you can migrate your data to a new Puppet Application Manager (PAM) instance.

- [Disaster recovery with PAM](#) on page 124

It is important to prepare your system and regularly capture full snapshots. This backs up your data and makes it easier to restore your system if disaster recovery is needed.

- [Troubleshooting PAM](#) on page 125

Use this guide to troubleshoot issues with your Puppet Application Manager installation.

## Welcome to Puppet Application Manager (PAM)

---

Puppet Application Manager is an administrative console where you can install, access, and manage your Puppet applications. It is also where you can go to access upgrades to new Puppet applications releases.

Useful links:

| Puppet Application Manager docs links                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Other useful places                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Before you install</b></p> <ul style="list-style-type: none"> <li><a href="#">Release notes</a></li> <li><a href="#">System requirements</a></li> </ul> <p><b>Install Puppet Application Manager</b></p> <ul style="list-style-type: none"> <li><a href="#">PAM standalone online installation</a> on page 93</li> <li><a href="#">PAM standalone offline installation</a> on page 96</li> <li><a href="#">PAM HA online installation</a> on page 86</li> <li><a href="#">PAM HA offline installation</a> on page 90</li> </ul> <p><b>Upgrading, disaster recovery, and troubleshooting</b></p> <ul style="list-style-type: none"> <li><a href="#">Upgrading PAM on a Puppet-supported cluster</a> on page 109</li> <li><a href="#">Backing up PAM using snapshots</a> on page 115</li> <li><a href="#">Disaster recovery with PAM</a> on page 124</li> <li><a href="#">Troubleshooting PAM</a> on page 125</li> </ul> | <p><b>Docs for related Puppet products</b></p> <ul style="list-style-type: none"> <li><a href="#">Continuous Delivery for PE</a></li> <li><a href="#">Comply</a></li> </ul> <p><b>Get support</b></p> <ul style="list-style-type: none"> <li><a href="#">Support</a></li> <li><a href="#">Upgrade your support plan</a></li> </ul> <p><b>Share and contribute</b></p> <ul style="list-style-type: none"> <li><a href="#">Engage with the Puppet community</a></li> <li><a href="#">Puppet Forge</a></li> <li><a href="#">Open source projects from Puppet on GitHub</a></li> </ul> <p><b>External resources</b></p> <ul style="list-style-type: none"> <li><a href="#">Getting started with Kubernetes Off-The-Shelf software (KOTS)</a></li> </ul> |

## PAM UI

The Puppet Application Manager (PAM) UI provides administration functionality where you can access and manage your Puppet applications.

### PAM console menu

Use the console menu at the top of the Puppet Application Manager UI to manage Puppet Application Manager itself. It has three tabs of interest to us:

- Use the **Dashboard** tab to:
  - Manage your applications
  - See version history
  - Set application configuration settings
  - Access support bundles for troubleshooting
  - Manage licenses
  - View files
  - Configure registry settings
- Use the **Cluster Management** tab to view current information on the nodes in your cluster. You can also use this tab to drain, and add nodes to your cluster.
- Use the **Snapshots** tab to create point-in-time backups of your deployment, which can be used to roll back to a previous state, or restore your installation into a new cluster for disaster recovery. For more information, see [Backing up PAM using snapshots](#) on page 115.

You can also use the console menu to **Add a new application** and to log out of Puppet Application Manager.

### Application monitoring graphs

When you have Prometheus installed, the **Dashboard** tab has an **Application** sub-tab that provides several simplified graphs for tracking overall health of the system.

- **Node CPU Usage (%)** shows when hosts are getting overwhelmed (high % usage).
- **Node Memory Usage (%)** shows when hosts are reaching full memory capacity that may result in processes being killed due to out-of-memory errors.

- **Node Available Storage (%)** shows when hosts are running out of storage. At 15%, pods may start to be evicted or reads/writes on databases are paused until more storage is made available.
- **Volume Available Storage (%)** shows when application persistent volumes are getting full (low %) that may lead to problems with a particular application. Note that -

**Note:** As of the 30 June 2021 Puppet Application Manager release, the monitoring/Prometheus-Kubernetes pods limit their storage use and are expected to never fall below 10% available storage.

Puppet Application Manager HA architectures include Prometheus and Grafana. Metrics about how the system is working are sent to Prometheus, and can be displayed with Grafana. Grafana credentials are printed during install, or can be retrieved later with the following command:

```
kubectl -n monitoring get secret grafana-admin -o go-template='{{index .data "admin-user" |base64decode}}:{{index .data "admin-password" |base64decode}}'
```

### Related information

[Backing up PAM using snapshots](#) on page 115

Snapshots are point-in-time backups of your Puppet Application Manager (PAM) deployment, which can be used to roll back to a previous state or restore your installation into a new cluster for disaster recovery.

## Release notes

---

### PAM release notes

These are the new features, enhancements, resolved issues, and deprecations for Puppet Application Manager.

**Restriction:** Because kURL does not support upgrading more than two Kubernetes minor release versions at once, if you're upgrading from an older version of PAM, you might need to follow a specific upgrade path to avoid failures. For example, PAM version 1.80.0 uses Kubernetes version 1.21.x, so you can upgrade up to PAM 1.91.3 (Kubernetes version 1.23.x), but not to PAM 1.94.0 (Kubernetes version 1.24.x). To determine the specific upgrade path for your installation, please check the [table of Kubernetes versions](#) for each version of PAM.

### 20 November 2025 (Puppet Application Manager 1.112.4-r3)

**Note:** Customers using the `pam_firewall` module must upgrade the module to version 1.0.5 prior to upgrading PAM to 1.112.4-r3.

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- kURL: v2025.11.15-0
- containerd: 1.7.29
- Flannel: 0.27.4
- Project Contour: 1.32.1
- Registry: 3.0.0
- ekco: 0.28.12
- Prometheus: 0.82.2-72.9.0
- Velero: 1.16.2
- Metrics Server: 0.8.0
- MinIO: 2025-10-15T17-29-55Z
- OpenEBS: 4.3.0
- Rook: 1.17.7

Resolved in this release:

- This release contains an update to `containerd` 1.7.29, to address CVE-2025-31133, CVE-2025-52565, and CVE-2025-52881.

### 10 July 2025 (Puppet Application Manager 1.112.4-r2)

**Note:** Customers using the `pam_firewall` module must upgrade the module to version 1.0.5 prior to upgrading PAM to 1.112.4-r2.

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.32.4.

**Important upgrade information:** The upgrade process takes place on all nodes, upgrading Kubernetes to version 1.32.4 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, please keep in mind that [kURL can only be upgraded two minor versions at a time](#) on page 129.

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.112.4
- kURL: v2025.06.25-0
- containerd: 1.7.26
- Flannel: 0.26.5
- ekco: 0.28.8
- Prometheus: 0.81.0-70.0.2
- MinIO: 2025-02-18T16-25-55Z
- Rook: 1.17.1

### 4 March 2025 (Puppet Application Manager 1.112.4-r1)

**Note:** Customers using the `pam_firewall` module must upgrade the module to version 1.0.5 prior to upgrading PAM to 1.112.4-r1.

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- kURL: v2025.02.14-0
- containerd: 1.7.25
- Flannel: 0.26.4
- Velero: 1.15.2
- Metrics Server: 0.7.2
- Prometheus: 0.80.0-69.2.0
- OpenEBS: 4.2.0
- MinIO: 2024-11-07T00-52-20Z

Resolved in this release:

- This release contains an update to `containerd` 1.7.25, which resolves a memory leak.

## 22 October 2024 (Puppet Application Manager 1.112.4)

**Note:** Customers using the `pam_firewall` module must upgrade the module to version 1.0.5 prior to upgrading PAM to 1.112.4.

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.30.5.

**Important upgrade information:** The upgrade process takes place on all nodes, upgrading Kubernetes to version 1.30.5 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, please keep in mind that [kURL can only be upgraded two minor versions at a time](#) on page 129.

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.112.4
- kURL: v2024.09.26-0
- containerd: 1.6.33
- Flannel: 0.25.6
- Project Contour: 1.30.0
- Velero: 1.14.0
- Metrics Server: 0.6.4
- ekco: 0.28.7
- Prometheus: 0.76.1-62.6.0
- Registry: 2.8.3
- OpenEBS: 4.1.0
- MinIO: 2024-08-26T15-33-07Z
- Rook: 1.12.8
- Goldpinger: 3.10.0-6.2.0

## 23 July 2024 (Puppet Application Manager 1.110.0)

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.110.0
- kURL: v2024.07.02-0
- containerd: 1.6.32
- Flannel: 0.25.4
- Project Contour: 1.29.0
- Velero: 1.13.2
- Metrics Server: 0.6.4
- ekco: 0.28.7
- Prometheus: 0.74.0-59.0.0
- Registry: 2.8.3
- OpenEBS: 4.0.0
- MinIO: 2024-05-10T01-41-38Z
- Rook: 1.12.8
- Goldpinger: 3.10.0-6.2.0

## 21 May 2024 (Puppet Application Manager 1.109.0)

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.28.9.

**Important upgrade information:** The upgrade process takes place on all nodes, upgrading Kubernetes to version 1.28.9 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, please keep in mind that [kURL can only be upgraded two minor versions at a time](#) on page 129.

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.109.0
- kURL: v2024.05.03-0
- containerd: 1.6.31
- Flannel: 0.25.1
- Project Contour: 1.28.3
- Velero: 1.13.2
- Metrics Server: 0.6.4
- ekco: 0.28.6
- Prometheus: 0.73.1-58.1.1
- Registry: 2.8.3
- OpenEBS: 4.0.0
- MinIO: 2024-04-06T05-26-02Z
- Rook: 1.12.8
- Goldpinger: 3.10.0-6.2.0

## 26 March 2024 (Puppet Application Manager 1.108.0)

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.108.0
- kURL: v2024.02.23-0
- containerd: 1.6.28
- Flannel: 0.24.2
- Project Contour: 1.27.0
- Velero: 1.12.3
- Metrics Server: 0.6.4
- ekco: 0.28.4
- Prometheus: 0.71.2-56.6.0
- OpenEBS: 3.10.0
- MinIO: 2024-02-17T01-15-57Z

### 13 February 2024 (Puppet Application Manager 1.107.0)

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.107.0
- kURL: v2024.01.09-0
- containerd: 1.6.26
- Flannel: 0.24.0
- Project Contour: 1.27.0
- Registry: 2.8.3
- Velero: 1.12.2
- ekco: 0.28.4
- Prometheus: 0.70.0-55.0.0
- OpenEBS: 3.10.0
- MinIO: 2024-01-01T16-36-33Z
- Rook: 1.12.8

### 7 November 2023 (Puppet Application Manager 1.103.3)

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.28.2.

**Important upgrade information:** The upgrade process takes place on all nodes, and first upgrades Kubernetes to version 1.27.6 before upgrading to version 1.28.2 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, please keep in mind that [kURL can only be upgraded two minor versions at a time](#) on page 129.

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.103.3
- kURL: v2023.10.26-0
- containerd: 1.6.24
- Flannel: 0.22.3
- Project Contour: 1.26.1
- Registry: 2.8.3
- Velero: 1.12.1
- OpenEBS: 3.9.0
- MinIO: 2023-10-16T04-13-43Z
- Rook: 1.12.6

## 26 September 2023 (Puppet Application Manager 1.102.2)

New in this release:

- **Migrated from Weave to Flannel.** Flannel has replaced Weave as the Kubernetes CNI on Puppet-supported clusters, as Weave is no longer supported. The installation has additional interactive prompts to support this change.

### Important upgrade information:

- IPv6 and dual-stack networks are not supported on Flannel.
- Pod-to-pod networking now depends on UDP port 8472 being open instead of ports 6783 and 6784.

- **Added a host preflight.** Added a host preflight in the installer to stop installation if the installer detects the presence of a default REJECT rule in the FORWARD chain of iptables.

**Important upgrade information:** This is a known issue with the Flannel installation. To check for a REJECT rule in the FORWARD chain of iptables, run:

```
iptables -vL FORWARD
```

If there are any REJECT rules, those rules must be removed prior to the upgrade. They can be restored afterwards.

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating standalone installations, ensure there is at least 10GB of free space in `/var/openebs` to allow for migration of MinIO in this release.

- KOTS: 1.102.2
- kURL: v2023.09.15-0
- containerd: 1.6.22
- Weave: REMOVED
- Flannel: 0.22.2
- Project Contour: 1.25.2
- Velero: 1.11.1
- Kubernetes Metrics Server: 0.6.4
- ekco: 0.28.3
- Prometheus: 0.68.0-51.0.0
- OpenEBS: 3.8.0
- MinIO: 2023-09-04T19-57-37Z
- Rook: 1.12.3

**Note:** If you are using the [firewall module](#) to manage your PAM install, you must update it to version 1.0.4 to support this PAM release.

### 18 July 2023 (Puppet Application Manager 1.100.3)

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.26.6.

**Important upgrade information:** The upgrade process takes place on all nodes, and first upgrades Kubernetes to version 1.25 before upgrading to version 1.26.6 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, please keep in mind that [kURL can only be upgraded two minor versions at a time](#) on page 129.

- **Component upgrades to address security issues.** This version upgrades, adds, and removes the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.100.3
- kURL: v2023.06.27-0
- Prometheus: 0.65.2-46.8.0
- OpenEBS: 3.7.0
- MinIO: 2023-06-19T19-52-50Z
- Rook: 1.11.8

**Note:** If you are using the [firewall module](#) to manage your PAM install, you must update it to version 1.0.3 to support this PAM release.

### 8 June 2023 (Puppet Application Manager 1.99.0)

New in this release:

- **Component upgrades to address security issues.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.99.0
- kURL: v2023.05.22-0
- containerd: 1.6.21
- Weave: 2.8.1-20230417
- Project Contour: 1.25.0
- Registry: 2.8.2
- Velero: 1.11.0
- ekco: 0.27.1
- Prometheus: 0.65.1-45.28.0
- OpenEBS: 3.6.0
- MinIO: 2023-05-18T00-05-36Z
- Rook: 1.11.5
- Goldpinger: 3.7.0-6.0.1

**Note:** For offline HA installs the Rook update in this release can cause significant downtime (around 4 hours) while downloading additional files. It is possible to [do some of this prior to upgrading](#) Puppet Application Manager from 1.97.0 to 1.99.0 to decrease the downtime.

## 25 April 2023 (Puppet Application Manager 1.97.0)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of OpenEBS to version 3.5.0, an upgrade of kURL to v2023.04.11-0, an upgrade of containerd to 1.6.20, an upgrade of Weave to version 2.8.1-20230324, an upgrade of Project Contour to version 1.24.3, an upgrade of ekco to 0.26.5, an upgrade of Velero to version 1.10.2, an upgrade of the Prometheus bundle to version 0.63.0-45.9.1, and upgrade of Kubernetes Metrics Server to version 0.6.3, an upgrade of KOTS to 1.97.0, an upgrade of MinIO to version 2023-03-24T21-41-23Z, and an upgrade of Goldpinger to 3.7.0-5.6.0.

**Note:** Before updating, ensure MinIO has 10GB of free space.

Deprecated in this release:

- **force-reapply-addons flag.** Starting with Puppet Application Manager 1.97.0, the `force-reapply-addons` flag is deprecated and generates a warning on use. This flag is only required when upgrading to a Puppet Application Manager version prior to 1.97.0.

## 28 February 2023 (Puppet Application Manager 1.94.0)

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.24.10.

**Important upgrade information:** The upgrade process takes place on all nodes, and first upgrades Kubernetes to version 1.24.10 on each. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, because [kURL can only be upgraded two minor versions at a time](#) on page 129, if you're on PAM version 1.80.0 or earlier, you must upgrade to PAM 1.81.1 before upgrading to PAM 1.94.0.

- This release also includes component upgrades to address security issues and general bug fixes.

## 10 January 2023 (Puppet Application Manager 1.91.3)

New in this release:

- **Component upgrades to address security issues and support RHEL 8.7.** This version upgrades the following:

**Note:** Before updating, ensure MinIO has 10GB of free space.

- KOTS: 1.91.3
- MinIO: 2022-10-20T00-55-09Z
- OpenEBS: 3.3.0
- Prometheus: 0.60.1-41.7.3
- ekco: 0.26.1
- Velero: 1.9.4
- Project Contour: 1.23.1
- kURL: v2022.12.12-0
- Weave: 2.8.1-20221122
- Goldpinger: 3.7.0-5.5.0

## 28 September 2022 (Puppet Application Manager 1.81.1)

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.23.9.

**Important upgrade information:** The upgrade process takes place on all nodes, and first upgrades Kubernetes to version 1.22 before upgrading to version 1.23.9. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

Additionally, because [kURL can only be upgraded two minor versions at a time](#) on page 129, if you're upgrading from PAM version 1.56.0 or earlier, you must upgrade to PAM 1.80.0 before upgrading to PAM 1.81.1.

For legacy installations, Kubernetes remains on version 1.19.15. If you're not sure which installation type you're running, see [How to determine your version of Puppet Application Manager](#).

## 16 August 2022 (Puppet Application Manager 1.80.0)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version upgrades containerd to 1.4.13, KOTS to 1.80.0, ekco to 0.19.6, and Goldpinger to 3.5.1-5.2.0.

Resolved in this release:

- Fixed an issue where legacy encryption keys didn't load properly during snapshot restores.

## 2 August 2022 (Puppet Application Manager 1.76.2)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of OpenEBS to version 3.2.0, an upgrade of Weave to version 2.8.1-20220720, an upgrade of Project Contour to version 1.21.1, and an upgrade of MinIO to version 2022-07-17T15-43-14Z.

**Note:** Before updating, ensure MinIO has 10GB of free space.

## 20 July 2022 (Puppet Application Manager 1.76.1)

New in this release:

- **Support for Red Hat Enterprise Linux version 8.6.** Beginning with version 1.76.1, PAM can be successfully installed on systems running Red Hat Enterprise Linux version 8.6.
- **More log data is now retained.** To ensure that you and our Support team have the data you need in debugging scenarios, the size of the pod logs has been increased from 10 files of 10MiB each to 10 files of 50MiB each. This change increases the storage used in `/var/log/pods` by 400MiB.
- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of Velero to version 1.9.0 and an upgrade of the Prometheus bundle to version 0.57.0-36.2.0.
- **Other component upgrades.** This version also includes an upgrade of Registry to version 2.8.1 and an upgrade of MinIO to version 2022-07-06T20-29-49Z.

**Note:** Before updating, ensure MinIO has 10GB of free space.

Resolved in this release:

- Velero pods no longer get stuck in a pending state when creating a snapshot to be saved to internal storage on a Puppet-supported cluster.

### 23 June 2022 (Puppet Application Manager 1.72.1)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of ekco to version 0.19.2 and an upgrade of kURL to v2022.06.17-0.

### 26 May 2022 (Puppet Application Manager 1.70.1)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of Project Contour to version 1.21.0, an upgrade of Velero to version 1.8.1, and an upgrade of the Prometheus bundle to version 0.56.2-35.2.0.

Resolved in this release:

- Image garbage collection in Kubernetes installer-created clusters (embedded clusters) no longer removes images outside of the application's dedicated registry namespace.
- The **Deploy** button is now present in newly updated versions after the configuration is updated from the previously deployed version.
- Legends are now shown properly for the performance graphs on the dashboard.

### 12 April 2022 (Puppet Application Manager 1.68.0)

New in this release:

- **Install a specific version of an application.** When installing a Puppet application using the automated installation method, you now have the option to specify the application's version by passing the `--app-version-label=<version>` flag to the `kubectl kots install` command. For more information, go to [Automate PAM and Puppet application online installations](#) on page 98.
- **Status reporting improvements.** The status reporting tools can now detect when an application is being upgraded.
- **Component upgrades to address CVEs.** To address various CVEs in Envoy, this version includes an upgrade of Project Contour to version 1.20.1.
- **Other component upgrades.** This version includes an upgrade of KOTS to version 1.68.0, which enables Kubernetes audit event logging by default and adds a 1 GB storage requirement for `/var/log/apiserver`.

Resolved in this release:

- During image garbage collection, images still in use by the cluster are no longer in danger of being deleted from the private registry in a Kubernetes installer-created cluster.

## 1 March 2022 (Puppet Application Manager 1.64.0)

Resolved in this release:

- Diffs are now shown correctly in the PAM UI.
- The OpenSSL package is no longer a prerequisite for successful installation on newer Red Hat Enterprise Linux 7 systems.
- You can now successfully install Puppet Application Manager on Red Hat Enterprise Linux 8 systems without the need to force-install the kurl-local audit-libs library.

## 17 February 2022 (Puppet Application Manager 1.62.0)

**Important:** Version 1.0.2 of the `puppetlabs/pam_firewall` module is now available. To avoid conflicts, upgrade the module **before** upgrading Puppet Application Manager to version 1.62.0.

New in this release:

- **Kubernetes version upgrade.** For standalone and HA installations, this version includes an upgrade of Kubernetes to version 1.21.8.

**Important upgrade information:** The upgrade process takes place on all nodes, and first upgrades Kubernetes to version 1.20 before upgrading to version 1.21.8. For a three-node cluster, you can expect the upgrade process to take around an hour. Confirmations are required during the upgrade process.

For legacy installations (installed before May 2021), this version includes an upgrade of Kubernetes to version 1.19.15.

**Tip:** See [How to determine your version of Puppet Application Manager](#) if you're not sure which installation type you're running.

- **Prometheus enabled on standalone architecture.** Beginning with version 1.62.0 Prometheus is enabled by default on all new and existing standalone Puppet Application Manager installations. Prometheus requires an additional 350m CPU and 500MiB of memory, so ensure your system is properly sized before upgrading. Prometheus is an optional component; if you need to disable it to conserve resources, see [Optional components](#) on page 128.
- **Automatic certificate rotation.** By default, the self-signed certificates used by Project Contour and Envoy expire after one year. This version includes an update that auto-rotates those certificates before they expire.
- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of containerd to version 1.4.12.
- **Other component upgrades.** This version includes an upgrade of KOTS to version 1.62.0.

Deprecated in this release:

- **Legacy architecture.** The legacy architecture, which was the version of Puppet Application Manager available for installation prior to May 2021, is now deprecated. (See [How to determine your version of Puppet Application Manager](#) if you need to confirm whether you're running the legacy architecture.) The legacy architecture utilizes Rook 1.0, which is incompatible with Kubernetes version 1.20 and newer versions. Kubernetes version 1.19 is no longer receiving security updates. Puppet will continue to update legacy architecture components other than Kubernetes until 30 June 2022. If security advisories against Kubernetes 1.19 arise, the remediation path is to migrate to one of the newer architectures by following the instructions in [Migrating PAM data to a new system](#) on page 118.

**Important:** Before beginning the migration process from a legacy deployment you must upgrade to PAM version 1.62.0 with the `force-reapply-addons` flag included in the upgrade command. Find upgrade instructions at [PAM legacy upgrades](#) on page 112 and [PAM offline legacy upgrades](#) on page 113.

## 30 November 2021 (Puppet Application Manager 1.56.0)

This release includes an upgrade of KOTS to version 1.56.0, which adds the following improvements:

- **Improved support bundles:** Adds an option to upload a support bundle directly from Puppet Application Manager.
- **Improved troubleshooting:** Adds detailed information on failing pods to the **Troubleshoot** tab.

### 6 October 2021 (Puppet Application Manager 1.52.1)

New in this release:

- **Improved statuses.** More granular status levels are now available from the **Application** tab.
- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of Kubernetes to 1.19.15.
- **Other component upgrades.** This version includes an upgrade of KOTS to version 1.52.1.

Resolved in this release:

- Generating a support bundle no longer results in unusually high memory use.
- Preflight check logs post to info level for progress messages and to error level for error messages.

### 25 August 2021 (Puppet Application Manager 1.49.0)

New in this release:

- **Component upgrades to address CVEs.** To address various CVEs, this version includes an upgrade of Kubernetes to 1.19.13, an upgrade of Project Contour to 1.18.0, and an upgrade of Velero to 1.6.2.
- **Goldpinger.** High availability architectures now include Goldpinger, which aids the debugging of network issues.
- **containerd upgrade.** This version includes an upgrade of containerd to version 1.4.6, and removes the need to use the `force-reapply-addons` option when upgrading.
- **Other component upgrades.** This version includes an upgrade of KOTS to version 1.49.0, an upgrade of ekco to 0.11.0, an upgrade of Prometheus to 0.49.0, and an upgrade of Rook to 1.5.12.

### 30 June 2021 (Puppet Application Manager 1.44.1)

New in this release:

- **Certificate auto-rotation for standalone architecture.** Certificates are now automatically rotated for the Kubernetes API and Puppet Application Manager UI in the standalone architecture. With this change, certificate auto-rotation is now supported in all Puppet Application Manager architectures.
- **Rook upgrades.** This version includes an upgrade of Rook in the high availability architecture to 1.5.11 and the version of Rook in the legacy architecture to 1.0.4-14.2.21. These upgrades address a vulnerability in Ceph components (CVE-2021-20288).
- **Prometheus upgrade.** This version includes an upgrade of Prometheus in the high availability and legacy architectures to 0.48.1. Additionally, Prometheus disk usage is now limited in order to preserve the storage space required for the usage charts on the **Application** tab.
- **Other component upgrades.** This version includes an upgrade of KOTS to version 1.44.1, an upgrade of Project Contour to version 1.15.1, and an upgrade of Weave to version 2.8.1.

Resolved in this release:

- Snapshots can now successfully use the **Other S3-Compatible Storage** option as the storage destination.

To apply this update, add the `force-reapply-addons` option during upgrade. For example:

```
curl <url> | bash -s force-reapply-addons
```

### 26 May 2021

New in this release:

- **runC.** The version of runC has been upgraded to v1.0.0-rc95 to address CVE-2021-30465.

Known issues in this release:

- Running the KOTS installer with the `airgap` and `kurl-registry-ip` flags results in an error.  
As a workaround (if you do not have any applications already installed in the cluster), delete the registry service, recreate the registry service IP and then re-run the installation script with the `kurl-registry-ip` flag.

### 10 May 2021 (Puppet Application Manager 1.40.0)

New in this release:

- Distinct architectures for standalone and high availability deployments of the Puppet Application Manager platform. Standalone supports lower system requirements and resolves inherent flaws in using Ceph on a single node. High availability uses an updated version of Rook for faster, more reliable distributed storage.

**Note:** It is not possible currently to upgrade to these architectures from existing installations. However, migrating applications between them is on the roadmap for a future release.

- The previous architecture is maintained as the legacy configuration. This version includes an upgrade of Kubernetes to 1.19.10; this upgrade process upgrades through Kubernetes 1.18, and happens on all nodes. It can take ~1 hour to do for a 3-node cluster, and requires confirmations during that period. It also includes an upgrade of Project Contour to version 1.14.1, adds Metrics Server 0.4.1, an upgrade of ekco to 0.10.1, and an upgrade of Prometheus to 2.26.0.

For more information on legacy upgrades, see [PAM legacy upgrades](#) on page 112.

### 15 April 2021 (Puppet Application Manager 1.38.0)

New in this release:

- **Snapshots.** Puppet Application Manager now supports full (instance-level) snapshots, which can be used for application rollbacks and disaster recovery. For more information, see [Backing up Puppet Application Manager using snapshots](#).
- **Component upgrades.** This version includes an upgrade of KOTS to version 1.38.0.

### 17 February 2021 (Puppet Application Manager 1.29.3)

New in this release:

- **Support for Ubuntu 20.04.** You can now run Puppet Application Manager on Ubuntu 20.04.
- **Component upgrades.** This version includes an upgrade of Prometheus to version 2.22.1 and Prometheus Operator to version 0.44.1, an upgrade of KOTS to version 1.29.3, an upgrade of Project Contour to version 1.12.0, and an upgrade of ekco to version 0.10.0.

### 3 February 2021 (Puppet Application Manager 1.29.2)

New in this release:

- **Component upgrades.** This version includes an upgrade of KOTS to version 1.29.2, an upgrade of Project Contour to version 1.11.0, and an upgrade of `containerd` to version 1.4.3.

Resolved in this release:

- During their initial preflight checks, new installations now pull images successfully and no longer report a `Failed to pull image` error.

### 7 December 2020

New in this release:

- **Support for Red Hat Enterprise Linux (RHEL) 8 and CentOS 8.** You can now run Puppet Application Manager on RHEL version 8 and CentOS version 8. To support this change, `containerd` is now used independently of Docker during the installation process.

- **Component upgrades.** This version includes an upgrade of Kubernetes to version 1.17.13.

### Related information

[Upgrading PAM on a Puppet-supported cluster](#) on page 109

Upgrade Puppet Application Manager (PAM) on a Puppet-supported cluster to take advantage of new features and bug fixes, and to upgrade your cluster to the latest version of Kubernetes when one is available.

[Backing up PAM using snapshots](#) on page 115

Snapshots are point-in-time backups of your Puppet Application Manager (PAM) deployment, which can be used to roll back to a previous state or restore your installation into a new cluster for disaster recovery.

## Known issues

These are the known issues for Puppet Application Manager (PAM).

### Restarting a PAM node on v1.103.3, v1.107.0, and v1.108.0 does not successfully bring back up all pods

A known Kubernetes issue impacts PAM versions v1.103.3, v1.107.0, and v1.108.0. Rebooting a node can result in pods incorrectly changing phases, causing them to remain at states such as 0/1 Completed and 0/1 Error. This can lead to pod crashloops and application outages. To work around this issue, delete the errant pods, which forces them to restart. Check the pods afterwards to ensure they come up as “Running”. This issue is resolved in PAM version 1.109.0

### PAM versions 1.72.1 and older cannot be installed on RHEL 8.6+ systems

A known issue in kURL prevents Puppet Application Manager versions 1.72.1 and older from successfully installing on Red Hat Enterprise Linux (RHEL) version 8.6 and newer versions. To work around this issue, install or upgrade to Puppet Application Manager version 1.76.1 or a newer version, which support RHEL version 8.6.

### Velero fails if network file system (NFS) snapshot storage is misconfigured

In Puppet Application Manager version 1.64.0 and newer versions, changes to the configuration of snapshot storage on a network file system (NFS) is appended to Velero containers, rather than replaced. This means that if NFS snapshot storage is misconfigured, attempts to fix the configuration do not correct the problem. This issue manifests as a failure of Velero to start up.

### OpenSSL package required for newer RHEL 7 systems with PAM 1.62.0

Attempts to install Puppet Application Manager version 1.62.0 or older on newer Red Hat Enterprise Linux (RHEL) 7 systems fail unless the OpenSSL package is present on the system before installation. To work around this issue, run `yum install openssl` and then re-run the PAM installation script.

### Package updates with yum or DNF fail after upgrading PAM

If you are unable to run `yum update` or `dnf upgrade` after a PAM upgrade, run one of the following commands to clean up a temporary module added by PAM:

```
yum module reset kurl.local
```

or

```
dnf module reset kurl.local
```

## Architecture overview

---

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

PAM can run on Puppet-supported or customer-supported Kubernetes clusters. Due to potential variations in the architecture of customer-supported clusters, the architecture overview provided on this page assumes PAM is running on Puppet-supported clusters. For more information on installing on a customer-supported Kubernetes cluster, see [Install Puppet applications using PAM on a customer-supported Kubernetes cluster](#) on page 84.

## Terminology

Throughout this documentation, we use a few terms to describe different roles nodes can take:

- **Primary** - A primary node runs core Kubernetes components (referred to as the Kubernetes control plane) as well as application workloads. At least three primaries are required to support high availability for Puppet Application Manager. These are also sometimes referred to as *masters*.
- **Secondary** - A secondary node runs application workloads. These are also sometimes referred to as *workers*.

Puppet Application Manager is built on the KOTS (Kubernetes off-the-Shelf) project, and we occasionally use its CLI tools (`kubect1`, `kots`) to manage the installation.

## Standalone architecture

Standalone is optimized for limited resources, storing data directly on disk. If you need to remove optional components like Prometheus and Grafana to decrease resource utilization, see [Optional components](#) on page 128. While additional compute capacity can be added through secondary nodes, this does not provide increased resilience as data is only stored on the node where a component service runs.

For information on migrating data from standalone to HA deployments, see [Migrating data between two systems with different architectures](#) on page 123.

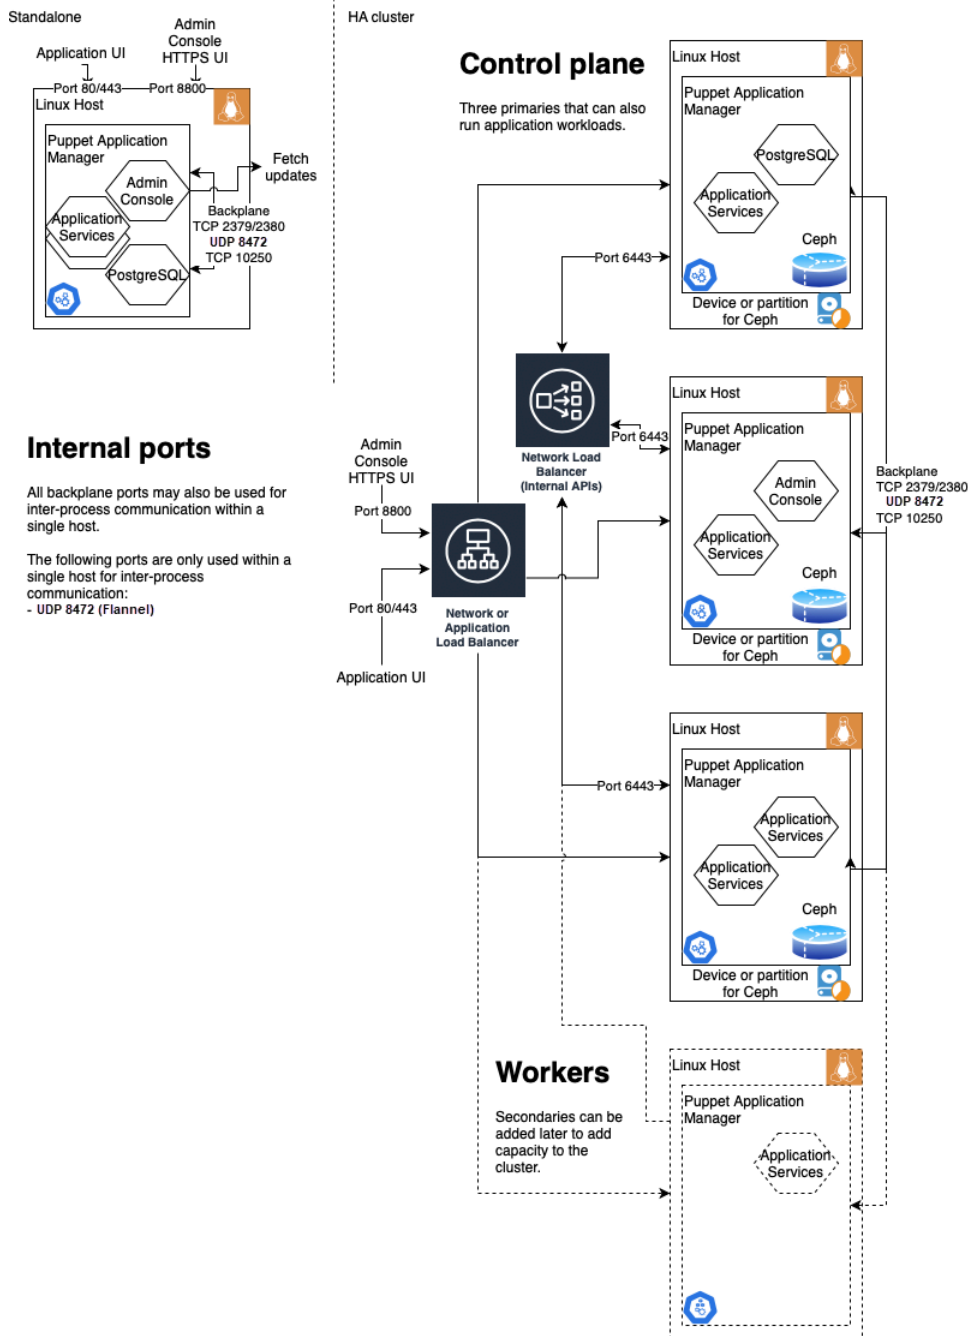
## HA architecture

A high availability (HA) architecture provides high availability for scheduling application services during failure and uses Ceph for distributed storage in case of node failure. Individual applications may still experience some loss of availability (up to 10 minutes) if individual services do not have replicas and need to be rescheduled. For more information, see [Reduce recovery time when a node fails](#) on page 127. An HA implementation requires a cluster of three primary nodes. Additional compute capacity can be added through secondary nodes.

The HA architecture installs Prometheus and Alertmanager. These are used to provide system monitoring in the Puppet Application Manager UI. Prometheus and Alertmanager are unauthenticated on ports 30900 and 30903, and you are recommended to control access to these ports via firewall rules. For information on how to remove Prometheus and Alertmanager, see [Optional components](#) on page 128.

## Puppet Application Manager architectures

The following diagram and lists outline some of the core components involved in standalone and HA architectures and how they communicate. For a detailed list of ports used by Puppet Application Manager, refer to the **Cluster port requirements** sections of the [PAM system requirements](#) on page 69. For firewall information, refer to [Web URL and port requirements for firewalls](#) on page 81.



**Standalone architecture**

**Puppet Application Manager**

Lives on a cluster within a Linux host.

The PAM application includes the admin console, application services, and PostgreSQL.

PAM communicates out of the Linux host to fetch updates.

**UI ports**

The application UI communicates on 80/443 to the Linux host.

The admin console HTTPS UI communicates on 8800 to the Linux host.

**Backplane and internal ports**

Backplane ports include 8472 (UDP) and 10250 (TCP).

Backplane ports can also be used within a single host for inter-process communication.

These ports are only used within a single host for inter-process communication (Flannel): 8472 (UDP)

#### **Additional default ports**

30900: Prometheus UI

30902: Grafana UI

30903: Alertmanager UI

### **HA cluster architecture**

#### **Control plane (primaries)**

Multiple primaries that can also run application workloads.

Structured as clusters within Linux hosts with a device or partition for Ceph.

Each primary hosts PAM and can run application services in addition to supporting either PostgreSQL or the admin console.

#### **Workers (secondaries)**

Can be added later to add capacity for running application workloads.

Structured as clusters within Linux hosts.

#### **Network or Application Balancer**

The balancer communicates out to the control plane (primaries) and workers (secondaries).

Receives admin console HTTPS UI communication over 8800.

Receives application UI communication over 80/443.

Network load balancer internal APIs communicate with primaries and secondaries over 6443.

To learn about setting up health checks for your load balancer, go to [Load balancer health checks](#) on page 108.

#### **Backplane and internal ports**

Backplane ports include 2379/2380 (TCP), 8472 (UDP), and 10250 (TCP).

Backplane ports can also be used within a single host for inter-process communication.

These ports are only used within a single host for inter-process communication (Flannel): 8472 (UDP)

#### **Additional default ports**

30900: Prometheus UI

30902: Grafana UI

30903: Alertmanager UI

### **UNSUPPORTED: Legacy architecture**

**Note:** The legacy architecture utilizes Rook 1.0, which is incompatible with Kubernetes version 1.20 and newer versions. Kubernetes version 1.19 is no longer receiving security updates. The legacy architecture reached the end of its support lifecycle on **30 June 2022**, and Puppet no longer updates legacy architecture components.

The Puppet Application Manager legacy architecture reflects an older configuration that used Ceph 1.0 which hosted data directly on the file system. Installing the legacy architecture is no longer supported.

For information on upgrading to a newer version of the legacy architecture, see [PAM legacy upgrades](#) on page 112 and [PAM offline legacy upgrades](#) on page 113.

For information on migrating data from a legacy architecture to a standalone or HA architecture, go to our Support Knowledge Base instructions:

- [Migrate to a supported PAM architecture for Continuous Delivery for PE](#)
- [Migrate to a supported PAM architecture for Comply](#)

## Related information

[Reduce recovery time when a node fails](#) on page 127

If a node running a non-replicated service like PostgreSQL fails, expect some service downtime.

[Install Puppet applications using PAM on a customer-supported Kubernetes cluster](#) on page 84

Use these instructions to install Puppet Application Manager and any Puppet applications on an existing Kubernetes cluster.

[PAM legacy upgrades](#) on page 112

The legacy architecture is no longer supported. However, if you have not yet migrated to a supported architecture, you can use this method to upgrade Puppet Application Manager (PAM).

[PAM offline legacy upgrades](#) on page 113

The legacy architecture is no longer supported. However, if you have not yet migrated to a supported architecture, you can use this method to upgrade Puppet Application Manager (PAM) on offline nodes.

[Troubleshooting PAM](#) on page 125

Use this guide to troubleshoot issues with your Puppet Application Manager installation.

## PAM system requirements

---

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

### Customer-supported cluster hardware requirements

The following Kubernetes distributions are supported:

- Google Kubernetes Engine
- AWS Elastic Kubernetes Service

If you use a different distribution, contact [Puppet Support](#) for more information on compatibility with PAM.

Application requirements:

| Application                                    | CPU   | Memory | Storage | Ports                                                            |
|------------------------------------------------|-------|--------|---------|------------------------------------------------------------------|
| Continuous Delivery for Puppet Enterprise (PE) | 3 CPU | 8 GB   | 280 GB  | Ingress, NodePort 8000<br><b>Note:</b> NodePort is configurable  |
| Puppet Comply®                                 | 7 CPU | 7 GB   | 35 GB   | Ingress, NodePort 30303<br><b>Note:</b> NodePort is configurable |

Make sure that your Kubernetes cluster meets the minimum requirements:

- Kubernetes version 1.24-1.26.
- A default storage class that can be used for relocatable storage.
- A standard Ingress controller that supports websockets (we have tested with Project Contour and NGINX).
- We currently test and support Google Kubernetes Engine (GKE) clusters.

**Cluster ports:** In addition to the NodePorts used by your Puppet applications, make sure that TCP port 443 is open for your ingress controller.

## Puppet-supported HA cluster hardware requirements

A high availability (HA) configuration uses multiple servers to provide availability in the event of a server failure.

A majority of servers must be available to preserve service availability. Below are suggested configurations for each application.

### Continuous Delivery for Puppet Enterprise (PE)

Three servers (referred to as primaries during installation) with the following minimum requirements:

| CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Open ports                                                                                  |
|-------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| 6 CPU | 10 GB  | <p>100 GB on an unformatted storage device.</p> <p>1 GB for <code>/var/log/apiserver</code> for Kubernetes audit logs.</p> <p>An additional 140 GB for <code>/var/lib</code>. You can use separate filesystems if necessary, but it is not a requirement to do so. For your reference, here is how the usage is roughly divided:</p> <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 10 GB for <code>/var/lib/rook</code> (plus buffer)</li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> </ul> <p><b>Note:</b> The storage backend prefers the file system inhabited by <code>/var/lib/rook</code> to remain below 70% utilization.</p> <p>SSDs (or similarly low-latency storage) are recommended for <code>/var/lib/etcd</code> and <code>/var/lib/rook</code>.</p> | <p><b>TCP:</b> 80, 443, 2379, 2380, 6443, 8000, 8800, and 10250</p> <p><b>UDP:</b> 8472</p> |

### Puppet Comply

Three servers (referred to as primaries during installation) with the following minimum requirements:

| CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Open ports                                                                                   |
|-------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| 7 CPU | 10 GB  | <p>100 GB on an unformatted storage device.</p> <p>1 GB for <code>/var/log/apiserver</code> for Kubernetes audit logs.</p> <p>An additional 140 GB for <code>/var/lib</code>. You can use separate filesystems if necessary, but it is not a requirement to do so. For your reference, here is how the usage is roughly divided:</p> <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 10 GB for <code>/var/lib/rook</code> (plus buffer)</li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> </ul> <p><b>Note:</b> The storage backend prefers the file system inhabited by <code>/var/lib/rook</code> to remain below 70% utilization.</p> <p>SSDs (or similarly low-latency storage) are recommended for <code>/var/lib/etcd</code> and <code>/var/lib/rook</code>.</p> | <p><b>TCP:</b> 80, 443, 2379, 2380, 6443, 8800, 10250, and 30303</p> <p><b>UDP:</b> 8472</p> |

**Continuous Delivery for Puppet Enterprise (PE) and Puppet Comply**

Three servers (referred to as primaries during installation) with the following minimum requirements:

| CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | Open ports                                                                                         |
|-------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| 8 CPU | 13 GB  | <p>150 GB on an unformatted storage device.</p> <p>1 GB for <code>/var/log/apiserver</code> for Kubernetes audit logs.</p> <p>An additional 140 GB for <code>/var/lib</code>. You can use separate filesystems if necessary, but it is not a requirement to do so. For your reference, here is how the usage is roughly divided:</p> <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 10 GB for <code>/var/lib/rook</code> (plus buffer)</li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> </ul> <p><b>Note:</b> The storage backend prefers the file system inhabited by <code>/var/lib/rook</code> to remain below 70% utilization.</p> <p>SSDs (or similarly low-latency storage) are recommended for <code>/var/lib/etcd</code> and <code>/var/lib/rook</code>.</p> | <p><b>TCP:</b> 80, 443, 2379, 2380, 6443, 8000, 8800, 10250, and 30303</p> <p><b>UDP:</b> 8472</p> |

For a detailed example of an HA configuration running Continuous Delivery for PE and Puppet Comply, see [Example of an HA cluster that supports CDPE and Comply](#).

### Networking requirements

Gigabit Ethernet (1GbE) and a latency of less than 10 milliseconds (ms) between cluster members is sufficient for most deployments. For more information on networking for specific Puppet Application Manager components, see the documentation for [Ceph](#), and [etcd](#).

### Cluster port requirements

Puppet Application Manager (PAM) uses the following ports in an HA cluster architecture:

| Category                                                    | Port | Protocol | Purpose                                                                          | Source         |
|-------------------------------------------------------------|------|----------|----------------------------------------------------------------------------------|----------------|
| <b>Puppet application ports</b>                             | 443  | TCP      | Web UI<br>Relies on Server Name Indication to route requests to the application. | Browser        |
| <b>Continuous Delivery for Puppet Enterprise (PE) ports</b> | 8000 | TCP      | Webhook service                                                                  | Source control |

| Category                   | Port       | Protocol | Purpose                                                                                                                     | Source                       |
|----------------------------|------------|----------|-----------------------------------------------------------------------------------------------------------------------------|------------------------------|
| <b>Puppet Comply ports</b> | 30303      | TCP      | Communication with Puppet Enterprise (PE)                                                                                   | PE instance                  |
| <b>Platform ports</b>      | 2379, 2380 | TCP      | High availability (HA) communication<br><br>Only needs to be open between the cluster's primary nodes.                      | etcd on the Kubernetes host. |
|                            | 6443       | TCP      | Kubernetes API<br><br>Might be useful to expose to workstations.                                                            | Admin workstation            |
|                            | 8472       | UDP      | Kubernetes networking - Flannel                                                                                             | Kubernetes host              |
|                            | 8800       | TCP      | PAM                                                                                                                         | Admin browser                |
|                            | 9001       | TCP      | Internal registry in offline installs only.<br><br>Requires configuring an Ingress to use this port.                        | Kubernetes host              |
|                            | 9090       | TCP      | Rook CSI RBD Plugin Metrics                                                                                                 | Kubernetes host              |
|                            | 10250      | TCP      | Kubernetes cluster management<br><br>Only communicates in one direction, from a primary to other primaries and secondaries. | Kubernetes host              |

Additionally, these ports are configured by default: 30900 (Prometheus UI), 30902 (Grafana UI), and 30903 (Alertmanager UI)

For Kubernetes-specific information, refer to [Networking Requirements in the Kurl documentation](#).

## IP address range requirements

**Important:** Puppet Application Manager must be installed on nodes with static IP assignments because IP addresses cannot be changed after installation.

Ensure that IP address ranges `10.96.0.0/22` and `10.32.0.0/22` are locally accessible. See [Resolve IP address range conflicts](#) for instructions.

**Note:** The minimum size for CIDR blocks used by PAM are:

- `/23` for pod and service CIDRs
- Default of `/22` is recommended to support future expansion

## Antivirus and antimalware considerations

Antivirus and antimalware software can impact PAM and its applications or prevent them from functioning properly.

To avoid issues, exclude the following directories from antivirus and antimalware tools that scan disk write operations:

- `/var/lib/rook`
- `/var/lib/kubelet`
- `/var/lib/containerd`

## Firewall modules

If you use the `puppetlabs/firewall` module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the `puppetlabs/pam_firewall` module before installing Puppet Application Manager.

If you've already installed PAM, apply the `pam_firewall` module and then restart the `kube-proxy` service to recreate its iptables rules by running the following on a primary:

```
systemctl restart kubelet
 kubectl -n kube-system delete pod -l k8s-app=kube-proxy
 kubectl -n kube-flannel delete pod -l app=flannel
```

For more information, see the PAM [firewall module](#).

## Supported operating systems

Puppet Application Manager and the applications it supports can be installed on these operating systems:

| Operating system                | Supported versions                                                                           |
|---------------------------------|----------------------------------------------------------------------------------------------|
| Amazon Linux                    | 2                                                                                            |
| CentOS                          | 7.4, 7.5, 7.6, 7.7, 7.8, 7.9<br>8.0, 8.1, 8.2, 8.3, 8.4                                      |
| Oracle Linux                    | 7.4, 7.5, 7.6, 7.7, 7.8, 7.9<br>8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 8.8                  |
| Red Hat Enterprise Linux (RHEL) | 7.4, 7.5, 7.6, 7.7, 7.8, 7.9<br>8.0, 8.1, 8.2, 8.3, 8.4, 8.5, 8.6, 8.7, 8.8<br>9.0, 9.1, 9.2 |

| Operating system                      | Supported versions      |
|---------------------------------------|-------------------------|
| Rocky Linux                           | 9.0, 9.1, 9.2           |
| Ubuntu (General availability kernels) | 18.04<br>20.04<br>22.04 |

## Puppet-supported standalone hardware requirements

Here are the suggested configurations for standalone installations.

### Continuous Delivery for Puppet Enterprise (PE)

| CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                          | Open ports                                                                       |
|-------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| 4 CPU | 8 GB   | 220 GB for <code>/var/lib</code> and <code>/var/opensbs</code> This is primarily divided among: <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> <li>• 100 GB for <code>/var/opensbs</code></li> </ul> | <b>TCP:</b> 80, 443, 2379, 2380, 6443, 8000, 8800, and 10250<br><b>UDP:</b> 8472 |

### Puppet Comply

| CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                          | Open ports                                                                        |
|-------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|
| 7 CPU | 7 GB   | 220 GB for <code>/var/lib</code> and <code>/var/opensbs</code> This is primarily divided among: <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> <li>• 100 GB for <code>/var/opensbs</code></li> </ul> | <b>TCP:</b> 80, 443, 2379, 2380, 6443, 8800, 10250, and 30303<br><b>UDP:</b> 8472 |

### Cluster port requirements

Puppet Application Manager (PAM) uses the following ports in a standalone architecture:

| Category                                                    | Port  | Protocol | Purpose                                                                                                                 | Source            |
|-------------------------------------------------------------|-------|----------|-------------------------------------------------------------------------------------------------------------------------|-------------------|
| <b>Puppet application ports</b>                             | 442   | TCP      | Web UI<br>Relies on Server Name Indication to route requests to the application.                                        | Browser           |
| <b>Continuous Delivery for Puppet Enterprise (PE) ports</b> | 8000  | TCP      | Webhook service                                                                                                         | Source control    |
| <b>Puppet Comply ports</b>                                  | 30303 | TCP      | Communication with Puppet Enterprise                                                                                    | PE instance       |
| <b>Platform ports</b>                                       | 6443  | TCP      | Kubernetes API<br>Might be useful to expose to workstations.                                                            | Admin workstation |
|                                                             | 8472  | UDP      | Kubernetes networking - Flannel                                                                                         | Kubernetes host   |
|                                                             | 8800  | TCP      | PAM                                                                                                                     | Admin browser     |
|                                                             | 9001  | TCP      | Internal registry in offline installs only.<br>Requires configuring an Ingress to use this port.                        | Kubernetes host   |
|                                                             | 10250 | TCP      | Kubernetes cluster management<br>Only communicates in one direction, from a primary to other primaries and secondaries. | Kubernetes host   |

Additionally, these ports are configured by default: 30900 (Prometheus UI), 30902 (Grafana UI), and 30903 (Alertmanager UI)

For Kubernetes-specific information, refer to [Networking Requirements in the Kurl documentation](#).

### IP address range requirements

**Important:** Puppet Application Manager must be installed on nodes with static IP assignments because IP addresses cannot be changed after installation.

Ensure that IP address ranges `10.96.0.0/22` and `10.32.0.0/22` are locally accessible. See [Resolve IP address range conflicts](#) for instructions.

**Note:** The minimum size for CIDR blocks used by PAM are:

- `/24` for pod and service CIDRs
- Default of `/22` is recommended to support future expansion

### Antivirus and antimalware considerations

Antivirus and antimalware software can impact PAM and its applications or prevent them from functioning properly.

To avoid issues, exclude the following directories from antivirus and antimalware tools that scan disk write operations:

- `/var/openebs`
- `/var/lib/kubelet`
- `/var/lib/containerd`

### Firewall modules

If you use the `puppetlabs/firewall` module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the `puppetlabs/pam_firewall` module before installing Puppet Application Manager.

If you've already installed PAM, apply the `pam_firewall` module and then restart the `kube-proxy` service to recreate its iptables rules by running the following on a primary:

```
systemctl restart kubelet
 kubectl -n kube-system delete pod -l k8s-app=kube-proxy
 kubectl -n kube-flannel delete pod -l app=flannel
```

For more information, see the PAM [firewall module](#).

## Detailed hardware requirements

For additional compute capacity, you can horizontally scale HA and standalone architectures by adding secondary nodes. During installation, only add secondaries after setting up all primaries.

You can add secondaries to HA and standalone architectures; however in standalone architectures, secondaries do not increase availability of the application, and data storage services are pinned to the host they start on and cannot be moved.

Here are the baseline requirements to run cluster services on primaries and secondaries. Any Puppet applications require additional resources on top of these requirements.

| Node type | CPU   | Memory | Storage                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Open ports                                                                                |
|-----------|-------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Primary   | 4 CPU | 7 GB   | <p>At least 50 GB on an unformatted storage device in addition to application-specific storage (below) for the Ceph storage backend. This can be satisfied by multiple devices if more storage is needed later, but should be balanced across primaries.</p> <p>1 GB for <code>/var/log/apiserver</code> for Kubernetes audit logs.</p> <p>An additional 140 GB for <code>/var/lib</code>. You can use separate filesystems if necessary, but it is not a requirement to do so. For your reference, here is how the usage is roughly divided:</p> <ul style="list-style-type: none"> <li>• 2 GB for <code>/var/lib/etcd</code></li> <li>• 10 GB for <code>/var/lib/rook</code> (plus buffer)</li> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> </ul> <p><b>Note:</b> Ceph storage backend prefers the file system inhabited by <code>/var/lib/rook</code> to remain below 70% utilization.</p> <p>SSDs (or similarly low-latency storage) are recommended for <code>/var/lib/etcd</code> and <code>/var/lib/rook</code>.</p> | <p><b>TCP:</b><br/>80, 443, 2379, 2380, 6443, 8800, and 10250</p> <p><b>UDP:</b> 8472</p> |
| Secondary | 1 CPU | 1.5 GB | <p>1 GB for <code>/var/log/apiserver</code> for Kubernetes audit logs.</p> <p>120 GB for <code>/var/lib</code>. You can use separate filesystems if necessary, but it is not a requirement to do so. For your reference, here is how the usage is roughly divided:</p> <ul style="list-style-type: none"> <li>• 32 GB for <code>/var/lib/kubelet</code></li> <li>• 80 GB for <code>/var/lib/containerd</code></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                                                                                           |

Applications are composed of multiple smaller services, so you can divide CPU and memory requirements across multiple servers. The listed ports can be accessed from all primaries and secondaries, but only need to be exposed on nodes you include in your load balancer. Apply application-specific storage to all primary nodes.

Application-specific requirements:

| Application                                    | CPU   | Memory | Storage | Ports          |
|------------------------------------------------|-------|--------|---------|----------------|
| Continuous Delivery for Puppet Enterprise (PE) | 3 CPU | 8 GB   | 50 GB   | 80, 443, 8000  |
| Puppet Comply                                  | 7 CPU | 7 GB   | 50 GB   | 80, 443, 30303 |

The minimum recommended size for a secondary node is 4 CPU and 8 GB of memory to allow some scheduling flexibility for individual services.

### Example of an HA cluster capable of running Continuous Delivery for PE and Comply

An HA cluster capable of running both Continuous Delivery for Puppet Enterprise (PE) and Puppet Comply requires 10 CPU and 15 GB of application-specific memory in addition to per-node baselines. You can create a cluster from 4 CPU, 8 GB nodes. Each primary uses all CPU and 7 GB of memory for cluster services, providing 0 CPU and 1 GB of memory for application workloads; each secondary uses 1 CPU and 1.5 GB of memory for cluster services, providing 3 CPU and 6.5 GB of memory for application workloads. Create the cluster as follows:

- Three primaries provide an excess of 3 GB of memory for application workloads. Each primary must have 150 GB of storage in an unformatted, unpartitioned storage device for Ceph and 140 GB of storage for `/var/lib`.
- Three secondaries provide an excess of 9 CPU and 19.5 GB of memory for application workloads. Each secondary must have 120 GB of storage for `/var/lib`.

This diagram illustrates the suggested three-node configuration for a cluster capable of running Continuous Delivery for Puppet Enterprise (PE) and Puppet

## Web URL and port requirements for firewalls

Puppet Application Manager interacts with external web URLs for a variety of installation, configuration, upgrade, and deployment tasks. Puppet Application Manager uses the following web URLs for internal and outbound network traffic.

| Category                                | URLs                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Puppet Application Manager and platform | <ul style="list-style-type: none"> <li>• get.replicated.com</li> <li>• registry.replicated.com</li> <li>• proxy.replicated.com</li> <li>• api.replicated.com</li> <li>• k8s.kurl.sh</li> <li>• kurl-sh.s3.amazonaws.com</li> <li>• replicated.app</li> <li>• registry-data.replicated.com</li> </ul> |
| Container registries                    | <ul style="list-style-type: none"> <li>• gcr.io</li> <li>• docker.io</li> <li>• index.docker.io</li> <li>• registry-1.docker.io</li> <li>• auth.docker.io</li> <li>• production.cloudflare.docker.com</li> <li>• quay.io</li> </ul>                                                                  |
| Puppet Enterprise                       | <ul style="list-style-type: none"> <li>• pup.pt</li> <li>• forgeapi.puppet.com</li> <li>• pm.puppetlabs.com</li> <li>• amazonaws.com</li> <li>• s3.amazonaws.com</li> <li>• rubygems.org</li> </ul>                                                                                                  |

For information about containers and firewalls, refer to the [Networking Requirements in the Kurl documentation](#).

### Firewall modules

If you use the [puppetlabs/firewall](#) module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the [puppetlabs/pam\\_firewall](#) module before installing Puppet Application Manager.

If you've already installed PAM, apply the `pam_firewall` module and then restart the `kube-proxy` service to recreate its iptables rules by running the following on a primary:

```
systemctl restart kubelet
 kubectl -n kube-system delete pod -l k8s-app=kube-proxy
 kubectl -n kube-flannel delete pod -l app=flannel
```

For more information, see the PAM [firewall module](#).

### Supported browsers

The following browsers are supported for use with the Puppet Application Manager UI:

| Browser         | Supported versions            |
|-----------------|-------------------------------|
| Google Chrome   | Current version as of release |
| Mozilla Firefox | Current version as of release |
| Microsoft Edge  | Current version as of release |
| Apple Safari    | Current version as of release |

### Component versions in PAM releases

These tables show the versions of components included in recent Puppet Application Manager (PAM) releases.

| Component         | PAM r3 | PAM r2 | PAM r1 | PAM 1.11 | PAM 1.12 | PAM 1.14 | PAM 1.16 | PAM 1.18 | PAM 1.20 | PAM 1.22 | PAM 1.24 | PAM 1.26 | PAM 1.28 | PAM 1.30 | PAM 1.32 | PAM 1.34 | PAM 1.36 | PAM 1.38 | PAM 1.40 | PAM 1.42 | PAM 1.44 | PAM 1.46 | PAM 1.48 | PAM 1.50 | PAM 1.52 | PAM 1.54 | PAM 1.56 | PAM 1.58 | PAM 1.60 |        |
|-------------------|--------|--------|--------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|----------|--------|
| Kubernetes        | 1.24.3 | 1.24.3 | 1.24.3 | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   | 1.24.3   |        |
| KOTS              | 1.11   | 1.11   | 1.11   | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     | 1.11     |        |
| kURL              | 2025.2 | 2025.2 | 2025.2 | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2   | 2025.2 |
| Weave             | N/A    | N/A    | N/A    | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      |        |
| Flannel           | 0.27   | 0.26   | 0.26   | 0.25     | 0.25     | 0.25     | 0.24     | 0.24     | 0.22     | 0.22     | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      | N/A      |        |
| Project Contour   | 1.32   | 1.30   | 1.30   | 1.30     | 1.29     | 1.28     | 1.27     | 1.27     | 1.26     | 1.25     | 1.25     | 1.25     | 1.25     | 1.25     | 1.24     | 1.24     | 1.23     | 1.22     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     | 1.21     |        |
| Registry          | 10.2   | 8.32   | 8.32   | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32     | 8.32   |
| Prometheus bundle | 0.82   | 0.81   | 0.80   | 0.79     | 0.78     | 0.77     | 0.76     | 0.75     | 0.74     | 0.73     | 0.72     | 0.71     | 0.70     | 0.69     | 0.68     | 0.67     | 0.66     | 0.65     | 0.64     | 0.63     | 0.62     | 0.61     | 0.60     | 0.59     | 0.58     | 0.57     | 0.56     | 0.55     | 0.54     |        |
| containers        | 7.29   | 7.26   | 7.25   | 6.33     | 6.31     | 6.31     | 6.28     | 6.26     | 6.24     | 6.21     | 6.21     | 6.20     | 5.11     | 4.13     | 4.13     | 4.13     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     | 4.12     |        |
| Velerl            | 16.2   | 15.2   | 15.2   | 14.0     | 13.2     | 13.2     | 12.3     | 12.2     | 12.1     | 11.1     | 11.1     | 11.0     | 10.2     | 9.5      | 9.4      | 9.1      | 9.1      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      | 9.0      |        |
| ekco              | 0.28   | 0.28   | 0.28   | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     | 0.28     |        |

|                   |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |
|-------------------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| Component         | 1.11.2 | 1.12.4 | 1.12.4 | 1.12.4 | 1.12.4 | 1.10.0 | 1.09.0 | 1.08.0 | 1.07.0 | 1.03.3 | 1.02.1 | 1.00.1 | 0.99.0 | 0.97.0 | 0.94.0 | 0.91.3 | 0.81.1 | 0.80.0 | 0.76.2 | 0.76.1 | 0.72.1 | 0.70.1 | 0.68.0 |
| Version           | r3     | r2     | r1     |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |
| Kubernetes        | 1.20.7 | 1.20.7 | 1.20.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.40.6 | 1.30.6 | 1.30.6 | 1.30.6 | 1.20.4 | 1.10.4 | 1.10.4 | 1.10.4 | 1.10.4 | 1.10.4 | 1.10.4 | 1.10.4 | 1.10.4 |
| Metrics Server    |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |
| Rook              | 1.17.7 | 1.17.1 | 1.12.8 | 1.12.8 | 1.12.8 | 1.12.8 | 1.12.8 | 1.12.8 | 1.12.6 | 1.12.3 | 1.11.8 | 1.11.5 | 1.11.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 | 1.12.5 |
| (HA only)         |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |
| MinIO             | 2025.2 | 2025.3 | 2024.8 | 2024.6 | 2024.5 | 2024.4 | 2024.3 | 2024.2 | 2024.1 | 2023.6 | 2023.4 | 2023.3 | 2023.2 | 2023.1 | 2022.9 | 2022.8 | 2022.7 | 2022.6 | 2022.5 | 2022.4 | 2022.3 | 2022.2 | 2022.1 |
| (Standalone only) |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |
| OpenEBS           | 4.2.0  | 4.2.0  | 4.1.0  | 4.0.4  | 4.0.3  | 10.0   | 10.0   | 9.0    | 8.0    | 7.0    | 6.0    | 5.0    | 3.0    | 3.0    | 3.0    | 2.0    | 3.0    | 2.0    | 2.0    | 2.0    | 2.0    | 2.0    | 2.0    |
| (Standalone only) |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |        |

### Looking up component versions

To view a list of the component versions included in your current version of PAM, run the following:

```
kubectl get installer -o jsonpath="{.items[].spec}" | jq
```

You can find more information about the current PAM version's component versions by navigating to the website appropriate to your installation type:

- **HA installations:** <https://kurl.sh/puppet-application-manager>
- **Standalone installations:** <https://kurl.sh/puppet-application-manager-standalone>

## Install PAM

You can install Puppet-supported Puppet Application Manager on a single node or in an HA configuration. Both online and offline install packages are available. You can also install it on an existing Kubernetes cluster.

Refer to the [Architecture overview](#) on page 65 for guidance on choosing which Puppet-supported Kubernetes cluster configuration is most appropriate for your needs.

**Important:** The Puppet-supported Puppet Application Manager cluster brings its own container runtime as part of the kURL installation.

- Do not install a container runtime from your operating system (OS) vendor or third-party.
- Do not install PAM on a node that has previously hosted a container runtime from your OS vendor or third-party.

Installing a different container runtime on a node, even if you installed and removed the packages before you installed PAM, causes failures that persist even after you've uninstalled the runtime.

For information on installing Puppet Application Manager on an existing Kubernetes cluster, see [Install Puppet applications using PAM on a customer-supported Kubernetes cluster](#) on page 84.

- [Install Puppet applications using PAM on a customer-supported Kubernetes cluster](#) on page 84
- Use these instructions to install Puppet Application Manager and any Puppet applications on an existing Kubernetes cluster.

- [PAM HA online installation](#) on page 86

The Puppet Application Manager (PAM) installation process creates a Kubernetes cluster for you and walks you through installing your Puppet application on the cluster.

- [PAM HA offline installation](#) on page 90

Use these instructions to install Puppet Application Manager (PAM) in an air-gapped or offline environment where the Puppet Application Manager host server does not have direct access to the internet.

- [PAM standalone online installation](#) on page 93

The Puppet Application Manager (PAM) installation process sets up the application manager (with a simple Kubernetes installation for container orchestration) for you and installs the application on the single-node cluster.

- [PAM standalone offline installation](#) on page 96

Use these instructions to install Puppet Application Manager (PAM) in an offline environment where the Puppet Application Manager host server does not have direct access to the internet.

- [Automate PAM and Puppet application online installations](#) on page 98

During a fresh online installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

- [Automate PAM and Puppet application offline installations](#) on page 100

During a fresh offline installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

- [Uninstall PAM](#) on page 104

Different uninstall procedures are required for Puppet-supported and customer-supported clusters

### Related information

[Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

## Install Puppet applications using PAM on a customer-supported Kubernetes cluster

Use these instructions to install Puppet Application Manager and any Puppet applications on an existing Kubernetes cluster.

### Before you begin

1. If you haven't already done so, [install kubectl](#).
2. Puppet Application Manager is expected to work on any certified Kubernetes distribution that meets the following requirements. We validated and support:

- Google Kubernetes Engine
- AWS Elastic Kubernetes Service

If you use a different distribution, contact [Puppet Support](#) for more information on compatibility with PAM.

3. Make sure your Kubernetes cluster meets the minimum requirements:

- Kubernetes version 1.24-1.26.
- A default storage class that can be used for relocatable storage.
- A standard Ingress controller that supports websockets (we have tested with Project Contour and NGINX).
- We currently test and support Google Kubernetes Engine (GKE) clusters.

**Note:** If you're using self-signed certificates on your Ingress controller, you must ensure that your job hardware nodes trust the certificates. Additionally, all nodes that use Continuous Delivery for PE webhooks must trust the certificates, or SSL checking must be disabled on these nodes.

**Important:** If you are installing Puppet Comply on Puppet Application Manager, the ingress controller must be configured to allow request payloads of up to 32 MB. Ingress controllers used by Amazon EKS commonly default to a 1 MB maximum — this causes all report submissions to fail.

The ingress must have a generous limit for total connection time. Setting the connection timeout to `infinity` in conjunction with an idle timeout is recommended.

- If you are setting up Puppet Application Manager behind a proxy server, the installer supports proxies configured via HTTP\_PROXY/HTTPS\_PROXY/NO\_PROXY environment variables.

**Restriction:** Using a proxy to connect to external version control systems is currently not supported.

Installation takes several (mostly hands-off) minutes to complete.

- Install the KOTS (Kubernetes off-the-shelf software) plugin on a workstation that has kubectl access to the cluster. Your kubectl configuration must have sufficient privileges to create cluster-level roles and permissions:

```
curl https://kots.io/install | bash
```

- If you are performing an offline install, ensure the required images are available in a local registry.

- Download the release assets matching the CLI version using the following command:

```
curl -LO https://github.com/replicatedhq/kots/releases/download/v
$(kubectl kots version | head -n1 | cut -d' ' -f3)/kotsadm.tar.gz
```

- Extract the images and push them into a private registry. Registry credentials provided in this step must have push access. These credentials are not stored anywhere or reused later.

```
kubectl kots admin-console push-images ./kotsadm.tar.gz
<private.registry.host>/puppet-application-manager \
--registry-username <rw-username> \
--registry-password <rw-password>
```

- Install Puppet Application Manager using images pushed in the previous step. Registry credentials provided in this step only need to have read access, and they are stored in a Kubernetes secret in the current namespace. These credentials are used to pull the images.

```
kubectl kots install puppet-application-manager \
--kotsadm-namespace puppet-application-manager \
--kotsadm-registry <private.registry.host> \
--registry-username <ro-username> \
--registry-password <ro-password>
```

**Note:** If you are setting up Puppet Application Manager behind a proxy server, add the `--copy-proxy-env` flag to this command to copy the proxy-related environment values from your environment.

- You can use similar commands to upload images from the application bundle to your registry to continue to use read-only access when pulling images. Use the same registry namespace (`puppet-application-manager`) to pull application images.

```
kubectl kots admin-console push-images ./<application-release>.airgap
<private.registry.host>/puppet-application-manager \
--registry-username <rw-username> \
--registry-password <rw-password>
```

- To perform an online install of Puppet Application Manager on your cluster, run the following commands from a workstation that has kubectl access to the cluster.

```
kubectl kots install puppet-application-manager --namespace <target
namespace>
```

This installs Puppet Application Manager on the cluster and sets up a port forward on the ClusterIP.

4. Navigate to `http://localhost:8800` and follow the prompts to be guided through the process of uploading a license for the application, configuring a local registry (for offline installs), checking to make sure your infrastructure meets system requirements, and configuring the application.

**Note:** If you are performing an offline install, download the application bundle and provide it when prompted.

**Tip:** Clusters like GKE often restrict ports to 30000-32767. The webhook for Continuous Delivery for PE defaults to port 8000. To update this port to something in the allowed range, when configuring the application, use the following steps:

- a. On the Puppet Application Manager **Dashboard** page, under **Config > Optional configuration**, select **View options for using a proxy or external load balancer**.
  - b. Enter a new value for **Webhook service port**.
5. To configure your installation further, click **Config**. On this tab, you can configure a public hostname, root user, and other settings. These are written as Kubernetes secrets in the deployment manifests. For information on how to configure your application, see the documentation for that application:
    - [Configure Continuous Delivery for PE in an online environment](#)
    - [Configure Comply in an online environment](#)
  6. To use cert-manager, in the **Customize endpoints** section, select **I have cert manager** and in the annotations section, add yours. For example:

```
kubernetes.io/ingress.class: nginx
cert-manager.io/cluster-issuer: letsencrypt-prod
```

7. When you are happy with your configuration, click **Save config** to deploy the application.

Follow the instructions for configuring and deploying your Puppet applications on Puppet Application Manager. For general information, go to [Install applications via the PAM UI](#) on page 105.

For more information on installing Continuous Delivery for PE online, see [Install Continuous Delivery for PE](#).

For more information on installing Comply online, see [Install Comply online](#).

### Related information

[Upgrade PAM on a customer-supported online cluster](#) on page 114

Upgrading Puppet Application Manager (PAM) on a customer-supported online Kubernetes cluster can be done with a single command.

[Upgrade PAM on a customer-supported offline cluster](#) on page 114

Upgrading Puppet Application Manager (PAM) on a customer-supported offline Kubernetes cluster requires a few simple kubectl commands.

## PAM HA online installation

The Puppet Application Manager (PAM) installation process creates a Kubernetes cluster for you and walks you through installing your Puppet application on the cluster.

### Before you begin

1. Review the [Puppet Application Manager system requirements](#).
2. Note that Swap is not supported for use with this version of Puppet Application Manager (PAM). The installation script attempts to disable Swap if it is enabled.

### 3. (Optional) If necessary, prepare additional steps related to SELinux and Firewalld:

The PAM installation script disables SELinux and Firewalld by default. If you want to keep SELinux enabled, append the `-s preserve-selinux-config` switch to the PAM install command. This may require additional configuration to adapt SELinux policy to the installation.

If you want to keep Firewalld enabled:

- a. Make sure Firewalld is installed on your system.
- b. To prevent the installation from disabling Firewalld, provide a patch file to the PAM install command using `-s installer-spec-file=patch.yaml`, where `patch.yaml` is the name of your patch file. For reference, here's an example patch file that enables Firewalld during installation, starts the service if it isn't running, and adds rules to open relevant ports:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 firewalldConfig:
 firewalld: enabled
 command: ["/bin/bash", "-c"]
 args: ["echo 'net.ipv4.ip_forward = 1' | tee -a /etc/sysctl.conf &&
sysctl -p"]
 firewalldCmds:
 - ["--permanent", "--zone=trusted", "--add-interface=flannel.1"]
 - ["--zone=external", "--add-masquerade"]
 # SSH port
 - ["--permanent", "--zone=public", "--add-port=22/tcp"]
 # HTTPS port
 - ["--permanent", "--zone=public", "--add-port=443/tcp"]
 # Kubernetes etcd port
 - ["--permanent", "--zone=public", "--add-port=2379-2830/tcp"]
 # Kubernetes API port
 - ["--permanent", "--zone=public", "--add-port=6443/tcp"]
 # Flannel Net port
 - ["--permanent", "--zone=public", "--add-port=8472/udp"]
 # CD4PE Webhook callback port (uncomment line below if needed)
 # - ["--permanent", "--zone=public", "--add-port=8000/tcp"]
 # KOTS UI port
 - ["--permanent", "--zone=public", "--add-port=8800/tcp"]
 # CD4PE Local registry port (offline only, uncomment line below if
needed)
 # - ["--permanent", "--zone=public", "--add-port=9001/tcp"]
 # Kubernetes component ports (kubelet, kube-scheduler, kube-
controller)
 - ["--permanent", "--zone=public", "--add-port=10250-10252/tcp"]
 # Reload firewall rules
 - ["--reload"]
 bypassFirewalldWarning: true
 disableFirewalld: false
 hardFailOnFirewalld: false
 preserveConfig: false
```

4. Ensure that IP address ranges `10.96.0.0/22` and `10.32.0.0/22` are locally accessible. See [Resolve IP address range conflicts](#) on page 126 for instructions.

**Note:** The minimum size for CIDR blocks used by Puppet Application Manager are:

- **Standalone** - /24 for pod and service CIDRs
- **HA** - /23 for pod and service CIDRs
- Default of /22 is recommended to support future expansion

5. If you are setting up Puppet Application Manager behind a proxy server, the installer supports proxies configured via `HTTP_PROXY/HTTPS_PROXY/NO_PROXY` environment variables.

**Restriction:** Using a proxy to connect to external version control systems is currently not supported.

6. Set all nodes used in your HA implementation to the UTC timezone.
7. If you use the [puppetlabs/firewall](#) module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the [puppetlabs/pam\\_firewall](#) module before installing Puppet Application Manager.

This installation process results in a Puppet Application Manager instance that is configured for high availability. Installation takes several minutes (mostly hands-off) to complete.

For more context about HA components and structure, refer to the **HA architecture** section of the [Architecture overview](#) on page 65.

1. Install and configure a load balancer (or two if you want to segment internal and external traffic - for more information, see [Architecture overview](#) on page 65). Round-robin load balancing is sufficient. For an HA cluster, the following is required:
  - A network (L4, TCP) load balancer for port 6443 across primary nodes. This is required for Kubernetes components to continue operating in the event that a node fails. The port is only accessed by the Kubernetes nodes and any admins using `kubectl`.
  - A network (L4, TCP) or application (L7, HTTP/S) load balancer for ports 80, and 443 across all primaries and secondaries. This maintains access to applications in event of a node failure. Include 8800 if you want external access to the Puppet Application Manager UI.

**Note:** Include port 8000 for webhook callbacks if you are installing Continuous Delivery for PE.

- From the command line of your first primary node, run the installation script:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager | sudo bash
```

**Tip:** If you're installing Puppet Application Manager behind a proxy server, using `sudo` might cause the installation to fail. Try running the command as root and replace `sudo bash` with `bash`.

**Note:** An unformatted, unpartitioned storage device is required.

By default this installation automatically uses devices (under `/dev`) matching the pattern `vd[b-z]`. Attach a device to each host. Only devices that match the pattern, and are unformatted, are used.

If necessary, you can override this pattern by providing a patch during installation; append `-s installer-spec-file=patch.yaml` to the installation command.

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 rook:
 blockDeviceFilter: "sd[b-z]"
```

- When prompted for a load balancer address, enter the address of the DNS entry for your load balancer.
- The installation script prints the address and password (only shown once, so make careful note of it) for Puppet Application Manager:

```

Kotsadm: http://<PUPPET APPLICATION MANAGER ADDRESS>:8800
Login with password (will not be shown again): <PASSWORD>

```

**Note:** If you lose this password or wish to change it, see [Reset the PAM password](#) on page 126 for instructions.

- When the installation script is complete, run `bash -l` to reload the shell.

**Tip:** If the installation script fails, run the following and upload the results to the Puppet Support team:

```
kubectl support-bundle https://kots.io
```

If you're installing as the root user, run the command directly:

```
/usr/local/bin/kubectl-support_bundle https://kots.io
```

- Add two additional primary nodes to the installation by following the instructions in the install script:

```
To add MASTER nodes to this installation, run the following script on your
other nodes:
curl -sSL
https://k8s.kurl.sh/puppet-application-manager-unstable/join.sh
| sudo bash -s kubernetes-master-address=...
```

If you want to keep SELinux enabled, append the `-s preserve-selinux-config` switch to the install command.

- Add the two new nodes to your load balancer.

5. Navigate to the Puppet Application Manager UI using the address provided by the installation script (`http://<PUPPET_APPLICATION_MANAGER_ADDRESS>:8800`) and follow the prompts.

The Puppet Application Manager UI is where you manage Puppet applications. You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets application system requirements.

Follow the instructions for configuring and deploying your Puppet applications on Puppet Application Manager. For more information, see [Install applications via the PAM UI](#) on page 105.

For more information on installing Continuous Delivery for PE online, see [Install Continuous Delivery for PE](#).

For more information on installing Comply online, see [Install Comply online](#).

### Related information

[Reset the PAM password](#) on page 126

As part of the installation process, Puppet Application Manager (PAM) generates a password for you. You can update this password to one of your choosing after installation.

[PAM system requirements](#) on page 69

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

[Resolve IP address range conflicts](#) on page 126

When installing Puppet Application Manager, IP address ranges `10.96.0.0/22` and `10.32.0.0/22` must not be used by other nodes on the local network.

[Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## PAM HA offline installation

Use these instructions to install Puppet Application Manager (PAM) in an air-gapped or offline environment where the Puppet Application Manager host server does not have direct access to the internet.

### Before you begin

1. Review the [Puppet Application Manager system requirements](#).
2. Note that Swap is not supported for use with this version of Puppet Application Manager (PAM). The installation script attempts to disable Swap if it is enabled.
3. (Optional) If necessary, prepare additional steps related to SELinux and FirewallD:

The PAM installation script disables SELinux and FirewallD by default. If you want to keep SELinux enabled, append the `-s preserve-selinux-config` switch to the PAM install command. This may require additional configuration to adapt SELinux policy to the installation.

If you want to keep FirewallD enabled:

- a. Make sure FirewallD is installed on your system.
- b. To prevent the installation from disabling FirewallD, provide a patch file to the PAM install command using `-s installer-spec-file=patch.yaml`, where `patch.yaml` is the name of your patch file. For reference, here's an example patch file that enables FirewallD during installation, starts the service if it isn't running, and adds rules to open relevant ports:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
```

```

metadata:
 name: patch
 spec:
 firewallldConfig:
 firewallld: enabled
 command: ["/bin/bash", "-c"]
 args: ["echo 'net.ipv4.ip_forward = 1' | tee -a /etc/sysctl.conf && sysctl -p"]
 firewallldCmds:
 - ["--permanent", "--zone=trusted", "--add-interface=flannel.1"]
 - ["--zone=external", "--add-masquerade"]
 # SSH port
 - ["--permanent", "--zone=public", "--add-port=22/tcp"]
 # HTTPS port
 - ["--permanent", "--zone=public", "--add-port=443/tcp"]
 # Kubernetes etcd port
 - ["--permanent", "--zone=public", "--add-port=2379-2830/tcp"]
 # Kubernetes API port
 - ["--permanent", "--zone=public", "--add-port=6443/tcp"]
 # Flannel Net port
 - ["--permanent", "--zone=public", "--add-port=8472/udp"]
 # CD4PE Webhook callback port (uncomment line below if needed)
 # - ["--permanent", "--zone=public", "--add-port=8000/tcp"]
 # KOTS UI port
 - ["--permanent", "--zone=public", "--add-port=8800/tcp"]
 # CD4PE Local registry port (offline only, uncomment line below if
 needed)
 # - ["--permanent", "--zone=public", "--add-port=9001/tcp"]
 # Kubernetes component ports (kubelet, kube-scheduler, kube-
 controller)
 - ["--permanent", "--zone=public", "--add-port=10250-10252/tcp"]
 # Reload firewall rules
 - ["--reload"]
 bypassFirewallldWarning: true
 disableFirewallld: false
 hardFailOnFirewallld: false
 preserveConfig: false

```

4. Ensure that IP address ranges 10.96.0.0/22 and 10.32.0.0/22 are locally accessible. See [Resolve IP address range conflicts](#) on page 126 for instructions.

**Note:** The minimum size for CIDR blocks used by Puppet Application Manager are:

- **Standalone** - /24 for pod and service CIDRs
- **HA** - /23 for pod and service CIDRs
- Default of /22 is recommended to support future expansion

5. Ensure that the nodes can resolve their own hostnames, through either local host mapping or a reachable DNS server.
6. Set all nodes used in your HA implementation to the UTC timezone.
7. If you use the [puppetlabs/firewall](#) module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the [puppetlabs/pam\\_firewall](#) module before installing Puppet Application Manager.
8. If you're restoring a backup from a previous cluster, make sure you include the `kurl-registry-ip=<YOUR_IP_ADDRESS>` installation option. For more information, see [Migrating PAM data to a new system](#) on page 118.

This installation process results in a basic Puppet Application Manager instance that is configured for optional high availability. Installation takes several minutes (mostly hands-off) to complete.

For more context about HA components and structure, refer to the **HA architecture** section of the [Architecture overview](#) on page 65.

1. Install and configure a load balancer (or two if you want to segment internal and external traffic - for more information, see [Architecture overview](#) on page 65). Round-robin load balancing is sufficient. For an HA cluster, the following is required:
  - A network (L4, TCP) load balancer for port 6443 across primary nodes. This is required for Kubernetes components to continue operating in the event that a node fails. The port is only accessed by the Kubernetes nodes and any admins using `kubectl`.
  - A network (L4, TCP) or application (L7, HTTP/S) load balancer for ports 80, and 443 across all primaries and secondaries. This maintains access to applications in event of a node failure. Include 8800 if you want external access to the Puppet Application Manager UI.

**Note:** Include port 8000 for webhook callbacks if you are installing Continuous Delivery for PE.

2. From a workstation with internet access, download the cluster installation bundle (note that this bundle is ~4GB):

```
https://k8s.kurl.sh/bundle/puppet-application-manager.tar.gz
```

3. Copy the installation bundle to your primary and secondary nodes and unpack it:

```
tar xzf puppet-application-manager.tar.gz
```

4. Run the installation command:

```
cat install.sh | sudo bash -s airgap
```

**Note:** An unformatted, unpartitioned storage device is required.

By default this installation automatically uses devices (under `/dev`) matching the pattern `vd[b-z]`. Attach a device to each host. Only devices that match the pattern, and are unformatted, are used.

If necessary, you can override this pattern by providing a patch during installation; append `-s installer-spec-file=patch.yaml` to the installation command.

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 rook:
 blockDeviceFilter: "sd[b-z]"
```

- a) When prompted for a load balancer address, enter the address of the DNS entry for your load balancer.
- b) The installation script prints the address and password (only shown once, so make careful note of it) for Puppet Application Manager:

```

Kotsadm: http://<PUPPET APPLICATION MANAGER ADDRESS>:8800
Login with password (will not be shown again): <PASSWORD>

```

**Note:** If you lose this password or wish to change it, see [Reset the PAM password](#) on page 126 for instructions.

5. Add two additional primary nodes to your offline installation using the instructions provided in the install script:

```
To add MASTER nodes to this installation, copy and unpack this bundle on
your other nodes, and run the following:
cat ./join.sh | sudo bash -s airgap
kubernetes-master-address=...
```

6. Add the two new nodes to your load balancer.
7. Navigate to the Puppet Application Manager UI using the address provided by the installation script (`http://<PUPPET_APPLICATION_MANAGER_ADDRESS>:8800`) and follow the prompts.

The Puppet Application Manager UI is where you manage Puppet applications. You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets application system requirements.

Follow the instructions for installing your Puppet applications on Puppet Application Manager. For more information, see [Install applications via the PAM UI](#) on page 105.

For more information on installing Continuous Delivery for PE offline, see [Install Continuous Delivery for PE in an offline environment](#).

For more information on installing Comply offline, see [Install Comply offline](#).

### Related information

[Reset the PAM password](#) on page 126

As part of the installation process, Puppet Application Manager (PAM) generates a password for you. You can update this password to one of your choosing after installation.

[PAM system requirements](#) on page 69

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

[Resolve IP address range conflicts](#) on page 126

When installing Puppet Application Manager, IP address ranges `10.96.0.0/22` and `10.32.0.0/22` must not be used by other nodes on the local network.

[Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## PAM standalone online installation

The Puppet Application Manager (PAM) installation process sets up the application manager (with a simple Kubernetes installation for container orchestration) for you and installs the application on the single-node cluster.

### Before you begin

1. Review the [Puppet Application Manager system requirements](#).
2. Note that Swap is not supported for use with this version of Puppet Application Manager (PAM). The installation script attempts to disable Swap if it is enabled.

### 3. (Optional) If necessary, prepare additional steps related to SELinux and Firewalld:

The PAM installation script disables SELinux and Firewalld by default. If you want to keep SELinux enabled, append the `-s preserve-selinux-config` switch to the PAM install command. This may require additional configuration to adapt SELinux policy to the installation.

If you want to keep Firewalld enabled:

- a. Make sure Firewalld is installed on your system.
- b. To prevent the installation from disabling Firewalld, provide a patch file to the PAM install command using `-s installer-spec-file=patch.yaml`, where `patch.yaml` is the name of your patch file. For reference, here's an example patch file that enables Firewalld during installation, starts the service if it isn't running, and adds rules to open relevant ports:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 firewalldConfig:
 firewalld: enabled
 command: ["/bin/bash", "-c"]
 args: ["echo 'net.ipv4.ip_forward = 1' | tee -a /etc/sysctl.conf && sysctl -p"]
 firewalldCmds:
 - ["--permanent", "--zone=trusted", "--add-interface=flannel.1"]
 - ["--zone=external", "--add-masquerade"]
 # SSH port
 - ["--permanent", "--zone=public", "--add-port=22/tcp"]
 # HTTPS port
 - ["--permanent", "--zone=public", "--add-port=443/tcp"]
 # Kubernetes etcd port
 - ["--permanent", "--zone=public", "--add-port=2379-2830/tcp"]
 # Kubernetes API port
 - ["--permanent", "--zone=public", "--add-port=6443/tcp"]
 # Flannel Net port
 - ["--permanent", "--zone=public", "--add-port=8472/udp"]
 # CD4PE Webhook callback port (uncomment line below if needed)
 # - ["--permanent", "--zone=public", "--add-port=8000/tcp"]
 # KOTS UI port
 - ["--permanent", "--zone=public", "--add-port=8800/tcp"]
 # CD4PE Local registry port (offline only, uncomment line below if
 needed)
 # - ["--permanent", "--zone=public", "--add-port=9001/tcp"]
 # Kubernetes component ports (kubelet, kube-scheduler, kube-
 controller)
 - ["--permanent", "--zone=public", "--add-port=10250-10252/tcp"]
 # Reload firewall rules
 - ["--reload"]
 bypassFirewalldWarning: true
 disableFirewalld: false
 hardFailOnFirewalld: false
 preserveConfig: false
```

4. Ensure that IP address ranges `10.96.0.0/22` and `10.32.0.0/22` are locally accessible. See [Resolve IP address range conflicts](#) on page 126 for instructions.
5. If you use the [puppetlabs/firewall](#) module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the [puppetlabs/pam\\_firewall](#) module before installing Puppet Application Manager.

This installation process results in a basic Puppet Application Manager instance. Installation takes several (mostly hands-off) minutes to complete.

1. From the command line of your node, run the installation script:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager-standalone | sudo
bash
```

**Tip:** If you're installing Puppet Application Manager behind a proxy server, using `sudo` might cause the installation to fail. Try running the command as root (use command `sudo su -`) and replace `sudo bash` with `bash`.

- a) When the installation script prints the Puppet Application Manager address and password, make a careful note of these credentials:

```

Kotsadm: http://<PUPPET APPLICATION MANAGER ADDRESS>:8800
Login with password (will not be shown again): <PASSWORD>

```

**Note:** If you lose this password or wish to change it, see [Reset the PAM password](#) on page 126 for instructions.

- b) When the installation script is complete, run `bash -l` to reload the shell.

**Tip:** If the installation script fails, run the following and upload the results to the Puppet Support team:

```
kubectl support-bundle https://kots.io
```

If you're installing as the root user, run the command directly:

```
/usr/local/bin/kubectl-support_bundle https://kots.io
```

2. Navigate to the Puppet Application Manager UI using the address provided by the installation script (`http://<PUPPET APPLICATION MANAGER ADDRESS>:8800`) and follow the prompts.

The Puppet Application Manager UI is where you manage Puppet applications. You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets application system requirements.

Follow the instructions for configuring and deploying your Puppet applications on Puppet Application Manager. For more information, see [Install applications via the PAM UI](#) on page 105.

For more information on installing Continuous Delivery for PE online, see [Install Continuous Delivery for PE](#).

For more information on installing Comply online, see [Install Comply online](#).

### Related information

[Reset the PAM password](#) on page 126

As part of the installation process, Puppet Application Manager (PAM) generates a password for you. You can update this password to one of your choosing after installation.

[PAM system requirements](#) on page 69

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

[Resolve IP address range conflicts](#) on page 126

When installing Puppet Application Manager, IP address ranges `10.96.0.0/22` and `10.32.0.0/22` must not be used by other nodes on the local network.

[Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## PAM standalone offline installation

Use these instructions to install Puppet Application Manager (PAM) in an offline environment where the Puppet Application Manager host server does not have direct access to the internet.

### Before you begin

1. Review the [Puppet Application Manager system requirements](#).
2. Note that Swap is not supported for use with this version of Puppet Application Manager (PAM). The installation script attempts to disable Swap if it is enabled.
3. (Optional) If necessary, prepare additional steps related to SELinux and Firewalld:

The PAM installation script disables SELinux and Firewalld by default. If you want to keep SELinux enabled, append the `-s preserve-selinux-config` switch to the PAM install command. This may require additional configuration to adapt SELinux policy to the installation.

If you want to keep Firewalld enabled:

- a. Make sure Firewalld is installed on your system.
- b. To prevent the installation from disabling Firewalld, provide a patch file to the PAM install command using `-s installer-spec-file=patch.yaml`, where `patch.yaml` is the name of your patch file. For reference, here's an example patch file that enables Firewalld during installation, starts the service if it isn't running, and adds rules to open relevant ports:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 firewalldConfig:
 firewalld: enabled
 command: ["/bin/bash", "-c"]
 args: ["echo 'net.ipv4.ip_forward = 1' | tee -a /etc/sysctl.conf &&
sysctl -p"]
 firewalldCmds:
 - ["--permanent", "--zone=trusted", "--add-interface=flannel.1"]
 - ["--zone=external", "--add-masquerade"]
 # SSH port
 - ["--permanent", "--zone=public", "--add-port=22/tcp"]
 # HTTPS port
 - ["--permanent", "--zone=public", "--add-port=443/tcp"]
 # Kubernetes etcd port
 - ["--permanent", "--zone=public", "--add-port=2379-2830/tcp"]
 # Kubernetes API port
 - ["--permanent", "--zone=public", "--add-port=6443/tcp"]
 # Flannel Net port
 - ["--permanent", "--zone=public", "--add-port=8472/udp"]
 # CD4PE Webhook callback port (uncomment line below if needed)
 # - ["--permanent", "--zone=public", "--add-port=8000/tcp"]
 # KOTS UI port
 - ["--permanent", "--zone=public", "--add-port=8800/tcp"]
 # CD4PE Local registry port (offline only, uncomment line below if
needed)
 # - ["--permanent", "--zone=public", "--add-port=9001/tcp"]
 # Kubernetes component ports (kubelet, kube-scheduler, kube-
controller)
```

```

- ["--permanent", "--zone=public", "--add-port=10250-10252/tcp"]
Reload firewall rules
- ["--reload"]
bypassFirewalldWarning: true
disableFirewalld: false
hardFailOnFirewalld: false
preserveConfig: false

```

4. Ensure that IP address ranges 10.96.0.0/22 and 10.32.0.0/22 are locally accessible. See [Resolve IP address range conflicts](#) on page 126 for instructions.
5. Ensure that the nodes can resolve their own hostnames, through either local host mapping or a reachable DNS server.
6. If you use the [puppetlabs/firewall](#) module to manage your cluster's firewall rules with Puppet, be advised that purging unknown rules from changes breaks Kubernetes communication. To avoid this, apply the [puppetlabs/pam\\_firewall](#) module before installing Puppet Application Manager.
7. If you're restoring a backup from a previous cluster, make sure you include the `kurl-registry-ip=<YOUR_IP_ADDRESS>` installation option. For more information, see [Migrating PAM data to a new system](#) on page 118.

This installation process results in a basic Puppet Application Manager instance. Installation takes several (mostly hands-off) minutes to complete.

1. From a workstation with internet access, download the cluster installation bundle (note that this bundle is ~4GB):

```
https://k8s.kurl.sh/bundle/puppet-application-manager-standalone.tar.gz
```

2. Copy the installation bundle to the host node and unpack it:

```
tar xzf puppet-application-manager-standalone.tar.gz
```

3. Run the installation command:

```
cat install.sh | sudo bash -s airgap
```

- a) The installation script prints the address and password (only shown once, so make careful note of it) for Puppet Application Manager:

```

Kotsadm: http://<PUPPET APPLICATION MANAGER ADDRESS>:8800
Login with password (will not be shown again): <PASSWORD>

```

**Note:** If you lose this password or wish to change it, see [Reset the PAM password](#) on page 126 for instructions.

4. Navigate to the Puppet Application Manager UI using the address provided by the installation script (`http://<PUPPET APPLICATION MANAGER ADDRESS>:8800`) and follow the prompts.

The Puppet Application Manager UI is where you manage Puppet applications. You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets application system requirements.

Follow the instructions for configuring and deploying your Puppet applications on Puppet Application Manager. For more information, see [Install applications via the PAM UI](#) on page 105.

For more information on installing Continuous Delivery for PE offline, see [Install Continuous Delivery for PE in an offline environment](#).

For more information on installing Comply offline, see [Install Comply offline](#).

### Related information

[Reset the PAM password](#) on page 126

As part of the installation process, Puppet Application Manager (PAM) generates a password for you. You can update this password to one of your choosing after installation.

[PAM system requirements](#) on page 69

You can install Puppet Application Manager (PAM) on a Puppet-supported cluster or add PAM to a customer-supported cluster. Before installing PAM, ensure that your system meets these requirements.

[Resolve IP address range conflicts](#) on page 126

When installing Puppet Application Manager, IP address ranges `10.96.0.0/22` and `10.32.0.0/22` must not be used by other nodes on the local network.

[Architecture overview](#) on page 65

Puppet Application Manager (PAM) runs on Kubernetes. We provide several supported configurations for different use cases.

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## Automate PAM and Puppet application online installations

During a fresh online installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

### Before you begin

Ensure that your system meets the [PAM system requirements](#) on page 69.

1. Install Puppet Application Manager. For detailed instructions, see [PAM HA online installation](#) on page 86.

2. Define the configuration values for your Puppet application installation, using Kubernetes YAML format.

```
apiVersion: kots.io/v1beta1
kind: ConfigValues
metadata:
 name: app-config
spec:
 values:
 accept_eula:
 value: has_accepted_eula
 annotations:
 value: "ingress.kubernetes.io/force-ssl-redirect: 'false'"
 hostname:
 value: "<HOSTNAME>"
 root_password:
 value: "<ROOT ACCOUNT PASSWORD>"
```

**Tip:** View the keyword names for all settings by clicking **View files > upstream > config.yaml** in Puppet Application Manager.

Replace the values indicated:

- Replace <HOSTNAME> with a hostname you want to use to configure an Ingress and to tell job hardware agents and web hooks how to connect to it. You might need to configure your DNS to resolve the hostname to your Kubernetes hosts.
- Replace <ROOT ACCOUNT PASSWORD> your chosen password for the application root account. The root account is used to administer your application and has full access to all resources and application-wide settings. This account must NOT be used for testing and deploying control repositories or modules.
- **Optional.** These configuration values disable HTTP-to-HTTPS redirection, so that SSL can be terminated at the load balancer. If you want to run the application over SSL only, change the `force-ssl-redirect` annotation to `true`.
- **Optional.** If your load balancer requires HTTP health checks, you can now enable Ingress settings that do not require Server Name Indication (SNI) for `/status`. To enable this setting, add the following to the config values statement:

```
enable_lb_healthcheck:
 value: "1"
```

**Note:** The automated installation automatically accepts the Puppet application end user license agreement (EULA). Unless Puppet has otherwise agreed in writing, all software is subject to the terms and conditions of the Puppet Master License Agreement located at <https://puppet.com/legal>.

3. Write your license file and the configuration values generated in step 1 to the following locations:
  - Write your license file to `./replicated_license.yaml`
  - Write your configuration values to `./replicated_config.yaml`
4. Add the Puppet application definition to Puppet Application Manager with the license file and configuration values, passing in the Puppet Application Manager password you set in step 4:

```
kubectl kots install <APPLICATION NAME> --namespace default --shared-
password <YOUR CHOSEN PASSWORD> --port-forward=false \
 --license-file ./replicated_license.yaml --config-values ./
replicated_config.yaml
```

**Note:** If you want to install a specific version of the application, include the `--app-version-label=<VERSION>` flag in the install command.

5. Wait five minutes to allow the software time to process the change.

6. Navigate to `http://<NODE_IP_ADDRESS>:8800` and log in with the Puppet Application Manager password.

Your configuration values are applied, and if preflight checks have passed, the application is deployed and in the process of starting up.

The application's status on the **Application** tab is shown as **Missing** for several minutes while deployment is underway. To monitor the deployment's progress, run `kubectl get pods --watch`.

When the deployment is complete, the application status changes to **Ready**.

7. Update your DNS or `/etc/hosts` file to include the hostname you chose during configuration.
8. Installation is now complete! Navigate to `https://<HOSTNAME>` and sign into Puppet application.

### Related information

[PAM HA online installation](#) on page 86

The Puppet Application Manager (PAM) installation process creates a Kubernetes cluster for you and walks you through installing your Puppet application on the cluster.

[Upgrade an automated online application installation](#) on page 106

If you installed a Puppet application following the automated online installation instructions, run a script to upgrade to the latest version.

## Automate PAM and Puppet application offline installations

During a fresh offline installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

### Before you begin

Ensure that your system meets the [PAM system requirements](#) on page 69.

### Related information

[PAM HA offline installation](#) on page 90

Use these instructions to install Puppet Application Manager (PAM) in an air-gapped or offline environment where the Puppet Application Manager host server does not have direct access to the internet.

[Upgrade an automated offline application installation](#) on page 107

If you installed a Puppet application following the automated offline installation instructions, run a script to upgrade to the latest version.

### Automate PAM and Puppet application offline installations on Puppet-supported clusters

1. Install Puppet Application Manager. For detailed instructions, see [PAM HA offline installation](#) on page 90.

- Define the configuration values for your Puppet application installation, using Kubernetes YAML format.

```
apiVersion: kots.io/v1beta1
kind: ConfigValues
metadata:
 name: app-config
spec:
 values:
 accept_eula:
 value: has_accepted_eula
 annotations:
 value: "ingress.kubernetes.io/force-ssl-redirect: 'false'"
 hostname:
 value: "<HOSTNAME>"
 root_password:
 value: "<ROOT ACCOUNT PASSWORD>"
```

**Tip:** View the keyword names for all settings by clicking **View files > upstream > config.yaml** in Puppet Application Manager.

Replace the values indicated:

- Replace <HOSTNAME> with a hostname you want to use to configure an Ingress and to tell job hardware agents and web hooks how to connect to it. You might need to configure your DNS to resolve the hostname to your Kubernetes hosts.
- Replace <ROOT ACCOUNT PASSWORD> your chosen password for the application root account. The root account is used to administer your application and has full access to all resources and application-wide settings. This account must NOT be used for testing and deploying control repositories or modules.
- Optional.** These configuration values disable HTTP-to-HTTPS redirection, so that SSL can be terminated at the load balancer. If you want to run the application over SSL only, change the `force-ssl-redirect` annotation to `true`.
- Optional.** If your load balancer requires HTTP health checks, you can now enable Ingress settings that do not require Server Name Indication (SNI) for `/status`. To enable this setting, add the following to the config values statement:

```
enable_lb_healthcheck:
 value: "1"
```

**Note:** The automated installation automatically accepts the Puppet application end user license agreement (EULA). Unless Puppet has otherwise agreed in writing, all software is subject to the terms and conditions of the Puppet Master License Agreement located at <https://puppet.com/legal>.

- Write your license file and the configuration values generated in the previous step to the following locations:
  - Write your license file to `./replicated_license.yaml`
  - Write your configuration values to `./replicated_config.yaml`
- Download the application bundle:

```
curl -L <APPLICATION BUNDLE URL> -o <APPLICATION BUNDLE FILE>
```

- Copy the application bundle to your primary and secondary nodes and unpack it:

```
tar xzf ./<APPLICATION BUNDLE FILE>
```

- Run the application install command on your primary node. Replace the `<YOUR CHOSEN PASSWORD>` , `<APPLICATION NAME>`, `<APPLICATION BUNDLE FILE>` values in the example below with your own values:

```
KOTS_PASSWORD=<YOUR CHOSEN PASSWORD>
kubectl kots install <APPLICATION NAME> --namespace default --shared-
password $KOTS_PASSWORD --license-file ./license.yaml --config-
values ./config.yaml --airgap-bundle ./<APPLICATION BUNDLE FILE> --port-
forward=false
wait several minutes for the application to deploy; if it doesn't show
up, preflights or another error might have occurred
```

**Note:** If you want to install a specific version of the application, include the `--app-version-label=<VERSION>` flag in the install command.

## Automate PAM and Puppet application offline installations on customer-supported clusters

### Before you begin

- If you haven't already done so, [install kubectl](#).
- Puppet Application Manager is expected to work on any certified Kubernetes distribution that meets the following requirements. We validated and support:
  - Google Kubernetes Engine
  - AWS Elastic Kubernetes Service

If you use a different distribution, contact [Puppet Support](#) for more information on compatibility with PAM.

- Make sure your Kubernetes cluster meets the minimum requirements:
  - Kubernetes version 1.24-1.26.
  - A default storage class that can be used for relocatable storage.
  - A standard Ingress controller that supports websockets (we have tested with Project Contour and NGINX).
  - We currently test and support Google Kubernetes Engine (GKE) clusters.

**Note:** If you're using self-signed certificates on your Ingress controller, you must ensure that your job hardware nodes trust the certificates. Additionally, all nodes that use Continuous Delivery for PE webhooks must trust the certificates, or SSL checking must be disabled on these nodes.

**Important:** If you are installing Puppet Comply on Puppet Application Manager, the ingress controller must be configured to allow request payloads of up to 32 MB. Ingress controllers used by Amazon EKS commonly default to a 1 MB maximum — this causes all report submissions to fail.

The ingress must have a generous limit for total connection time. Setting the connection timeout to `infinity` in conjunction with an idle timeout is recommended.

- If you are setting up Puppet Application Manager behind a proxy server, the installer supports proxies configured via `HTTP_PROXY/HTTPS_PROXY/NO_PROXY` environment variables.

**Restriction:** Using a proxy to connect to external version control systems is currently not supported.

1. Define the configuration values for your Puppet application installation, using Kubernetes YAML format.

```
apiVersion: kots.io/v1beta1
kind: ConfigValues
metadata:
 name: app-config
spec:
 values:
 accept_eula:
 value: has_accepted_eula
 annotations:
 value: "ingress.kubernetes.io/force-ssl-redirect: 'false'"
 hostname:
 value: "<HOSTNAME>"
 root_password:
 value: "<ROOT ACCOUNT PASSWORD>"
```

**Tip:** View the keyword names for all settings by clicking **View files > upstream > config.yaml** in Puppet Application Manager.

Replace the values indicated:

- Replace <HOSTNAME> with a hostname you want to use to configure an Ingress and to tell job hardware agents and web hooks how to connect to it. You might need to configure your DNS to resolve the hostname to your Kubernetes hosts.
- Replace <ROOT ACCOUNT PASSWORD> your chosen password for the application root account. The root account is used to administer your application and has full access to all resources and application-wide settings. This account must NOT be used for testing and deploying control repositories or modules.
- **Optional.** These configuration values disable HTTP-to-HTTPS redirection, so that SSL can be terminated at the load balancer. If you want to run the application over SSL only, change the `force-ssl-redirect` annotation to `true`.
- **Optional.** If your load balancer requires HTTP health checks, you can now enable Ingress settings that do not require Server Name Indication (SNI) for `/status`. To enable this setting, add the following to the config values statement:

```
enable_lb_healthcheck:
 value: "1"
```

**Note:** The automated installation automatically accepts the Puppet application end user license agreement (EULA). Unless Puppet has otherwise agreed in writing, all software is subject to the terms and conditions of the Puppet Master License Agreement located at <https://puppet.com/legal>.

2. Write your license file and the configuration values generated in the previous step to the following locations:
  - Write your license file to `./replicated_license.yaml`
  - Write your configuration values to `./replicated_config.yaml`
3. Download the application bundle:

```
curl -L <APPLICATION BUNDLE URL> -o <APPLICATION BUNDLE FILE>
```

#### 4. Create and run the following script, supplying values specific to your installation for the variables:

```
#!/bin/bash
REGISTRY=<YOUR_CONTAINER_REGISTRY>
APP_K8S_NAMESPACE=<DESIRED_NAMESPACE_IN_TARGET_CLUSTER>
APP_BUNDLE=<PATH_TO_AIRGAP_BUNDLE_FROM_STEP_3>
PAM_PASSWORD=<DESIRED_PAM_CONSOLE_PASSWORD>
LICENSE_FILE=<PATH_TO_LICENSE_FILE_FROM_STEP_1>
CONFIG_FILE=<PATH_TO_CONFIG_FILE_FROM_STEP_2>

curl https://kots.io/install | bash
curl -LO https://github.com/replicatedhq/kots/releases/download/v$(kubect
kots version | head -n1 | cut -d' ' -f3)/kotsadm.tar.gz

kubectl kots admin-console push-images ./kotsadm.tar.gz ${REGISTRY}
kubectl kots admin-console push-images ${APP_BUNDLE} ${REGISTRY}
kubectl kots install puppet-application-manager --namespace
 ${APP_K8S_NAMESPACE} --shared-password ${PAM_PASSWORD} --license-
file ${LICENSE_FILE} --config-values ${CONFIG_FILE} --airgap-bundle
 ${APP_BUNDLE} --disable-image-push --kotsadm-registry ${REGISTRY} --port-
forward=false --skip-preflights
```

**Tip:** If the script fails, it might be because:

- The `push-images` commands require that the local machine where the script is running has push access to the registry.
- The `install` command requires read access to the registry from the target cluster.
- Offline HA installs of GKE can't run preflights; therefore `--skip-preflights` must be included.

## Uninstall PAM

Different uninstall procedures are required for Puppet-supported and customer-supported clusters

At this time it's not possible to cleanly uninstall PAM from Puppet-supported clusters.

If you need to start with a fresh PAM install, you'll need to provision a new host.

### Uninstall PAM on customer-supported clusters

To uninstall Puppet Application Manager from customer-supported clusters, use:

```
kubectl delete namespace <pam-namespace>
kubectl delete clusterrolebinding kotsadm-rolebinding
kubectl delete clusterrole kotsadm-role
```

### Related information

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## Working with Puppet applications

You can install and upgrade Puppet applications using the Puppet Application Manager UI.

- [Install applications via the PAM UI](#) on page 105

The process of adding an application once you've installed Puppet Application Manager is simple.

- [Update a license for online installations](#) on page 106

If you have performed online installation of an application, you can use the Puppet Application Manager UI to update your license.

- [Update a license for offline installations](#) on page 106

If you have performed offline installation of an application, you can use the Puppet Application Manager UI to update your license.

- [Upgrade an automated online application installation](#) on page 106

If you installed a Puppet application following the automated online installation instructions, run a script to upgrade to the latest version.

- [Upgrade an automated offline application installation](#) on page 107

If you installed a Puppet application following the automated offline installation instructions, run a script to upgrade to the latest version.

## Install applications via the PAM UI

The process of adding an application once you've installed Puppet Application Manager is simple.

**Important:** Ensure you are using the following Puppet application versions if you want to add more than one Puppet application via the Puppet Application Manager UI:

| Application                               | Version        |
|-------------------------------------------|----------------|
| Continuous Delivery for Puppet Enterprise | 4.6.0 or later |
| Comply                                    | 1.0.4 or later |

For information on installing Puppet applications via the command line, see [Automate PAM and Puppet application online installations](#) on page 98 and [Automate PAM and Puppet application offline installations](#) on page 100.

To install a Puppet application using the Puppet Application Manager UI:

1. Log into the Puppet Application Manager UI, and click **Add a new application**.
  - If you have not added a Puppet application before you are prompted to upload a license.
  - If you have already added a Puppet application, click **Add a new application**.
2. Upload your `replicated_license.yaml` file when requested.

**Note:** Once the license file is installed, if offline installations are enabled, you are presented with an option to proceed with an offline setup.

Add the following information to install an offline application:

- **Hostname** - the hostname you want to use to configure an Ingress and to tell job hardware agents and web hooks how to connect to it. You might need to configure your DNS to resolve the hostname to your Kubernetes hosts.

**Important:** The hostname must be unique for each application you install.

- **Username/Password** - The username and password for the application root account. The root account is used to administer your application and has full access to all resources and application-wide settings. This account must NOT be used for testing and deploying control repositories or modules.
- **Registry namespace** - the registry namespace for the application, e.g. *CD4PE* or *Comply*.
- **Airgap bundle** - upload the relevant application bundle tarball. Click **Continue**.

3. Add any additional required information that is presented on the **Config** page. Configure any other settings on the page relevant to your installation, such as external databases, customized endpoints, a load balancer, or TLS certificates. Click **Save Config** when you are done.

Saving your new configuration settings prompts the creation of a new application version.

4. Click **Go to new version**, which redirects you to the **Version history** tab. The newly created version is shown in the **All versions** section of the page.
5. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while your system is checked to make sure your cluster meets minimum system requirements. When the preflight check is complete:
  - If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the status is **Ready to Deploy**, move on to the next step.
6. Once the version is ready to deploy, click **Deploy**. On the **Application** tab, monitor the application for readiness. The application's status is shown as **Missing** for several minutes while deployment is underway. To monitor the deployment's progress, run `kubectl get pods --watch`.  
When the deployment is complete, the application status changes to **Ready**.
  7. Navigate to `https://<HOSTNAME>` (using the hostname you entered on the **Config** screen) and sign into your application.

## Update a license for online installations

If you have performed online installation of an application, you can use the Puppet Application Manager UI to update your license.

To update the license for an online application:

1. Log in to Puppet Application Manager, click the **License** tab, and then **Sync License**.
2. On the **Version history** tab, click **Deploy**.

Puppet Application Manager adds “License Change” as the deployment cause on the **Version history** tab.

## Update a license for offline installations

If you have performed offline installation of an application, you can use the Puppet Application Manager UI to update your license.

To update the license for an offline application:

1. Ask your Puppet sales representative to email you an updated license file.
2. Log in to Puppet Application Manager, click the **License** tab.
3. Drag and drop or upload the updated license file provided by your Puppet sales representative.
4. On the **Version history** tab, click **Deploy**.

Puppet Application Manager adds “License Change” as the deployment cause on the **Version history** tab.

## Upgrade an automated online application installation

If you installed a Puppet application following the automated online installation instructions, run a script to upgrade to the latest version.

**Important:** Ensure that you are following an approved upgrade path for the application you want to upgrade. For more information, check the relevant application documentation.

1. From the command line of your primary (control plane) node, get the *application slug* for the application you want to upgrade:

```
kubectl kots --namespace <NAMESPACE> get apps
```

Replace <NAMESPACE> with the name of the namespace in which you installed PAM (usually default).

2. Run the upgrade script:

```
kubectl kots upstream upgrade <APPLICATION SLUG> --namespace <NAMESPACE>
--deploy
```

Replace <APPLICATION SLUG> with the relevant application slug for the application you want to upgrade.

Replace <NAMESPACE> with the name of the namespace in which you installed PAM (usually default).

3. Wait five minutes to allow the software time to process the change.
4. Navigate to `http://<NODE IP ADDRESS>:8800` and log in with the Puppet Application Manager password.

If preflight checks have passed, the upgraded application is deployed and in the process of starting up. To monitor the deployment's progress, run:

```
kubectl get pods --watch
```

### Related information

[Automate PAM and Puppet application offline installations](#) on page 100

During a fresh offline installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

## Upgrade an automated offline application installation

If you installed a Puppet application following the automated offline installation instructions, run a script to upgrade to the latest version.

**Important:** Ensure that you are following an approved upgrade path for the application you want to upgrade. For more information, check the relevant application documentation.

1. Download the application bundle you want to upgrade to. Copy to your primary node.
2. From the command line of your primary (control plane) node, get the *application slug* for the application you want to upgrade:

```
kubectl kots --namespace <NAMESPACE> get apps
```

Replace <NAMESPACE> with the name of the namespace in which you installed PAM (usually default).

3. Run the upgrade script:

```
kubectl kots upstream upgrade <APPLICATION SLUG> --airgap-bundle ./
<APPLICATION BUNDLE FILE> --kotsadm-namespace <REGISTRY NAMESPACE> --
namespace <NAMESPACE> --deploy
```

- Replace <APPLICATION SLUG> with the relevant application slug for the application you want to upgrade.
  - Replace <APPLICATION BUNDLE FILE> with the name of the application bundle file.
  - Replace <REGISTRY NAMESPACE> with your Registry namespace where images are uploaded.
  - Replace <NAMESPACE> with the name of the namespace in which you installed PAM (usually default).
4. Wait five minutes to allow the software time to process the change.

5. Navigate to `http://<NODE IP ADDRESS>:8800` and log in with the Puppet Application Manager password.

If preflight checks have passed, the upgraded application is deployed and in the process of starting up. To monitor the deployment's progress, run:

```
kubectl get pods --watch
```

### Related information

[Automate PAM and Puppet application offline installations](#) on page 100

During a fresh offline installation of Puppet Application Manager (PAM) and a Puppet application, you have the option to configure the software automatically rather than completing the installation script interview.

## Maintenance and tuning

---

Follow these guidelines when you're tuning or performing maintenance on a node running Puppet Application Manager (PAM).

### How to look up your Puppet Application Manager architecture

If you're running PAM on a Puppet-supported cluster, you can use the following command to determine your PAM architecture version:

```
kubectl get installer --sort-by=.metadata.creationTimestamp -o
 jsonpath='{.items[-1:].metadata.name}' ; echo
```

Depending on which architecture you used when installing, the command returns one of these values:

- **HA architecture:** `puppet-application-manager`
- **Standalone architecture:** `puppet-application-manager-standalone`
- **Legacy architecture:** Any other value, for example, `puppet-application-manager-legacy`, `cd4pe`, or `comply`

### Rebooting PAM nodes

Where possible, avoid rebooting or shutting down a PAM node. Shutting down an HA PAM node incorrectly could result in storage volume corruption and the loss of data.

For tasks such as package updates or security patches, where you must perform a reboot or shut down, follow the procedure below to gracefully shut down the node and ensure that it is drained correctly.

To reboot a node:

1. Shut down services using Ceph-backed storage:

```
/opt/ekco/shutdown.sh
```

2. If you're using a high availability (HA) cluster, drain the node:

```
kubectl drain <NODE NAME> --ignore-daemonsets --delete-local-data
```

3. Reboot the node.

### Load balancer health checks

To set up health checks for the load balancer that your Puppet Application Manager (PAM) applications are running behind, set up rules for these applications and services.

| Application/service                                                                         | URL/port                                                                   | Notes                                                                                                                                                                                         |
|---------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Puppet application. For example, Continuous Delivery for Puppet Enterprise or Puppet Comply | <code>https://&lt;CDPE HOSTNAME&gt;:443/status</code>                      | Although Puppet applications might expose other ports (Continuous Delivery for PE exposes ports 443, 80, and 8000), 443 is the HTTPS endpoint, and is the best port to use for health checks. |
| Puppet Application Manager (PAM)                                                            | <code>https://&lt;KUBERNETES PRIMARY IP&gt;:8800/healthz</code>            |                                                                                                                                                                                               |
| External load balancer endpoint                                                             | Port 6443 or <code>https://&lt;KUBERNETES PRIMARY IP&gt;:6443/livez</code> | For information on setting up a TCP probe on an external load balancer endpoint, consult the <a href="#">kURL load balancer documentation</a> .                                               |
| Local container registry (for offline installations)                                        | <code>https://&lt;KUBERNETES PRIMARY IP&gt;:9001</code>                    |                                                                                                                                                                                               |

## Load balancing

The following load balancer requirements are needed for a HA install:

- A network (L4, TCP) load balancer for port 6443 across primary nodes. This is required for Kubernetes components to continue operating in the event that a node fails. The port is only accessed by the Kubernetes nodes and any admins using `kubectl`.
- A network (L4, TCP) or application (L7, HTTP/S) load balancer for ports 80, and 443 across all primaries and secondaries. This maintains access to applications in event of a node failure. Include 8800 if you want external access to the Puppet Application Manager UI.

**Note:** Include port 8000 for webhook callbacks if you are installing Continuous Delivery for PE.

**Important:** If you are using application load balancing, be aware that Ingress items use Server Name Indication (SNI) to route requests, which may require additional configuration with your load balancer. If your load balancer does not support SNI for health checks, enable **Enable load balancer HTTP health check** in the Puppet Application Manager UI **Config** page .

## Upgrading PAM on a Puppet-supported cluster

Upgrade Puppet Application Manager (PAM) on a Puppet-supported cluster to take advantage of new features and bug fixes, and to upgrade your cluster to the latest version of Kubernetes when one is available.

There are four possible upgrade types for Puppet Application Manager installations:

- **Online** - For standalone or HA installations with a connection to the internet.
- **Offline** - For air-gapped standalone or HA installations without a connection to the internet.
- **Online legacy** - For standalone or HA installations created prior to April 2021 with a connection to the internet.
- **Offline legacy** - For air-gapped standalone or HA installations created prior to April 2021 without a connection to the internet.

**Restriction:** You cannot use the upgrade process to move from a legacy deployment to a non-legacy deployment, or from standalone to HA, or vice versa. If you wish to change architecture types, see [Migrating PAM data to a new system](#) on page 118.

## How to look up your Puppet Application Manager architecture

If you're running PAM on a Puppet-supported cluster, you can use the following command to determine your PAM architecture version:

```
kubectl get installer --sort-by=.metadata.creationTimestamp -o
 jsonpath='{.items[-1:].metadata.name}' ; echo
```

Depending on which architecture you used when installing, the command returns one of these values:

- **HA architecture:** puppet-application-manager
- **Standalone architecture:** puppet-application-manager-standalone
- **Legacy architecture:** Any other value, for example, puppet-application-manager-legacy, cd4pe, or comply

## Upgrade PAM online

Upgrade Puppet Application Manager (PAM) to take advantage of new features and bug fixes, and to upgrade your cluster to the latest version of Kubernetes when one is available.

### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

If you are upgrading from a version of PAM that used Weave (versions 1.100.3 and earlier) to a version of PAM that uses Flannel (versions 1.102.2 and later), pod-to-pod networking now depends on UDP port 8472 being open instead of ports 6783 and 6784.

**Note:** Starting with Puppet Application Manager 1.97.0, the `force-reapply-addons` flag is deprecated and generates a warning on use. If you are upgrading to a version prior to 1.97.0, you need to add the `force-reapply-addons` flag to the `bash` command using the `-s` flag.

1. On your first primary node, rerun the installation script, passing in any arguments you included when installing for the first time:

For standalone deployments, use:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager-standalone | sudo
 bash
```

For HA deployments, use:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager | sudo bash
```

2. If a new version of Kubernetes is available, the installer notes upgrade scripts to run on other nodes in an HA cluster.

The installer also pauses before draining nodes as part of the Kubernetes upgrade. The node draining process can take several minutes to complete, during which time application workloads are stopped or migrated to other systems. This migration may cause several minutes of downtime while databases are rescheduled.

### Related information

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## Upgrade PAM offline

Users operating in environments without direct access to the internet must use the links below to upgrade to the latest version of Puppet Application Manager (PAM).

### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

If you are upgrading from a version of PAM that used Weave (versions 1.100.3 and earlier) to a version of PAM that uses Flannel (versions 1.102.2 and later), pod-to-pod networking now depends on UDP port 8472 being open instead of ports 6783 and 6784.

For offline installs, you can upgrade ONE Kubernetes minor version at a time automatically. If you need to upgrade two minor versions of Kubernetes, the installer stops and prompts you to download the intermediate version. For more information about installing Kubernetes in airgapped environments, please see [About kURL Cluster Updates | Replicated Docs](#).

**Note:** Starting with Puppet Application Manager 1.97.0, the `force-reapply-addons` flag is deprecated and generates a warning on use. If you are upgrading to a version prior to 1.97.0, you need to add the `force-reapply-addons` flag in **Step 3** to the bash command after `-s airgap`.

To upgrade Puppet Application Manager:

1. From a workstation with internet access, download the latest version of the installation bundle that is relevant for your installation type:

For standalone installations, enter the following command (note that this bundle is ~4GB):

```
curl -LO https://k8s.kurl.sh/bundle/puppet-application-manager-standalone.tar.gz
```

For HA installations, enter the following command (note that this bundle is ~4GB):

```
curl -LO https://k8s.kurl.sh/bundle/puppet-application-manager.tar.gz
```

2. Copy the installation bundle to your primary and secondary nodes and unpack it:

For standalone installations, use:

```
tar xzf puppet-application-manager-standalone.tar.gz
```

For HA installations, use:

```
tar xzf puppet-application-manager.tar.gz
```

3. Manually load the images from the installation bundle:

```
cat tasks.sh | bash -s load-images
```

4. On your primary node, rerun the installation script, passing in any arguments you included when installing for the first time:

```
cat install.sh | sudo bash -s airgap
```

**Note:** This script issues a prompt to run the `task.sh` and `upgrade.sh` scripts on your secondary nodes. Use the versions of these scripts from the downloaded bundle in step 2.

**Note:** In cases where Kubernetes needs to be upgraded more than one minor version, the installer stops with a message about where to download the intermediate version for Kubernetes.

5. If a new version of Kubernetes is available, the installer systems provide upgrade scripts to run on other nodes in an HA cluster. The installer also pauses before draining nodes as part of the Kubernetes upgrade. Node draining is performed as part of a Kubernetes upgrade.

The node draining process can take several minutes to complete, during which time application workloads are stopped or migrated to other systems. This migration may cause several minutes of downtime while databases are rescheduled.

When the deployment is complete, sign into Puppet Application Manager- `http://<PUPPET_APPLICATION_MANAGER_ADDRESS>:8800` - and verify that the new version number is displayed in the bottom left corner of the web UI.

## PAM legacy upgrades

The legacy architecture is no longer supported. However, if you have not yet migrated to a supported architecture, you can use this method to upgrade Puppet Application Manager (PAM).

### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

**Legacy architecture is no longer supported:** The legacy architecture utilizes Rook 1.0, which is incompatible with Kubernetes version 1.20 and newer versions. Kubernetes version 1.19 is no longer receiving security updates. The legacy architecture reached the end of its support lifecycle on **30 June 2022**, and Puppet no longer updates legacy architecture components. For information on migrating data from a legacy architecture to a standalone or HA architecture, go to our Support Knowledge Base instructions:

- [Migrate to a supported PAM architecture for Continuous Delivery for PE](#)
- [Migrate to a supported PAM architecture for Comply](#)

**Restriction:** It is not possible to upgrade from an online legacy install to a new offline install configuration. Similarly, upgrades from an offline legacy configuration to a new online install are not supported.

To upgrade a legacy version of Puppet Application Manager on nodes with internet access:

1. On your node (or control plane node if you have a HA deployment), rerun the installation script, passing in any arguments you included when installing for the first time:

- For standalone installs:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager-legacy | sudo bash -s force-reapply-addons
```

- For HA installs:

```
curl -sSL https://k8s.kurl.sh/puppet-application-manager-legacy | sudo bash -s ha force-reapply-addons
```

- If a new version of Kubernetes is available, the systems provide upgrade scripts to run on each node in your cluster.

Node draining is performed as part of a Kubernetes upgrade. The node draining process can take several minutes to complete.

**Note:** During the Kubernetes upgrade process, nodes are not able to properly route network connections. If you have a HA deployment, make sure you have load balancers or a multi-node fail-over process in place, or schedule downtime before upgrading.

## PAM offline legacy upgrades

The legacy architecture is no longer supported. However, if you have not yet migrated to a supported architecture, you can use this method to upgrade Puppet Application Manager (PAM) on offline nodes.

### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

**Legacy architecture is no longer supported:** The legacy architecture utilizes Rook 1.0, which is incompatible with Kubernetes version 1.20 and newer versions. Kubernetes version 1.19 is no longer receiving security updates. The legacy architecture reached the end of its support lifecycle on **30 June 2022**, and Puppet no longer updates legacy architecture components. For information on migrating data from a legacy architecture to a standalone or HA architecture, go to our Support Knowledge Base instructions:

- [Migrate to a supported PAM architecture for Continuous Delivery for PE](#)
- [Migrate to a supported PAM architecture for Comply](#)

**Restriction:** It is not possible to upgrade from an online legacy install to a new offline install configuration. Similarly, upgrades from an offline legacy configuration to a new online install are not supported.

To upgrade Puppet Application Manager on nodes without a connection to the internet:

- From a workstation with internet access, download the latest version of the cluster installation bundle (note that this bundle is ~4GB):

```
https://k8s.kurl.sh/bundle/puppet-application-manager-legacy.tar.gz
```

- Copy the installation bundle to your primary and secondary Puppet Application Manager nodes and unpack it:

```
tar xzf puppet-application-manager-legacy.tar.gz
```

- Rerun the installation script. Don't forget to pass in any additional arguments you included when installing for the first time you installed the product:

For standalone installs use:

```
cat install.sh | sudo bash -s airgap force-reapply-addons
```

For HA installs use:

```
cat install.sh | sudo bash -s airgap ha force-reapply-addons
```

**Note:** During the upgrade process, follow any prompts to run commands on your other cluster nodes.

When the deployment is complete, sign into Puppet Application Manager and verify that the new version number is displayed in the bottom center of the web UI.

## Upgrading PAM on a customer-supported cluster

Upgrade Puppet Application Manager (PAM) on your own Kubernetes cluster to take advantage of new features and bug fixes.

There are two possible upgrade types for customer-supported Puppet Application Manager deployments:

- **Online** - For installations with a connection to the internet.
- **Offline** - For air-gapped installations without a connection to the internet.

### Upgrade PAM on a customer-supported online cluster

Upgrading Puppet Application Manager (PAM) on a customer-supported online Kubernetes cluster can be done with a single command.

#### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

To upgrade Puppet Application Manager on a customer-supported online cluster:

1. Upgrade kubectl KOTS:

```
curl https://kots.io/install | bash
```

2. Issue the following KOTS command:

```
kubectl kots admin-console upgrade --namespace <target namespace>
```

**Tip:** Run the `kubectl kots admin-console upgrade -h` command for more usage information.

### Upgrade PAM on a customer-supported offline cluster

Upgrading Puppet Application Manager (PAM) on a customer-supported offline Kubernetes cluster requires a few simple kubectl commands.

#### Before you begin

Make sure you have captured an up-to-date snapshot of your PAM installation, which you can use to fall back the current version if there is an issue with the upgrade process. Learn more about snapshots at [Backing up PAM using snapshots](#) on page 115.

To upgrade Puppet Application Manager on a customer-supported offline cluster, perform the following steps from a workstation that has kubectl access to the cluster:

1. Upgrade kubectl KOTS:

```
curl https://kots.io/install | bash
```

2. Ensure the required images are available in your local registry. Download the release assets matching the CLI version using the following command:

```
curl -LO https://github.com/replicatedhq/kots/releases/download/v$(kubectl kots version | head -n1 | cut -d' ' -f3)/kotsadm.tar.gz
```

3. Extract the images and push them to your private registry. Registry credentials provided in this step must have push access. These credentials are not stored anywhere or reused later.

```
kubectl kots admin-console push-images ./kotsadm.tar.gz
 <private.registry.host>/puppet-application-manager \
 --registry-username <rw-username> \
 --registry-password <rw-password>
```

4. After you push the images to your private registry, execute the upgrade command with registry read-only credentials:

```
kubectl kots upgrade puppet-application-manager \
 --kotsadm-namespace puppet-application-manager \
 --kotsadm-registry <private.registry.host> \
 --registry-username <ro-username> \
 --registry-password <ro-password> \
 --namespace <target namespace>
```

## Backing up PAM using snapshots

Snapshots are point-in-time backups of your Puppet Application Manager (PAM) deployment, which can be used to roll back to a previous state or restore your installation into a new cluster for disaster recovery.

### Related information

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## Full and partial snapshots

There are two options available when you're creating a snapshot for your Puppet Application Manager (PAM) deployment, full snapshots (also known as instance snapshots) and partial (or application) snapshots. For full disaster recovery, make sure you've configured and scheduled regular full snapshots stored on a remote storage solution such as an S3 bucket or NFS share.

Full snapshots offer a comprehensive backup of your PAM deployment, because they include the core PAM application together with the Puppet applications you've installed in your PAM deployment. You can use a full snapshot to restore your PAM deployment and all of your installed Puppet applications to a previous backup. For example, you could use a full snapshot to revert an undesired configuration change or a failed upgrade, or to migrate your PAM deployment to another Puppet-supported cluster.

Partial snapshots are available from the PAM console, but are limited in their usefulness. To restore from a partial snapshot, you must already have an installed and functioning version of PAM. A functioning PAM installation is needed because the option to restore a partial snapshot can only be accessed from the **Snapshots** section of the PAM admin console.

Partial snapshots only back up the Puppet application you specified when you configured the snapshot, for example, Continuous Delivery for Puppet Enterprise, or Puppet Comply. They do not back up the underlying PAM deployment. Partial snapshots are sometimes useful if you want to roll back to a previous version of a specific Puppet application that you've installed on your PAM deployment, but are far less versatile than full snapshots. To make sure that you have all disaster recovery options available to you, use a full snapshot wherever possible.

## Configure snapshots

Before using snapshots, select a storage location, set a snapshot retention period, and indicate whether snapshots are created manually or on a set schedule.

**Important:** Disaster recovery requires that the store backend used for backups is accessible from the new cluster. When setting up snapshots in an *offline* cluster, make sure to record the registry service IP address with the following command:

```
kubectl -n kurl get svc registry -o jsonpath='{.spec.clusterIP}'
```

Be sure to record the value returned by this command as it is required when creating a new cluster to restore to as part of [Disaster recovery with PAM](#) on page 124.

1. In the upper navigation bar of the Puppet Application Manager UI, click **Snapshots > Settings & Schedule**.
2. The snapshots feature uses <https://velero.io>, an open source backup and restore tool. Click **Check for Velero** to determine whether Velero is present on your cluster, and to install it if needed.

3. Select a destination for your snapshot storage and provide the required configuration information. You can choose to set up snapshot storage in the PAM UI or on the command line. Supported destinations are listed below. We recommend using an external service or NFS, depending on what is available to you:

- Internal storage (default)
- Amazon S3
- Azure Blob Storage
- Google Cloud Storage
- Other S3-compatible storage
- Network file system (NFS)
- Host path

### Amazon S3 storage

If using the PAM UI, provide the following information:

| Field                  | Description                                                                                                                      |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Bucket                 | The name of the AWS bucket where snapshots are stored.                                                                           |
| Region                 | The AWS region the bucket is available in.                                                                                       |
| Path                   | <b>Optional.</b> The path within the bucket where all snapshots are stored.                                                      |
| Use IAM instance role? | If selected, an IAM instance role is used instead of an access key ID and secret.                                                |
| Access key ID          | <b>Required only if not using an IAM instance role.</b> The AWS IAM access key ID that can read from and write to the bucket.    |
| Access key secret      | <b>Required only if not using an IAM instance role.</b> The AWS IAM secret access key that is associated with the access key ID. |

If using the command line, run the appropriate command:

**Not** using an IAM instance role:

```
kubectl kots velero configure-aws-s3 access-key --access-key-id <string>
--bucket <string> --path <string> --region <string> --secret-access-key
<string>
```

Using an IAM instance role:

```
kubectl kots velero configure-aws-s3 instance-role --bucket <string> --
path <string> --region <string>
```

### Azure Blob Storage

If using the PAM UI, provide the following information:

| Field                                                                         | Description                                                                                                                           |
|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Note:</b> Only connections via service principals are currently supported. |                                                                                                                                       |
| Bucket                                                                        | The name of the Azure Blob Storage container where snapshots are stored.                                                              |
| Path                                                                          | <b>Optional.</b> The path within the container where all snapshots are stored.                                                        |
| Subscription ID                                                               | <b>Required only for access via service principal or AAD Pod Identity.</b> The subscription ID associated with the target container.  |
| Tenant ID                                                                     | <b>Required only for access via service principal .</b> The tenant ID associated with the Azure account of the target container.      |
| Client ID                                                                     | <b>Required only for access via service principal .</b> The client ID of a Service Principle with access to the target container.     |
| Client secret                                                                 | <b>Required only for access via service principal .</b> The Client Secret of a Service Principle with access to the target container. |

4. Click **Update storage settings** to save your storage destination information.

Depending on your chosen storage provider, saving and configuring your storage provider might take several minutes.

5. Optional: To automatically create new snapshots on a schedule, select **Enable automatic scheduled snapshots** on the **Full snapshots (instance)** tab. (If desired, you can also set up a schedule for capturing partial (application-only) snapshots.)

You can schedule a new snapshot creation for every hour, day, or week, or you can create a custom schedule by entering a cron expression.

6. Finally, set the retention schedule for your snapshots by selecting the time period after which old snapshots are automatically deleted. The default retention period is one month.

**Note:** A snapshot's retention period cannot be changed once the snapshot is created. If you update the retention schedule, the new retention period applies only to snapshots created after the update is made.

7. Click **Update schedule** to save your changes.

Snapshots are automatically created according to your specified schedule and saved to the storage location you selected. You can also create an unscheduled snapshot at any time by clicking **Start a snapshot** on the **Dashboard** or on the **Snapshots** page.

## Roll back changes using a snapshot

When necessary, you can use a snapshot to roll back to a previous version of your Puppet Application Manager set-up without changing the underlying cluster infrastructure.

To roll back changes:

1. In console menu of the Puppet Application Manager UI, click **Snapshots > Full Snapshots (Instance)**.
2. Select the snapshot you wish to roll back to from the list of available snapshots and click **Restore from this**

**backup** 

3. Follow the instructions to complete either a partial restore or a full restore.

A full restore is useful if you need to stay on an earlier version of an application and want to disable automatic version updates. Otherwise, a partial restore is the quicker option.

## Migrating PAM data to a new system

By using a snapshot, you can migrate your data to a new Puppet Application Manager (PAM) instance.

### Data migration prerequisites

In order to perform a data migration, your system must be configured as follows:

- On the original system, Puppet Application Manager (PAM) must be configured to support **Full Snapshots (Instance)**. For instructions on configuring snapshots, see [Backing up PAM using snapshots](#) on page 115.
- Velero must be configured to use an external snapshot destination accessible to both the old and new clusters, such as S3 or NFS.
- Both the old and new clusters must have the same connection status (online or offline). Migrating from offline to online clusters or vice versa is not supported.
- For offline installs, both the old and new clusters must use the same version of PAM.
- Upgrade to the latest version of PAM on both the old and new clusters before you begin.

### Migrating data between two systems with the same architecture

To perform data migration between two systems using the same architecture (from standalone to standalone, or from HA to HA), you must create a new cluster to migrate to, then follow the process outlined below to recover your instance from a snapshot.

## Before you begin

Review the requirements in [Data migration prerequisites](#) on page 118.

**Important:** If you are migrating from a legacy architecture, go to our Support Knowledge Base instructions for migrating to a supported architecture for your Puppet application:

- [Migrate to a supported PAM architecture for Continuous Delivery for PE](#)
- [Migrate to a supported PAM architecture for Comply](#)

1. On the original system, find the version of kURL your deployment is using by running the following command. Save the version for use in step 3.

```
kubectl get configmap -n kurl kurl-current-config -o
 jsonpath="{.data.kurl-version}" && echo
```

2. Get the installer spec section by running the command appropriate to your PAM installation type:

**Tip:** See [How to determine your version of Puppet Application Manager](#) if you're not sure which installation type you're running.

- HA installation: `kubectl get installers puppet-application-manager -o yaml`
- Standalone installation: `kubectl get installers puppet-application-manager-standalone -o yaml`
- Legacy installation: `kubectl get installers puppet-application-manager-legacy -o yaml`

The command's output looks similar to the following. The spec section is shown in bold in the example below. Save your spec section for use in step 3.

```
kubectl get installers puppet-application-manager-standalone -o yaml
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 annotations:
 kubectl.kubernetes.io/last-applied-configuration: |

 {"apiVersion":"cluster.kurl.sh/v1beta1","kind":"Installer","metadata":
{"annotations":{},"creationTimestamp":null,"name":"puppet-application-
manager-standalone","namespace":"default"},"spec":{"containerd":
{"version":"1.4.12"},"contour":{"version":"1.18.0"},"ekco":
{"version":"0.16.0"},"kotsadm":{"applicationSlug":"puppet-
application-manager","version":"1.64.0"},"kubernetes":
{"version":"1.21.8"},"metricsServer":{"version":"0.4.1"},"minio":
{"version":"2020-01-25T02-50-51Z"},"openebs":
{"isLocalPVEnabled":true,"localPVStorageClassName":"default","version":"2.6.0"},"prom
{"version":"0.49.0-17.1.1"},"registry":{"version":"2.7.1"},"velero":
{"version":"1.6.2"},"weave":
{"podCidrRange":"/22","version":"2.8.1"}}, "status":{}}
 creationTimestamp: "2021-06-04T00:05:08Z"
 generation: 4
 labels:
 velero.io/exclude-from-backup: "true"
 name: puppet-application-manager-standalone
 namespace: default
 resourceVersion: "102061068"
 uid: 4e7f1196-5fab-4072-9399-15d18dcc5137
spec:
 containerd:
 version: 1.4.12
 contour:
 version: 1.18.0
 ekco:
 version: 0.16.0
 kotsadm:
 applicationSlug: puppet-application-manager
 version: 1.64.0
 kubernetes:
 version: 1.21.8
 metricsServer:
 version: 0.4.1
 minio:
 version: 2020-01-25T02-50-51Z
 openebs:
 isLocalPVEnabled: true
 localPVStorageClassName: default
 version: 2.6.0
 prometheus:
 version: 0.49.0-17.1.1
 registry:
 version: 2.7.1
 velero:
 version: 1.6.2 © 2024 Puppet, Inc., a Perforce company
 weave:
```

- On a new machine, create a file named `installer.yaml` with the following contents, replacing `<SPEC>` and `<KURL VERSION>` with the information you gathered in the previous steps.

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 <SPEC>
 kurl:
 installerVersion: "<KURL VERSION>"
```

**Important:** If you are running PAM version 1.68.0 or newer, the kURL installer version might be included in the `spec` section. If this is the case, omit the `kurl:` section from the bottom of the `installer.yaml` file. There must be only one `kurl:` entry in the file.

**Tip:** Spacing is critical in YAML files. Use a YAML file linter to confirm that the format of your file is correct.

Here is an example of the contents of the `installer.yaml` file:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
spec:
 containerd:
 version: 1.4.12
 contour:
 version: 1.18.0
 ekco:
 version: 0.16.0
 kotsadm:
 applicationSlug: puppet-application-manager
 version: 1.64.0
 kubernetes:
 version: 1.21.8
 metricsServer:
 version: 0.4.1
 minio:
 version: 2020-01-25T02-50-51Z
 openebs:
 isLocalPVEnabled: true
 localPVStorageClassName: default
 version: 2.6.0
 prometheus:
 version: 0.49.0-17.1.1
 registry:
 version: 2.7.1
 velero:
 version: 1.6.2
 weave:
 podCidrRange: /22
 version: 2.8.1
 kurl:
 installerVersion: "v2022.03.11-0"
```

- Build an installer using the `installer.yaml` file. Run the following command:

```
curl -s -X POST -H "Content-Type: text/yaml" --data-binary
"@installer.yaml" https://kurl.sh/installer |grep -o "[^/]*$"
```

The output is a hash. Carefully save the hash for use in step 5.

5. Install a new cluster. To do so, you can either:

- a. Point your browser to `https://kurl.sh/<HASH>` (replacing `<HASH>` with the hash you generated in step 4) to see customized installation scripts and information.
- b. Follow the appropriate PAM documentation.
  - **For online installations:** Follow the steps in [PAM HA online installation](#) on page 86 or [PAM standalone online installation](#) on page 93, replacing the installation script with the following:

```
curl https://kurl.sh/<HASH> | sudo bash
```

- **For offline installations:** Follow the steps in [PAM HA offline installation](#) on page 90 or [PAM standalone offline installation](#) on page 96, replacing the installation script with the following:

```
curl -LO https://k8s.kurl.sh/bundle/<HASH>.tar.gz
```

When setting up a new *offline* cluster as part of disaster recovery, add `kurl-registry-ip=<IP>` to the install options, replacing `<IP>` with the value you recorded when setting up snapshots.

**Note:** If you do not include the `kurl-registry-ip=<IP>` flag, the registry service will be assigned a new IP address that does not match the IP on the machine where the snapshot was created. You must align the registry service IP address on the new offline cluster to ensure that the restored configuration can pull images from the correct location.

**Important:** Do not install any Puppet applications after the PAM installation is complete. You'll recover your Puppet applications later in the process.

6. To recover using a snapshot saved to a **host path**, ensure user/group 1001 has full access on all nodes by running:

```
chown -R 1001:1001 /<PATH/TO/HOSTPATH>
```

7. Configure the new cluster to connect to your snapshot storage location. Run the following to see the arguments needed to complete this task:

```
kubectl kots -n default velero configure-{hostpath,nfs,aws-s3,other-s3,gcp} --help
```

8. Run `kubectl kots get backup` and wait for the list of snapshots to become available. This might take several minutes.
9. Start the restoration process by running `kubectl kots restore --from-backup <BACKUP NAME>`. The restoration process takes several minutes to complete. When the PAM UI is available, use it to monitor the status of the application.

**Note:** When restoring, wait for all restores to complete before making any changes. The following command waits for pods to finish restoring data from backup. Other pods may not be ready until updated configuration is deployed in the next step:

```
kubectl get pod -o json | jq -r '.items[] | select(.metadata.annotations."backup.velero.io/backup-volumes" | .metadata.name' | xargs kubectl wait --for=condition=Ready pod --timeout=20m
```

This command requires the [jq](#) CLI tool to be installed. It is available in most Linux OS repositories.

10. After the restoration process completes, save your config and deploy:

- a) From the PAM UI, click **Config**.
- b) (Optional) If the new cluster's hostname is different from the old one, update the **Hostname**.
- c) Click **Save Config**.
- d) Deploy the application. You must save your config and deploy even if you haven't made any changes.

**Note:** If you have installed Continuous Delivery for PE and changed the hostname, you need to update the webhooks that connect Continuous Delivery for PE with your source control provider. For information on how to do this, see [Update webhooks](#).

## Migrating data between two systems with different architectures

To perform data migration between two systems using different PAM architectures (from standalone to HA, or from HA to standalone), you must create a new cluster to recover to, then follow the process outlined below to recover your instance from a snapshot.

### Before you begin

Review the requirements in [Data migration prerequisites](#) on page 118.

**Important:** If you are migrating from a legacy architecture, go to our Support Knowledge Base instructions for migrating to a supported architecture for your Puppet application:

- [Migrate to a supported PAM architecture for Continuous Delivery for PE](#)
- [Migrate to a supported PAM architecture for Comply](#)

1. On the original system, find the version of kURL your deployment is using by running the following command. Save the version for use in step 2.

```
kubectl get configmap -n kurl kurl-current-config -o
jsonpath="{.data.kurl-version}" && echo
```

2. Set up a new cluster to house the recovered instance, following the system requirements for your applications.

**Important:** Do not install any Puppet applications after the PAM installation is complete. You'll recover your Puppet applications later in the process.

- Install PAM using the version of kURL you retrieved earlier:

- For online installs:

```
curl -sSL https://k8s.kurl.sh/version/<VERSION STRING>/puppet-
application-manager | sudo bash <--s options>
```

- For offline installs:

```
curl -O https://k8s.kurl.sh/bundle/version/<VERSION STRING>/puppet-
application-manager.tar.gz
```

- When setting up a new *offline* cluster as part of disaster recovery, add `kurl-registry-ip=<IP>` to the install options, replacing `<IP>` with the value you recorded when setting up snapshots.

**Note:** If you do not include the `kurl-registry-ip=<IP>` flag, the registry service will be assigned a new IP address that does not match the IP on the machine where the snapshot was created. You must align the registry service IP address on the new offline cluster to ensure that the restored configuration can pull images from the correct location.

- To recover using a snapshot saved to a **host path**, ensure user/group 1001 has full access on all nodes by running:

```
chown -R 1001:1001 /<PATH/TO/HOSTPATH>
```

- Configure the new cluster to connect to your snapshot storage location. Run the following to see the arguments needed to complete this task:

```
kubectl kots -n default velero configure-{hostpath,nfs,aws-s3,other-s3,gcp} --help
```

- Run `kubectl kots get backup` and wait for the list of snapshots to become available. This might take several minutes.
- Start the restoration process by running `kubectl kots restore --from-backup <BACKUP NAME>`. The restoration process takes several minutes to complete. When the PAM UI is available, use it to monitor the status of the application.

**Note:** When restoring, wait for all restores to complete before making any changes. The following command waits for pods to finish restoring data from backup. Other pods may not be ready until updated configuration is deployed in the next step:

```
kubectl get pod -o json | jq -r '.items[] |
 select(.metadata.annotations."backup.velero.io/backup-volumes")
 | .metadata.name' | xargs kubectl wait --for=condition=Ready pod --
 timeout=20m
```

This command requires the [jq](#) CLI tool to be installed. It is available in most Linux OS repositories.

- After the restoration process completes, save your config and deploy:
  - From the PAM UI, click **Config**.
  - (Optional) If the new cluster's hostname is different from the old one, update the **Hostname**.
  - Click **Save Config**.
  - Deploy the application. You must save your config and deploy even if you haven't made any changes.

**Note:** If you have installed Continuous Delivery for PE and changed the hostname, you need to update the webhooks that connect Continuous Delivery for PE with your source control provider. For information on how to do this, see [Update webhooks](#).

## Disaster recovery with PAM

It is important to prepare your system and regularly capture full snapshots. This backs up your data and makes it easier to restore your system if disaster recovery is needed.

### Prepare your system to support future disaster recovery

To make sure your system is equipped to help you recover from a potential system failure, you must:

- Configure Puppet Application Manager (PAM) to support **Full Snapshots (Instance)**. For instructions on configuring snapshots, see [Backing up PAM using snapshots](#) on page 115.
- Configure Velero to use an external snapshot destination that is accessible to both your current cluster and future new clusters, such as S3 or NFS.

- Disaster recovery requires that the store backend used for backups is accessible from the new cluster. When setting up snapshots in an *offline* cluster, use the following command to record the registry service IP address:

```
kubectl -n kurl get svc registry -o jsonpath='{.spec.clusterIP}'
```

Make a record of the value returned by this command, because you'll need it to create a new cluster to restore to as part of disaster recovery.

- Run the latest version of PAM. Disaster recovery is only available on systems running PAM version 1.44.1 or newer.

## Disaster recovery process

To perform a disaster recovery, you must create a new cluster to recover to and then recover your instance from a snapshot.

1. Find the version of kURL your original deployment was using.

If you have access to the original cluster, you can use this command:

```
kubectl get configmap -n kurl kurl-current-config -o jsonpath="{.data.kurl-version}" && echo
```

If you aren't able to run the command, you remember your PAM version, and you were on version 1.68.0 or later, you can look up the kURL version in the [Component versions in PAM releases](#) on page 82 table.

If you don't remember your PAM version or you were on a version earlier than 1.68.0, you need to contact your technical account manager or Support for assistance.

2. If you have access to the original cluster, follow the steps for [Migrating data between two systems with the same architecture](#) on page 118.

If your original cluster is completely offline and inaccessible, you'll need to contact your technical account manager or Support for assistance.

### Restriction:

Your old and new clusters must have the same connection status (online or offline). Disaster recovery from an offline to an online cluster (or vice versa) is not supported.

Additionally, for offline installs, both the old and new clusters must use the same PAM version.

## Troubleshooting PAM

Use this guide to troubleshoot issues with your Puppet Application Manager installation.

### How to look up your Puppet Application Manager architecture

If you're running PAM on a Puppet-supported cluster, you can use the following command to determine your PAM architecture version:

```
kubectl get installer --sort-by=.metadata.creationTimestamp -o jsonpath='{.items[-1:].metadata.name}' ; echo
```

Depending on which architecture you used when installing, the command returns one of these values:

- **HA architecture:** puppet-application-manager
- **Standalone architecture:** puppet-application-manager-standalone
- **Legacy architecture:** Any other value, for example, puppet-application-manager-legacy, cd4pe, or comply

## Resolve IP address range conflicts

When installing Puppet Application Manager, IP address ranges `10.96.0.0/22` and `10.32.0.0/22` must not be used by other nodes on the local network.

**Note:** The minimum size for CIDR blocks used by Puppet Application Manager are:

- **Standalone** - /24 for pod and service CIDRs
- **HA** - /23 for pod and service CIDRs
- Default of /22 is recommended to support future expansion

To resolve IP address range conflicts, create a `patch.yaml` file and add the `installer-spec-file=patch.yaml` argument when running the installation script (see below):

1. If you use IP addresses internally that overlap `10.32.0.0/22`, add the following to your `patch.yaml` file (`10.40.0.0/23` used here as an example range):

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 flannel:
 podCIDR: 10.40.0.0/23
 podCIDRRange: "/22"
```

2. If you use IP addresses internally that overlap `10.96.0.0/22`, add the following to your `patch.yaml` file (`10.100.0.0/23` used here as an example range):

```
spec:
 ...
 kubernetes:
 serviceCIDR: 10.100.0.0/23
 serviceCIDRRange: "/23"
```



**CAUTION:** The `podCIDR` and `serviceCIDR` ranges must not overlap.

3. Once your `patch.yaml` file is set up, add the `installer-spec-file=patch.yaml` argument when you run the installation script:

```
cat install.sh | sudo bash -s airgap installer-spec-file=patch.yaml
```

**Remember:** Add the `installer-spec-file=patch.yaml` argument any time you re-run the installation script, such as when reinstalling to upgrade to a new version.

### Related information

[Using sudo behind a proxy server](#) on page 129

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## Reset the PAM password

As part of the installation process, Puppet Application Manager (PAM) generates a password for you. You can update this password to one of your choosing after installation.

1. To reset the Puppet Application Manager password, run the following command as the root user:

```
kubectl kots reset-password default
```

The system prompts you to enter a new password of your choosing.

2. If the command fails with an unknown command "kots" for "kubectl" error, it's because /usr/local/bin is not in the path. To address this error, either update the path to include /usr/local/bin, or run this command:

```
/usr/local/bin/kubectl-kots reset-password default
```

## Update the PAM TLS certificate

A self-signed TLS certificate secures the connection between your browser and Puppet Application Manager (PAM). Once the initial Puppet Application Manager setup process is complete, you can upload new certificates by enabling changes to the installation's Kubernetes secrets.

Use this process if you chose not to add a TLS certificate when installing Puppet Application Manager, or if you need to update your existing TLS certificate.

1. Enable changes to your installation's kotsadm-tls Kubernetes secret by running:

```
kubectl -n default annotate secret kotsadm-tls acceptAnonymousUploads=1
```

2. Restart the kurl-proxy pod to deploy the change by running:

```
kubectl delete pods $(kubectl get pods -A | grep kurl-proxy | awk '{print $2}')
```

3. Once the kurl-rpxy pod restarts and is back up and running, navigate to `https://<HOSTNAME>:8800/tls` and upload your new TLS certificate.

## Reduce recovery time when a node fails

If a node running a non-replicated service like PostgreSQL fails, expect some service downtime.

How much downtime depends on the following factors:

- Timeout for communication between Kubernetes services (at least one minute to mark the node as unreachable).
- Timeout for the ekco service to determine that pods need to be rescheduled. The default is five minutes after node is marked unreachable.
- Time to restart services (at least two minutes, possibly up to five minutes, if there are complex dependencies).

The ekco service can be configured to reschedule pods more quickly by configuring the installation with a `patch.yaml` similar to the following:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 ekco:
 nodeUnreachableToleration: 1m
```

Apply the patch during an install or upgrade by including `installer-spec-file=patch.yaml` as an install option.

**Important:** This patch needs to be included during all future upgrades to avoid resetting the option.

## PAM components

Puppet Application Manager (PAM) uses a range of mandatory and optional components.

### Support services

- Database: PostgreSQL (single instance) - <https://www.postgresql.org/>
- Object storage: previously MinIO - <https://min.io>, now Ceph - <https://ceph.io>
- tlser for basic TLS cert management - <https://github.com/puppetlabs/tlser>
- kurl\_proxy for HTTPS proxying outside the Ingress (ports besides 80/443): [https://github.com/replicatedhq/kots/tree/v1.36.1/kurl\\_proxy](https://github.com/replicatedhq/kots/tree/v1.36.1/kurl_proxy)

### Kubernetes components

- Networking (CNI): Flannel - <https://github.com/flannel-io/flannel>
- Storage (CSI): Rook - <https://rook.io>, Ceph - <https://ceph.io>
- Ingress: Project Contour - <https://projectcontour.io>
- Kubernetes Cluster: kURL - <https://kurl.sh>
- Embedded kURL Cluster Operator: ekco - <https://github.com/replicatedhq/ekco>
- Admin Console: KOTS - <https://kots.io>
- Snapshots: Velero - <https://velero.io>, Restic - <https://restic.net>
- Monitoring: Prometheus - <https://prometheus.io>
- Registry: Docker Registry - <https://docs.docker.com/registry/>

### Optional components

Prometheus (+Grafana) and Velero (+Restic) are optional components:

- Prometheus+Grafana uses 112m/node + 600m CPU, 200MiB/node + 1750MiB RAM
- Velero+Restic uses 500m/node + 500m CPU, 512MiB/node + 128MiB RAM

If you do not need these optional components, they can be omitted from the initial install and further upgrades with a patch similar to the following:

```
apiVersion: cluster.kurl.sh/v1beta1
kind: Installer
metadata:
 name: patch
spec:
 prometheus:
 version: ''
 velero:
 version: ''
```

**Important:** This patch needs to be included during upgrades to avoid adding the components later.

If you want to remove optional components that are already installed, use the following command:

```
kubectl delete ns/monitoring ns/velero
```

## Generate a support bundle

When seeking support, you might be asked to generate and provide a support bundle. This bundle collects a large amount of logs, system information and application diagnostics.

To create a support bundle:

1. In Puppet Application Manager UI, click **Troubleshoot** > **Generate a support bundle**.

2. Select a method for generating the support bundle:
  - **Generate the bundle automatically.** Click **Analyze <APPLICATION NAME>** (<APPLICATION NAME> is replaced in the UI by the name of the Puppet application you have installed), and Puppet Application Manager generates the bundle for you and uploads it to the **Troubleshoot** page.
  - **Generate the bundle manually.** Click the prompt to generate a custom command for your installation, then run the command on your cluster. Follow the prompts to upload the bundle to Puppet Application Manager.
3. Review the collected data before forwarding it to Puppet, as it may contain sensitive information that you wish to redact.
4. Return to the **Troubleshoot** page, download the newly created support bundle, and send it to your Puppet Support contact.

## Create a support bundle from the command line

If installation of the Puppet Application Manager, or upload of an app, on an embedded kURL cluster fails, it may not be possible to access the UI to generate a support bundle.

You can generate a support bundle by using the default [kots.io](https://kots.io) spec. To do this, run the following command:

```
kubectl support-bundle https://kots.io
```

On an offline server, you can copy the default [kots.io](https://kots.io) spec by using the following command:

```
curl -o spec.yaml https://kots.io -H 'User-agent:Replicated_Troubleshoot/v1beta1'
```

The spec can then be uploaded to the server. Use the local spec by running:

```
kubectl support-bundle /path/to/spec.yaml
```

If the Puppet Application Manager UI is working and the app is installed, you can use:

```
kubectl support-bundle http://<server-address>:8800/api/v1/troubleshoot/<app-slug>
```

If the app is not installed but the Puppet Application Manager UI is running:

```
kubectl support-bundle http://<server-address>:8800/api/v1/troubleshoot
```

If you do not already have the support-bundle kubectl plugin installed, install it by using the command below:

```
curl https://krew.sh/support-bundle | bash
```

Or by installing [krew2](#) and running:

```
kubectl krew install support-bundle
```

## Using sudo behind a proxy server

Many of the commands you run to install or configure Puppet Application Manager (PAM) require root access. In the PAM documentation, commands that require root access use `sudo` to elevate privileges. If you're running PAM behind a proxy, `sudo` might not work correctly. If you're having trouble running commands with `sudo`, and you're behind a proxy, try switching to the `root` user and running the command without `sudo`.

## kURL can only be upgraded two minor versions at a time

Because kURL does not support upgrading more than two Kubernetes minor release versions at once, if you're upgrading from an older version of PAM, you might need to follow a specific upgrade path to avoid failures. For example, PAM version 1.80.0 uses Kubernetes version 1.21.x, so you can upgrade up to PAM 1.91.3 (Kubernetes

version 1.23.x), but not to PAM 1.94.0 (Kubernetes version 1.24.x). To determine the specific upgrade path for your installation, please check the [table of Kubernetes versions](#) for each version of PAM.

Attempting to upgrade too far at once returns the following error message: The currently installed kubernetes version is <CURRENT\_VERSION>. The requested version to upgrade to is <INVALID\_TARGET\_VERSION>. Kurl can only be upgraded two minor versions at time. Please install <VALID\_TARGET\_VERSION> first.

## Installing

---

To begin using Puppet Comply, you must first complete the initial setup process.

- [System requirements](#) on page 130

Refer to these system requirements to allow your Puppet Comply application to connect to Puppet Enterprise (PE).

- [Hosting the CIS-CAT Pro Assessor bundle internally](#) on page 131

There are three separate assessor bundles: one each for Linux, Mac, and Windows. If you choose to host your CIS-CAT Pro Assessor internally, you must download the appropriate assessor bundle for your operating system.

- [Set up Comply](#) on page 132

To start using Puppet Comply, you must complete the setup process, using both Puppet Application Manager (PAM) and Puppet Enterprise (PE).

- [Uninstall Comply and remove PAM](#) on page 142

If you are running Comply only in a Puppet-supported cluster, and you want to delete Comply, Puppet Application Manager (PAM), and any related files, uninstall Comply by deleting the Comply application and purging the Kubernetes cluster.

- [Uninstall Comply without removing PAM](#) on page 142

If you are running Comply in a customer-supported cluster, or if you don't want to remove Puppet Application Manager (PAM) for other reasons, you can remove Comply from the customer-supported cluster without taking further action.

- [Remove the CIS-CAT Pro Assessor from a node](#) on page 143

In rare cases, you might want to remove the CIS-CAT Pro Assessor from a node. For example, you can remove the assessor to exclude the node from compliance scans if the node is no longer relevant or is expiring. You can also remove the assessor if the node has issues that are causing the assessor to malfunction.

## System requirements

---

Refer to these system requirements to allow your Puppet Comply application to connect to Puppet Enterprise (PE).

### Open port requirements

Comply is deployed on a Kubernetes cluster that requires the following ports:

| Port                | Protocol | Purpose          | Source       | Destination           |
|---------------------|----------|------------------|--------------|-----------------------|
| <i>PE ports</i>     |          |                  |              |                       |
| 8140                | TCP      | Preflight checks | Comply       | Puppet primary server |
| 8143                | TCP      | PE integration   | Comply       | PE Orchestrator       |
| 8081                | TCP      | PE integration   | Comply       | PuppetDB              |
| 4433                | TCP      | PE integration   | Comply       | PE RBAC               |
| <i>Comply ports</i> |          |                  |              |                       |
| 443                 | TCP      | Comply access    | User browser | Comply UI             |

| Port  | Protocol | Purpose                                | Source           | Destination   |
|-------|----------|----------------------------------------|------------------|---------------|
| 443   | TCP      | Sending reports                        | Scan target node | Comply server |
| 30303 | TCP      | Assessor downloads and sending reports | Scan target node | Comply        |

**Tip:** Port 30303 is not required if you bring your own ingress. You can also set a custom Comply port in the **Comply port** field on the **Config** tab in Puppet Application Manager if you do not want to use port 30303.

### Supported Puppet Enterprise versions

The following versions of Puppet Enterprise (PE) are supported for use with Comply:

| PE version         |
|--------------------|
| 2023.0 and later   |
| 2021.7.2 and later |

For more information about PE versions, see [Puppet Enterprise lifecycle policy](#).

### Java Runtime Environment requirements

If you install the Comply module with the default setting of `true` for the `manage_java` option, the correct version of Java Runtime Environment (JRE) is installed automatically, and no further action is required.

**Restriction:** You cannot use the `manage_java` option on some operating systems, such as Ubuntu 16.04 and Mac OS X.

If you are independently managing JRE, ensure that the appropriate version is installed on the host system where the CIS-CAT Pro Assessor resides. JRE v1.8 or later is required. For the latest information about JRE requirements, see the [CIS-CAT Pro Assessor Configuration Guide](#).

## Hosting the CIS-CAT Pro Assessor bundle internally

There are three separate assessor bundles: one each for Linux, Mac, and Windows. If you choose to host your CIS-CAT Pro Assessor internally, you must download the appropriate assessor bundle for your operating system.

Comply supports a limited number of concurrent downloads of the CIS-CAT Pro Assessor to each Puppet-managed node. In lab testing, a maximum of approximately 120 concurrent downloads was achieved. For large-scale environments, hosting the CIS-CAT Pro Assessor on an internal server can help facilitate simultaneous downloading of the assessor to a large number of nodes during the installation or upgrade of SCM.

To host the assessor file internally, complete the following steps:

1. Download the appropriate assessor bundle for your operating system. The assessor bundles are located at:
  - `https://<COMPLY_FQDN>/files/assessor/linux`
  - `https://<COMPLY_FQDN>/files/assessor/mac`
  - `https://<COMPLY_FQDN>/files/assessor/windows`
2. In the Puppet Enterprise (PE) console, click **Node Groups** > **PE Infrastructure** > **PE Agent** > **Classes**.
3. In the **Add new class** field, select the Comply class.
4. In the **Parameter name** field, select `scanner_source`.

- Set the value of the scanner source to the URL where you want to host the assessor. For example, the URL can have the following structure, where `server-hosting-assessor-ip` specifies the IP address of the server that hosts the assessor and `os` specifies either `mac`, `linux`, or `windows`:

```
http://server-hosting-assessor-ip/assessor/os/assessor.zip
```

- Commit the changes.

### Related reference

[Guidelines for running Comply at scale](#) on page 154

You can run Puppet Comply on a maximum of 100,000 nodes. Before you run Comply at scale, review the guidelines for configuring the environment and running the scan. The process of running Comply at scale was tested by Puppet in a controlled environment. Because many factors affect performance, results in your system environment might vary.

## Set up Comply

To start using Puppet Comply, you must complete the setup process, using both Puppet Application Manager (PAM) and Puppet Enterprise (PE).

**Important:** Before you set up Comply, ensure that you have installed Puppet Application Manager (PAM) and Puppet Enterprise (PE) and have reviewed the [system requirements](#).

Setting up Comply involves the following steps:

- Configure Comply in Puppet Application Manager (PAM) in an online or offline environment. You can use the default ingress or, if you prefer, a custom NGINX ingress.
- Configure Comply TLS certificates in Puppet Enterprise (PE). You can configure these for the default ingress or, if you prefer, a custom NGINX ingress.
- Install the `comply` module.
- Classify the nodes you want to scan in PE.
- Deploy Comply.
- Add your PE credentials to Comply.

- [Configure Comply in an online environment](#) on page 133

The Comply configuration process creates a Kubernetes cluster and installs the application on that cluster.

- [Configure Comply in an offline environment](#) on page 133

Configure Puppet Comply in an air-gapped or offline environment where the Comply host server does not have direct access to the internet.

- [Configure Comply TLS certificates](#) on page 134

You need to generate certificates for Comply in Puppet Enterprise (PE) to enable automatic upgrades of the CIS-CAT Pro Assessor and for tasks to upload reports.

- [Configure Comply for a custom NGINX ingress \(online environment\)](#) on page 135

The Comply configuration process requires some extra configuration parameters if you use a custom NGINX ingress.

- [Configure Comply for a custom NGINX ingress \(offline environment\)](#) on page 136

Configure Puppet Comply in an air-gapped or offline environment where the Comply host server does not have direct access to the internet.

- [Configure Comply TLS certificates for a custom NGINX ingress](#) on page 137

You need to generate certificates for Comply in Puppet Enterprise (PE) to enable automatic upgrades of the CIS-CAT assessor and for tasks to upload reports.

- [Install the Comply module](#) on page 139

Install the Comply module from Puppet Forge.

- [Classify the nodes you want to scan](#) on page 140

In Puppet Enterprise (PE), classify the nodes you want to scan. You can scan a maximum of 5000 nodes in a batch.

- [Specify your own Java binary](#) on page 140

Specify your own Java binary to use rather than the version embedded with the CIS-CAT Pro Assessor.

- [Deploy Comply](#) on page 140

Now that you have completed the setup process, you can deploy Comply.

- [Add your PE credentials to Comply](#) on page 141

To allow Comply to communicate with PE, you must add your PE credentials to Comply.

- [Configure inventory refresh interval](#) on page 142

Configure how often to poll Puppet Enterprise for changes to the inventory, including changes in node and fact information. By default, polling occurs every 24 hours.

- [Configure data retention policy](#) on page 142

Configure how long to retain scan data. By default, Comply retains scan data indefinitely.

## Configure Comply in an online environment

The Comply configuration process creates a Kubernetes cluster and installs the application on that cluster.

### Before you begin

Follow the instructions to [install Puppet Application Manager](#).

1. In Puppet Application Manager, upload your Comply license and follow the prompts.

You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets Comply system requirements.

**Note:** The license file is issued by Puppet. If you do not have a license file, contact your Puppet representative. You must also agree to our [license agreement](#). If your license terms update, for example the expiry date or number of licensed nodes, upload your updated license file to Puppet Application Manager.

2. To configure your installation, click **Config**.

- a) In the **Hostname** field, enter the fully qualified domain name (FQDN) that you want to use to access Comply.

For example, this could be the name of the node you have installed Comply on. If you choose to use an FQDN that is different from the name of this node, you must configure your domain name system (DNS) to resolve the FQDN to the IP address of the Comply node.

- b) Configure any other settings on the page relevant to your installation. For example, you can determine how often the Comply inventory retrieves node and fact information from Puppet Enterprise. The default refresh interval for the Comply inventory is 24 hours, but you can specify a different value in the **Inventory Refresh Interval** section.

- c) When you have finished making any necessary changes to the configuration, click **Continue**.

3. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while your system is checked to make sure your cluster meets minimum system requirements. When the preflight check is complete:

- If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the status is **Ready to Deploy**, move on to the next step.

[Generate Comply certificates in PE.](#)

## Configure Comply in an offline environment

Configure Puppet Comply in an air-gapped or offline environment where the Comply host server does not have direct access to the internet.

### Before you begin

Follow the instructions to [install Puppet Application Manager](#).

Obtain the Comply bundle for air-gapped and offline environments by taking the following actions:

1. Locate the email that you received with the Comply licensing information. The email should include a password and a custom URL from which to download the bundle. If you no longer have the email, open a ticket with Puppet support so that you can obtain a custom URL and reset your password.
2. Navigate to the download portal (for example, <https://get.replicated.com/airgap/#/kots/comply/>) and log in with the password.
3. Select **Embedded cluster**.
4. Click **Download airgap bundle**.

1. In Puppet Application Manager (PAM), upload your Comply license and follow the prompts.

You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets Comply system requirements.

**Note:** The license file is issued by Puppet. If you do not have a license file, contact your Puppet representative. You must also agree to our [license agreement](#). If your license terms update, for example the expiry date or number of licensed nodes, upload your updated license file to Puppet Application Manager.

2. When prompted, upload the `.airgap` bundle for the most recent version of Comply.
3. To configure your installation, click **Config**.
  - a) In the **Hostname** field, enter the fully qualified domain name (FQDN) that you want to use to access Comply.
 

For example, this could be the name of the node you have installed Comply on. If you choose to use an FQDN that is different from the name of this node, you must configure your domain name system (DNS) to resolve the FQDN to the IP address of the Comply node.
  - b) Configure any other settings on the page relevant to your installation. For example, you can determine how often the Comply inventory retrieves node and fact information from Puppet Enterprise. The default refresh interval for the Comply inventory is 24 hours, but you can specify a different value in the **Inventory Refresh Interval** section.
  - c) When you have finished making any necessary changes to the configuration, click **Continue**.
4. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while Comply checks your system to make sure your cluster meets minimum system requirements. When the preflight check is complete:

- If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the status is **Ready to Deploy**, move on to the next step.

[Generate Comply certificates in PE.](#)

## Configure Comply TLS certificates

You need to generate certificates for Comply in Puppet Enterprise (PE) to enable automatic upgrades of the CIS-CAT Pro Assessor and for tasks to upload reports.

Certificates are required when setting up Comply for the following interactions:

- **Interactions between Comply and PE.** Interactions between Comply and PE require correct configuration of the CA certificate. Any issues with the CA certificate with regard to communication between Comply and PE result in an error on the Comply UI.
- **Agent runs.** If you have set up the Comply module to download the assessor from the Comply server (as opposed to being hosted locally), the assessor is downloaded using Mutual Transport Layer Security (MTLS) with the client certificate from the node. The Comply mtls-proxy component requires the configured TLS and CA certificate.
- **Scan task runs.** Running a scan sends reports back into Comply via an HTTP POST. This POST goes through the mtls-proxy and uses MTLS with the client certificate from the node.

Configuring Comply TLS certificates involves first generating the certificates in Puppet Enterprise (PE) and then setting up MTLs in PAM. MTLs enables a secure authenticated connection between your nodes and Comply.

For information on troubleshooting problems with certificates, see [Troubleshooting TLS issues in Comply](#) on page 172.

1. SSH into your PE primary server and generate the certificates:

```
puppetserver ca generate --certname <COMPLY-HOSTNAME>
```

This command does the following:

- Saves the private key to `/etc/puppetlabs/puppet/ssl/private_keys/<COMPLY-HOSTNAME>.pem`
  - Saves the certificate to `/etc/puppetlabs/puppet/ssl/certs/<COMPLY-HOSTNAME>.pem`
2. Log in to Puppet Application Manager, click the **Version history** tab, and click **Check for update**.
  3. Click the **Config** tab, and scroll down to **Transport layer security (TLS) certificates to interact with PE**.
  4. Ensure **Use a NodePort** is selected. If you want to change the Comply port from the default (30303), add the new port number in the **Comply port for PE nodes** field.

**Note:** To host the assessor on your own supported cluster via NGINX ingress, see [Configure Comply for a custom NGINX ingress \(online environment\)](#) on page 135 and [Configure Comply TLS certificates for a custom NGINX ingress](#) on page 137.

5. Enter the hostname of your PE instance in the **PE hostname** field to enable validation of the keys and certificates added in the next step.
6. Upload the signed certificate public key, the private key files, and the CA certificate, with the following locations:
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/certs/<COMPLY-HOSTNAME>.pem` into the **TLS certificate** field.
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/private_keys/<COMPLY-HOSTNAME>.pem` into the **TLS private key** field.
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/ca/ca.crt.pem` into the **CA certificate** field.
7. Click **Save Config**.
8. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while your system is checked to make sure your cluster meets minimum system requirements.

The **Config: Check if we can connect to PE using provided certificates** preflight passes if the certificates are configured correctly.

- If the preflight check status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the preflight check status is **Ready to Deploy**, proceed with the next step.
9. Click **Deploy**.

[Install the comply module.](#)

## Configure Comply for a custom NGINX ingress (online environment)

The Comply configuration process requires some extra configuration parameters if you use a custom NGINX ingress.

### Before you begin

Follow the instructions to [install Puppet Application Manager](#).

1. In Puppet Application Manager, upload your Comply license and follow the prompts.

You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets Comply system requirements.

**Note:** The license file is issued by Puppet. If you do not have a license file, contact your Puppet representative. You must also agree to our [license agreement](#). If your license terms update, for example the expiry date or number of licensed nodes, upload your updated license file to Puppet Application Manager.

2. To configure your installation, click **Config**.

- a) In the **Hostname** field, enter the fully qualified domain name (FQDN) that you want to use to access Comply.

For example, this could be the name of the node you have installed Comply on. If you choose to use an FQDN that is different from the name of this node, you must configure your domain name system (DNS) to resolve the FQDN to the IP address of the Comply node.

- b) In the **Configure access** section, add the following annotations to configure the Ingress if you use cert-manager.

```
kubernetes.io/ingress.class: nginx
cert-manager.io/cluster-issuer: letsencrypt-prod
```

- c) Configure any other settings on the page relevant to your installation. For example, you can determine how often the Comply inventory retrieves node and fact information from Puppet Enterprise. The default refresh interval for the Comply inventory is 24 hours, but you can specify a different value in the **Inventory Refresh Interval** section.
  - d) When you have finished, click **Continue**.
3. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while your system is checked to make sure your cluster meets minimum system requirements. When the preflight check is complete:
    - If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the status is **Ready to Deploy**, move on to the next step.

[Configure Comply TLS certificates for a custom NGINX ingress](#) on page 137.

## Configure Comply for a custom NGINX ingress (offline environment)

Configure Puppet Comply in an air-gapped or offline environment where the Comply host server does not have direct access to the internet.

### Before you begin

Follow the instructions to [install Puppet Application Manager](#).

Obtain the Comply bundle for air-gapped and offline environments by taking the following actions:

1. Locate the email that you received with the Comply licensing information. The email should include a password and a custom URL from which to download the bundle. If you no longer have the email, open a ticket with Puppet support so that you can obtain a custom URL and reset your password.
2. Navigate to the download portal (for example, <https://get.replicated.com/airgap/#/kots/comply/>) and log in with the password.
3. Select **Embedded cluster**.
4. Click **Download airgap bundle**.

1. In Puppet Application Manager (PAM), upload your Comply license and follow the prompts.

You'll be guided through the process of setting up SSL certificates, uploading a license, and checking to make sure your infrastructure meets Comply system requirements.

**Note:** The license file is issued by Puppet. If you do not have a license file, contact your Puppet representative. You must also agree to our [license agreement](#). If your license terms update, for example the expiry date or number of licensed nodes, upload your updated license file to Puppet Application Manager.

2. When prompted, upload the `.airgap` bundle for the most recent version of Comply.
3. To configure your installation, click **Config**.
  - a) In the **Hostname** field, enter the fully qualified domain name (FQDN) that you want to use to access Comply.
 

For example, this could be the name of the node you have installed Comply on. If you choose to use an FQDN that is different from the name of this node, you must configure your domain name system (DNS) to resolve the FQDN to the IP address of the Comply node.
  - b) In the **Configure access** section, add the following annotations to configure the Ingress if you use cert-manager:
 

```
kubernetes.io/ingress.class: nginx
cert-manager.io/cluster-issuer: letsencrypt-prod
```
  - c) Configure any other settings on the page relevant to your installation. For example, you can determine how often the Comply inventory retrieves node and fact information from Puppet Enterprise. The default refresh interval for the Comply inventory is 24 hours, but you can specify a different value in the **Inventory Refresh Interval** section.
  - d) When you have finished making any necessary changes to the configuration, click **Continue**.
4. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while Comply checks your system to make sure your cluster meets minimum system requirements. When the preflight check is complete:

- If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

- If the status is **Ready to Deploy**, move on to the next step.

[Configure Comply TLS certificates for a custom NGINX ingress](#) on page 137.

## Configure Comply TLS certificates for a custom NGINX ingress

You need to generate certificates for Comply in Puppet Enterprise (PE) to enable automatic upgrades of the CIS-CAT assessor and for tasks to upload reports.

### Before you begin

Make sure you have [set up Comply in Puppet Application Manager \(PAM\)](#) and you have followed the instructions in [Configure Comply for a custom NGINX ingress \(online environment\)](#) on page 135 or [Configure Comply for a custom NGINX ingress \(offline environment\)](#) on page 136 as appropriate to your implementation.

This process involves generating certificates in Puppet Enterprise (PE) and setting up Mutual Transport Layer Security (MTLS) in Puppet Application Manager (PAM). MTLS enables a secure authenticated connection between your nodes and Comply.

Certificates are required when setting up Comply for the following interactions:

- **Interactions between Comply and PE.** Interactions between Comply and PE require correct configuration of the CA certificate. Any issues with the CA certificate with regard to communication between Comply and PE result in an error on the Comply UI.

- **Agent runs.** If you have set up the Comply module to download the assessor from the Comply server (as opposed to being hosted locally) then the assessor is downloaded using MTLS with the client certificate from the node. The Comply mtls-proxy component requires the configured TLS and CA certificate.
- **Scan task runs.** Running a scan sends reports back into Comply via an HTTP POST. This POST goes through the mtls-proxy and uses MTLS with the client certificate from the node.

Configuring Comply TLS certificates involves first generating the certificates in Puppet Enterprise (PE) and then setting up MTLS in PAM. MTLS enables a secure authenticated connection between your nodes and Comply.

For information on troubleshooting problems with certificates, see [Troubleshooting TLS issues in Comply](#) on page 172.

1. SSH into your PE primary server and generate the certificates:

```
puppetserver ca generate --certname <COMPLY-HOSTNAME>
```

This command does the following:

- Saves the private key to `/etc/puppetlabs/puppet/ssl/private_keys/<COMPLY-HOSTNAME>.pem`
  - Saves the certificate to `/etc/puppetlabs/puppet/ssl/certs/<COMPLY-HOSTNAME>.pem`
2. Log in to Puppet Application Manager, click the **Version history** tab, and click **Check for update**.
  3. Click the **Config** tab, and scroll down to **Transport layer security (TLS) certificates to interact with PE**.
  4. Select **Use an ingress with a hostname**.
  5. Enter the fully-qualified domain name in the **PE TLS hostname** field using the same fully-qualified domain name that you used to generate the TLS certificates.

**Important:** The fully-qualified domain name in the **PE TLS hostname** field MUST be different from that used in the **Hostname** field in the **Required set-up** area.

6. Check that the following annotations are in the **SSL Passthrough Annotation** field and add them if not:

```
kubernetes.io/ingress.class: nginx
nginx.ingress.kubernetes.io/ssl-passthrough: "true"
```

**Note:** If you are not using NGINX, replace these annotations with those specific to your chosen ingress controller. For example, add [OpenShift as an MLTS proxy for PE certificates](#).

7. Enter the hostname of your PE instance in the **PE hostname** field to enable validation of the keys and certificates added in the next step.
8. Copy the signed certificate public key, the private key files, and the CA certificate, to the following locations:
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/certs/<COMPLY-HOSTNAME>.pem` to the **TLS certificate** field.
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/private_keys/<COMPLY-HOSTNAME>.pem` to the **TLS private key** field.
  - Paste the contents of `/etc/puppetlabs/puppet/ssl/ca/ca.crt.pem` to the **CA certificate** field.
9. Click **Save Config**.
10. Monitor the new version's preflight checks. The **Running Checks** indicator is shown on the screen while Comply checks your system to make sure your cluster meets minimum system requirements. When the preflight check is complete:
  - If the status is **Checks Failed**, click **View preflights**. Correct the issues and click **Re-run**. Repeat this step as needed.

**Important:** Do not move on until all preflight checks pass.

If the status is **Ready to Deploy**, move on to the next step.

11. The nginx-ingress controller configuration is configured with the `--enable-ssl-passthrough` setting disabled by default. This feature is required to enable passthrough in Ingress controller, allowing for the connection to be accepted by the application backends.

a) To edit the configuration inline with the running configuration, execute:

```
kubectl edit deployment -n <namespace> <ingress-controller>
```

b) Find the `spec:` configuration section in the `ingress-nginx-deployment.yaml` and append `--enable-ssl-passthrough` under `containers:` `-args`. For example:

```
apiVersion: apps/v1
kind: Deployment
metadata:
 [...]
spec:
 containers:
 - args:
 [...]
 - --enable-ssl-passthrough
```

[Install the comply module.](#)

## Install the Comply module

Install the Comply module from Puppet Forge.

### Before you begin

Make sure you have [generated the Comply certificates in PE](#).

Modules are self-contained, shareable bundles of code and data. The Comply module contains a Puppet Bolt task — the tool that runs the CIS assessor on your nodes.

The Comply module lives on Puppet Forge, a repository of thousands of modules. If you're new to PE and Comply, see [Managing environment content with a Puppetfile](#) for more information on the Puppetfile and installing modules.

1. Go to the [comply module on the Forge](#).

Follow the instructions in the **r10k or Code Manager** drop-down menu to add the module declaration to your Puppetfile. You also need to add its dependencies. For example:

```
Puppet comply module
mod 'puppetlabs/comply', '2.14.0'

dependencies for comply
mod 'puppet/archive', '7.1.0'
mod 'puppetlabs/chocolatey', '8.0.0'
mod 'puppetlabs/inifile', '6.1.1'
mod 'puppetlabs/ruby_task_helper', '0.6.1'
mod 'puppetlabs/stdlib', '9.6.0'
mod 'puppetlabs/powershell', '6.0.0'
mod 'puppetlabs/registry', '5.0.1'
mod 'puppetlabs/pwshlib', '1.1.1'
```

If you don't specify options, Code Manager installs the latest version and does not update it automatically. To always have the latest version installed, specify `:latest` and it updates automatically when a new version is released. Make sure you are always running the latest version of Comply if you intend to use the `:latest` keyword to update the Comply module. To install a specific version of the module that does not update automatically, specify the version number as a string.

**Important:** If you choose a specific version of the module, it **must** be the same as the Comply version. For example, version 2.3.0 of the module must be installed for Comply 2.3.0.

- SSH into your PE primary server and deploy the code:

```
puppet-code deploy --all
```

Classify the nodes you want to scan in Puppet Enterprise (PE).

## Classify the nodes you want to scan

In Puppet Enterprise (PE), classify the nodes you want to scan. You can scan a maximum of 5000 nodes in a batch.

### Before you begin

Make sure you have [installed the comply module](#).

*Classification* is when you create a node group, add nodes to the group, and assign *classes* to the group — in this case, the `comply` class. Classes are the blocks of Puppet code used to configure nodes and assign resources to them. If you are new to Puppet, see [Grouping and classifying nodes](#) for more information.

**Tip:** For guidelines about scanning thousands of nodes in a single batch, see [Guidelines for running Comply at scale](#) on page 154.

- In the PE console, click **Node groups**.
- Create a new node group or select an existing node group that you want to scan.
- On the **Classes** tab — in the **Add new class** field — select the `comply` class.
- Click **Add class**.
- In your new `comply` class, select the `scanner_source` **Parameter**.

**Note:** *Parameters* allow a class to request external data.

- Change the default parameter value for the Puppet-supported cluster to: `https://<COMPLY-HOSTNAME>:30303/assessor`.
- Click **Add to node group**, and then commit the changes.
- Run Puppet **twice**.

[Deploy Comply](#).

### Related information

[Create a custom profile](#) on page 158

Create a custom profile based on an existing benchmark.

## Specify your own Java binary

Specify your own Java binary to use rather than the version embedded with the CIS-CAT Pro Assessor.

Comply allows you to specify your own Java binary to use rather than the version embedded with the CIS-CAT Pro Assessor. If you specify your own binary, it must comply with the Java [requirements specified in the CIS-CAT Pro Assessor documentation](#).

Set the `assessor_java_path` module parameter to specify your own Java binary.

```
assessor_java_path => 'C:/Program Files/Java/jre1.8.0_451/bin/java.exe'
```

When this is configured, the embedded CIS Java binary is not copied to, or removed from, the node.

If the `assessor_java_path` parameter is removed, the embedded CIS java binary is used instead.

## Deploy Comply

Now that you have completed the setup process, you can deploy Comply.

**Before you begin**

Make sure you have [classified the nodes you want to scan in Puppet Enterprise \(PE\)](#).

1. Navigate to Puppet Application Manager. Once the version is ready to deploy, click **Deploy**. On the **Application** tab, monitor the application for readiness.

The application's status is shown as **Missing** for several minutes while deployment is underway.

**Tip:** You can monitor the deployment's progress by running `kubectl get pods --watch`.

When the deployment is complete, the application status changes to **Ready**. Comply is now deployed.

2. Navigate to `https://<COMPLY-HOSTNAME>` using the name of the Hostname FQDN you created earlier and sign into Comply using the default values:

- **username:** comply
- **password:** compliance

You are then prompted to create a new password.

[Add your PE credentials to Comply.](#)

**Add your PE credentials to Comply**

To allow Comply to communicate with PE, you must add your PE credentials to Comply.

**Before you begin**

Make sure you have [deployed Comply](#).

Adding your PE credentials authenticates Comply with Role Based Access Control (RBAC). Your PE account requires the following permissions:

| Type             | Action                       | Instance                                                              |
|------------------|------------------------------|-----------------------------------------------------------------------|
| Console          | View                         | -                                                                     |
| Job Orchestrator | Start, stop and view jobs    | -                                                                     |
| Node Groups      | View                         | All                                                                   |
| Nodes            | View node data from PuppetDB | -                                                                     |
| Tasks            | Run Tasks                    | Task: <code>comply::backup_assessor</code><br>Permitted on: All nodes |
| Tasks            | Run Tasks                    | Task: <code>comply::ciscat_scan</code><br>Permitted on: All nodes     |
| User Roles       | View                         | All                                                                   |
| User Roles       | Create                       | All                                                                   |

For more information on permissions, see [User permissions and user roles](#).

1. In Comply — located at `https://<COMPLY-HOSTNAME>/` — click **Settings**.
2. Click **Puppet Enterprise instance**.
3. Enter your PE **hostname**, **username**, and **password**.
4. Click **Submit**.

**Tip:** You can refresh the PE node and fact information by clicking **Refresh data**.

You'll now see a list of your classified nodes on the **Nodes** page.

You have completed the Comply setup process! You can now start running CIS scans on your nodes. If you're new to Comply, try out the [beginner's guide](#).

## Configure inventory refresh interval

Configure how often to poll Puppet Enterprise for changes to the inventory, including changes in node and fact information. By default, polling occurs every 24 hours.

1. In Comply — located at `https://<COMPLY-HOSTNAME>/` — click **Settings**.
2. Click **Configure your install**.
3. Select the units of time to set the inventory refresh interval in, either hours or minutes, under **Inventory refresh interval**.
4. Specify the inventory refresh interval. By default, the inventory refresh interval is set to 24 hours.

## Configure data retention policy

Configure how long to retain scan data. By default, Comply retains scan data indefinitely.

1. In Comply — located at `https://<COMPLY-HOSTNAME>/` — click **Settings**.
2. Click **Configure your install**.
3. Select **Enable data retention policy** to specify the retention period for your scan data.
4. Specify the scan data retention duration in weeks. Scan data older than this is deleted.

## Uninstall Comply and remove PAM

---

If you are running Comply only in a Puppet-supported cluster, and you want to delete Comply, Puppet Application Manager (PAM), and any related files, uninstall Comply by deleting the Comply application and purging the Kubernetes cluster.



**CAUTION:** By completing this procedure, you reset the Replicated installation.

1. From the command line of the node where Comply is installed, run the following command to delete the Comply application:

```
kubectl delete $(kubectl api-resources --verbs=delete -o name | paste -sd ", " -) -A -l app.kubernetes.io/part-of=comply
```

2. On the same node, run the following command to uninstall the embedded Kubernetes cluster:

```
curl https://k8s.kurl.sh/comply/tasks.sh | sudo bash -s reset
```

**Tip:** This command resets the Replicated installation with a purge.

3. Reboot your node to clear the `kube-ipvs0` device.

## Uninstall Comply without removing PAM

---

If you are running Comply in a customer-supported cluster, or if you don't want to remove Puppet Application Manager (PAM) for other reasons, you can remove Comply from the customer-supported cluster without taking further action.

From the command line of the node where Comply is installed, run the following command to delete the Comply application:

```
kubectl delete $(kubectl api-resources --verbs=delete -o name | paste -sd
", " -) -A -l app.kubernetes.io/part-of=comply
```

## Remove the CIS-CAT Pro Assessor from a node

In rare cases, you might want to remove the CIS-CAT Pro Assessor from a node. For example, you can remove the assessor to exclude the node from compliance scans if the node is no longer relevant or is expiring. You can also remove the assessor if the node has issues that are causing the assessor to malfunction.

**Note:** To remove old versions of the CIS-CAT Pro Assessor without uninstalling it from the node, use the `remove_assessor` task in the Comply module. For more information on running tasks in Comply, see [Running tasks in PE](#).

To remove the assessor:

1. Declassify the node by taking the following actions:
  - a. In the Puppet Enterprise console, click **Node groups**.
  - b. Select the node group where the node is classified with a Comply class.
  - c. If the node is pinned to a rule, click the **Rules** tab. Select the node name and click **Unpin**. If the node is not pinned to a rule, remove the class from the entire node group by clicking the **Classes** tab. Then, select the `scanner_source` parameter and click **Remove**.

2. On the command line of the node where the assessor is installed, run the appropriate command.

On a Linux operating system, run the following command:

```
rm -rf /opt/puppetlabs/comply/
```

On a Microsoft Windows operating system, run the following command:

```
Remove-Item -path C:\ProgramData\PuppetLabs\comply -recurse
```

3. Update the facts in Puppet Enterprise by running Puppet.
4. Retrieve the latest inventory from Puppet Enterprise by taking the following actions:
  - a. In Comply, click **Settings**.
  - b. Click **Refresh data**.

The assessor folder is removed along with the assessor JAR file and any backup copies of the JAR file. Because the node is declassified in Puppet Enterprise, Puppet will not reinstall the assessor during future runs. Declassified nodes are no longer visible to Comply and will be skipped in future compliance scans.

Optionally, if you want to resume scans on the node, you must classify the node so that it will be visible to Comply and the assessor will be reinstalled. For instructions, see [Classify the nodes you want to scan](#).

## Upgrading

New versions of Puppet Comply are released regularly. Upgrading to the current version helps you take advantage of the latest features, fixes, and improvements.

**Important:** The CIS-CAT Pro Assessor setup process is embedded in the `comply` module. If you are upgrading to the latest version of the CIS-CAT Pro Assessor, upgrade the `comply` module **before** you upgrade the Comply application. Note that you cannot run scans until you complete **both** of these upgrades.


## Upgrade to Comply 2.20.0

Comply 2.20.0 automatically upgrades the CIS-CAT Pro Assessor to the latest version when upgrading Comply. However, by adjusting your configuration you can choose to stay on a previous version of the assessor. Comply supports the latest plus the two previous versions of the assessor.

**Important:** When upgrading to 2.20.0, you must ensure that your PE account has permissions to View and Create User Roles. For more information, visit <https://www.puppet.com/docs/comply/2.x/configure-comply-with-pe.html>.

### Tips:

- If you are upgrading Comply in an environment with thousands of nodes, see [Guidelines for running Comply at scale](#) on page 154.
- If you are upgrading Comply to a version that includes a new assessor, you can expedite the process of installing the assessor on all nodes. In the PAM **Config** tab, in the **CIS-CAT Pro Assessor upgrade** section, select the checkbox to automatically start two Puppet runs after an assessor upgrade. To help prevent performance issues, enable this option only in small to medium Puppet deployments. If you enable this option, you can verify that a PE job was run: In Comply, select **Activity Feed > Assessor Upgrades** and click the assessor version to see the PE job number and detailed results of the upgrade.

1. Log in to PAM, click the **Version history** tab, and click **Check for update**.
2. Navigate to **CIS-CAT Pro Assessor version** and ensure that the correct version of the CIS-CAT Pro Assessor is selected.
3. Click **Save Config**.
4. Upgrade the comply module:
  - a) Update your Puppetfile with the latest version of the `comply` module and its dependencies.
  - b) To stay on a previous version of the CIS-CAT Pro Assessor, configure the module's **scanner\_version** and **scanner\_checksum** class parameters to the desired version of the assessor. To find the checksum for your desired version of the assessor, visit [CIS-CAT Pro Assessor history](#) on page 8. The version configured must match the version selected in step 5. To upgrade to the latest version of the CIS-CAT Pro Assessor, remove those parameters and the module defaults to the latest version.
  - c)  **CAUTION:** Only the latest version of the CIS-CAT Pro Assessor has the latest security fixes. Customers on older versions of the CIS-CAT Pro Assessor might be vulnerable to security issues.
  - d) Deploy code by running the `puppet-code deploy --all` command.
5. Navigate back to PAM. After the pre-flight checks are successfully completed, click **Go to updated version**, and then click **Deploy**.

If the upgrade of an assessor on a node fails, the node is marked in red on the **Inventory** page. Failures may be due to network issues. If that is the case, Comply attempts to upgrade the node when connectivity returns. An hourly background task runs to check if nodes are upgraded. If a node is not upgraded and remains red on the **Inventory** page, run Puppet. If the upgrade continues to fail, see the Puppet agent logs for more information as described in [Agent logs](#).

## Upgrade from Comply 2.2.2 to 2.3.0

Comply 2.3.0 automatically upgrades the CIS-CAT assessor to the latest version every time you upgrade Comply.

### Before you begin

Make sure you have generated certificates in Puppet Enterprise (PE) and set up Mutual Transport Layer Security (MTLS) in Puppet Application Manager (PAM). MTLS enables a secure authenticated connection between your nodes and Comply. For more information, see [Configure Comply TLS certificates](#) on page 134.

1. If you want Comply to update the CIS-CAT Assessor automatically, select **Automatically kick off PE jobs on assessor upgrade** on the **Config** page in Puppet Application Manager.

If you select this option, on upgrade Comply kicks off 2 PE agent runs: the first to download the new assessor, and the second update the facts in PE.

**Tip:** Because this option starts PE jobs automatically on upgrading Comply, systems administrators, especially of larger implementation, may wish to consider leaving this option unchecked. Assessor upgrade then takes place automatically when the next two PE jobs are run.

Comply requires the latest version of the assessor on the node in order to perform runs. A background task runs to check if nodes have been upgraded every 15 minutes if this option is selected and every hour if it is not selected. If a node does not upgrade and remains red on the **Inventory** page, run the Puppet agent. If the upgrade continues to fail, see the Puppet agent logs for more information.

2. Click **Save Config**.
3. If you have not already configured the `comply` class `scanner_source` parameter, you can do so at this point. Otherwise proceed to the next step. Navigate to Puppet Enterprise (PE), and update the default **value** of the `comply` class `scanner_source` parameter to one of the following assessor distribution files:

- If using the Puppet supported cluster: `https://<COMPLY-HOSTNAME>:30303/assessor`
- If using NGINX Ingress `scanner_source`: `https://<PE-TLS-FQDN>/assessor`

For more information, see [Classify the nodes you want to scan in PE](#).

Click **Add to node group**, and then commit the changes.

4. Upgrade the `comply` module.
  - a) Update your Puppetfile with the latest version of the `comply` module and its dependencies.
  - b) Deploy code by running the `puppet-code deploy --all` command.

**Warning:** When upgrading the `comply` module, running the agent before Comply is updated may cause an error.

5. Navigate back to Puppet Application Manager. After pre-flight checks have completed successfully, click **Go to updated version**, and then click **Deploy**.

**Note:** If the upgrade of an assessor on a node fails, the node is marked in red on the **Inventory** page. Failures may be due to network issues. If that is the case, Comply attempts to upgrade the node once connectivity returns. An hourly background task runs to check if nodes have been upgraded or not. If a node does not upgrade and remains red on the **Inventory** page, run the Puppet agent. If the upgrade continues to fail, see the Puppet agent logs for more information.

## Upgrade Comply in an online environment

Check for downloads and deploy updates from the **Version history** tab in the Puppet Application Manager (PAM) UI.

### Before you begin

Upgrade the `comply` module.

1. In the PAM UI, click **Version history**.
2. Click **Check for updates**.

Configure an automatic update check by clicking **Configure automatic updates**. You can check for updates hourly, every four hours, daily, weekly, or at a custom interval.

3. If an update is available, PAM downloads it for you and performs preflight checks on your system to make sure your cluster meets system requirements for the new version. Review the outcome of these checks by clicking **View preflight**.
4. When you're ready to upgrade to the new version of Comply, click **Deploy**.

## Upgrade Comply in an offline environment

---

If your environment does not have direct access to the internet, follow the documented procedure to upgrade Comply to the latest version.

### Before you begin

Upgrade the `comply` module.

1. Navigate to the portal provided to you by Puppet in the licence email, for example, `https://get.replicated.com/airgap/#/kots/comply/`, and log in with the password.
2. Select **Embedded cluster** and click **Download airgap bundle**.
3. Log into Puppet Application Manager:

```
https://<PLATFORM-ADMIN-CONSOLE-ADDRESS>:8800
```

4. Select **Version history** and upload the new version of the `.airgap` file that you downloaded in step 2.
5. Click **Deploy**.

## Using the Comply API

---

The Comply API allows you to automate actions, retrieve Comply data, and share Comply data with other groups and tooling. To use the API, you first must create a personal access token, after which you can access API endpoints.

Puppet's open & integrated approach ensures data sharing with enterprise tools (Such as risk management, systems of records etc.), improves productivity and cross-team collaboration leveraging the same data to ensure transparency.

- [Manage personal access tokens](#) on page 146

Authentication tokens allow a user to enter their credentials once, then receive an alphanumeric token to use to access different services or parts of the system infrastructure.

- [Authenticate public APIs](#) on page 147

The Comply API supports authentication via personal access tokens associated with Comply accounts. Authentication tokens are tied to the permissions granted to the user through role-based access control (RBAC), and they provide the user with the appropriate access to application programming interfaces (APIs).

- [REST API](#) on page 148

Fetch data and automate your workflows with the Comply REST API. Use the API to promote pipelines based on your own criteria and automate workspace setup to make onboarding new users more streamlined.

## Manage personal access tokens

---

Authentication tokens allow a user to enter their credentials once, then receive an alphanumeric token to use to access different services or parts of the system infrastructure.

- [Create and manage personal access tokens as a user](#) on page 147

As a Comply user, you can create a personal access token for yourself from your user profile in the Comply UI. The endpoints you can access with your token are dependent on your user role's permissions at the time you call the Comply API.

- [Manage personal access tokens as an admin](#) on page 147

As a Comply admin, you can view details of personal access tokens created by all Comply users. You can revoke tokens on behalf of other Comply users, but you cannot create a token for another user.

## Create and manage personal access tokens as a user

As a Comply user, you can create a personal access token for yourself from your user profile in the Comply UI. The endpoints you can access with your token are dependent on your user role's permissions at the time you call the Comply API.

To create a new access token:

1. In Comply, click your username at the bottom of the left navigation pane to access your user profile.
2. Specify a name and expiration date for your new API access token.
3. Click **Create personal access token**.
4. Copy your new token code and keep it in a secure location. Your code is only displayed once, so you must copy it immediately in order to access the API using that token.

Within your user profile you can also view all existing tokens for your user. If you no longer need a particular token, you can revoke it. To revoke a token:

1. Click **Revoke** on the token you wish to revoke.
2. In the pop-up, check the box for **Yes, revoke [token name]**.
3. Click **Revoke**.

You cannot edit existing tokens. If you need to edit a pre-existing token, revoke it and create a new token with the same name.

## Manage personal access tokens as an admin

As a Comply admin, you can view details of personal access tokens created by all Comply users. You can revoke tokens on behalf of other Comply users, but you cannot create a token for another user.

To revoke another user's token:

1. In Comply, click **Tokens**.
2. The **Tokens** page displays a list of all Comply users. Click a username to view their personal access tokens.
3. Find the token you want to revoke, and click **Revoke**.
4. In the pop-up, check the box for **Yes, revoke [token name]**.
5. Click **Revoke**.

## Authenticate public APIs

The Comply API supports authentication via personal access tokens associated with Comply accounts. Authentication tokens are tied to the permissions granted to the user through role-based access control (RBAC), and they provide the user with the appropriate access to application programming interfaces (APIs).

### Before you begin

You need the personal access token you created in [Create and manage personal access tokens as a user](#) on page 147. It is required for authentication.

Using personal access token to authenticate the Comply public API.

Add the personal access token you copied when creating the token and add this to the Authorization header of the API call. For example:

```
curl --insecure --header \
 "Authorization: <token value>" \
 "https://<hostname>/comply/api/v1/user"
```

## REST API

---

Fetch data and automate your workflows with the Comply REST API. Use the API to promote pipelines based on your own criteria and automate workspace setup to make onboarding new users more streamlined.

The Comply REST API uses the OpenAPI standard, *which you can refer to at any point at <https://<COMPLY-HOSTNAME>/openapi.json>*, where COMPLY-HOSTNAME is your organization's Comply server hostname. Comply uses [version 3.0.1](#) of the standard.

- [REST API tutorial](#) on page 148

Learn how to interact with the diverse set of endpoints and harness the power of REST architecture in integrating with and extending the functionality of Comply. This tutorial guides you through running an ad-hoc scan in Comply using the API.

- [Extract Compliance results using the Comply API](#) on page 151

If you have a personal API access token and the correct user permissions, you can use the Comply API to extract compliance results from Puppet Comply and share those results with your organization's other third party tooling.

- [Export Comply data using the Comply API](#) on page 151

If you have a personal API access token and the correct user permissions, you can use the Comply Exports API to create, retrieve, download, and delete exports of data from Comply.

- [Synchronize inventory with Puppet Enterprise using the Comply API](#) on page 151

If you have a personal API access token and the correct user permissions, you can use the Comply Inventory API to initiate a PE inventory sync.

### REST API tutorial

Learn how to interact with the diverse set of endpoints and harness the power of REST architecture in integrating with and extending the functionality of Comply. This tutorial guides you through running an ad-hoc scan in Comply using the API.

#### Before you begin

Make sure that you have read the [Authentication](#) section and have a personal access token. You need the personal access token you created in [Create and manage personal access tokens as a user](#) on page 147. It is required for authentication.

You also need to have `comply-operator` or `comply-admin` permissions on the token to trigger an ad-hoc scan.

1. Optional: Get the host's information. This step can be skipped if you do not want to run a scan on a particular host.

```
curl --request GET \
 --url https://<COMPLY-HOSTNAME>/api/public/v1/hosts \
 --header 'Authorization: Bearer <TOKEN>'
```

This returns something similar to:

```
{
 "data": [
 {
 "id": 1,
 "name": "00.example.puppet.net",
 "environment": "development",
 "assessor_version": "4.41.0",
 "operating_system": {
 "id": 9,
 "name": "AlmaLinux 8"
 }
 },
 ...
 {
 "id": 10,
 "name": "09.example.puppet.net",
 "environment": "production",
 "assessor_version": "4.41.0",
 "operating_system": {
 "id": 45,
 "name": "Debian 12"
 }
 }
],
 "total": 100,
 "next": "/api/public/v1/hosts?offset=10&limit=10",
 "offset": 0,
 "limit": 10
}
```

From this you can extract the name of the hosts.

## 2. Get the operating system information.

```
curl --request GET \
 --url 'https://<COMPLY_HOSTNAME>/api/public/v1/operating-systems' \
 --header 'Authorization: Bearer <TOKEN>'
```

This returns something similar to:

```
[
 {
 "id": 8,
 "name": "CentOS 7",
 "compatible_benchmarks": [
 {
 "id": 8,
 "name": "4.0.0 CIS CentOS Linux 7"
 }
]
 },
 ...
 {
 "id": 3,
 "name": "macOS 13.0",
 "compatible_benchmarks": [
 {
 "id": 3,
 "name": "2.0.0 CIS Apple macOS 13.0 Ventura"
 }
]
 }
]
```

From this you can extract the name of the operating system.

## 3. Run the ad-hoc scan on nodes with specified hostnames, operating systems, and environment and node groups. Pass these as a filter object of the request body.

```
curl --request POST \
 --url https://<COMPLY_HOSTNAME>/api/public/v1/scan \
 --header 'Authorization: Bearer <TOKEN>' \
 --header 'Content-Type: application/json' \
 --data '{
 "scan-type": "desired",
 "filters": {
 "operating_systems": [
 "redhat 7",
 "redhat 8"
]
 }
 }'
```

This triggers a scan on all Red Hat 7 and Red Hat 8 nodes in the Puppet Enterprise inventory.

4. Optional: You can run a scan on particular hosts using the hostnames gathered from Step 1.

```
curl --request POST \
 --url https://<COMPLY_HOSTNAME>/api/public/v1/scan \
 --header 'Authorization: Bearer <TOKEN>' \
 --header 'Content-Type: application/json' \
 --data '{
 "scan-type": "desired",
 "filters": {
 "hostnames": [
 "00.example.puppet.net"
]
 }
 }'
```

This triggers a scan on the 00.example.puppet.net host.

## Extract Compliance results using the Comply API

If you have a personal API access token and the correct user permissions, you can use the Comply API to extract compliance results from Puppet Comply and share those results with your organization's other third party tooling.

You can extract both summary and raw data results for one, many, or all nodes up to 100,000 nodes. You can also filter results by hostname, node group, environment or operating system.

For information on how to access and use the Comply API, reference the documentation at <https://<COMPLY-HOSTNAME>/openapi.json>, where COMPLY-HOSTNAME is your organization's Comply server hostname. This documentation is in OpenAPI 3.0 format, so you can view it more clearly by importing it into an API documentation tool that supports the format, such as Swagger UI or Postman.

## Export Comply data using the Comply API

If you have a personal API access token and the correct user permissions, you can use the Comply Exports API to create, retrieve, download, and delete exports of data from Comply.

You can interact with the Exports API using the following methods:

- `POST /v1/export-job`: Create an export job. The export job will be queued and subsequently processed by Comply. If the export job completes successfully, you may download the new export.
- `GET/v1/export-job/{id}`: Retrieve information about an export job, including its current status and the location of the associated export, if one has been generated successfully.
- `DELETE/v1/export/{id}`: Delete an export permanently. This action also deletes the associated export job.
- `GET/v1/export/{id}/file`: Download an export file as a ZIP archive.
- `GET/v1/exports`: Retrieve a list of all exports.

You can view inventory sync jobs requested by the Inventory API and their status in the **Settings** page of the Comply console.

For information on how to access and use the Comply API, reference the documentation at <https://<COMPLY-HOSTNAME>/openapi.json>, where COMPLY-HOSTNAME is your organization's Comply server hostname. This documentation is in OpenAPI 3.0 format, so you can view it more clearly by importing it into an API documentation tool that supports the format, such as Swagger UI or Postman.

## Synchronize inventory with Puppet Enterprise using the Comply API

If you have a personal API access token and the correct user permissions, you can use the Comply Inventory API to initiate a PE inventory sync.

You can interact with the Inventory API using the following methods:

- `POST/v1/inventory-sync-job`: Create an inventory sync job. The inventory sync job starts immediately if there is not already an inventory sync job running in Comply.
- `GET/v1/inventory-sync-job/{id}`: Retrieve information about an inventory sync job.

For information on how to access and use the Comply API, reference the documentation at <https://<COMPLY-HOSTNAME>/openapi.json>, where COMPLY-HOSTNAME is your organization's Comply server hostname. This documentation is in OpenAPI 3.0 format, so you can view it more clearly by importing it into an API documentation tool that supports the format, such as Swagger UI or Postman.

## Managing access for Comply users

Comply integrates with Puppet Enterprise (PE) for role-based access control (RBAC). You can create or import new Comply users and assign them to roles in the PE Console. There are three default roles provided for Comply users: comply-admin, comply-operator, and comply-viewer. Users must be assigned to one of these roles in order to log into Comply.

### Adding new Comply users and roles

In order to add a new local user in Comply, log into the Puppet Enterprise (PE) console associated with your Comply instance. Your user in PE must have the ability to create and edit user roles. Follow the instructions found in the PE documentation at [https://www.puppet.com/docs/pe/2023.2/rbac\\_user\\_roles\\_intro.html#create\\_a\\_new\\_user](https://www.puppet.com/docs/pe/2023.2/rbac_user_roles_intro.html#create_a_new_user) to add a new user and assign them to one of the three provided default Comply roles.

For more information on configuring Comply with PE, visit [Add your PE credentials to Comply](#) on page 141.

### Importing existing users to Comply

RBAC integrates with LDAP for easy import of existing remote users. You can find instructions on how to connect to LDAP, import users, and assign them to roles at <https://www.puppet.com/docs/pe/2023.2/rbac-ldap.html>.

**Note:** Importing SAML users from Microsoft ADFS and Okta is not currently supported in Comply.

### Default Comply roles

There are three default roles provided for Comply users. Each role is assigned different permissions and has a different view of the Comply console, meaning that some options in Comply are greyed out or unavailable for users with certain roles.

The following table explains the permissions included by default for each role:

| Category            | Action                          | Puppet Comply Role |                 |               |
|---------------------|---------------------------------|--------------------|-----------------|---------------|
|                     |                                 | comply-admin       | comply-operator | comply-viewer |
| <b>Dashboard</b>    | View compliance dashboard       | #                  | #               | #             |
| <b>Node Results</b> | View node results list          | #                  | #               | #             |
|                     | Export node results data to CSV | #                  | #               |               |
|                     | View node detail                | #                  | #               | #             |
| <b>Rule Detail</b>  | View rule detail                | #                  | #               | #             |
|                     | Create an exception             | #                  | #               |               |
| <b>Scan Reports</b> | View scans list                 | #                  | #               | #             |

|                          |                                                      |   |   |   |
|--------------------------|------------------------------------------------------|---|---|---|
|                          | View scan report                                     | # | # | # |
|                          | View scan report: rule performance                   | # | # | # |
|                          | View scan report: node performance                   | # | # | # |
|                          | Run an ad hoc scan                                   | # | # |   |
| <b>Generated Reports</b> | View the list of exported data                       | # | # | # |
|                          | Download exported data                               | # | # | # |
| <b>Inventory</b>         | View inventory list                                  | # | # | # |
|                          | Update desired compliance (in bulk and individually) | # | # |   |
| <b>Scan Schedules</b>    | View scan schedules list                             | # | # | # |
|                          | Create a scan schedule                               | # | # |   |
|                          | View a scan schedule detail                          | # | # | # |
|                          | Edit a scan schedule                                 | # | # |   |
|                          | Manage the nodes linked to a scan schedule           | # | # |   |
|                          | Pause, end, restart a scan schedule                  | # | # |   |
|                          | Delete a scan schedule                               | # | # |   |
| <b>Custom Profiles</b>   | Create a custom profile                              | # | # |   |
|                          | View custom profiles list                            | # | # | # |
|                          | View custom profile details                          | # | # | # |
|                          | Create a custom profile                              | # | # |   |
|                          | Edit a custom profile                                | # | # |   |
|                          | Delete a custom profile                              | # | # |   |
|                          | Export custom profiles to csv                        | # | # |   |
| <b>Exceptions</b>        | View exceptions list                                 | # | # | # |
|                          | View exceptions detail                               | # | # | # |

|                      |                                                     |   |   |
|----------------------|-----------------------------------------------------|---|---|
|                      | Create an exception                                 | # | # |
|                      | Edit an exception                                   | # | # |
|                      | Resolve an exception<br>(one, many, all nodes)      | # | # |
|                      | Delete an exception                                 | # | # |
| <b>Activity Feed</b> | View activity feed scans tab                        | # | # |
|                      | View activity feed assessor upgrade tab             | # | # |
|                      | View activity feed assessor upgrade summary page    | # | # |
| <b>License</b>       | View license page                                   | # |   |
|                      | Sync license                                        | # |   |
| <b>Settings</b>      | View settings page                                  | # |   |
|                      | Edit settings page<br>(refresh data, remove/add PE) | # |   |
| <b>Upgrade</b>       | See alert advising there is an upgrade available    | # |   |

## Guidelines for running Comply at scale

You can run Puppet Comply on a maximum of 100,000 nodes. Before you run Comply at scale, review the guidelines for configuring the environment and running the scan. The process of running Comply at scale was tested by Puppet in a controlled environment. Because many factors affect performance, results in your system environment might vary.

### System requirements and configuration for large-scale environments

To support environments with more than 10,000 nodes, your Comply installation needs a total of at least 16GB of memory and 100GB of storage space available.

Depending on your node count, scan frequency, and desired retention period, you may also need to adjust the "Comply PostgreSQL capacity" and "Comply PostgreSQL memory" values under the "Additional Support" configuration section. The table below contains recommended values based on certain landmark node counts. Note that these values assume one scan per week and the default data retention period of 14 weeks.

**Table 1:**

| 14 weeks data retention at 1 scan per week | PostgreSQL capacity | PostgreSQL memory |
|--------------------------------------------|---------------------|-------------------|
| 1,000 nodes                                | Default             | Default           |
| 10,000 nodes                               | 25Gi                | Default           |

| 14 weeks data retention at 1 scan per week | PostgreSQL capacity | PostgreSQL memory |
|--------------------------------------------|---------------------|-------------------|
| 50,000 nodes                               | 75Gi                | 4Gi               |
| 75,000 nodes                               | 105Gi               | 8Gi               |
| 100,000 nodes                              | 135Gi               | 12Gi              |

For higher node counts, add 6Gi to PostgreSQL capacity for each additional 5,000 nodes. For longer retention periods, divide calculated storage requirement by default retention period (14) to determine per week storage requirement and then multiply by desired retention period.

**Note:** Comply PostgreSQL capacity cannot be modified after initial installation without contacting Puppet support.

### Configure the scan process

To help optimize the scan process, follow the guidelines:

- In Puppet orchestrator, set the `task_concurrency` parameter to a value appropriate for your environment and number of nodes. This value sets the maximum number of task or plan actions that can run concurrently in the orchestrator. If you set the parameter to 250 and run a scan of 5000 nodes, the orchestrator will be fully consumed until the scans are completed on all 5000 nodes. (For more information about optimizing performance, see [Tune task and plan performance in Puppet Enterprise \(PE\)](#).)
- Schedule scans to coincide with periods of minimal workflow to help ensure adequate network throughput.
- Plan adequate time for the initial inventory ingestion from Puppet Enterprise (PE). In lab testing, the ingestion of 100,000 nodes took 20 minutes.
- If you have a large number of nodes, consider configuring ad hoc and scheduled scans in smaller batches of up to 10,000 nodes.

### Upgrade Comply in a large-scale environment

Before you upgrade Comply in an environment with thousands of nodes, review the limitations and consider the best strategy for your environment.

During the standard upgrade process, a new version of the CIS-CAT Pro Assessor is downloaded to each Puppet-managed node. However, Comply supports a limited number of concurrent downloads of the assessor. In lab testing, a maximum of about 120 concurrent downloads was achieved. Thus, if you initiate an upgrade of thousands of nodes, not all nodes are updated on the first run.

You can resolve the issue in one of the following ways:

- Run Puppet manually on a maximum of 120 nodes. Repeat the process until all nodes are updated.
- Configure Comply to host the assessor file on an internal web server and then upgrade Comply. If you choose this option you need to ensure that you host the correct assessor bundle based on your operating system.

To host the assessor file internally and upgrade Comply, complete the following steps:

1. If you have not already, download the appropriate assessor bundle for your operating system. The assessor bundles are located at:
  - `https://<COMPLY_FQDN>/files/assessor/linux`
  - `https://<COMPLY_FQDN>/files/assessor/mac`
  - `https://<COMPLY_FQDN>/files/assessor/windows`
2. In the Puppet Enterprise (PE) console, click **Node Groups** > **PE Infrastructure** > **PE Agent** > **Classes**.
3. In the **Add new class** field, select the Comply class.
4. In the **Parameter name** field, select `scanner_source`.

- Set the value of the scanner source to the URL where the assessor will be hosted. For example, the URL can have the following structure, where `server-hosting-assessor-ip` specifies the IP address of the server that will host the assessor and `os` specifies either `mac`, `linux`, or `windows`:

```
http://server-hosting-assessor-ip/assessor/os/assessor.zip
```

- Commit the changes.
- In the PE console, click **Run > Puppet**.
- Complete the upgrade process by selecting the relevant nodes and running the job.

### Optimize scanning and reporting at scale

You can compare the results of your scanning and reporting processes against the results obtained in lab testing. If performance is not adequate in your environment, determine the cause of bottlenecks and address the issues.

Comply has been tested and is able to process reports from up to 100,000 nodes in a single scan. Processing this number of reports can take up to 120 minutes depending on system resources. However, total scan time may be significantly longer based on Puppet orchestrator concurrency limits as well as the amount of time the CIS-CAT Pro Assessor takes to run on individual nodes. If you have a large number of nodes, consider configuring ad hoc and scheduled scans in smaller batches of up to 10,000 nodes.

Node results raw data exports can take up to 36 minutes for 90,000 nodes, or up to 4 minutes per batch of 10,000 nodes. Allow additional time if generating several exports of over 10,000 nodes concurrently.

The assessor run times are affected by the host type. In general, scans on Microsoft Windows systems take longer than scans on `*nix` systems. Run times can vary significantly, depending on many other factors. For example, run times are longer for nodes with many user accounts and for nodes with many types of software installed. Results obtained in the lab represent an optimal use case.

To help understand performance issues, you can analyze log files. For more information, see [Access logs](#) on page 170.

For more information on configuration, workflow, and best practices, visit <https://www.puppet.com/docs/patterns-and-tactics/latest/patterns-and-tactics.html>.

## Desired compliance

---

Set your desired compliance. This is the benchmark and profile that you assign to a particular node and that is scanned on that node by default. Generally, you set compliance only once for your nodes.

By default, Comply automatically assigns an appropriate benchmark for each operating system, along with a Level 1 profile, to nodes that have not been set based on fact information from PE. Accepting this option is the quickest way to get up and running with desired compliance. You can also customize the default desired compliance settings for an operating system.

The **Benchmark** and **Profile** columns on the **Inventory** page show you the benchmark and profile for each node. You can view further details for a node by clicking on the row assigned to the node. To change a node's desired compliance, follow the instructions in [Manually set desired compliance](#) on page 156. You can also follow the manual instructions to assign a different benchmark and profile to a node, or to assign a custom profile.

**Restriction:** Only one benchmark and profile can be assigned to each node.

### Manually set desired compliance

---

If you don't want to use the benchmark and profile that Comply assigns automatically to your node, you can set the benchmark and profile that you prefer from the **Inventory** page.

1. On the **Inventory** page of Comply, click the node for which you want to specify desired compliance.  
In the **Information** window that appears on the right, you can see facts about the node and whether desired compliance has been set.
2. Choose the CIS Benchmark and profile that you want to assign to the node using the drop-down menus.  
The benchmark and profile you set here is the desired compliance option for future scans.  
If you have created a custom profile, you can set it as the desired compliance by clicking **Use an associated custom profile?**
3. Click **Update**.

Now that you have applied desired compliance, you can run scans based on your selection.

#### Related information

[Create a custom profile](#) on page 158

Create a custom profile based on an existing benchmark.

[Run an ad hoc scan](#) on page 161

Run your desired compliance scan or an ad hoc scan on your nodes.

## Bulk set desired compliance

---

You can also bulk assign desired compliance to a batch of nodes if the nodes are running on the same operating system, and the latest version of the CIS-CAT Pro Assessor is installed on each node.

1. On the **Inventory** page of Comply, click the check-boxes of the nodes for which you want to specify desired compliance.
2. In the toolbar at the top of the page, click **Actions > Set desired compliance**.
3. Choose the CIS Benchmark and profile that you want to assign to the node using the drop-down menus.
4. If you have created a custom profile, you can set it as the desired compliance by clicking **Use an associated custom profile?**
5. Click **Update**.

The **Benchmark** and **Profile** columns tell you the desired compliance set for each node. You can view the node's information, including its assigned benchmark and profile, by clicking on the node.

Now that you have applied desired compliance, you can run scans based on your selection.

## Set desired compliance by operating system

---

If you prefer to assign the desired compliance to an operating system, you can set a benchmark and profile for this operating system from the **Settings** page.

1. On the **Settings** page of Comply, click **Manage desired compliance**.
2. Click the operating system you want to assign a benchmark and profile to. Any nodes with this operating system that are added have these benchmark and profiles automatically assigned to them.
3. Choose the CIS Benchmark and profile that you want to assign to the node using the drop-down menus.
4. If you have created a custom profile, you can set it as the desired compliance by clicking **Use an associated custom profile?**
5. Check the check box to assign these settings to all of the selected operating system's current nodes, overriding any individually assigned benchmarks and profiles. Leave this unchecked to keep the individual desired compliance settings for any nodes that already exist on this operating system.
6. Click **Save**.

When nodes are added to an operating system, these benchmarks and profiles are automatically assigned to them.

Now that you have applied desired compliance, you can run scans based on your selection.

## Custom profiles

---

A custom profile is a benchmark profile that you customize to fit your organization's internally defined standards. You can base a custom profile on an existing benchmark and profile combination, and then specify which rules to apply.

For example, assume that your Center for Internet Security (CIS) Benchmark includes a rule that prohibits users from reusing any of the last 24 passwords that they specified. However, your organization enforces a stricter password policy. In this case, you could create a custom profile that enforces all other benchmark rules but excludes the CIS password rule. In this way, you would achieve more realistic compliance scores.

Custom profiles are typically created for long-term use. During an audit, you can note that a custom profile is applied to meet your organization's requirements.

The Comply API allows you to retrieve lists of profiles and information about specific profiles using the [Profiles endpoints](#).

### Create a custom profile

---

Create a custom profile based on an existing benchmark.

1. Navigate to **Custom profiles**.
2. Click **Create custom profile**.
3. Select a **Benchmark** and **Profile**.
4. Deselect rules in the profile that you **do not** want to scan and click **Next**.
5. Enter the name of the profile and, optionally, a description.
6. Click **Save custom profile**.

Your custom profile appears as an option when you assign the associated benchmark to a node.

Navigate to **Nodes** to set your custom profile as the desired compliance for your nodes or perform an ad hoc scan by selecting your custom profile on the **Scans** page.

To apply a custom profile to several nodes simultaneously, go to the **Inventory** page and select the nodes. From the **Actions** menu, select **Set desired compliance**. In the **Benchmark**, **Profile**, and **Custom profile** fields, specify the desired compliance and click **Update**.

**Restriction:** The selected nodes must be running on the same operating system, and the latest version of the CIS-CAT Pro Assessor must be installed on each node.

#### Related information

[Manually set desired compliance](#) on page 156

If you don't want to use the benchmark and profile that Comply assigns automatically to your node, you can set the benchmark and profile that you prefer from the **Inventory** page.

[Run an ad hoc scan](#) on page 161

Run your desired compliance scan or an ad hoc scan on your nodes.

### Delete a custom profile

---

When a custom profile is no longer necessary, you can delete the profile.



**CAUTION:** After you delete a custom profile, you cannot restore it.

1. In the **Custom profiles** table, select one or more profiles to delete.
2. In the **Actions** drop-down menu, select **Delete selected**.
3. When you are prompted to confirm the choice, click **Delete**.

Any nodes that were assigned to the deleted custom profile are unassigned. During future scans, the nodes will not be checked against the deleted custom profile. Any nodes that were assigned to the deleted custom profile will be reassigned to their default profile.

Previously run scan reports will continue to show results that reflect the custom profile. However, the custom profile appears in red and will be flagged with a warning triangle. The hover help will indicate that the custom profile no longer exists.

## Export a custom profile

---

You can export a list of all, many, or one of your custom profiles along with their details.

1. In the **Custom profiles** table, select one or more custom profiles to export.
2. Select **Actions**, then **Export raw data**.
3. Write a name and description for the export.
4. Click **Submit**.
5. You can view your exported custom profiles in the **Exported data** tab.

## Exceptions

---

Each Center for Internet Security (CIS) Benchmark specifies many controls, commonly known as rules. In some cases, you might find it useful to create a temporary exception to a rule and apply the exception to one node, several nodes, or all nodes.

For example, assume that your environment includes legacy nodes that are installed on an operating system that is not CIS compliant, and you plan to decommission those nodes. You create an exception that specifies the rule, the affected nodes, the expiration date, the reason for the exception, and the name of the approver. On the next scan, the rule is not applied to the specified nodes, and the compliance score accurately reflects the exception. Later, after the nodes are decommissioned, the exception expires on your specified date. If an audit occurs, a record of the exception remains available on the **Exceptions** page.

## Create an exception

---

When you create an exception to a rule, you prevent the rule from being applied to one or more nodes. If you run a scan while the exception is active, the compliance score of the rule is excluded from the overall compliance score of any specified nodes.

**Tip:** Exceptions are typically temporary with a specified expiration date and time. However, you can create an exception with no expiration date or time.

1. Click **Scans > Scan reports** and select a scan to which you want to add an exception.
2. On the **Scan report** page, on the **Rules** tab, locate the rule for which you want to create an exception. Click **View report**.
3. On the **Scan report: Rule performance** page, next to the rule name, click **View rule detail**.

4. On the **Rule detail** page, click **Create exception** and follow the exception creation workflow:
  - a. Select a profile and, optionally, a custom profile. Click **Next**.
  - b. Select one or more nodes to which the exception will apply. Click **Set expiry**.
  - c. Optionally, set an expiration date, time, and time zone. Click **Add details and review**.
  - d. Provide a name and reason for the exception.
  - e. Optionally, for audit or tracking purposes, you can specify the name of the person who approved the exception and the associated ticket number, if applicable.
  - f. Click **Save exception and exit**.

**Tip:** Alternatively, you can create an exception by going to the Comply navigation pane, clicking **Exceptions** and then clicking **How do I create an exception?**

Optionally, to see how the exception affects the compliance score, run a scan.

## View an exception

---

To view one or more exceptions, go to the Puppet Comply navigation pane and click **Exceptions**.

You can filter exceptions by **Active**, **Resolved**, or **Expired**. For each exception, you can view the associated benchmark and profile. You can also see the rule, the number of nodes affected, and the expiration information.

When viewing exceptions, select an exception and then click **View exception detail** for a detailed view of the exception. Here you can find the nodes for which the exception is active. You can also edit the exception details by clicking **Edit details**, or resolve the exception by clicking **Resolve**. For more information on resolving exceptions, visit [Resolve an exception](#) on page 160.

The **Exceptions** page also includes the **How do I create an exception?** button. You can click the button for instructions on how to create an exception.

## Resolve an exception

---

To stop using an exception before its expiration date, resolve the exception for all nodes or a subset of nodes. After an exception is resolved, the rule scan results again count towards the overall compliance score for the impacted nodes.

1. Go to the Puppet Comply navigation pane and click **Exceptions**.
2. Specify the exception to resolve, and then click **View exception detail**.
3. To resolve the exception for all nodes, click **Resolve**.
  - Provide a reason for resolution, and an approver if applicable, and then click **Submit**.
4. To resolve the exception for only some nodes, select the checkboxes for the nodes on which you would like to resolve the exception, and then select **Resolve selected** from the **Actions** dropdown menu.
  - Provide a reason for resolution, and an approver if applicable, and then select **Submit**.

## Delete an exception

---

In general, exceptions should not be deleted because an auditor might want to see a record of the exception. However, you might want to delete an exception in rare cases. For example, if you create an exception by mistake, create an exception incorrectly, or you no longer require a record of the exception, you can delete it.



**CAUTION:** After you delete an exception, you cannot restore it.

1. Go to the Puppet Comply navigation pane and click **Exceptions**.
2. Specify the exception to delete, then select **View exception detail**.

3. Select **Delete**.
4. Provide a reason for deletion, and an approver if applicable, then select **Delete**.

## CIS scans

---

Run your desired compliance scan or an ad hoc scan on your nodes.

*Ad hoc*, or manual, scans give you the freedom to run a CIS scan on your network straight away. Alternatively, you can schedule scans to run on a regular basis and at times when network traffic is low.

Scan reports provide you with in depth information on not just the timing of scans but the nodes affected and the rules that passed or failed.

- [Run an ad hoc scan](#) on page 161

Run your desired compliance scan or an ad hoc scan on your nodes.

- [Scheduled scans](#) on page 162

In addition to manual ad hoc scans, you can also schedule scans. A scan can be scheduled to run either once or at user-defined intervals.

- [CIS scan reports](#) on page 165

The **Scan reports** page displays all scans run within the defined scan data retention period.

## Run an ad hoc scan

---

Run your desired compliance scan or an ad hoc scan on your nodes.

**Note:** Should you prefer to run ad hoc scans using the Comply API, you can use one of the following [endpoints](#):

- **Custom Scan API.** Run an ad hoc scan with custom CIS profiles.
- **Scan API.** Run an ad hoc scan with a provided CIS profile.

1. In Comply, click **Scan reports** and then **Run an ad hoc scan**.
2. In the drop-down menu, select **Desired compliance** or **Custom**.  
If you have not set desired compliance, follow the instructions in [Setting desired compliance](#).
3. If you selected **Custom**, select a benchmark from the **Benchmark** drop-down menu, then select an option from the **Profile** drop-down menu. To use a custom profile for this scan, select the **Use an associated custom profile?** option and choose the relevant option from the **Custom profile** drop-down menu.
4. Click **Next** to see the nodes selected for scanning. Use the drop-down menus to filter nodes by operating system, environment, or node group.

To scan only a subset of nodes, deselect any nodes that you want to exclude.

**Debug mode:** By default, assessor logs are set to WARN level. To troubleshoot an issue, you can set the logging level to DEBUG for the scan by clicking **Run in debug mode**. The assessor logs can then be retrieved from the individual node.

On Linux and macOS platforms the assessor log is located at:

```
/opt/puppetlabs/comply/Assessor-CLI/logs/assessor-cli.log
```

On Windows the assessor log is located at:

```
C:/ProgramData/PuppetLabs/comply/Assessor-CLI/logs/assessor-cli.log
```

Note that scanning in debug mode increases the size of the assessor log file significantly.

## 5. Click **Scan**.

You are taken to the **Activity feed**, which lists each scan. Scans are run as a task in PE. Click the scan name to see the scan report, or click the job ID to be taken to PE.

## 6. Optionally, to review the results of your scan, navigate to the **Compliance Dashboard** page.

See [Scan results](#) for a description of the scan data.

### Related information

[Enforce CIS benchmarks](#) on page 169

Puppet Comply provides visibility into your compliance status, but it cannot fix your failing nodes. Instead, you can use Puppet's Compliance Enforcement Modules (CEM).

[Custom profiles](#) on page 158

A custom profile is a benchmark profile that you customize to fit your organization's internally defined standards. You can base a custom profile on an existing benchmark and profile combination, and then specify which rules to apply.

[Desired compliance](#) on page 156

Set your desired compliance. This is the benchmark and profile that you assign to a particular node and that is scanned on that node by default. Generally, you set compliance only once for your nodes.

## Scheduled scans

---

In addition to manual ad hoc scans, you can also schedule scans. A scan can be scheduled to run either once or at user-defined intervals.

From the **Scheduled scans** page you can:

- View scan schedules
- Create and edit scans
- Pause and resume scans
- Add or remove nodes in scheduled scans
- End and delete scans.

Comply gives you many options for scan schedules. You can schedule a one-off scan or a scan that repeats periodically:

- You can set scans to run on individual days of the week, every day, weekdays only, or weekends only.
- You can set scans to run on any day of the month. Be aware that if you, for example, set a scan to run on day 31, the scan does not run in months that do not have 31 days.

### Scheduled scans table

The table on the **Scheduled scans** page provides the following information:

- **Name** - The name assigned to the scan.
- **Nodes** - The number of nodes to be scanned.
- **Start** - The date and time scheduled for the scan to start according to the selected time zone.
- **End** - The date and time scheduled for the scan to finish according to the selected time zone.
- **Frequency** - How often the scan is scheduled to run.
- **Last run** - Date and time when the last scan of this type started.
- **Status** - Whether the scan is not currently running (*Not started*), has completed processing (*Ended*), or is currently running (*Active*).

### View details about a scan schedule

You can view details about a scan schedule in the **Scheduled scan information** window.

To view details about a scan schedule:

1. On the **Scheduled scans** page, click the schedule you wish to see details about from the **Scheduled scans table**.
2. In the **Scheduled scan information** window, click **View detail**.

On the **Scheduled scan detail** page, you can view details about the node selection method, scan schedule, scan history, and affected nodes.

## Pause and resume a scan schedule

You can pause or resume a scan schedule by using the **Scheduled scan information** window.

To pause and resume a scan schedule:

1. On the **Scheduled scans** page, click the schedule you want to pause or resume from the **Scheduled scans table**.
2. In the **Scheduled scan information** window, click **View detail**.
3. To pause a scheduled scan, click **Pause**.
4. To resume a paused scan, click **Resume**.

You can also end a scan by clicking **End**.

**Restriction:** After you click **End**, the scan cannot be restarted.

## Edit a scan schedule

You can edit a scan schedule by using the **Scheduled scan information** window.

To edit a scan schedule:

1. On the **Scheduled scans** page, click the schedule you wish to edit from the **Scheduled scans table**.
2. In the **Scheduled scan information** window, click **View detail**.
3. Click **Edit details**.
4. In the **Edit scan schedule** window, update the details of the scan. For example, you can change the schedule type, the frequency, the start date and time, and the end date.

**Restriction:** You cannot add or remove nodes from the scan schedule.

5. To save your changes, click **Save**.

## Delete a scan schedule

You can delete a scan schedule in any state (active, paused, or ended) by using the **Scheduled scan information** window.



**CAUTION:** After you delete a scan schedule, the scan cannot be restarted. However, scan reports that were generated by the scan remain available.

To delete a scan schedule:

1. On the **Scheduled scans** page, click the schedule you want to delete from the **Scheduled scans table**.
2. In the **Scheduled scan information** window, click **View detail**.
3. Click **Delete**.
4. When you are prompted to confirm the deletion, click **Delete scheduled scan**.

## Add or remove nodes in scheduled scans

You can add or remove the nodes in scan schedule by using the **Scheduled scan information** window.

To add or remove nodes in a scan schedule:

1. On the **Scheduled scans** page, click the schedule you wish to edit from the **Scheduled scans table**.
2. In the **Scheduled scan information** window, click **View detail**.

3. Click **Manage nodes** or **Edit Selection** from the list of nodes at the bottom of the page, depending on the node selection method. If the scheduled scan has a static node selection method, **Manage nodes** appears. If the scheduled scan has a dynamic node selection method, **Edit Selection** appears.
4. In the **Edit node selection** (or **Edit node group selection**) window, review and change the nodes in the scheduled scan.
  - a) If you chose **Create a static node list**, specify the nodes to scan by selecting the relevant checkboxes. Use the drop-down menus to filter nodes by operating system, desired compliance, or node group.
  - b) If you chose **Target nodes dynamically**, specify the node groups to scan from the **Select node groups** drop-down menu. You may select more than one node group to include in the scheduled scan. Click **Apply** when you are done selecting which node groups to include.
5. To save your changes, click **Save and exit**.

## Create a one-off scan schedule

You can schedule a compliance scan to run once.

To create a one-off scan schedule:

1. On the **Scheduled scans** page, click **Create scheduled scan**.
2. Choose node selection method. Specify method used to select nodes and click **Next**. The selection methods are:
  - **Create a static node list**: Choose this selection method to explicitly pick each node to include in the scheduled scan. After you schedule a scan, if you onboard new nodes and want those nodes to be included in the scan, you must manually update the scheduled scan by adding the new nodes to the list.
  - **Target nodes dynamically**: Choose this selection method to target nodes from your specified node groups. Scheduled scans are automatically updated to add or exclude nodes based on changes in node group membership.
3. Select the nodes to include in the scheduled scan.

**Note:** If desired compliance is not set on any of the nodes, you are prompted to select a suggested desired compliance profile from the existing options or exclude those nodes from the scan.

- a) If you chose **Create a static node list**, specify the nodes to scan by selecting the relevant checkboxes. Use the drop-down menus to filter nodes by operating system, desired compliance, or node group.
- b) If you chose **Target nodes dynamically**, specify the node groups to scan from the **Select node groups** drop-down menu. You may select more than one node group to include in the scheduled scan. Click **Apply** when you are done selecting which node groups to include.
- c) Click **Next** when you have finished selecting nodes.
4. Set the schedule:
  - a) Select **Once only**.
  - b) Specify the date for your scan schedule.
  - c) Specify the time and time zone for starting the scan schedule.
  - d) Click **Next**.
5. Add details and review. Enter a name in the **Scheduled scan name** field and, optionally, enter information in the **Description** field. Click **Create scheduled scan**.  
The scan schedule information is listed on the table on the **Scheduled scans** page.

## Create a repeating scan schedule

You can schedule a compliance scan to run regularly.

To create a scan schedule that repeats:

1. On the **Scheduled scans** page, click **Create scheduled scan**.

2. Choose node selection method. Specify method used to select nodes and click **Next** . The selection methods are:
  - **Create a static node list:** Choose this selection method to explicitly pick each node to include in the scheduled scan. After you schedule a scan, if you onboard new nodes and want those nodes to be included in the scan, you must manually update the scheduled scan by adding the new nodes to the list.
  - **Target nodes dynamically:** Choose this selection method to target nodes from your specified node groups. Scheduled scans are automatically updated to add or exclude nodes based on changes in node group membership.

3. Select the nodes to include in the scheduled scan.

**Note:** If desired compliance is not set on any of the nodes, you are prompted to select a suggested desired compliance profile from the existing options or exclude those nodes from the scan.

- a) If you chose **Create a static node list**, specify the nodes to scan by selecting the relevant checkboxes. Use the drop-down menus to filter nodes by operating system, desired compliance, or node group.
  - b) If you chose **Target nodes dynamically**, specify the node groups to scan from the **Select node groups** drop-down menu. You may select more than one node group to include in the scheduled scan. Click **Apply** when you are done selecting which node groups to include.
  - c) Click **Next** when you have finished selecting nodes.
4. Set the schedule:
    - a) Select **Repeating**.
    - b) Select the desired **Frequency** of the scan:
      - If you select the **Weekly** option, you can specify a day of the week for a weekly scan, or you can specify that the scan runs daily, only on weekdays, or only on weekends.
      - If you select the **Monthly** option, specify a day of the month for the scan.

**Remember:** Your scan does not run if the month does not include the day you selected. For example, if you choose day 31, your scan cannot run in those months that do not have 31 days.

- c) Choose the time and time zone when you want the scan to run.
  - d) Optionally, add an **End Date** for the scan schedule.
 

If you choose an end date for your scan schedule, no more scans run on or after that date and no scan reports are produced.
  - e) Click **Next**.
5. Add details and review. Enter a name in the **Scheduled scan name** field and, optionally, enter information in the **Description** field. Click **Create scheduled scan**.
 

The scan schedule information is listed on the table on the **Scheduled scans** page.

## CIS scan reports

---

The **Scan reports** page displays all scans run within the defined scan data retention period.

The **Scan reports** page provides the following information:

- **Name** - the name given to the scan when it was run.
- **Scan type** - ad hoc or scheduled.
- **Nodes scanned** - the total number of nodes included in the scan.
- **Compliance** - the percentage of nodes that passed compliance in the scan.
- **Time started** - the date and time stamp when the scan was initiated.

Click the row assigned to any scan to go to its **Scan report** page, which provides details.

For more information on defining the scan data retention period, see [Scan data retention policy](#) on page 169.

## Related information

[CIS scan report details](#) on page 166

The **Scan report** page provides detailed information on a selected CIS scan.

[Scan results](#) on page 166

View the results of your CIS scans and find out whether your nodes are compliant.

## CIS scan report details

The **Scan report** page provides detailed information on a selected CIS scan.

The metrics bar at the top of the **Scan report** page is divided into two sections: **Compliance scan status** and **Puppet Enterprise job status**.

The **Compliance scan status** section provides a brief overview of the number of nodes that have passed and failed compliance, the error percentage, the rules that couldn't be evaluated across nodes, and the scan initiation date and time.

The **Puppet Enterprise job status** section shows the number of nodes that ran the CIS scanner job successfully, the number that failed to run the scanner job, and the number of nodes that showed an error for the scanner job.

On the **Scans** page, you can click **Run an ad hoc scan** to kick off a new scan.

More detailed information on the success and failure of rules is given on the **Scan report** page on the **Nodes** tab. The **Rules** tab provides detail on the performance of individual rules in the scan.

### Rules tab

On the **Scan report** page, the table on the **Rules** tab lists all the rules that were assessed as part of the latest scan. The table provides information on the rule profile and the number of nodes on which the rule failed. To view details about a rule, locate the row that is associated with that rule and click **View report**.

### Nodes tab

The table on the **Nodes** tab lists all nodes that were part of the latest scan. The table provides information on the node profile, and the percentage of rules in compliance on each node. To view details about a node, locate the row that is associated with that node and click **View report**.

## Scan results

---

View the results of your CIS scans and find out whether your nodes are compliant.

### Scan report metrics bar

On the **Scan report** metrics bar, the **Compliance scan status** section displays the compliance score with any applicable exceptions. The metrics bar also displays the percentage of nodes that passed, failed, or could not be evaluated, and the scan initiation date and time. The **Puppet Enterprise job status** section displays the status of scan jobs in Puppet Enterprise.

### Compliance Dashboard

The **Compliance Dashboard** provides a breakdown of your latest CIS scan.

The dashboard can be filtered by environment, operating system, and node group in order to show subsets of your infrastructure. It has several widgets containing information about your nodes and exceptions:

- **Number of scanned nodes:** the number of nodes included in your last scan. Clicking this widget takes you to the **Node results** page.
- **Number of nodes in inventory:** clicking this widget takes you to the **Inventory** page.

- **Active exceptions:** clicking this widget takes you to the **Exceptions** page.
- **Nodes added:** the number of nodes added in the specified time frame. Use the arrow buttons to change the time frame.
- **Overall compliance score over time:** a graph of changes in your compliance score. Use the drop-down list to change the specified time frame.
- **Compliance score with exceptions applied** and **Compliance score without exceptions applied:** clicking this widget takes you to the **Node results** page.
- **Nodes without desired compliance applied:** clicking this widget takes you to a filter on the **Inventory** page.
- **Expiring exceptions:** clicking this widget takes you to a filter on the **Exceptions** page.
- **Out-of-date assessors:** clicking this widget takes you to a filter on the **Inventory** page.
- **Unscanned nodes:** clicking this widget takes you to the **Run a desired compliance scan** page.
- **5 least compliant nodes:** clicking a node name within this widget takes you to the **Node detail** page for that node.

## Node compliance

From the **Node results** page, click a node name to navigate to the **Node detail** page and see the results of the latest scan on that node:

- The **Scan status** pane shows a status breakdown for the latest scan, including the total number of rules and the number of rules that passed, failed, reported an error, or had an unknown status. You can hover over the statuses in the legend to see percentages in the donut chart. The chart and legend reflect only the statuses that are subject to scoring. Non-scoring statuses (for example, cases in which a recommendation is not applicable or cannot be automatically assessed) are excluded. Statuses are described in the following table:

| Value          | Included in scoring? | Description                                                                                                                                                                                                                                                                                                                                                                       |
|----------------|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Pass           | Yes                  | The target system or component state satisfied all the conditions of any checks or rules for the recommendation.                                                                                                                                                                                                                                                                  |
| Fail           | Yes                  | The target system or component state did not satisfy at least one condition of any checks or rules for the recommendation.                                                                                                                                                                                                                                                        |
| Error          | Yes                  | The assessor checking engine encountered a system error and could not complete the test. The status of the target's compliance is not certain.                                                                                                                                                                                                                                    |
| Unknown        | Yes                  | The assessor was unable to collect, interpret, or evaluate against any check or rule conditions associated with the recommendation.                                                                                                                                                                                                                                               |
| Other          | No                   | The Other status includes all statuses that do not fall into the categories of Pass, Fail, Error, or Unknown. For details about the statuses that are included in the Other category, see the following rows.                                                                                                                                                                     |
| Manual         | No                   | This recommendation cannot be fully automated and requires manual evaluation. This status occurs when, in the CIS Benchmarks, a recommendation is deemed important but cannot be fully and reliably verified without a manual check by an organization. This status corresponds to the Extensible Configuration Checklist Description Format (XCCDF) term, <i>Informational</i> . |
| Not Applicable | No                   | Rules, checks, or both were not applicable to the target. This situation typically occurs when the benchmark and platform are mismatched.                                                                                                                                                                                                                                         |

|               |    |                                                                                                                                                      |
|---------------|----|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Not Checked   | No | The recommendation was not evaluated as there are no rule or check properties.                                                                       |
| Not Selected  | No | This recommendation was not part of the profile selected for the configuration assessment.                                                           |
| Informational | No | This is the same result that is displayed as Manual on the HTML report. The recommendation cannot be fully automated and requires manual evaluation. |

- The **Rule scan results** table lists each rule that was checked and the status of that rule from the latest scan. The table also shows the date and time of the last successful scan for each rule.

## Rule results

From the **Node results** page, in the **Node results** table, click a node. Then, in the **Rule scan results** table, click a rule. The **Rule detail** page includes the following information:

- The **Scan status** pane shows the total number of nodes scanned and detailed results. You can hover over the results to see percentages in the donut chart. The compliance score in the chart and legend reflects only the statuses that are subject to scoring. Non-scoring statuses (for example, cases in which a recommendation is not applicable or cannot be automatically assessed) are excluded.
- A tabbed section displays information about each rule:
  - **Fix** — the steps you can take to fix the rule if it is failing on a node.
  - **Description** — information on what is being checked.
  - **Rationale** — the reason why it is important to check that rule.
- The **Node results** table lists each node the rule has been checked against and shows the current status, including when the node was last checked and when it last passed that rule. The table shows the profile, the environment in which the scan took place (for example, production or test), and any exceptions that apply.
- The **Exceptions** tab displays any exceptions that are relevant to the selected rule.

## Exporting results

To export your results as a .csv file, select **Export CSV** at the top right of the **Node results** tab, and then choose whether to export raw data or a report summary. After exporting, you can download past reports from the **Generated reports** tab in the left menu.

The raw data export contains detailed scan results for each rule, including the rule's name, ID, and status, whether the rule has an exception against it, and details about the exception if applicable.

Rather than raw data, the summary export provides an exception score and an adjusted compliance score for each rule. The exception score is the latest overall compliance score for all nodes. This score accounts for any temporary compliance rule exceptions in place, and any rules with exceptions are excluded from the overall compliance score. The adjusted compliance score does not account for any temporary compliance rule exceptions, instead providing a true compliance score for all nodes.

## Scan rule report

You can view a report about scan results for a single rule. The **Scan report: Rule performance** page lists the nodes on which the rule was run and the results.

From the **Scans** page, click a scan report. Ensure that the **Rules** tab is displayed. Locate a rule in the table and click **View report**.

The data includes:

- Overall compliance status for the nodes on which the rule was run
- The date and time when the scan was started
- The scan status for each node, including an indication of whether exceptions apply

## Scan node report

You can view a report about scan results for a single node. The **Scan report: Node performance** page lists the rules that were run on the node and the results.

From the **Scans** page, click a scan report. Ensure that the **Nodes** tab is displayed. In the table, locate a node and click **View report**.

The data includes:

- Overall compliance status for the node
- The date and time when the scan was started
- The scan status for each rule, including an indication of whether exceptions apply

**Tip:** If you created one or more exceptions to a rule, you must then run a scan to ensure that the compliance score correctly reflects the exceptions.

## Scan data retention policy

By default, no retention period is defined and Comply retains scan data indefinitely. You can, however, [enable this feature](#) in **Settings > Configure your install** in Comply. We recommend that you run no more than 1 complete scan per week and retain data for no more than 14 weeks at a time. If you do plan to retain historical data for more than 14 complete scans, increase **Comply PostgreSQL capacity** in PAM by approximately 3GB per additional scan.

# Enforce CIS benchmarks

---

Puppet Comply provides visibility into your compliance status, but it cannot fix your failing nodes. Instead, you can use Puppet's Compliance Enforcement Modules (CEM).

Available to [premium content subscribers](#), CEM consists of two modules — `cem_linux` and `cem_windows`. These are supported Puppet modules developed specifically to bring your Puppet Enterprise (PE) managed nodes under CIS (Center for Internet Security) compliance.

By default, CEM enforces the latest CIS Level 1 benchmarks on your nodes, automating hundreds of operating system settings — the default profile depends on your operating system. You can also customize these configurations to suit your organization's policies.

**Tip:** Starting with CEM for Linux 1.4.0, CEM also enforces the Security Technical Implementation Guides (STIG) developed by the US Defense Information Systems Agency (DISA). The DISA STIG standards, widely used by US government agencies, can be enforced by CEM on Red Hat Enterprise Linux 7 and 8 operating systems.

To get started with CEM, see [Introducing the Compliance Enforcement Modules](#) on page 173.

# Troubleshooting

---

Use this section to troubleshoot issues with your Puppet Comply installation.

## Reset your Comply password

---

If you forget your password, you can reset it in the user admin console.

1. SSH into your Comply node and run the following commands to retrieve the admin username and password:

```
kubectl exec $(kubectl get pod -l app.kubernetes.io/name=comply-auth
-o jsonpath="{.items[0].metadata.name}") -- /bin/bash -c 'env | grep
KEYCLOAK_ADMIN='
```

```
kubectl exec $(kubectl get pod -l app.kubernetes.io/name=comply-auth
-o jsonpath="{.items[0].metadata.name}") -- /bin/bash -c 'env | grep
KEYCLOAK_ADMIN_PASSWORD='
```

2. Navigate to `https://<COMPLY-HOSTNAME>/auth/admin` using the FQDN of your Comply node.
3. Login using the credentials from step 1.
4. Select **Puppet Realm** from the dropdown menu.
5. Navigate to **Users**.
6. Use an asterisk(\*) as a search term to show all users and select the user account you want to update.
7. Click **Edit** to edit the selected user account.
8. Select the **Credentials** tab and then reset the password.

## Access logs

---

If you run into issues with Puppet Comply, you can download the relevant log files. The Comply logs are stored in Puppet Application Manager.

1. Log into Puppet Application Manager — `https://<PUPPET-APPLICATION-MANAGER-ADDRESS>:8800`.
2. Select the **Troubleshoot** tab, and click **Analyse Comply**.
3. Download the bundle of log files.

## Resolve the Comply domain

---

If the Puppet Comply gatekeeper is unable to resolve the Comply domain, try the following troubleshooting steps.

When you assign a hostname to Comply, it needs to be resolved by the pods in your Kubernetes cluster. A preflight check verifies the domain you specified in the configuration is resolvable. You must ensure that the nodes can resolve their own hostnames, through either local host mapping or a reachable DNS server.

1. To verify your whether your hostname is resolvable, run the following commands:

```
kubectl exec $(kubectl get pod -l app=kotsadm -o
jsonpath="{.items[0].metadata.name}") -- /bin/sh -c 'curl --SI
<hostname>'
```

If the hostname was resolved, the command returns an exit code 0 with no output.

If the hostname cannot be resolved, the command returns an exit code 6. Proceed to step 2 to add DNS entries.

2. To add DNS entries for CoreDNS, run the following command to open the CoreDNS configuration maps:

```
kubectl -n kube-system edit configmaps coredns
```

3. Add a `hosts` entry below `kubernetes`. This is where you configure the DNS entry for Comply. For example:

```
kubernetes cluster.local in-addr.arpa ip6.arpa {
 pods insecure
 fallthrough in-addr.arpa ip6.arpa
 ttl 30
}
hosts {
 10.23.24.25 comply.mycompany.net comply // IP_address canonical_hostname
 [aliases...]
 fallthrough
}
prometheus :9153
```

4. Run the command from step 1 to verify whether the DNS entry was updated:

```
kubectl exec $(kubectl get pod -l app=kotsadm -o
 jsonpath="{.items[0].metadata.name}") -- /bin/sh -c 'curl --SI
 <hostname>'
```

5. Re-run the preflight checks.

## Resolve a failed assessor upgrade

---

If an upgrade of the assessor has failed on one of your nodes, try the following troubleshooting step.

If the upgrade of an assessor on a node fails, the node is marked in red on the **Inventory** page. Failures may be due to network issues. If that is the case, Comply attempts to upgrade the node once connectivity returns. An hourly background task runs to check if nodes have been upgraded or not. If a node does not upgrade and remains red on the **Inventory** page, run the Puppet agent. If the upgrade continues to fail, see the Puppet agent logs for more information.

## Resolve a failed scan

---

If an inappropriate version of Java Runtime Environment (JRE) is installed on the host system where the CIS-CAT Pro Assessor resides, you might see an error message about a failed scan.

The error message is similar to the following example:

```
Error: Scan did not complete successfully 'java _Server -b
/opt/puppetlabs/comply/Assessor-CLI/benchmarks/
CIS_Red_Hat_Enterprise_Linux_7_Benchmark_v3.1.1-xccdf.xml -D
cisat.license.filepath=/opt/puppetlabs/comply/Assessor-CLI/license/
license.xml', 'Exception in thread "main"
java.lang.UnsupportedClassVersionError: org/cisecurity/assessor/cli/
Assessor : Unsupported major.minor version 52.0 at
 java.lang.ClassLoader.defineClass1(Native Method) at
 java.lang.ClassLoader.defineClass(ClassLoader.java:808) at
 java.security.SecureClassLoader.defineClass(SecureClassLoader.java:142) at
 java.net.URLClassLoader.defineClass(URLClassLoader.java:443) at
 java.net.URLClassLoader.access$100(URLClassLoader.java:65) at
 java.net.URLClassLoader$1.run(URLClassLoader.java:355)
at java.net.URLClassLoader$1.run(URLClassLoader.java:349) at
 java.security.AccessController.doPrivileged(Native Method)
at java.net.URLClassLoader.findClass(URLClassLoader.java:348) at
 java.lang.ClassLoader.loadClass(ClassLoader.java:430)
at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:326) at
 java.lang.ClassLoader.loadClass(ClassLoader.java:363) at
 sun.launcher.LauncherHelper.checkAndLoadMain(LauncherHelper.java:482)
```

```
N.B. The java.lang.UnsupportedClassVersionError: org/cisecurity/assessor/
cli/Assessor : Unsupported major.minor version
52.0
```

To resolve the issue, ensure that JRE v1.8 or later is installed. For the latest information about JRE requirements, see the [CIS-CAT Pro Assessor Configuration Guide](#).

## Troubleshooting TLS issues in Comply

---

Incorrect configuration of TLS certificates when setting up Comply to work with PE can cause agents and/or scan tasks to fail.

There are two main certificate configuration errors that cause problems:

- If the CA certificate configured for Comply in Puppet Application Manager does not match the CA certificate the Puppet Enterprise certificate is signed with, then a trust store issue is returned upon setting up PE or trying to run a scan.
- A hostname issue can arise if a server identity check does not pass. The `dnsName` in the `subjectAltName` of the server certificate must match the hostname Comply is communicating with. The hostname configured in the Comply settings page for Puppet Enterprise must match one of the `dnsName` entries in the PE certificate.

### Troubleshooting agent issues

Agents can fail if the certificate is configured incorrectly in two ways:

- The hostname could be the issue - the Comply TLS certificate must have the `dnsName` with which the agent is trying to contact the Comply server. This is relevant only if you set up the Comply module to download the assessor from the Comply server. The hostname the Comply server is contacted with is the hostname in the configured `scanner_source` parameter URL.
- The trust store may also be the problem if the Comply server TLS certificate and the client certificate are not signed by the same CA.

### Troubleshooting scan task issues

Scan tasks can also fail if the certificate is configured incorrectly in two ways:

- If the hostname is incorrectly configured - The Comply TLS certificate must have the `dnsName` with which the agent is trying to contact the Comply server. The hostname used is passed through by the task and can be seen in the PE UI by checking the task parameters. Verify the task parameters to ensure that the hostname is correct.
- If the Comply server TLS certificate and the client certificate are not signed by the same CA a trust store issue occurs and this too can cause the scan task to fail.

## Troubleshoot TLS from a node

You can run a simple `cURL` command to troubleshoot TLS issues from a node.

To troubleshoot TLS from a node:

1. SSH into a node and change to the SSL subdirectory:

```
cd /etc/puppetlabs/puppet/ssl/
```

- Issue the following cURL command replacing values in angle brackets with values relevant to your installation:

```
curl -G --key private_keys/<local host key> --cacert certs/ca.pem --
cert ../certs/<local host cert> https://<comply-fqdn>:30303/assessor --
output /tmp/assessor.zip
```

**Note:** If you are using your own ingress, issue the following command:

```
curl -G --key private_keys/<local host key> --cacert certs/ca.pem --
cert ../certs/<local host cert> https://<PE TLS hostname>/assessor --
output /tmp/assessor.zip
```

If this command fails, there is an issue with certificates. The error message can help to identify if there is a CA, client certificate or hostname issue.

- If there is a hostname error, the output resembles the following error message:

```
curl: (60) SSL: no alternative certificate subject name matches target
host name '<comply-fqdn>'
```

- If it is a CA issue, the output is likely to be similar to the example below:

```
curl: (60) SSL certificate problem: unable to get local issuer
certificate
```

**Important:** This may be because the CA of the `https://<comply-fqdn>:30303/assessor` (or `https://<PE TLS hostname>/assessor`) does not match the CA passed to the cURL command or because there is a mismatch with the CA of the client certificate.

## Introducing the Compliance Enforcement Modules

The Puppet Compliance Enforcement Modules (CEM) were developed to enforce the secure configuration of IT infrastructures and thus protect operations and data. You can use CEM with Puppet Enterprise (PE) or open source Puppet. You can enforce the Center for Internet Security (CIS) compliance rules, which embody internationally recognized standards. You can also enforce the Security Technical Implementation Guides (STIGs) developed by the US Defense Information Systems Agency (DISA). DISA STIG standards are implemented by many US government agencies.

**Important:** In May 2024, CEM was renamed to Security Compliance Enforcement (SCE). For information about the SCE modules, including new features and fixes, see the [Release notes for Linux](#) and the [Release notes for Windows](#).

After you install and configure CEM, PE or open source Puppet runs on any classified nodes without user intervention to enforce compliance. By default, CEM enforces CIS rules for the Level 1 profile. However, you can enforce a variety of security standards and levels, depending on the operating system of the nodes where your servers and workstations are installed. For a list of supported standards for Linux nodes, see [Prepare to install the module](#) on page 193. For a list of supported standards for Microsoft Windows nodes, see [Prepare to install the module](#) on page 222.

The following sections provide instructions for installing CEM and customizing the configuration settings, if necessary, to meet your organization's requirements.

Separate documentation is provided for Linux nodes and for Windows nodes:

- To manage Linux nodes, see [CEM for Linux](#) on page 174.
- To manage Windows nodes, see [CEM for Windows](#) on page 212.

# CEM for Linux

---

You can deploy the Compliance Enforcement Module (CEM) for Linux to help ensure that your servers on Linux operating systems comply with security recommendations. You can enforce the controls that are specified by the Center for Internet Security (CIS). Alternatively, you can apply the standards published in the US Defense Information Systems Agency (DISA) Security Technical Implementation Guides (STIGs).

To get started, review the basic concepts and then follow the instructions to deploy CEM in your environment. See [Getting started](#) on page 190.

- [Release notes](#) on page 174

Review the release notes to learn about updates and resolved issues in the Compliance Enforcement Module (CEM) for Linux.

- [Getting started](#) on page 190

Learn the basic concepts associated with the Compliance Enforcement Modules and then review the steps for deploying CEM in your environment.

- [Installing CEM](#) on page 192

Before you install CEM, complete the preparation steps: review the system requirements, install and configure Puppet Enterprise (PE) or open source Puppet, and purchase CEM. To help avoid issues, install and evaluate CEM in a *test* environment before you install CEM in a production environment.

- [Upgrading CEM](#) on page 196

You can upgrade CEM for Linux to take advantage of the latest features, fixes, and improvements. To help ensure a smooth upgrade process, complete the preparation tasks first.

- [Configuring CEM](#) on page 198

Configuration of CEM is optional. If you installed CEM and assigned the `cem_linux` class to one or more node groups, the Center for Internet Security (CIS) Server Level 1 profile is enforced automatically during the next Puppet run. However, if the default values leave your infrastructure in an undesirable state, or if you want to customize compliance to meet your organization's requirements, you can configure CEM.

- [Auditing and querying issues identified during scans](#) on page 209

In some cases, a CIS or DISA STIG compliance scan might identify an issue that you want to investigate and fix. To get started, you can run an audit or query.

- [Reference: Benchmarks and controls](#) on page 210

For help with configuring CEM, review the [CEM Linux Reference](#) on Puppet Forge.

## Release notes

---

Review the release notes to learn about updates and resolved issues in the Compliance Enforcement Module (CEM) for Linux.

### v1.9.1

Released 8 February 2024

CEM for Linux v1.9.1 improves operational efficiency by correcting an issue where Puppet runs fail prematurely in certain scenarios. Furthermore, this release enhances security safeguards by correcting issues that could have prevented three Security Technical Implementation Guide (STIG) controls from being enforced as intended.

#### Fixed

- Fixed an issue that could potentially cause Puppet runs to fail prematurely on Red Hat Enterprise Linux (RHEL) 8 systems. The issue is related to the Factor `fact_cem_mount_info`, which can fail to resolve when home directories cannot be determined for a user on a system. The fix helps to ensure that the correct default directory is used, and the fact resolves successfully.

- Fixed the default value for STIG Control V-230270 to help prevent kernel profiling by unauthorized users. The kernel parameter `kernel.perf_event_paranoid` is now set to a default value of 2 to help prevent attackers from gaining system information.
- Fixed an issue that could cause incorrect failure reports for STIG Control V-230281. The control helps to ensure that the RHEL operating system removes previously installed software components when updated versions are installed. In CEM, the specified default value of `true` was changed to `True` to ensure that the control works as designed.
- Fixed the default value for STIG Control V-230494. The control helps to protect the security of systems by disabling the Asynchronous Transfer Mode (ATM) protocol. The default value was changed from `ATM` to `atm` so that the control works as designed and helps to protect systems from exploitation.

## v1.9.0

Released 14 December 2023

CEM for Linux v1.9.0 introduces extended coverage with support for two new operating systems -- Red Hat Enterprise Linux (RHEL) 9 and Oracle Linux 9. This release also improves operational efficiency by introducing an automated audit process and instant access to audited controls for result review.

### Added

- Introduced support for the RHEL 9 operating system. You can now enforce the Center for Internet Security (CIS) RHEL 9 Benchmarks, Levels 1 and 2, to help ensure a secure configuration for your RHEL 9 nodes.
- Introduced support for the Oracle Linux 9 operating system. You can now protect your Oracle Linux 9 nodes by enforcing the CIS Benchmarks at Levels 1 and 2.
- To improve the efficiency of the auditing process, you can now run a single Puppet Bolt® plan that includes 40 audit tasks. The Bolt plan, `run_audit`, can be run on one or more specified nodes to verify their configuration. Afterward, the Bolt log file provides a list of audited controls and detailed results. You can still run [audit tasks](#) separately, as supported in earlier releases.

### Changed

- To help ensure compatibility between CEM for Linux and Puppet 8, the range of supported versions for the Puppet Labs® `firewall` module was changed. The [firewall module](#) must be at version 5.0 or later, but earlier than 6.0.

## v1.8.0

Released 24 October 2023

### Added

- Introduced support for enforcing Center for Internet Security (CIS) Benchmarks on the Rocky Linux operating system. With CEM for Linux v1.8.0, you can enforce the CIS Rocky Linux 8 Benchmarks, Levels 1 and 2. In this way, you can help to secure your Rocky Linux system and reduce the manual overhead associated with solution configuration.
- Added controls to help enhance security on Red Hat Enterprise Linux (RHEL) 8, AlmaLinux 8, and Oracle Linux 8 systems. The following controls can now be enforced:
  - Control 2.3.2 helps to ensure that remote shell (rsh) clients are not installed. The rsh program presents a security vulnerability because users who run rsh commands risk exposing their user credentials.
  - Control 2.3.3 helps to ensure that the Linux `talk` client is not installed. Talk software makes it possible for users to send and receive messages using unencrypted protocols.
  - Control 5.6.1.4 helps to ensure that inactive user accounts are locked automatically within 30 days of password expiration. Inactive accounts present a vulnerability because users are not logging in to check for failed login attempts and other anomalies.

- Added controls to help enhance security on Oracle Linux 7 systems. The following controls can now be enforced:
  - Control 1.4.1 ensures that a bootloader password is set. A user who restarts a system must enter a password before setting command-line boot parameters. This restriction helps to prevent unauthorized users from undermining system security.
  - Control 1.4.2 ensures that permissions are specified for the `grub` configuration file, which contains information about boot settings and passwords for unlocking boot options. This restriction helps to prevent unauthorized users from viewing and changing boot parameters.
  - Control 1.6.1.2 ensures that Security-Enhanced Linux (SELinux) is enabled at boot time and is not overridden by `grub` boot parameters.
  - Control 4.1.1.3 ensures that processes are audited even if they are running before the `auditd` service is started. Audit events must be recorded to detect potential malicious activity.
  - Control 4.1.2.4 helps to ensure that the audit backlog limit is sufficient to prevent audit records from being lost. If records from the `auditd` service are lost, malicious activity could go undetected.

#### Fixed

- Fixed an issue related to data protection in log files. The CIS Oracle Linux 8 Benchmark v2.0.0 includes Control 4.2.3, which requires the configuration of access permissions for log files. Previously, the control enforced access permissions only for non-hidden log files. With the fix, access permissions are enforced for both non-hidden and hidden log files.

#### v1.7.1

Released 29 September 2023

#### Fixed

- Fixed the command that automatically generates the [Reference](#) section on Puppet Forge. The section now includes information about the Security Technical Implementation Guide (STIG) controls that are enforced by CEM for Linux.

#### v1.7.0

Released 28 September 2023

#### Added

- An update was implemented in the `cem_mount_info` fact to help ensure that Puppet runs successfully. (The `cem_mount_info` fact uses the `homedir_mounts` function to determine the home directory mount paths from all mounted file systems. Previously, if the `homedir_mounts` function failed to return a value, the fact would have no value for home directories, and subsequent Puppet runs using the fact could fail. Now, if the `homedir_mounts` function fails to detect a value, the function verifies whether the `/home` directory is mounted and is a valid directory. If so, the function returns a value of `/home`.)

#### Changed

- This release includes updates that are designed to enhance security on Red Hat Enterprise Linux (RHEL) systems. By upgrading to CEM for Linux v1.7.0, you can enforce the latest US Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standards and reduce the manual overhead associated with solution configuration:
  - For the RHEL 8 operating system, the DISA STIG standard was upgraded from Red Hat Enterprise Linux 8 STIG, Version 1, Release 8, to Version 1, Release 11. For information about controls that were added or changed, see [Control updates introduced for Red Hat Enterprise Linux 8 STIG, Version 1, Release 11](#) on page 210.
  - For the RHEL 7 operating system, the DISA STIG standard was upgraded from Red Hat Enterprise Linux 7 STIG, Version 3, Release 8, to Version 3, Release 12. For information about controls that were added or changed, see [Control updates introduced for Red Hat Enterprise Linux 7 STIG, Version 3, Release 12](#) on page 211.

- The CEM for Linux documentation now provides more detailed upgrade instructions, including preparation steps that you can take to help ensure a smooth upgrade. See [Upgrading CEM](#) on page 196.

#### Fixed

- Fixed an issue related to communication security on RHEL 8 systems. STIG Control V-230507 is designed to disable the use of Bluetooth communication technologies because communications transmitted over Bluetooth can be intercepted. Previously, Bluetooth was not fully disabled. In this release, Bluetooth is fully disabled.
- Fixed an issue that could cause erroneous scan failures for users enforcing STIG Control V-204563. The STIG control ensures that audit records are generated when the `kmod` command is run to manage kernel modules. The audit rule for kernel module loading was fixed to prevent the erroneous scan failures.
- Fixed an issue that could cause scan failures for users enforcing STIG Control V-204579. The control is designed to enforce a timeout on RHEL 7 systems after a session is terminated or after 15 minutes of command-line inactivity by the user. The content of the timeout script, `999-tmout.sh`, was updated to align with the STIG control.
- Fixed an issue that could cause scan failures for users enforcing STIG Control V-204605. The control helps to ensure that the date and time of the previous logon is displayed at the next logon. In the `cem_linux::utils::postlogin` utility class, you can specify the `prune_pam_lastlog` parameter to remove the silent option from the `pam_lastlog.so` module entries. In this way, you can help to prevent erroneous scan failures.
- Fixed an issue that caused a failure to set kernel parameter values. This issue occurred when kernel parameter values contained valid shell operators, such as `|`. Specifically, when the STIG profile attempted to set `kernel.core_patterns = |/bin/false`, no value was set:

```
kernel.core_patterns =
```

- Fixed an issue that could cause non-idempotent Puppet agent runs in which the resources related to `grub` bootloader arguments are reset during each run. With the fix, `grub` bootloader kernel arguments that are managed by the `cem_grub_args` custom resource type are now correctly persisted to the `grub` configuration file.
- Fixed an issue that prevented some user-specified configuration options from being applied. The issue affected only some parameters on some controls.

### v1.6.3

Released 10 August 2023

#### Fixed

- Fixed an issue that caused Puppet run failures. The issue occurred when CEM for Linux set `grub2` bootloader arguments on systems that had "non linux entry" kernel entries on the `grubby --info=ALL` command. The issue affected users on all supported Linux operating systems.

### v1.6.2

Released 8 August 2023

#### Added

- The CEM for Linux code is updated to help ensure compatibility with Puppet 8. Compatibility with earlier Puppet releases remains unchanged.

#### Changed

- A default setting was changed to help ensure that audit logs are encrypted before being offloaded to a remote system. This change affects users who implement the US Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standard on Red Hat Enterprise Linux (RHEL) 8 operating systems. Because the default setting for STIG Control V-230478 is now `true`, the GNU Privacy Guard (GnuPG) utility is installed by default and the `auditd` service is enabled to help protect audit logs from unauthorized access.

#### Fixed

- Fixed an issue related to audit backlogs on RHEL 8 operating systems. The issue arose because Center for Internet Security (CIS) Control 4.1.1.4 was not enforced. The control helps to ensure that a sufficient number of audit records are retained in the backlog on system startup so that users can detect potential malicious activity. Starting with CEM for Linux v1.6.2, users can choose whether to enable CIS Control 4.1.1.4. If enabled, the control is enforced.
- Fixed an issue that affected users who enforce the DISA STIG standard on RHEL 8 operating systems. STIG Control V-230307 is designed to implement the `nodev` mount option to prevent unauthorized access from untrusted file systems. Previously, the control did not work as designed because the `nodev` option was not applied. The issue is fixed to help ensure that the control works as expected.
- Fixed an issue related to the Security-Enhanced Linux (SELinux) architecture, which controls system access. In previous releases, when a user changed the SELinux security setting, the change went into effect after a system restart. Starting with CEM for Linux v1.6.2, a user can specify that the updated setting is enforced immediately, and the change goes into effect. A user can also specify that the updated setting is enforced only after a system restart.

### v1.6.1

Released 16 May 2023

#### Fixed

- Fixed an issue that occasionally caused compliance scan failures for users who applied Center for Internet Security (CIS) controls in a Red Hat Enterprise Linux (RHEL) 7 environment. The issue occurred because CIS Control 4.1.15 'Ensure system administrator command executions sudo are collected' did not apply `auditd` configuration entries. The issue is fixed to ensure that scans run as expected.
- Fixed an issue that prevented CIS Control 4.1.3.7 'Ensure unsuccessful file access attempts are collected' from being enforced. The issue affected users on RHEL 8, AlmaLinux 8, and Oracle Linux 8 systems. The control is now enforced to ensure that information is collected about unsuccessful attempts to access files. These unsuccessful attempts can indicate that an unauthorized user or process is trying to gain access to the system.
- Fixed an error that could cause system operations to halt unexpectedly for users on RHEL 7. The issue was caused by an incorrect default setting for CIS Control 4.1.2.3 'Ensure system is disabled when audit logs are full.' The incorrect default setting could cause system operations to halt when disk space for audit logs runs low. The control is now correctly enforced as recommended by CIS: when disk space runs low, an email is sent to the address specified by the `'action_mail_acct'` parameter.
- Implemented a fix to help ensure that the use of privileged programs is monitored on all partitions on RHEL 7 operating systems. Previously, CIS Control 4.1.3.10 'Ensure use of privileged commands is collected' was not fully enforced. The control was only partially enforced because the `find` command did not descend into all partitions. The fix helps to ensure that all partitions are monitored so that unauthorized use of privileged programs is detected.
- Fixed an issue that occasionally caused the `10-cem_priv_commands.rules` file to be blank. This issue affected users on RHEL operating systems. The issue is fixed to ensure that CEM operates as expected when the user specifies a value of `true` for the `audit_privileged_commands` option:
  - If privileged commands exist, the `10-cem_priv_commands.rules` file displays the audit rule for that control.
  - If privileged commands do not exist or the privileged commands are ignored, a `10-cem_priv_commands.rules` file is not created.
- Fixed an issue that caused a failure to remove the Automatic Bug Reporting Tool (ABRT). In CEM for Linux v1.6.0, when the `remove_automatic_bug_report_tools` class was specified to remove ABRT, the tool remained installed. The problem is resolved to ensure that ABRT packages can be removed based on the class specification without further user intervention.

### v1.6.0

Released 28 March 2023

**Added**

- Enforcement of Center for Internet Security (CIS) Benchmarks on three new operating systems:
  - CIS Benchmark for AlmaLinux 8, v2.0.0, levels 1 and 2
  - CIS Benchmark for Oracle Linux 7, v3.1.1, levels 1 and 2
  - CIS Benchmark for Oracle Linux 8, v2.0.0, levels 1 and 2
- For users who enforce the US Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standard on Red Hat Enterprise Linux (RHEL) operating systems, CEM now supports several controls that are designed to secure graphical user interfaces (GUIs). The following new GUI controls are enforced:
  - Control V-204397 - The Red Hat Enterprise Linux operating system must uniquely identify and must authenticate users using multifactor authentication via a graphical user logon.
  - Control V-204398 - The Red Hat Enterprise Linux operating system must initiate a screensaver after a 15-minute period of inactivity for graphical user interfaces.
  - Control V-204399 - The Red Hat Enterprise Linux operating system must prevent a user from overriding the screensaver lock-delay setting for the graphical user interface.
  - Control V-204400 - The Red Hat Enterprise Linux operating system must prevent a user from overriding the session idle-delay setting for the graphical user interface.
  - Control V-204402 - The Red Hat Enterprise Linux operating system must initiate a session lock for the screensaver after a period of inactivity for graphical user interfaces.
  - Control V-204403 - The Red Hat Enterprise Linux operating system must prevent a user from overriding the screensaver idle-activation-enabled setting for the graphical user interface.
  - Control V-204404 - The Red Hat Enterprise Linux operating system must initiate a session lock for graphical user interfaces when the screensaver is activated.
  - Control V-204432 - The Red Hat Enterprise Linux operating system must not allow an unattended or automatic logon to the system via a graphical user interface.
  - Control V-204433 - The Red Hat Enterprise Linux operating system must not allow an unrestricted logon to the system.
  - Control V-204456 - The Red Hat Enterprise Linux operating system must be configured so that the x86 Ctrl-Alt-Delete key sequence is disabled in the Graphical User Interface.
  - Control V-214937 - The Red Hat Enterprise Linux operating system must prevent a user from overriding the screensaver lock-enabled setting for the graphical user interface.
  - Control V-219059 - The Red Hat Enterprise Linux operating system must disable the graphical user interface automounter unless required.
  - Control V-251718 - The graphical display manager must not be the default target on RHEL 8 unless approved.
  - Control V-230530 - The x86 Ctrl-Alt-Delete key sequence in RHEL 8 must be disabled if a graphical user interface is installed.
  - Control V-230553 - The graphical display manager must not be installed on RHEL 8 unless approved.
  - Control V-244519 - RHEL 8 must display a banner before granting local or remote access to the system via a graphical user logon.
  - Control V-244535 - RHEL 8 must initiate a session lock for graphical user interfaces when the screensaver is activated.
  - Control V-244536 - RHEL 8 must disable the user list at logon for graphical user interfaces.
  - Control V-244538 - RHEL 8 must prevent a user from overriding the session idle-delay setting for the graphical user interface.
  - Control V-244539 - RHEL 8 must prevent a user from overriding the screensaver lock-enabled setting for the graphical user interface.
  - Control V-230347 - RHEL 8 must enable a user session lock until that user re-establishes access using established identification and authentication procedures for graphical user sessions.
  - Control V-230351 - RHEL 8 must be able to initiate directly a session lock for all connection types using smartcard when the smartcard is removed.
  - Control V-230352 - RHEL 8 must automatically lock graphical user sessions after 15 minutes of inactivity.

- Control V-230354 - RHEL 8 must prevent a user from overriding the session lock-delay setting for the graphical user interface.
- Control V-230329 - Unattended or automatic logon via the RHEL 8 graphical user interface must not be allowed.
- Control V-230226 - RHEL 8 must display the Standard Mandatory DoD Notice and Consent Banner before granting local or remote access to the system via a graphical user logon.
- Control V-204393 - The Red Hat Enterprise Linux operating system must display the Standard Mandatory DoD Notice and Consent Banner before granting local or remote access to the system via a graphical user logon.
- Control V-204394 - The Red Hat Enterprise Linux operating system must display the approved Standard Mandatory DoD Notice and Consent Banner before granting local or remote access to the system via a graphical user logon.
- Control V-204396 - The Red Hat Enterprise Linux operating system must enable a user session lock until that user re-establishes access using established identification and authentication procedures.
- The ability to enable controls that are designed to secure a GNOME desktop environment. Users who enforce the DISA STIG standard on a RHEL 7 or 8 operating system can set the top-level `manage_gnome` option to `true` to enforce all relevant controls in a GNOME desktop environment. Users who enforce CIS controls in RHEL 7 or 8 or CentOS environments must also specify `manage_gnome=true` to enable controls for GNOME. For CIS users, the following GNOME controls are added:
  - 1.8.1 - Ensure GNOME Display Manager is removed
  - 1.8.2 - Ensure GDM login banner is configured
  - 1.8.3 - Ensure last logged in user display is disabled
  - 1.8.4 - Ensure XDCMP is not enabled
  - 1.8.5 - Ensure automatic mounting of removable media is disabled

### Changed

- The implementation for DISA STIG Control V-204600 was changed to ensure that the `StrictModes=yes` setting is explicitly applied to the Secure Shell (SSH) configuration. This setting helps to protect SSH keys. If file system permissions for SSH keys are too lax, SSH authentication fails. The change affects users who enforce DISA STIG in a RHEL 7 environment.

### Fixed

- An issue that caused invalid permissions to be assigned to the audit directory. This issue affected users who enforced the DISA STIG standard on RHEL operating systems. The issue was related to STIG Control V-230471, which helps to ensure that only users with sufficient permissions can specify which events will be audited. The correct permissions are now set on the audit directory.
- An issue that prevented proper enforcement of STIG Controls V-204614 and V-204615, which are designed to prevent man-in-the-middle attacks. The controls are enforced on RHEL systems to specify whether redirect messages sent with the Internet Control Message Protocol (ICMP) using Internet Control version 4 (IPv4) can be accepted. This issue is corrected to ensure that the ICMP redirect messages are disabled by default for both IPv4 and IPv6.
- A CEM configuration issue for users who enforce the DISA STIG standard in a RHEL 8 environment. In CEM for Linux v1.5.0, the base class `cem_linux::utils::packages::linux::sudo` was not included for STIG use in the `data/RedHat/RedHat/8.yaml` file. The omission resulted in additional configuration steps for some users. The YAML file is now updated to include the missing `sudo` class.
- A scan issue for users who enforce the DISA STIG standard on RHEL operating systems. Previously, Controls V-204427 (5.4.12) and V-204428 (5.4.13) were reported as failing in Puppet Comply. Failures were reported even when CEM appeared to correctly enforce the controls. These controls help to protect systems against unauthorized access by limiting the number of failed login attempts. Scans for the controls now report accurate results.

### v1.5.2

Released 9 March 2023

### Fixed

- An issue that prevented the CEM for Linux v1.5.1 reference topics from being displayed on Puppet Forge. With the v1.5.2 release, the [Reference](#) tab is restored. The topics on the **Reference** tab for v1.5.2 are applicable to both the v1.5.2 and v1.5.1 releases.

## v1.5.1

Released 7 March 2023

### Changed

- A change was introduced to simplify configuration in Red Hat Enterprise Linux (RHEL) 8 environments where the US Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standard is enforced. The update applies to DISA STIG control V-230339, which is designed to limit login attempts and thus help to prevent brute-force attacks. The default directory where login failure records are kept can now be changed.

### Fixed

- An issue that prevented the OpenSSH server process from starting on RHEL 8 systems. The issue affected users who enforced the DISA STIG standard. When a value of `FUTURE` was set for cryptographic policies to help prevent malware attacks, the OpenSSH process failed with the following error message:

```
Extra argument FUTURE.
```

- An issue that prevented users from running the system account task from the Puppet Enterprise (PE) console. Previously, attempts to run the task resulted in error messages about missing metadata. With the fix applied, users can run the `cem_linux::system_account` task from the PE console to view system accounts.
- An issue that caused an error message pertaining to the `audit.rules` file. The issue was seen on RHEL operating systems after an upgrade to CEM for Linux v1.5.0. The following error message was issued:

```
Could not stat /etc/audit/rules.d/audit.rules
```

To resolve the issue, the CEM for Linux module was updated to reference all existing files in `/etc/audit/rules.d/` directory.

- An issue that caused a failure to audit `sudo` log files. When events pertaining to a `sudo` log file are collected, system administrators can review the events to detect whether unauthorized commands were run. The issue, which affected users on RHEL 8 systems, was caused by a failure to enforce Center for Internet Security (CIS) Control 4.1.3.3. The control is now enforced.
- A failure to enable GNU Privacy Guard (GPG) checks for downloaded packages on RHEL 8 operating systems. CIS Benchmark Control 1.2.3 ("Ensure `gpgcheck` is globally activated") is designed to ensure that downloaded packages from the RPM package management tool are checked. However, these checks failed to occur because the `repo_files` parameter associated with the CIS control does not specify the YUM files that are used to manage RHEL packages. The fix ensures that GPG checks will be enabled on a per-repository basis for each file that is listed in the `repo_files` parameter.
- An issue that can cause configuration problems in RHEL 8 environments where DISA STIG standards are enforced. Because of the issue, the following top-level parameters for the Grub2 bootloader could not be set: `cem_linux::regenerate_grub2_config`, `cem_linux::set_grub2_password`, `cem_linux::grub2_superuser`, and `cem_linux::grub2_superuser_password`. The issue was resolved to ensure that the parameters can be set, and the values are applied.
- A configuration issue that affects the security of messages when DISA STIG standards are applied in a RHEL 8 environment. The issue pertains to STIG control V-230245, which is designed to ensure that unauthorized persons cannot access system messages. Security is enforced by setting a permissions mode for access to the `/var/log/messages` file. The issue occurred because the resource data for STIG control V-230245 specified a value of `directory` instead of `file`. The issue was fixed to ensure that permissions are set for the `messages` file. The fix also ensures that a `/var/log/messages` directory is not created inadvertently.
- An issue that can cause configuration problems for users who attempt to enforce the DISA STIG standard in a RHEL 8 environment. The issue was caused by extraneous text in the `cem_linux/manifests/utils/`

bootloader/grub2/fips.pp file. The extraneous text, a Universal Unique Identifier of 6484, is now removed.

### v1.5.0

Released 14 February 2023

#### Added

- Enforcement of the DISA STIG standard on Red Hat Enterprise Linux (RHEL) 8 operating systems:
  - The Security Technical Implementation Guide (STIG) standard was developed by the US Defense Information Systems Agency (DISA). DISA STIG compliance is required for some infrastructures managed by the US government. On RHEL 8, the following DISA STIG standard is supported: Red Hat Enterprise Linux 8 STIG, Version 1, Release 8.
  - For the RHEL 8 operating system, STIG can be enabled by adding the following Hiera data to the control repository:

```
cem_linux::benchmark: 'stig'
```

- STIG supports Mission Assurance Category (MAC) levels 1, 2, and 3 and their associated “public,” “sensitive,” and “classified” profiles. STIG controls can be configured with their vulnerability ID (*v-*nnn**) or rule ID (*sv-*nnn**).
- For a list of supported STIG controls and configurations, see the [CEM Linux Reference](#).
- New top-level configuration option, `disable_package_gpgcheck`. By enabling this option, you disable GNU Privacy Guard (GPG) checks of downloaded packages. Disabling GPG checks can be helpful in rare cases if you enable more stringent system encryption standards, such as the Federal Information Processing Standards (FIPS). These standards can introduce stricter criteria than are normally available for GPG package signatures. If GPG and more stringent criteria are applied simultaneously, package downloads can fail. Specify the `disable_package_gpgcheck=true` setting only when necessary. Enabling this option can make your infrastructure less secure.

#### Fixed

- An error that occasionally prevented system startups and that caused failures of the Internet Control Message Protocol (ICMP) was resolved. The error was identified in the Puppet manifest file `disable_icmp_redirects.pp`, which specifies whether messages sent with ICMP can be redirected. In the file, extraneous text is now commented out.
- An issue with the `cem::utils::boot_fstab_entry` class was fixed to help ensure that Puppet runs would not overwrite user-specified settings.

### v1.4.3

Released 15 December 2022

#### Fixed

- Fixed an issue that resulted in catalog compilation errors on the Red Hat Enterprise Linux (RHEL) 7 operating system. The issue, caused by a duplicate instance of control V-204450 in the YAML file, resulted in error messages like the following:

```
Error: Could not retrieve catalog from remote server: Error 500 on SERVER:
Server
Error: Evaluation Error: Error while evaluating a Function Call, CEM:
cem_create_resources: failed
resource: class
```

- Fixed an issue related to Center for Internet Security (CIS) Control 4.1.3.6 in a RHEL 8 environment. When Control 4.1.3.6 is enabled, privileged programs are monitored to determine whether unauthorized users are trying to gain access. However, when Control 4.1.3.6 was enabled on systems using the Postfix mail transfer agent, two

`setuid` binary files (`postdrop` and `postqueue`) were not being added to the `auditd` monitor list. The issue is corrected so that scans can run successfully.

- Fixed an issue related to CIS Control 1.5.3 in a RHEL environment. Control 1.5.3 is designed to ensure that address space layout randomization (ASLR) is enabled. ASLR randomly arranges the address space of data areas in processes to help protect system security. Enforcement for Control 1.5.3 was available in a RHEL 7 environment but was missing from RHEL 8. The control is now available to RHEL 8 users.
- Fixed an issue that affects users of the RHEL 7 operating system and pertains to control V-204444. When enabled, the control helps to prevent non-privileged users from initiating privileged functions such as disabling, circumventing, or altering security safeguards. However, after a system administrator specified the privileged users (resources) for control V-204444 and scans were run, the Puppet agent overwrote the resource list on each run. The issue is fixed to ensure that the resource list is not overwritten.
- Fixed an issue that caused scan failures for CIS Control 4.1.16, ‘Ensure kernel module loading and unloading is collected,’ in a RHEL 7 environment. When this control is enforced, the process of loading and unloading kernel modules is monitored to help detect unauthorized access to the system. Users of RHEL 7 found that Control 4.1.16 failed scans even when the control was correctly configured. The issue is corrected so that scans can run successfully.

## v1.4.2

Released 8 November 2022

### Added

- Added the ability to configure multiple `rsyslog` remote hosts to CEM for Linux. In previous releases, only single remote hosts were fully configurable. This software update simplifies the process of using the `rsyslog` software utility to forward logs to remote servers.
- Added an audit script for the V-204392 control, which is included in a Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standard. The DISA STIG control helps to ensure that file permissions, ownership, and group membership of system files and commands match vendor values. You can use the new audit script to troubleshoot issues related to the control.

### Changed

- Updated the Advanced Intrusion Detection Environment (AIDE) utility class to add support for the `cron` scheduling utility. As a result, AIDE scans can be scheduled by using a `cron` task rather than `systemd` timers.
- Updated CEM for Linux to ensure that the `nullok` option cannot be included in the `system-auth` file. The `nullok` option determines whether users can access a service with a blank password. This software update is designed to prevent unauthorized access to the system.

### Fixed

- Fixed an issue that prevented certificates from being checked for Public Key Infrastructure (PKI) authentication. This software update affects users who are enforcing DISA STIG controls on a Red Hat Enterprise Linux (RHEL) operating system.
- Fixed an issue to help ensure that any new password must contain at least 8 characters that differ from the previous password. This software update affects users who are enforcing DISA STIG controls on a RHEL operating system.
- Fixed an issue related to the Center for Internet Security (CIS) Red Hat Enterprise Linux 8 Benchmark 2.0.0, Control 3.3.2: Ensure ICMP redirects are not accepted. This software update helps to ensure that the control is enforced so that Internet Control and Error Message Protocol redirects are prevented.
- Fixed an issue that caused the `cem_semanage` fact to run and log errors on an unsupported operating system, RHEL 6. `semanage` is a Security-Enhanced Linux (SELinux) management tool.
- Fixed an issue that caused catalog compilation errors when users selected the Network Time Protocol (NTP) synchronization service.

## v1.4.1

Released 24 October 2022

**Fixed**

- Fixed an issue that prevented the `cem_mount_info` fact from resolving on Puppet Enterprise (PE) versions 2019.x.x. The issue prompted the following error message:

```
Facter: error while resolving custom facts...
```

To resolve the issue, you can install CEM Linux v1.4.1. To help avoid the issue, you can install the latest version of PE.

**v1.4.0**

Released 20 October 2022

**Added**

- Support for the DISA STIG standard on Red Hat Enterprise Linux (RHEL) 7:
  - For the first time, CEM supports a Security Technical Implementation Guide (STIG) standard developed by the US Defense Information Systems Agency (DISA). DISA STIG compliance is required for some infrastructures managed by the US government. On RHEL 7, the following DISA STIG standard is supported: Red Hat Enterprise Linux 7 STIG, Version 3, Release 8.
  - For the RHEL 7 operating system, STIG can be enabled by adding the following Hiera data to the control repository:

```
cem_linux::benchmark: 'stig'
```

- STIG supports Mission Assurance Category (MAC) levels 1, 2, and 3 and their associated “public,” “sensitive,” and “classified” profiles. STIG controls can be configured with their vulnerability ID (`V-nnn`) or rule ID (`SV-nnn`).
- To support STIG controls that require information audits, new Puppet Bolt tasks were added.
- The following new Facter facts were added: `cem_mount_info`, `cem_nfs_exports`, `cem_semanage`, and `cem_sssd_domains`.
- For a list of supported STIG controls and configurations, see the [CEM Linux Reference](#).

**Changed**

- The product documentation was revised to improve usability and retrievability:
  - The changelog was migrated from Puppet Forge to the central location for Puppet documentation, Puppet Docs. The changelog was renamed to [Release notes](#) on page 174.
  - The readme file was transformed into a series of topics with a structure similar to other Puppet documentation. The CEM topics can now be viewed on Puppet Docs, starting with [Introducing the Compliance Enforcement Modules](#) on page 173.
  - The [Reference](#), [Tasks](#), and [Dependencies](#) documentation, which is generated automatically, remains on Puppet Forge.
- To facilitate the implementation of DISA STIG standards, new parameters were introduced for some resources. The new parameters do not affect default configurations and are backward compatible with previous user configurations. All parameters are documented in the [CEM Linux Reference](#).

**Fixed**

- Fixed an issue that caused the `auditd` service to restart multiple times. The problem was caused by an incorrect sequence when setting a rule for immutable configuration.
- Fixed an issue that caused catalog compilations to fail although the specified configuration was valid. The failures occurred when certain time-server options were specified for the chrony implementation of the Network Time Protocol.

**v1.3.2**

Released 8 September 2022

**Added**

- The `Ensure core dump storage is disabled` and `Ensure core dump backtraces are disabled` controls are now enforced on Red Hat Enterprise Linux (RHEL) 8 systems.
- Added a new enforcement mode, `disabled`, so that you can disable Security Enhanced Linux (SELinux) in your environment.

**Changed**

- The `Ensure audit log is disabled` when `audit logs are full` control is updated to halt the machine when the audit log is full. This change helps to ensure better compliance with Center for Internet Security (CIS) recommendations.
- To simplify configuration, the `ntp` and `chrony` classes were combined into the `timesync` class.

**Fixed**

- The `Disable USB Storage` control is updated to work as designed.
- The regular expression for matching Linux username patterns is updated to accept capital letters.
- Rules in the `/etc/auditd/rules.d` directory are now loaded by using the `augenrules --load` command. This fix helps to ensure that all rule files within the directory are loaded into the kernel.
- Fixed the per-resource ordering process by using the correct metaparameter `before` instead of `subscribe`.
- Fixed a parsing error for `chrony` that caused catalog compilation failures.
- Fixed a command injection vulnerability that could occur when unsanitized user input was used in the `command`, `onlyif`, or `unless` parameters of an `exec` resource.
- Fixed an issue with the permissions of Secure Shell (SSH) host private keys to ensure that the permissions are sufficiently restrictive.
- Fixed the `cem_systemctl` feature to return a result of `false` without error messages in Puppet run logs when the feature is evaluated on Microsoft Windows machines.
- Fixed an issue with the `cem_mta` fact that caused errors in RHEL 6.

**v1.3.1**

Released 18 August 2022

**Fixed**

- Controls that configure `journald` now properly configure the `journald.conf` file.
- The `cem_coredump` fact will no longer attempt to resolve on nodes that do not support `systemctl`.
- The `cem_grub_cfg` fact will now identify the correct GRUB2 configuration file on Red Hat Enterprise Linux (RHEL).
- The Center for Internet Security (CIS)-specific parameters `enable_systemd_journal` and `enable_nopasswd_sudo_prune` now function correctly.
- Fixed how Ruby code is loaded during Continuous Delivery for Puppet Enterprise impact analysis. This update fixes a bug that caused impact analysis to fail after upgrading CEM for Linux to v1.3.0.
- Fixed invalid default parameter values that caused catalog compilation failures when enforcing the control `ensure_password_creation_requirements_are_configured`.
- Fixed a duplicate resource defaults statement that caused catalog compilation failures when selecting `ntp` as the time synchronization service.

**v1.3.0**

Released 3 August 2022

**Changed**

- The core architecture for the module has changed. These changes should be transparent to the user. However, using Hiera automatic parameter lookup to set configurations directly on classes in the `cem_linux::benchmarks::controls::*` namespace will no longer work. This configuration method

was not supported previously, and with the new architecture those classes have been removed and replaced with module Hiera data.

- For more information on the new architecture, see [Configuring CEM](#) on page 198.
- The [Reference: Benchmarks and controls](#) on page 210 was revised to improve usability. Sample configurations are provided for each supported control.

#### Fixed

- Added proper containment to the `cem_coredump` fact so that it will no longer run on operating systems that do not support it.
- Fixed how Network Time Protocol (NTP) options are handled. This fix resolves failures that occurred when using certain timeserver options.

#### v1.2.0

Released 24 May 2022

#### Added

- Added the Center for Internet Security (CIS) Level 2 Server profile for Red Hat Enterprise Linux (RHEL) 7.

#### Changed

- Updated the CIS RHEL 8 benchmark to version 2.0.0.
- Removed support for CentOS 8 because the operating system has reached End of Life (EOL). CEM for Linux has never supported CentOS Stream, and with non-stream CentOS 8 being EOL, support for it was removed entirely.

#### Fixed

- Fixed an issue that prevented the `coredump` configuration setting from being properly enforced. Now, you can use the module to configure core dumps.
- Fixed an issue related to file system mount points, which were not properly remounted after changes in mount-option enforcement. This issue prevented certain configuration changes from being applied.

#### v1.1.4

Released 25 March 2022

#### Changed

- Updated the `audit_user_homedir` task to prevent the task from modifying permissions on top-level directories: `/boot`, `/boot/`, `/etc`, `/lib`, `/lib64`, `/proc`, `/proc/`, `/home`, `/opt`, `/tmp`, `/var`, and `/srv/`. The `audit_user_homedir` task can still modify permissions on subdirectories within the listed directories, except for `/boot` and `/proc`.
- In the `audit_user_homedir` task, added `rtkit` to the list of ignored usernames. Because `rtkit` is a system user, CIS states that the home directory permissions for `rtkit` should not be audited.

#### v1.1.3

Released 24 March 2022

#### Fixed

- Fixed a bug in the `audit_user_homedir` task to prevent the inadvertent modification of permissions on `bin` directories: `/bin`, `/sbin`, `/usr/bin`, and `/usr/sbin`.

#### v1.1.2

Released 16 March 2022

#### Added

- Added a section to the [CEM Reference](#) about configuring `chrony/ntp` time servers.

**Changed**

- Expanded the range of versions in the `metadata.json` file so that users can install the latest modules to meet dependency requirements.

**Fixed**

- Fixed a bug in the `cem_linux::utils::timesync` configuration option that caused Puppet run failures when Network Time Protocol (NTP) was selected for time synchronization.
- Fixed a bug that caused a Puppet run failure during attempts to use a template to provide the Message of the Day (MOTD).
- Fixed a bug relating to unsupported options in the `auditd` config template on Red Hat Enterprise Linux (RHEL) 7. The bug caused startup failures for the `auditd` service.

**v1.1.1**

Released 25 January 2022

**Fixed**

- Fixed an issue related to non-idempotent resources when managing permissions for the `Grub2` bootloader configuration. This issue affected Red Hat Enterprise Linux (RHEL) systems that did not use Extensible Firmware Interface (EFI) mode.

**v1.1.0**

Released 14 December 2021

**Added**

- Enforcement for Center for Internet Security (CIS) Red Hat Enterprise Linux (RHEL) 8 Server Level 2 recommendations.
- Updates related to bootloader configurations. Configurations, including password settings, can now be managed through the CEM module on systems that use the `grub2` bootloader. You can also opt in to automatically regenerate the bootloader config files after changes are made. For details, see the [CEM for Linux readme file](#).
- Permissions management for log files in the `/var/log` directory is now available in the module. Previously, you had to run a Puppet Bolt task to manage permissions for log files. Because this feature is now supported natively, the Puppet Bolt task `cem_linux::logfile_permissions` was removed.
- Added a new fact, `cem_grub_cfg`. This fact contains information related to general `grub` configuration on the machine.

**Changed**

- Replaced the `camptocamp-systemd` module with the supported `puppet-systemd` module. To help ensure compatibility, you must update your Puppetfile to use the `puppet-systemd` module v3.5.0 or later.
- The `cem_uefi_boot` fact was changed to `cem_efi` and more information was added to the fact. The new name is more representative because the fact now includes boot and other information.

**Restriction**

- When you scan a node with Puppet Comply after applying CEM, some recommendations that are enforced by CEM might be reported as having failed the scan. This issue is due to bugs in the CIS-CAT Pro Assessor that is used by Comply. For more information, see the [readme file](#).

**v1.0.0**

Released 28 September 2021

This is the initial public release of CEM for Linux.

**Known issues and limitations**

The current release includes known issues and limitations. In most cases, workarounds are provided.

## Comply scan issues

During a Comply scan, you might see errors about Center for Internet Security (CIS) recommended guidelines that are not enforced. These error messages are triggered by bugs in the CIS-CAT Pro Assessor that is bundled with Comply. CEM **does** correctly enforce these settings.

The following Comply scan errors might be reported:

- Red Hat Enterprise Linux (RHEL) Benchmark v2.0.0:
  - 1.4.2 - Ensure permissions on bootloader are configured
    - On EFI systems, the script that was run by the CIS-CAT Pro Assessor did not locate the correct `grub` file path. Permissions are set correctly by CEM. No action is required.
  - 1.4.1 - Ensure bootloader password is set
    - On EFI systems, the script that was run by the CIS-CAT Pro Assessor did not locate the correct `grub` file path. It is not mandatory to set a bootloader password. However, if you want to set a password to protect your system against unauthorized startup, follow the instructions in [Set a bootloader password](#) on page 203.
  - 4.1.2.3 Ensure system is disabled when audit logs are full
    - This is set to `halt` by CEM. The CIS-CAT Pro Assessor incorrectly shows this as a scan failure. No action is required.
  - 5.2.18 Ensure SSH MaxSessions is set to 10 or less
    - This is set to 10 by default. The CIS-CAT Pro Assessor incorrectly shows this as a scan failure. The scanner is looking for `<=4` instead of `<=10`. No action is required.

## General issues and limitations

- If you are using CEM for Linux on a RHEL 9 or Oracle Linux 9 operating system and you are enforcing a CIS Benchmark that requires `auditd` rules to be loaded, you must manually load the `auditd` rules by using the `"augenrules --load"` command. In this way, you can meet the requirements of `auditd` controls that require `auditd` rules and help to prevent scan failures.
- The [Reference](#) section on Puppet Forge does not include Security Technical Implementation Guide (STIG) information for the CEM for Linux v1.7.0 release.
- CEM for Linux does not support version 9.0.0 or later of the Puppet standard library of resources for modules (`puppetlabs-stdlib`). Several functions were removed from `puppetlabs-stdlib` starting with version 9.0.0, and the omission of these functions can cause warnings and catalog compilation failures in CEM for Linux. To help prevent issues, ensure that your Puppetfile specifies a version of `puppetlabs-stdlib` that is earlier than 9.0.0.
- If you run CEM for Linux on Oracle Linux 7, Oracle Linux 8, or AlmaLinux 8, two Puppet runs might be required to ensure that the required intentional changes, corrective changes, or both are made on the nodes in the infrastructure. Subsequent Puppet runs are idempotent.
- CEM for Linux v1.6.0 and later provides *limited* support for the System Security Services Daemon (SSSD), which manages access to remote directory services and authentication mechanisms. However, the following US Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) controls are not implemented:
  - Control V-230274 - RHEL 8 must implement certificate status checking for multifactor authentication.
  - Control V-230372 - RHEL 8 must implement smart card logon for multifactor authentication for access to interactive accounts.
  - Control V-230376 - RHEL 8 must prohibit the use of cached authentications after one day.

- The DISA STIG standard for RHEL 8 includes controls that help to manage and secure removable media file systems. Support for the related controls is currently outside the scope of CEM for Linux. As a result, the following controls are not enforced:
  - Control V-230303 - RHEL 8 must prevent special devices on file systems that are used with removable media.
  - Control V-230304 - RHEL 8 must prevent code from being executed on file systems that are used with removable media.
  - Control V-230305 - RHEL 8 must prevent files with the `setuid` and `setgid` bit set from being executed on file systems that are used with removable media.
- The DISA STIG standard for RHEL 8 includes controls related to the file access policy (`fapolicy`) module, which helps to ensure that only authorized applications can be run on a system. The `fapolicy` module must be configured with caution because improper configuration can result in a non-functional system. Because of the potential risk to system operations, controls related to the `fapolicy` module are currently outside the scope of CEM for Linux. The following controls are not enforced:
  - Control V-230523 - The RHEL 8 `fapolicy` module must be installed.
  - Control V-244545 - The RHEL 8 `fapolicy` module must be enabled.
  - Control V-244546 - The RHEL 8 `fapolicy` module must be configured to employ a deny-all, permit-by-exception policy to allow the execution of authorized software programs.
- Multifactor controls and configurations are outside the scope of CEM for Linux. However, you can set up multifactor authentication for an infrastructure that is protected by CEM for Linux by implementing a network authentication system. For example, you can set up one-time password authentication on the client side by following the instructions in [Setting up multi-factor authentication on Linux systems](#).
- If you are enforcing the DISA STIG standard on the RHEL 7 operating system, the V-204392 auditing control is not working as designed. The control is missing a script that audits file permissions, ownership, and group membership of system files and commands. As a workaround, you can audit file permissions manually.
- Starting with v1.3.0, CEM for Linux implements a new architecture. If you upgrade CEM from v1.2.0 or earlier to v1.3.0 or later, and you encounter errors, try restarting the `pe-puppetserver` service or restarting or reloading Puppet Server. For instructions, see [Restarting Puppet Server](#).
- You cannot use the `iolog_dir` option to specify a directory for `sudo` log files. If you attempt to use the `iolog_dir` option in the `sudors` file to specify a log directory other than the default, errors are reported by the Augeas program. Augeas is a tool used for configuration editing in CEM.
- CEM cannot create file system partitions. This limitation can cause certain scanner checks to fail.
- CEM cannot set permissions on removable media partitions. To set the required permissions on these partitions, ensure that `nodev`, `nosuid`, `noexec` exists in the options portion of `/etc/fstab` for the partition.
- Support for the eXecute Disable/No eXecute (XD/NX) hardware feature is dependent on the host kernel and cannot be configured by CEM. If you plan to enable XD/NX support, ensure that you are using up-to-date kernels. If you plan to enable XD/NX support on newer kernels, be aware that CEM cannot manage this feature.
- To comply with CIS recommendations, you must prevent root users from logging onto the system console. Because this action requires knowledge of the site, you must configure this control manually by removing entries in `/etc/securetty` for consoles that are not in secure locations.
- CEM does not enforce `authselect` controls for CIS 2.0.0 5.4.x on Red Hat Enterprise Linux 8. Enforcement requires site knowledge and can break network authentication. CIS recommends that these controls not be enforced in specific environments. For example, the controls should not be enforced if the node is joined to an Active Directory domain or to Red Hat Identity Management. If you enforce `authselect` controls, you must ignore all other controls that affect authentication and use a predefined `authselect` profile or manage a custom profile. For more information, see [Configure system authentication with the authselect utility](#) on page 205. CEM includes a Puppet Bolt task, `audit_authselect`, to audit these controls.
- You can configure the `ensure_nodev_option_set_on_home_partition` control only if the `/home` setting is mounted on its own partition. Puppet does not create a partition for `/home`.
- If your system is running on Red Hat Enterprise Linux 8:
  - The `ensure_nis_server_is_not_installed` control is dependent on `ensure_rpcbind_is_not_installed_or_the__rpcbind_services_are_masked`. If you enforce `ensure_nis_server_is_not_installed`, you must also enforce `ensure_rpcbind_is_not_installed_or_the__rpcbind_services_are_masked`.

- The `ensure_nfs_utils_is_not_installed_or_the_nfs_server_service_is_masked` control is dependent on `ensure_rpcbind_is_not_installed_or_the_rpcbind_services_are_masked`. If you do not enforce `ensure_rpcbind_is_not_installed_or_the_rpcbind_services_are_masked`, you must also not enforce `ensure_nfs_utils_is_not_installed_or_the_nfs_server_service_is_masked`.
- The `ensure_the_running_and_on_disk_configuration_is_the_same` control is always enforced if `auditd` is managed by CEM.
- The `ensure_users_must_provide_password_for_escalation` control is disabled by default. You might want to enable this control to help ensure CIS compliance. However, a potential risk exists: It is possible that removing `NOPASSWD:` from `sudoers` files could invalidate the syntax of those files and break system authentication. If you accept the risk and want to enable this control, set the top-level configuration option `enable_nopasswd_sudo_prune` to `true`.
- If your system is running on Red Hat Enterprise Linux 7 or CentOS 7:
  - The `ensure_rpcbind_is_not_installed_or_the_rpcbind_services_are_masked` control is dependent on `ensure_nfsutils_is_not_installed_or_the_nfs_server_service_is_masked`. If you enforce `ensure_rpcbind_is_not_installed_or_the_rpcbind_services_are_masked`, you must also enforce `ensure_nfsutils_is_not_installed_or_the_nfs_server_service_is_masked`.
  - The `disable_wireless_interfaces` control requires that you install the `NetworkManager` package and that the service is running.

## Getting started

---

Learn the basic concepts associated with the Compliance Enforcement Modules and then review the steps for deploying CEM in your environment.

- [Basic concepts](#) on page 190

The basic concepts include an overview of CEM and the security standards that it enforces. You can also learn about `Hiera`, a key-value store that is used to configure CEM.

- [Next steps](#) on page 192

After you are familiar with the basic concepts, you can take the next steps.

## Basic concepts

The basic concepts include an overview of CEM and the security standards that it enforces. You can also learn about `Hiera`, a key-value store that is used to configure CEM.

### Compliance Enforcement Modules

CEM is software that automatically enforces security standards on IT infrastructures. After CEM is installed and configured, you can run Puppet Enterprise (PE) or open source Puppet on your specified nodes, and CEM automatically enforces security controls.

### Hiera

To configure CEM, you can use the *Hiera* key-value store. *Hiera* stores configuration data in a hierarchical structure in key-value pairs. For an introduction, see [About Hiera](#).

### Center for Internet Security (CIS)

The *Center for Internet Security, Inc.*, is a nonprofit organization that strives to protect IT infrastructures through collaboration and innovation. Contributors to the organization include security experts from government, business,

and academia who develop and maintain internationally recognized security standards. For more information, see [Center for Internet Security](#).

## CIS Benchmarks

CIS develops and maintains *CIS Benchmarks*, which are configuration recommendations for product families. For example, if your nodes run on the Red Hat Enterprise Linux (RHEL) 8 operating system, you can enforce the CIS Red Hat Enterprise Linux 8 Benchmark v2.0.0, Level 1 or 2. For more information, see [CIS Benchmarks](#).

## CIS profiles and levels

Each CIS Benchmark has a *profile*, which consists of a level and an applicability.

The *level* refers to the degree of protection:

- *Level 1* is intended to be practical and prudent, providing a clear security benefit without inhibiting the use of the technology.
- *Level 2* extends the Level 1 profile to provide additional protection for systems in which security is paramount. Level 2 can affect a system's performance and usability while promoting enhanced security.

The *applicability* refers to the affected system component. For example, if a benchmark has a profile of `Level 1 - Server`, the benchmark provides Level 1 (basic) security protections for servers.

## CIS controls

Each CIS Benchmark consists of *controls*, which are also called *recommendations* or *rules*. Each control is a security safeguard. For example, a control might disable the use of Bluetooth communication technologies on the protected system because Bluetooth transmissions can be intercepted. Or a control might specify that passwords must consist of at least 14 characters to help prevent unauthorized access.

To learn about the CIS controls that are enforced by CEM, go to the [CEM Linux Reference](#) section on Puppet Forge and click a benchmark to see the list of enforced controls. The control description starts with its config ID and name, for example:

```
5.4.2 - Ensure lockout for failed password attempts is configured
```

The anatomy of a CIS control is as follows:

- **Parameters:** Configuration options, along with data types and default values.
- **Supported Levels:** The supported levels for the control, for example, Level 1 or 2.
- **Supported Profiles:** The applicability of the control. For example, a control with a `server` profile is applicable to server components.
- **Hiera Configuration Example:** Snippet of code that can be used to configure the control in Hiera.
- **Alternate Config IDs:** The alternate config IDs for a control. If you configure the control in Hiera, you can use any of the listed config IDs. However, you cannot mix and match types within a configuration; you must use a single type of config ID.
- **Resource:** The name of the Puppet resource that enforces the control.

To enforce a CIS control on a node, you add Hiera data to your control repository. A control repository is the location where CEM configuration data is stored. For example, to enforce a CIS Benchmark with a profile of `Level 1 - Server` and specify only one control (for file system integrity), you would enter the following values:

```
control-repo/data/nodes/<node name>.yaml
cem_linux::benchmark: 'cis'
cem_linux::config:

 profile: 'server'
 level: '1'

 only:
```

```
- 'ensure_filesystem_integrity_is_regularly_checked'
```

## Security Technical Implementation Guides (STIGs)

The *Security Technical Implementation Guides (STIGs)* specify practices for configuring hardware and software to help prevent security vulnerabilities. Each STIG is a security standard for a vendor-specific component, such as RHEL 8. For details, see [Security Technical Implementation Guides \(STIGs\)](#).

## Defense Information Systems Agency (DISA)

The *Defense Information Systems Agency (DISA)* is a unit within the US Department of Defense. DISA provides IT and communications support to individuals and institutes working for the Department of Defense. DISA maintains and updates hundreds of STIGs, known as *DISA STIGs*.

### DISA STIGs

*DISA STIGs* are security standards. To learn about DISA STIG standards that you can enforce with CEM, go to the [CEM Linux Reference](#) on Puppet Forge and click a STIG standard to see the list of enforced controls. Each control is identified by its vulnerability ID, for example: V-204486.

To learn about the purpose of a STIG control, go to the [STIG Viewer](#) and search for the operating system associated with the control. For example, you can find control V-204486 in the [Red Hat Enterprise Linux 7 Security Technical Implementation Guide](#).

The anatomy of a DISA STIG control is as follows:

- **Parameters:** Configuration options, along with data types and default values.
- **Supported Levels:** The level of required protection, specified in terms of Mission Assurance Category levels (MACs):
  - `mac-1` specifies a critically important system. The loss of availability or integrity of MAC 1 systems is considered unacceptable.
  - `mac-2` specifies a system that handles support for deployed or contingency forces. The loss of a MAC 2 system can be tolerated only briefly.
  - `mac-3` specifies a system that handles information required for day-to-day operations. The loss of a MAC 3 system can be tolerated without largely impacting mission or operational readiness.
- **Supported Profiles:** The supported profiles for the control, such as `public`, `classified`, or `sensitive`.
- **Hiera Configuration Example:** Snippet of Hiera that can be used to configure the control.
- **Resource:** The name of the Puppet resource that enforces the control.

## Next steps

After you are familiar with the basic concepts, you can take the next steps.

1. To prepare for the installation and install CEM, follow the instructions in [Installing CEM](#) on page 192.
2. To configure CEM, follow the instructions in [Configuring CEM](#) on page 198.

## Installing CEM

Before you install CEM, complete the preparation steps: review the system requirements, install and configure Puppet Enterprise (PE) or open source Puppet, and purchase CEM. To help avoid issues, install and evaluate CEM in a *test* environment before you install CEM in a production environment.

- [Prepare to install the module](#) on page 193

To help ensure the successful deployment of CEM, complete the preparation steps.

- [Install and evaluate the module in a test environment](#) on page 194

In some cases, compliance controls can negatively impact services that run on nodes. To help avoid possible issues, install and evaluate CEM in a *test* environment before running CEM in a production environment.

- [Install the module in the production environment](#) on page 195

To deploy CEM, you must install the module in the production environment and then classify the nodes on which you want to enforce secure configuration compliance.

- [Uninstall the module](#) on page 195

To stop using CEM, you can uninstall the `cem_linux` module. Alternatively, to stop using CEM on one or more nodes, declassify the nodes to remove their association with the `cem_linux` class.

## Prepare to install the module

To help ensure the successful deployment of CEM, complete the preparation steps.

1. Review the following table to ensure that CEM can meet your organization's requirements. CEM for Linux supports the following operating systems, frameworks, and standards:

| Operating system                  | Framework or standard                                                                                                                                                      | Profile             |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| Red Hat Enterprise Linux (RHEL) 7 | CIS Benchmarks v3.1.1                                                                                                                                                      | Level 1, 2 – Server |
| RHEL 7                            | The following Defense Information Systems Agency (DISA) Security Technical Implementation Guide (STIG) standard:<br>Red Hat Enterprise Linux 7 STIG, Version 3, Release 12 |                     |
| RHEL 8                            | The following DISA STIG standard:<br>Red Hat Enterprise Linux 8 STIG, Version 1, Release 11                                                                                |                     |
| CentOS Linux 7                    | CIS Benchmarks v3.1.2                                                                                                                                                      | Level 1, 2 – Server |
| RHEL 8                            | CIS Benchmarks v2.0.0                                                                                                                                                      | Level 1, 2 – Server |
| RHEL 9                            | CIS Benchmarks v1.0.0                                                                                                                                                      | Level 1, 2 – Server |
| AlmaLinux 8                       | CIS Benchmarks v2.0.0                                                                                                                                                      | Level 1, 2 – Server |
| Oracle Linux 7                    | CIS Benchmarks v3.1.1                                                                                                                                                      | Level 1, 2 – Server |
| Oracle Linux 8                    | CIS Benchmarks v2.0.0                                                                                                                                                      | Level 1, 2 – Server |
| Oracle Linux 9                    | CIS Benchmarks v1.0.0                                                                                                                                                      | Level 1, 2 – Server |
| Rocky Linux 8                     | CIS Benchmarks v1.0.0                                                                                                                                                      | Level 1, 2 – Server |

2. Review the dependencies to ensure that your infrastructure will meet the requirements. Go to Puppet Forge and review the [Dependencies](#) tab.
3. To manage nodes with Puppet Enterprise (PE), install any PE version in the 2021.7 or 2023 release stream. For instructions, see [Installing](#).
4. If you installed PE, follow the instructions in [Configuring Puppet Enterprise](#).
5. To manage nodes with open source Puppet, install Puppet 7 or 8. For instructions, see [Install Puppet](#).
6. If you installed open source Puppet, follow the instructions in [Configure Puppet settings](#).
7. If you are using Puppet 7, verify that the agent and server are at v7.8.0 or later. You can also use any level of Puppet 8.
8. To purchase CEM software, contact a Puppet by Perforce sales representative. For more information, see [Contact Us](#).

- Optionally, to help avoid issues during deployment to a production environment, you can initially install and evaluate CEM in a test environment. For instructions, see [Install and evaluate the module in a test environment](#) on page 194.

## Install and evaluate the module in a test environment

In some cases, compliance controls can negatively impact services that run on nodes. To help avoid possible issues, install and evaluate CEM in a *test* environment before running CEM in a production environment.

- Learn about the CIS Benchmark or STIG standard that you plan to enforce:

- For a list of supported CIS Benchmarks and STIG standards, see [Prepare to install the module](#) on page 193.
- If you plan to enforce CIS Benchmarks, you can find a list of benchmarks and associated controls in the [CEM Linux Reference](#) on Puppet Forge. You can also download benchmark information from the [CIS Benchmarks List](#). If you are using CEM with Puppet Comply, you can view details about the benchmarks in Comply.
- If you plan to enforce STIG controls, you can find a list of guides and associated controls in the [CEM Linux Reference](#) on Puppet Forge. For more detailed information about STIG controls, go to the [STIG Viewer](#).

- Make a list of any CIS or STIG controls that you plan to enable, disable, or configure to meet your organization's requirements.

For example, if a control specifies that a password must be changed every 60 days, but your organization requires a password change every 30 days, you can change the expected value for the associated control.

**Tip:** For the sake of simplicity, some users review the controls and enable only a *limited subset* to meet their organization's requirements.

- Identify a test environment. Many users follow the instructions in [Environments](#). You can also use any alternative method that works for you:

- For Puppet Enterprise, create a test node group and then assign the `cem_linux` class to that node group.
- For open source Puppet, follow the instructions in [Classifying nodes](#). Ensure that the CEM module is included on the test nodes.

- Download CEM from Puppet Forge. CEM is available as a subscription. For more information, see the [Premium content page](#).

- If the host server is connected to the internet, install the module by following the instructions in [Installing modules from the Forge by using an internet connection](#).

- If the host server is not connected to the internet, install the module by following the instructions in [Installing modules from the Forge in an air-gapped environment](#).

- Verify that the CEM module is successfully installed in the test environment.

**Tip:** If the installation was successful, you can find `cem_linux` in the following directory:

```
/etc/puppetlabs/code/environments/<environment_name>/modules/cem_linux
```

- If you plan to implement a CIS Benchmark at Level 2, ensure that the level is set to 2 in the control repository. (CIS Benchmarks are set to Level 1 by default, but Level 2 must be specified manually.) You can simplify configuration by using the Hiera key-value store as described in [Getting started with Hiera](#). For instructions about specifying the CIS Benchmark level, see [Basic configuration example](#) on page 202.

- If you plan to implement STIG controls, configure STIG. Follow the instructions in [Configure DISA STIG](#) on page 204.

- Implement any other configuration updates that you identified in Step 2. Take the following actions:

- Specify the updates as described in [Find and set configuration options](#) on page 198.
- Ensure that the updates are deployed to the test environment. For example, if you are using Hiera and Code Manager, you must update the Hiera YAML files, commit the changes to the appropriate branch of your control repository, and trigger [Code Manager](#). If you are using open source Puppet, you would follow the same procedure but trigger `r10k`.

11. To detect and resolve any errors, take the following actions:

- a. Look for errors in Puppet runs in your test environment.
- b. If you detect errors, review and update your configuration. For help with configuration options, see the [CEM Linux Reference](#).
- c. If the configuration is correct but errors persist, enable debug logging on the Puppet primary server and review the `puppetserver.log` file. For more information, see [Puppet Server logging](#).

**Tip:** In the log list, CEM errors are prefixed with `CEM` or `cem`.

- d. Optionally, for additional insight into errors, enable tracing and debugging when you run the Puppet agent.
- e. If you are unable to resolve the errors, take one of the following actions:
  - PE users can post a question in the [#compliance](#) Slack channel in the [Puppet Community](#) or open a ticket with [Puppet Support](#).
  - Open source Puppet users can post a question in the [#compliance](#) Slack channel in the [Puppet Community](#), open a ticket on the [cem\\_issues webpage in GitHub](#), or open a ticket with [Puppet Support](#). As an open source Puppet user, your options vary depending on the support package that you purchased with CEM.

## Install the module in the production environment

To deploy CEM, you must install the module in the production environment and then classify the nodes on which you want to enforce secure configuration compliance.

### Before you begin

If you have not done so, install and evaluate CEM in a *test* environment before running CEM in a production environment. In this way, you can help to detect and avoid possible issues. For instructions, see [Install and evaluate the module in a test environment](#) on page 194.

1. If you have not done so, download CEM from Puppet Forge. The module is available as a premium content subscription. For more information, see the [Premium content page](#).
2. If the host server is connected to the internet, install the module by following the instructions in [Installing modules from the Forge by using an internet connection](#).
3. If the host server is not connected to the internet, install the module by following the instructions in [Installing modules from the Forge in an air-gapped environment](#).
4. If you installed Puppet Enterprise (PE), specify the nodes on which you want PE to run and enforce compliance. To specify the nodes, open the PE console and assign the `cem_linux` class to one or more node groups. For instructions about creating and classifying node groups, see [Grouping and classifying nodes](#).
5. If you installed open source Puppet, specify the nodes on which you want Puppet to run and enforce compliance. Follow the instructions in [Classifying nodes](#).

## Uninstall the module

To stop using CEM, you can uninstall the `cem_linux` module. Alternatively, to stop using CEM on one or more nodes, declassify the nodes to remove their association with the `cem_linux` class.



**CAUTION:** If you uninstall CEM or declassify one or more nodes, Puppet Enterprise (PE) or open source Puppet no longer runs on the affected nodes to enforce secure configuration compliance. When you uninstall CEM or declassify a node, the node does not revert to its pre-CEM state.

Take one of the following actions:

- To uninstall the CEM module, follow the instructions in [Uninstalling modules](#).
- On PE, to declassify nodes, you can remove the nodes from the node group that is associated with the `cem_linux` class. For instructions, see [Remove nodes from a node group](#). Alternatively, to stop using CEM on an entire node group, you can remove the `cem_linux` class from the node group. For instructions, see [Remove classes from a node group](#).

- On open source Puppet, ensure that `cem_linux` is no longer included in the manifest that applies the `cem_linux` class.

## Upgrading CEM

---

You can upgrade CEM for Linux to take advantage of the latest features, fixes, and improvements. To help ensure a smooth upgrade process, complete the preparation tasks first.

- [Prepare to upgrade the module](#) on page 196

Before you upgrade the module, learn about the target release. Familiarize yourself with any updates that were implemented to comply with Center for Internet Security (CIS) Benchmarks or with Defense Information Systems Agency (DISA) Security Technical Implementation Guidelines (STIGs). Then, test the upgrade in a non-production environment and troubleshoot any issues.

- [Upgrade the module](#) on page 197

After you complete the preparation steps, you are ready to upgrade CEM for Linux.

### Prepare to upgrade the module

Before you upgrade the module, learn about the target release. Familiarize yourself with any updates that were implemented to comply with Center for Internet Security (CIS) Benchmarks or with Defense Information Systems Agency (DISA) Security Technical Implementation Guidelines (STIGs). Then, test the upgrade in a non-production environment and troubleshoot any issues.

1. To learn about a new CEM release, review the [Release notes](#) on page 174, which document major updates, such as the introduction of additional operating systems, new or changed CIS Benchmarks, and new or changed STIG standards. You can also learn about defect fixes, minor changes, and known issues.
2. Learn about the CIS Benchmark or STIG standard that you plan to enforce in the new release. Review the associated controls and configuration options:
  - For a list of supported CIS Benchmarks and STIG standards, see [Prepare to install the module](#) on page 193.
  - If you plan to enforce CIS Benchmarks, you can find a list of benchmarks and associated controls in the [CEM Linux Reference](#) on Puppet Forge. You can also download benchmark information from the [CIS Benchmarks List](#). If you are using CEM with Puppet Comply, you can view details about the benchmarks in Comply.
  - If you plan to enforce STIG controls, you can find a list of benchmarks and associated controls in the [CEM Linux Reference](#) on Puppet Forge. For more detailed information about STIG controls, go to the [STIG Viewer](#). If CEM was updated to accommodate updates in a STIG standard, look for a list of changes in the [Reference](#) section on the documentation website.
3. Identify a test environment. Many users follow the instructions in [Environments](#). You can also use any alternative method that works for you.

Testing is important because CEM can make hundreds of changes to a system, and many of those changes are critical to components. By testing and troubleshooting issues in advance, you can help to prevent issues later.

4. In the test environment, upgrade CEM. To help simplify the process, you can use a [Puppetfile](#):
  - a. In the Puppetfile, update the version in the CEM module declaration. For example, to upgrade CEM to version 1.8.0, specify the CEM declaration as shown:

```
mod 'puppetlabs/cem_linux', '1.8.0'
```

- b. Commit the change to the appropriate branch or test environment.
  - c. Deploy the change by using Code Manager or r10k. For instructions about using Code Manager and r10k, see [Managing code with Code Manager](#).
5. Verify that the CEM module is successfully upgraded in the test environment.

6. Determine whether the CEM configuration must be updated. If so, make a list of the required updates.

In many cases, you can upgrade CEM *without* additional configuration. However, if the relevant CIS Benchmark or STIG standard was updated in the release, the CEM configuration typically requires updates because controls might have been added or removed, or their numbers or titles might have changed. For example, if your configuration uses a "normalized number" control ID ("c1\_1\_1" or similar), pay close attention to any controls whose number changed because you must update the corresponding control ID in your configuration. If the reference section indicates that the number for control "1.1.1" changed to "1.1.2," you must replace "c1\_1\_1" in your configuration with "c1\_1\_2."

To find any documented control updates for CEM for Linux, see [Reference: Benchmarks and controls](#) on page 210.

7. Implement any required configuration updates. This step is crucial to help prevent CEM errors caused by misconfiguration.

You can simplify configuration by using the Hiera key-value store as described in [Getting started with Hiera](#). Take the following actions:

- a. Update the configuration of the test environment. For instructions, see [Find and set configuration options](#) on page 198.
  - b. Ensure that the updates are deployed to the test environment. For example, if you are using Hiera and Code Manager, you must update the Hiera YAML files, commit the changes to the appropriate branch of your control repository, and trigger Code Manager.
8. To detect and resolve any errors, take the following actions:
    - a. Look for errors in Puppet runs in your test environment.
    - b. If you detect errors, review and update your configuration. For help with configuration options, see the [CEM Linux Reference](#).
    - c. If the configuration is correct but errors persist, enable debug logging on the Puppet primary server and review the `puppetserver.log` file. For more information, see [Puppet Server logging](#).
- Tip:** In the log list, CEM errors are prefixed with `CEM` or `cem`.
- d. Optionally, for additional insight into errors, enable tracing and debugging when you run the Puppet agent.
  - e. If you are unable to resolve the errors, post a question in the [#compliance](#) Slack channel in the [Puppet Community](#) or open a ticket with [Puppet Support](#).
9. To plan the upgrade in the *production* environment, consider the following questions:
    - Should the upgrade occur in stages or on all nodes simultaneously?
    - What is the risk of the upgrade, and is further testing required to mitigate the risk?
    - In the rare event that the upgrade is not successful, what is the plan to roll back the changes, given the fact that CEM cannot revert changes automatically?

**Tip:** For assistance with these questions, contact [Puppet Support](#). Support engineers are prepared to assist you before, during, and after the upgrade.

## Upgrade the module

After you complete the preparation steps, you are ready to upgrade CEM for Linux.

1. Update the CEM declaration in the Puppetfile. Specify the CEM version number to which you are upgrading. For example, to upgrade the `cem_linux` module to version 1.8.0, specify the CEM declaration as shown:

```
mod 'puppetlabs/cem_linux', '1.8.0'
```

For instructions about modifying the Puppetfile, see [Declare Forge modules in the Puppetfile](#).

2. Commit the change to the appropriate branch.
3. Deploy the change by using Code Manager or r10k.

- If you determined during the preparation process that configuration updates are required in CEM, implement the configuration updates. For instructions, see [Configuring CEM](#) on page 198.

**Tip:** Starting with v1.3.0, CEM for Linux implements a new architecture. If you upgrade CEM from v1.2.0 or earlier to v1.3.0 or later, and you encounter errors, try restarting the `pe-puppetserver` service or restarting or reloading Puppet Server. For instructions, see [Restarting Puppet Server](#).

## Configuring CEM

Configuration of CEM is optional. If you installed CEM and assigned the `cem_linux` class to one or more node groups, the Center for Internet Security (CIS) Server Level 1 profile is enforced automatically during the next Puppet run. However, if the default values leave your infrastructure in an undesirable state, or if you want to customize compliance to meet your organization's requirements, you can configure CEM.

For example, if a CIS control sets the maximum password age at 365 days, but your organization requires a password change every 90 days, you can configure CEM accordingly.

You must also configure CEM if you plan to enforce DISA STIG standards rather than a CIS Benchmark. Follow the instructions in [Configure DISA STIG](#) on page 204.



**CAUTION:** CIS and STIG controls are developed and maintained by security experts, and CEM implements the controls as code to help secure your configuration. CEM can make hundreds of changes to a system, and many of those changes are critical to components. Because every system environment is different, some of the default control settings might not be appropriate in all environments. For this reason, when you configure CEM or update a CEM configuration, test the configuration in a limited environment on one or two nodes and evaluate the results. Resolve any issues before implementing the configuration in a production environment. For a new installation, see [Install and evaluate the module in a test environment](#) on page 194. For an upgrade, see [Prepare to upgrade the module](#) on page 196.

For all types of configuration tasks, you can use the Hiera key-value store in your control repository. For more information, see [About Hiera](#) and [Getting started with Hiera](#).

For general information about CEM configuration options, see [Overview of configuration options](#) on page 198. For detailed information about CEM configuration options, see the [CEM Linux Reference](#).

For configuration examples, see [How to configure the module: Examples and guidelines](#) on page 201.

- [Overview of configuration options](#) on page 198

Configuration options include top-level options, benchmark options, and Center for Internet Security (CIS)-specific options.

- [How to configure the module: Examples and guidelines](#) on page 201

Configuration examples are provided to help you understand how CEM is used in a production environment.

Guidelines are provided to help optimize your configuration.

## Overview of configuration options

Configuration options include top-level options, benchmark options, and Center for Internet Security (CIS)-specific options.

### Find and set configuration options

You can find the configuration options in the [CEM Linux Reference](#) on Puppet Forge. Locate the CIS Benchmark or the DISA STIG standard that you want to enforce, and then view the associated controls and configuration options.

**Tip:** When you use Hiera to specify CEM configuration options, the configuration looks different from the configuration for other Puppet products. The reason is that the `config` variable receives a hash, and the key values in the hash control the CEM variables in the configuration.

### CIS controls

Each CIS Benchmark has a list of associated controls, and additional information is provided for each control. The control description starts with a config ID and name, for example:

```
5.4.1 - Ensure password creation requirements are configured
```

The anatomy of a CIS control is as follows:

- **Parameters:** Configuration options for a control, along with the data type and default value.
- **Supported Levels:** The supported levels for a CIS control.
- **Supported Profiles:** The applicability of the control. For example, a control with a profile of `server` is applicable to server components.
- **Hiera Configuration Example:** Snippet of Hiera that can be used to configure a control.
- **Alternate Config IDs:** The alternate config IDs for a control. Any of these config IDs, along with the full control name, can be used as a key in the `control_config` hash.
- **Resource:** The name of the Puppet resource that enforces the control.

#### Guidelines for specifying CIS config IDs

You can specify CIS controls in the `control_config` hash by referencing the full control name, the control number, the normalized control name, or the normalized control number. You *cannot* mix and match these forms and must pick a single config ID form to use for your config. Full control names and control numbers are copied verbatim from the benchmarks and are case-sensitive. Normalized control names have lowercase letters and contain only alphanumeric characters and underscores. Normalized control numbers are always prefixed with a `c` and contain only numeric characters separated by underscores.

Example of alternate config IDs:

- Full control name: `(L1) Ensure 'Enforce password history' is set to '24 or more password(s)'`
- Control number: `1.1.1`
- Normalized control name: `ensure_enforce_password_history_is_set_to_24_or_more_passwords`
- Normalized number: `c1_1_1`

#### DISA STIG controls

You can view DISA STIG controls in the [CEM Linux Reference](#) on Puppet Forge. Each DISA STIG control is identified by its vulnerability ID, for example: V-204486. To learn about the purpose of a STIG control, go to the [STIG Viewer](#) and search for the operating system associated with the control, and then the control.

For example, if your nodes are running on RHEL 7, you can find relevant information in the [Red Hat Enterprise Linux 7 Security Technical Implementation Guide](#). That webpage also provides a description of Control V-204486.

The anatomy of a DISA STIG control is as follows:

- **Parameters:** Configuration options, along with data types and default values.
- **Supported Levels:** The level of required protection, specified in terms of Mission Assurance Category levels (MACs):
  - `mac-1` specifies protection for a critically important system. The loss of availability or integrity of MAC 1 systems is considered unacceptable.
  - `mac-2` specifies protection for a system that handles support for deployed or contingency forces. The loss of a MAC 2 system can be tolerated only briefly.
  - `mac-3` specifies protection for a system that handles information required for day-to-day operations. The loss of a MAC 3 system can be tolerated without largely impacting mission or operational readiness.
- **Supported Profiles:** The supported profiles for the control, such as `public`, `classified`, or `sensitive`.
- **Hiera Configuration Example:** Hiera code snippet that can be used to configure the control.
- **Resource:** The name of the Puppet resource that enforces the control.

## Top-level configuration options

In Hiera, the top-level configuration options are found directly under the `cem_linux` namespace. If you must customize CEM to meet your organization's requirements, you can specify values for the top-level configuration options.

In Hiera, the top-level options are prefixed with `cem_linux:`. The following list describes the options:

- `benchmark` - `Enum[ 'cis', 'stig' ]` - the compliance framework to use. CEM for Linux supports only `cis` and `stig`. Default: `cis`.
- `config` - `Optional[Hash]` - the location for all non-top-level configuration options. Default: `undef`.
- `allow_on_kubernetes_node` - `Boolean` - If `cem_linux` detects that it is running on a Kubernetes cluster node or host, CEM does not enforce controls, and it logs a warning. In this way, CEM helps to prevent the accidental enforcement of incorrect compliance settings that can render Kubernetes non-functional. Default: `false`.
- `manage_gnome` - `Boolean` - When set to `true`, controls are enforced to secure a GNOME desktop environment. When set to `false`, the controls are not enforced. Default: `false`.
- `disable_package_gpgcheck` - `Boolean` - When set to `true`, GNU Privacy Guard (GPG) checks of downloaded packages are disabled. Disabling GPG checks can be helpful in rare cases if you enable more stringent system encryption standards, such as the Federal Information Processing Standards (FIPS). These tighter standards can introduce stricter criteria than are normally applied to GPG package signatures. If GPG and more stringent criteria are applied simultaneously, package downloads can fail. Specify `disable_package_gpgcheck=true` only when necessary because this setting can make your infrastructure less secure. Default: `false`.
- `regenerate_grub2_config` - `Boolean` - Some configurations in CEM for Linux modify the Grub2 bootloader configuration. To regenerate the Grub2 configuration after applying a change, set this parameter to `true`. If you do not set this parameter to `true`, you must manually regenerate the Grub2 configuration. Default: `false`.
- `set_grub2_password` - `Boolean` - Set the password for the Grub2 bootloader. If you set this value to `true`, you must also set the `grub2_superuser` and `grub2_superuser_password` parameters, or configure the specific bootloader password control by using the `control_configs` option. Default: `false`.
- `grub2_superuser` - `Optional[String[1]]` - The superuser for the Grub2 bootloader if you set the `set_grub2_password` parameter to `true`. Default: `Undef`.
- `grub2_superuser_password` - `Optional[Sensitive[String]]` - The superuser password for the Grub2 bootloader if you set the `set_grub2_password` parameter to `true`. This value is sensitive in terms of security and should be stored in a `Sensitive` data type. Default: `Undef`.

## Hiera example

The following example configures CEM for Linux to regenerate the Grub2 bootloader config on a node using the CIS benchmark:

```
cem_linux::benchmark: 'cis'
cem_linux::allow_on_kubernetes_node: false
cem_linux::regenerate_grub2_config: true
cem_linux::config:
 ...
```

## Benchmark configuration options

Each Center for Internet Security (CIS) Benchmark and each Security Technical Implementation Guide (STIG) is associated with a set of configuration options. You can use Hiera to specify values for the benchmark configuration options.

The benchmark configuration options are available as key-value pairs within the `cem_linux::config:` hash. The following options are available:

- `only:` - `Optional[Array[String]]` — takes an array of control class names (`manifests/benchmarks/<benchmark>/controls/*.pp`). Classes specified here are included in the catalog. This option takes precedence over `ignore:`. Default: `undef`.

- `ignore`: - `Optional[Array[String]]` — takes an array of control class names (`manifests/benchmarks/<benchmark>/controls/*.pp`). The classes specified here are not included in the catalog. If `only` is specified, this option does nothing. Default: `undef`.
- `control_configs` - `Optional[Hash]` — where all rule-specific configurations live. Default: `undef`.

### CIS-specific configuration options

To meet your organization's requirements, you can specify CIS-specific configuration options, such as settings related to firewalls and log files. Use Hiera to specify these options.

The CIS-specific configuration options are available as key-value pairs within the `cem_linux::config` hash:

- `profile`: - `Optional[Enum['server', 'workstation']]` — the name of the benchmark profile. The only value supported by CEM is `server`. Default: `server`.
- `level`: - `Optional[Enum['1', '2']]` — the name of the profile level. Default: `1`.
- `firewall_type`: - `Optional[Enum['iptables', 'firewalld', 'unmanaged']]` — the preferred firewall provider. If set to `unmanaged`, CEM will not enforce firewall-related rules. Default: `firewalld`.
- `enable_systemd_journal` - `Optional[Boolean]` - Whether to enable the `systemd-journal` logging service. The default value is `false`. If this option is enabled, the `systemd-journal-remote` package will be installed and the `systemd-journal-upload.service` service will be enabled. However, several configuration parameters are required to ensure that the `systemd-journal-upload.service` functions correctly:

```
cem_linux::config:
 control_configs:
 'ensure_systemd_journal_remote_is_configured':
 address: '<IP address or FQDN of the remote host>'
 server_key_file: '<path to the server key file>'
 server_certificate_file: '<path to the server certificate file>'
 trusted_certificate_file: '<path to the trusted certificate file>'
```

## How to configure the module: Examples and guidelines

Configuration examples are provided to help you understand how CEM is used in a production environment. Guidelines are provided to help optimize your configuration.

- [Basic configuration example](#) on page 202

When you specify a compliance framework, CEM is configured to provide rule enforcement and configuration for that framework. For example, to enforce the Center for Internet Security (CIS) Server Level 1 benchmark for a node, you must classify the node with the `cem_linux` class, set the `benchmark` parameter to `cis`, and run Puppet.

- [Advanced configuration example](#) on page 202

Building on the basic configuration example, the advanced configuration example customizes the Advanced Intrusion Detection Environment (AIDE) configuration file in Hiera.

- [Automatically regenerate and enforce bootloader configurations](#) on page 203

In rare cases, it might be useful to enable automatic regeneration of the bootloader configuration, and you might want to set a bootloader password. By setting a bootloader password, you can provide an extra layer of security for your infrastructure.

- [Configure DISA STIG](#) on page 204

The US Defense Information Systems Agency (DISA) has developed Security Technical Implementation Guide (STIG) standards that are designed to secure information systems and software.

- [Configure authentication rules with PAM](#) on page 204

You can use a pluggable authentication module (PAM) to set authentication rules. To configure PAM settings, specify control values in Hiera.

- [Configure system authentication with the authselect utility](#) on page 205

The `authselect` utility can be used to configure system authentication on a Red Hat Enterprise Linux (RHEL) host. If you installed CEM on a RHEL 8 operating system, `authselect` options are available, but should be avoided in almost all cases. The `authselect` utility is disabled by default because enablement of `authselect` can break authentication methods, and use of the utility requires extensive configuration.

- [Apply CIS Benchmarks to a new Puppet Enterprise installation](#) on page 206

To ensure that the Compliance Enforcement Module (CEM) for Linux can enforce Center for Internet Security (CIS) Benchmarks or STIG standards after a new installation of Puppet Enterprise (PE), you must update the CEM configuration. The configuration update helps to ensure that CEM can work on a PE primary server without issues that might be caused by default log rotation rules.

- [Configure custom logrotate rules](#) on page 206

To help ensure that logs are pruned on a regular basis to conserve system space, you can specify `logrotate` rules.

- [Configure sudo without a password](#) on page 207

You can give users and user groups the ability to run some or all commands as root without a password.

- [Configure user SSH keys](#) on page 207

To use the Secure Shell (SSH) protocol for communication between computers, you must configure SSH keys. You can also configure SSH keys for individual users.

- [Configure SSH permissions for users and groups](#) on page 207

You can configure Secure Shell (SSH) protocol settings at a granular level to specify permissions for users and groups.

- [Configure the firewall type](#) on page 208

To help protect your infrastructure, CEM enforces a firewall solution, `firewalld`, by default. `firewalld` is zone-based software that is designed to monitor traffic and take appropriate action. To change the firewall type or specify that CEM does not manage a firewall, you can update the firewall configuration.

- [Configure rules that rely on site-specific information](#) on page 208

Some Center for Internet Security (CIS) rules require information that is specific to a customer site. You can use Puppet Bolt tasks to configure these rules. For more information about Puppet Bolt, see [Welcome to Puppet Bolt](#).

### Basic configuration example

When you specify a compliance framework, CEM is configured to provide rule enforcement and configuration for that framework. For example, to enforce the Center for Internet Security (CIS) Server Level 1 benchmark for a node, you must classify the node with the `cem_linux` class, set the `benchmark` parameter to `cis`, and run Puppet.

In the following example, CEM enforces the CIS Level 1 server controls `Ensure AIDE is installed` and `Ensure filesystem integrity is regularly checked` on a CentOS 7 node:

1. Add the following Hiera data to your control repository, `control_repo`:

```
control-repo/data/nodes/<node name>.yaml
cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 only:
 - 'ensure_aide_is_installed'
 - 'ensure_filesystem_integrity_is_regularly_checked'
```

2. Classify the node with the `cem_linux` class.
3. Run Puppet.

This example is for CIS configuration. For information about configuring STIG controls, see [Configure DISA STIG](#) on page 204.

Some CIS recommendations require you to run a Puppet Bolt task. To determine which task to run, review the output of the Puppet debug logs.

### Advanced configuration example

Building on the basic configuration example, the advanced configuration example customizes the Advanced Intrusion Detection Environment (AIDE) configuration file in Hiera.

1. Add the following code to the node's Hiera file:

```
control-repo/data/nodes/<node name>.yaml
cem_linux::benchmark: 'cis'
cem_linux::config:
```

```

profile: 'server'
level: '1'
only:
 - 'ensure_aide_is_installed'
 - 'ensure_filesystem_integrity_is_regularly_checked'
control_configs:
 ensure_aide_is_installed:
 conf_rules:
 - 'PERMS = p+u+g+acl+xattrs'
 - 'CONTENT_EX = sha256+ftype+p+u+g+n+acl+xattrs'
 conf_checks:
 - '/root/\.* PERMS'
 - '/root/ CONTENT_EX'

```

2. Classify the node with the `cem_linux` class.
3. Run Puppet.
4. Run the Puppet Bolt task that is specified in the debug log.

The AIDE configuration file now reflects the changes in Hiera.

This example is for CIS configuration. For information about configuring STIG controls, see [Configure DISA STIG](#) on page 204.

### Automatically regenerate and enforce bootloader configurations

In rare cases, it might be useful to enable automatic regeneration of the bootloader configuration, and you might want to set a bootloader password. By setting a bootloader password, you can provide an extra layer of security for your infrastructure.

**Restriction:** The only bootloader supported by CEM for Linux is `grub2`.

CEM for Linux enforces various bootloader configurations as required by the selected compliance framework and benchmark. However, because changes to bootloader configurations can be potentially dangerous, a minimalistic approach to configuration changes is used by CEM for Linux.

Several CIS recommendations modify the bootloader config. If you run CEM for Linux with the full range of default settings, these changes will be applied, but the bootloader config will not be regenerated. While changes are pending on the node, bootloader operations remain the same until the configurations are regenerated. The exception to this is the bootloader password, which is **not set** by default. To learn how to configure CEM for Linux to automatically regenerate the bootloader config and set the bootloader password, see the following topics:

- [Regenerate bootloader configs automatically](#) on page 203  
You can regenerate bootloader configs automatically by editing a Hiera `.yaml` file.
- [Set a bootloader password](#) on page 203  
You can set a bootloader password by editing a Hiera `.yaml` file.

### Regenerate bootloader configs automatically

You can regenerate bootloader configs automatically by editing a Hiera `.yaml` file.

Edit the `.yaml` file to specify the `regenerate` setting:

```

control-repo/data/nodes/<node name>.yaml

cem_linux::regenerate_grub2_config: true

```

### Set a bootloader password

You can set a bootloader password by editing a Hiera `.yaml` file.

You can set a bootloader password as shown:

```

control-repo/data/nodes/<node name>.yaml

cem_linux::regenerate_grub2_config: true

```

```

cem_linux::set_grub2_password: true
cem_linux::grub2_superuser: 'root'
cem_linux::grub2_superuser_password: 'password'
lookup_options:
 cem_linux::grub2_superuser_password:
 convert_to: 'Sensitive'

```

**Restriction:** The `cem_linux::grub2_superuser_password` key **must** be of type `Sensitive[String]`. Setting a lookup option for that key to convert it to `Sensitive` is the best way to ensure that the value is `Sensitive[String]`.



**CAUTION:** Do not store plain-text passwords in Hiera. To help protect passwords, use a hierarchy entry such as `hiera-eyaml`.

## Configure DISA STIG

The US Defense Information Systems Agency (DISA) has developed Security Technical Implementation Guide (STIG) standards that are designed to secure information systems and software.

### Before you begin

Verify that the DISA STIG standards are available for your operating system. See [Prepare to install the module](#) on page 193.

To configure DISA STIG, do not use the `profile` and `level` parameters, which are associated with the Center for Internet Security (CIS). Instead, specify the `mac` parameter to determine the Mission Assurance Category (MAC) level and the `confidentiality` parameter to determine the confidentiality level. The values that you specify will depend on the type of information that your system processes. For detailed information about specifying parameters, see the DISA STIG documentation and any relevant US Department of Defense instructions.

To configure DISA STIG, add Hiera data to your control repository, `control-repo`, as shown in the following example:

```

control-repo/data/nodes/<node name>.yaml
cem_linux::benchmark: 'stig'
cem_linux::config:
 # @param [Optional[Enum['1', '2', '3']]] mac
 # Which STIG benchmark Mission Assurance Category (MAC) level to
 # enforce.
 mac: '3'
 # @param [Optional[Enum['classified', 'sensitive', 'public']]]
 # confidentiality
 # Which STIG benchmark confidentiality level to enforce.
 confidentiality: 'public'

```

## Configure authentication rules with PAM

You can use a pluggable authentication module (PAM) to set authentication rules. To configure PAM settings, specify control values in Hiera.

For example, assume that you want to enforce a minimum length of 30 characters for passwords. Because you are implementing the CIS Oracle Linux 8 Benchmark 2.0.0, you go to the [CEM Linux Reference](#) and look for the relevant control in that benchmark. The control is 5.5.1, “Ensure password creation requirements are configured,” which specifies a default minimum password length of 14. On the `minlen` parameter, you replace the default value of 14 with a new value of 30, as shown in the example:

```

cem_linux::config:
 control_configs:
 "Ensure password creation requirements are configured":
 manage_pwquality: true
 manage_pam_auth: true
 minlen: 30

```

```

minclass: 4
faillock_args: ["preauth", "silent", "audit", "deny=5",
"unlock_time=900"]
pwhistory_args: ["use_authtok", "remember=5", "retry=3"]

```

### Configure system authentication with the authselect utility

The `authselect` utility can be used to configure system authentication on a Red Hat Enterprise Linux (RHEL) host. If you installed CEM on a RHEL 8 operating system, `authselect` options are available, but should be avoided in almost all cases. The `authselect` utility is disabled by default because enablement of `authselect` can break authentication methods, and use of the utility requires extensive configuration.

### Guidelines and restrictions

If `authselect` is used by your system to configure authentication, and you instead want CEM to fully manage authentication, CEM must be configured to manage authentication by using pluggable authentication modules (PAM) directly. For instructions, see [Configure authentication rules with PAM](#) on page 204.



**CAUTION:** When using CEM with `authselect` to configure system authentication, you must set the `use_authselect` option to `true`. Failure to do so will cause CEM to attempt to manage system authentication by using PAM instead of `authselect` and can compromise the security of your system's authentication configuration.

Before you configure system authentication with `authselect`, review the following restrictions:

- If a node is joined to an Active Directory domain or to Red Hat Identity Management (idM), do not enable the `authselect` utility. Enabling the `authselect` utility on these nodes *will break* your authentication configurations.
- To fully configure system authentication with CEM, you must use PAM instead of `authselect`.

### Authselect options

The following `authselect` options are available for RHEL 8:

- `use_authselect`: - Optional[Boolean] - Whether to use `authselect` to manage most authentication options. Defaults to `false`.
- `authselect_profile`: - Optional[String] - Profile for `authselect` configuration options. If using the `authselect` utility, you must specify an `authselect` profile. Defaults to `undef`.

### Enabling the authselect utility

Both of the `authselect` options *must* be set directly in the `cem_linux::config` hash for the `authselect` utility to work properly. All `authselect` configurations are managed via the `ensure_custom_authselect_profile_is_used` control, regardless of whether you use a custom profile.

To enable the `authselect` utility:

1. Set the config option `use_authselect` to `true`.
2. Specify an `authselect` profile with the config option `authselect_profile`.

By default, `cem_linux` uses standard PAM rules to configure the authentication controls specified by CIS. However, if you are enforcing CIS compliance on RHEL 8, CIS guidelines call for the `authselect` utility to be used. The following configuration example shows how to enable the `authselect` utility on a node by using the minimal system default profile:

```

control-repo/data/nodes/<node name>.yaml

cem_linux::config:
 use_authselect: true
 authselect_profile: 'minimal'

```

## Custom authselect profiles

If you are enforcing CIS compliance on a RHEL 8 system and you want to enable additional features for your authselect profile, you can create a custom profile.

To create a custom authselect profile in `cem_linux`, prefix the profile name in `authselect_profile` with `custom/`. If the custom profile does not exist on the node, the profile will be created automatically. The following example shows how to create and use a custom profile, `my_custom_profile`, which is based on the minimal system profile with additional features enabled:

```
control-repo/data/nodes/<node name>.yaml

cem_linux::config:
 use_authselect: true
 authselect_profile: 'custom/my_custom_profile'
 control_configs:
 ensure_custom_authselect_profile_is_used:
 custom_profile_base: 'minimal'
 profile_features:
 - with-faillock
 - with-mkhomedir
```

For more information about authselect features, see the [authselect documentation](#) in the [Red Hat Customer Portal](#).

## Apply CIS Benchmarks to a new Puppet Enterprise installation

To ensure that the Compliance Enforcement Module (CEM) for Linux can enforce Center for Internet Security (CIS) Benchmarks or STIG standards after a new installation of Puppet Enterprise (PE), you must update the CEM configuration. The configuration update helps to ensure that CEM can work on a PE primary server without issues that might be caused by default log rotation rules.

Add the following Hiera data to your control repository, `control_repo`.

```
cem_linux::config:
 control_configs:
 ensure_logrotate_is_configured:
 rules:
 puppetserver:
 path:
 - '/var/log/puppetlabs/puppetserver/puppetserver.log'
 - '/var/log/puppetlabs/puppetserver/pcp-broker.log'
 - '/var/log/puppetlabs/puppetserver/puppetserver-access.log'
 - '/var/log/puppetlabs/puppetserver/puppetserver-daemon.log'
 - '/var/log/puppetlabs/puppetserver/puppetserver-status.log'
 - '/var/log/puppetlabs/puppetserver/code-manager-access.log'
 - '/var/log/puppetlabs/puppetserver/file-sync-access.log'
 - '/var/log/puppetlabs/puppetserver/masterhttp.log'
 create_owner: 'puppet'
 create_group: 'puppet'
```

## Configure custom logrotate rules

To help ensure that logs are pruned on a regular basis to conserve system space, you can specify logrotate rules.

The following example creates custom logrotate rules for the primary Puppet server's puppetserver logs.

```
control-repo/data/nodes/<your puppetserver>.yaml

cem_linux::config:
 control_configs:
 ensure_logrotate_is_configured:
 rules:
 puppetserver:
```

```

path:
- '/var/log/puppetlabs/puppetserver/puppetserver.log'
- '/var/log/puppetlabs/puppetserver/pcp-broker.log'
- '/var/log/puppetlabs/puppetserver/puppetserver-access.log'
- '/var/log/puppetlabs/puppetserver/puppetserver-daemon.log'
- '/var/log/puppetlabs/puppetserver/puppetserver-status.log'
- '/var/log/puppetlabs/puppetserver/code-manager-access.log'
- '/var/log/puppetlabs/puppetserver/file-sync-access.log'
- '/var/log/puppetlabs/puppetserver/masterhttp.log'
create_owner: 'puppet'
create_group: 'puppet'

```

### Configure sudo without a password

You can give users and user groups the ability to run some or all commands as root without a password.

The following example configures the `admins` group to grant sudo access without a password:

```

cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 control_configs:
 ensure_sudo_is_installed:
 package_ensure: 'installed'
 options:
 user_group:
 %admins:
 options:
 - 'NOPASSWD:'

```

### Configure user SSH keys

To use the Secure Shell (SSH) protocol for communication between computers, you must configure SSH keys. You can also configure SSH keys for individual users.

The following example shows how to configure keys for two users on Red Hat Enterprise Linux 8. This example uses options of the `cem_linux::utils::packages::linux::ssh` class that are not currently documented in the reference, but are fully supported for use. To use undocumented options, or options that are supported but do not fit directly under a control, such as options documented under the `cem_options` key, you can place the options under any control associated with your specified backing resource. Controls and their backing resources can be found in the [CEM Linux Reference](#).

In the example, SSH root login is permitted, and keys are configured for `testuser1` and `testuser2`:

```

cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 control_configs:
 ensure_ssh_root_login_is_disabled:
 permit_root_login: 'yes'
 user_ssh_keys:
 testuser1:
 home_dir: /home/testuser1
 ssh_key: ssh-rsa A...ZcTFw== rsa-key-20201022
 testuser2:
 home_dir: /home/testuser2
 ssh_key: ssh-rsa A...ZcTFw== rsa-key-20201022

```

### Configure SSH permissions for users and groups

You can configure Secure Shell (SSH) protocol settings at a granular level to specify permissions for users and groups.

The following example configures SSH to grant permissions to some users and groups and deny permissions to other users and groups:

```
cem_linux::benchmark: 'cis'
cem_linux::config:
 control_configs:
 ensure_permissions_on_etcsshsshd_config_are_configured:
 allow_users:
 - testuser1
 - the_dude
 allow_groups:
 - testgroup1
 - goonies
 deny_users:
 - testuser2
 - the_emperor
 deny_groups:
 - testgroup2
 - legion_of_doom
```

### Configure the firewall type

To help protect your infrastructure, CEM enforces a firewall solution, `firewalld`, by default. `firewalld` is zone-based software that is designed to monitor traffic and take appropriate action. To change the firewall type or specify that CEM does not manage a firewall, you can update the firewall configuration.

**Restriction:** Firewalls that are based on the `nftables` framework are not supported. Use the `firewalld` or `iptables` setting instead.

The following examples show how to configure a firewall type.

The default setting is `firewalld`:

```
cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 firewall_type: 'firewalld'
```

You can also specify a value of `iptables`:

```
cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 firewall_type: 'iptables'
```

You can also specify a value of `unmanaged`. If you specify `unmanaged`, CEM does not enforce a state on any firewall resource:

```
cem_linux::benchmark: 'cis'
cem_linux::config:
 profile: 'server'
 level: '1'
 firewall_type : 'unmanaged'
```

### Configure rules that rely on site-specific information

Some Center for Internet Security (CIS) rules require information that is specific to a customer site. You can use Puppet Bolt tasks to configure these rules. For more information about Puppet Bolt, see [Welcome to Puppet Bolt](#).

By using Puppet Enterprise (PE), you can run Puppet Bolt tasks and plans to audit or configure specific parts of a node. To run Puppet Bolt tasks, open the PE console and select the **Tasks** menu. Then, select **cem\_linux**.

You can also use open source Puppet to run Puppet Bolt tasks and plans. If you are using open source Puppet, run Puppet Bolt tasks from the command line:

1. Install [Puppet Development Kit \(PDK\)](#) and [Bolt](#).
2. In the root of the CEM directory, run the `pdk bundle exec rake 'spec_prep'` command. This command downloads the required dependencies as RSpec fixtures, and then creates a symbolic link from the module directory to the fixtures directory.
3. Run the tasks on one or more hosts. For example:

```
bolt task run
 comply_enforcement_module::audit_unowned_files_and_directories -t
 $nodefqdn --modulepath spec/fixtures/modules
```

You must add the `--modulepath spec/fixtures/modules` option to Puppet Bolt commands. Otherwise, Puppet Bolt is not able to find the tasks and plans.

## Auditing and querying issues identified during scans

In some cases, a CIS or DISA STIG compliance scan might identify an issue that you want to investigate and fix. To get started, you can run an audit or query.

CEM for Linux supports Puppet Bolt *tasks* that you can use to run an audit or query. For more information, go to Puppet Forge and review the [Tasks](#) list.

**Tip:** Some tasks have associated parameters. In the Tasks list, you can click the green downward arrow to display any associated parameters.

You can run tasks individually by following the instructions in [Running tasks](#).

You can run all audit tasks *simultaneously* by using a Bolt plan, `run_audit`. To run the Bolt plan, use a command with the following structure:

```
bolt plan run cem_linux::run_audit -targets <IP_address_of_target_node>
--tasks_dir <relative_path_of_cem_tasks_directory> --user <user_name>
--password <user_password>
```

On this command, the following parameters are optional:

| Optional parameter     | Details                                                                                                                                  |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <code>tasks_dir</code> | Specify this parameter only if your tasks are not in the default directory:<br><br>./tasks                                               |
| <code>user</code>      | Specify this parameter only if you are running the command on a remote node, and you require a username and password to access the node. |
| <code>password</code>  | Specify this parameter only if you require a password to access a remote node.                                                           |

To run all audit tasks simultaneously from your *local* computer, issue the following command:

```
bolt plan run cem_linux::run_audit --targets localhost
```

To run all audit tasks simultaneously on a *remote* node, issue a command similar to one of the following examples:

- If the IP address of the remote node is 34.83.231.128, your username is expert23, and your password is RainInSpain!, issue the following command:

```
bolt plan run cem_linux::run_audit --targets 34.83.231.128
--user expert23 --password RainInSpain!
```

- However, if a username and password are not required to access the remote node, you would issue the following command:

```
bolt plan run cem_linux::run_audit --targets 34.83.231.128
```

- If your tasks are not in the default directory (./tasks) but are in the /cem/tasks directory, you would issue the following command:

```
bolt plan run cem_linux::run_audit --targets 34.83.231.128
--tasks_dir /cem/tasks
```

You can also specify *multiple* target nodes. For instructions, see [TargetSpec](#).

After you run the plan, review the output in the Bolt log file on the computer where you ran the command. The following sample output is for an individual task, `audit_check_ipv6`, that was run as part of a Bolt plan. In this case, the audit was successful:

```
Task ->
 Task_name: audit_check_ipv6,
 Task_details: Audit IPV6 for RHEL8
Task: audit_check_ipv6
Starting: task cem_linux::audit_check_ipv6 on 34.83.231.128
Finished: task cem_linux::audit_check_ipv6 with 0 failures in 6.27 sec

Task Result: [{"target": "34.83.231.128", "action": "task",
"object": "cem_linux::audit_check_ipv6", "status": "success",
"value": {"_output": "\nIPv6 is enabled on the system\n\n"}}]
```

For more information about running Bolt plans, see [Orchestrating workflows with plans](#).

## Reference: Benchmarks and controls

---

For help with configuring CEM, review the [CEM Linux Reference](#) on Puppet Forge.

In the [CEM Linux Reference](#), locate the relevant CIS Benchmark or DISA STIG standard and then the control.

For information about control updates that were implemented between CEM releases, see the following topics:

- [Control updates introduced for Red Hat Enterprise Linux 8 STIG, Version 1, Release 11](#) on page 210  
The Compliance Enforcement Module (CEM) for Linux v1.7.0 introduces enforcement for an updated Security Technical Implementation Guide (STIG) standard: Red Hat Enterprise Linux STIG - Version 1, Release 11. The transition from the previously supported version (Version 1, Release 8) resulted in module updates.
- [Control updates introduced for Red Hat Enterprise Linux 7 STIG, Version 3, Release 12](#) on page 211  
The Compliance Enforcement Module (CEM) for Linux v1.7.0 introduces enforcement for an updated Security Technical Implementation Guide (STIG) standard: Red Hat Enterprise Linux STIG 7 - Version 3, Release 12. The transition from the previously supported version (Version 3, Release 8) resulted in module updates.

### Control updates introduced for Red Hat Enterprise Linux 8 STIG, Version 1, Release 11

The Compliance Enforcement Module (CEM) for Linux v1.7.0 introduces enforcement for an updated Security Technical Implementation Guide (STIG) standard: Red Hat Enterprise Linux STIG - Version 1, Release 11. The transition from the previously supported version (Version 1, Release 8) resulted in module updates.

- **Added**

- The following controls are added to the module:
  - V-255924 - RHEL 8 SSH server must be configured to use only FIPS-validated key exchange algorithms.
  - V-256974 - RHEL 8 must be configured to allow sending email notifications of unauthorized configuration changes to designated personnel.
  - V-257258 - RHEL 8 must terminate idle user sessions.

- **Changed**

- The following controls are changed in the module:
  - V-230244 - RHEL 8 must be configured so that all network connections associated with SSH traffic are terminated at the end of the session or after 10 minutes of inactivity, except to fulfill documented and validated mission requirements.

Previously, the system could send up to 600 messages to a client without receiving a response, and the session was still considered active. In the new release, the `ClientAliveCountMax` threshold is changed from 600 to 1. As a result, if the system sends a single message without receiving a response, the session is considered inactive and is terminated. This termination reduces the window of opportunity for unauthorized personnel to take control of a management system that is unattended.

- V-230251 - The RHEL 8 SSH server must be configured to use only Message Authentication Codes (MACs) employing FIPS 140-2 validated cryptographic hash algorithms.

The control is updated with the following new default setting to help prevent unauthorized access to data:

```
oMACS=hmac-sha2-512,hmac-sha2-256,hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com
```

- V-230252 - The RHEL 8 operating system must implement DoD-approved encryption to protect the confidentiality of SSH server connections.

The control is updated with the following new default setting to help prevent unauthorized users from altering data:

```
CRYPTO_POLICY='-oCiphers=aes256-ctr,aes192-ctr,aes128-ctr,aes256-gcm@openssh.com,aes128-gcm@openssh.com'
```

- V-230525 - A firewall must be able to protect against or limit the effects of Denial of Service (DoS) attacks by ensuring RHEL 8 can implement rate-limiting measures on impacted network interfaces.

The control is updated to ensure that `nftables` is the backend used by `firewalld`.

- A Puppet Bolt task was updated to support the functionality in the following control:

V-256973 - RHEL 8 must ensure cryptographic verification of vendor software packages.

The Bolt task `query_rpm_gpg_keys` can now be used to report the data required by this control. To ensure compliance with this control, run the Bolt task `query_rpm_gpg_keys` and verify the output.

- **Removed**

- The following control is removed from the module:
  - V-230289 - The RHEL 8 SSH daemon must not allow compression or must only allow compression after successful authentication.

## Control updates introduced for Red Hat Enterprise Linux 7 STIG, Version 3, Release 12

The Compliance Enforcement Module (CEM) for Linux v1.7.0 introduces enforcement for an updated Security Technical Implementation Guide (STIG) standard: Red Hat Enterprise Linux STIG 7 - Version 3, Release 12. The transition from the previously supported version (Version 3, Release 8) resulted in module updates.

- **Added**
  - The following controls are added to the module:
    - V-255925 - The Red Hat Enterprise Linux operating system SSH server must be configured to use only FIPS-validated key exchange algorithms.
    - V-255926 - The Red Hat Enterprise Linux operating system must have the screen package installed.
    - V-255927 - The Red Hat Enterprise Linux operating system must restrict access to the kernel message buffer.
    - V-255928 - The Red Hat Enterprise Linux operating system must be configured to prevent overwriting of custom authentication configuration settings by the `authconfig` utility.
    - V-256969 - The Red Hat Enterprise Linux operating system must disable the login screen user list for graphical user interfaces.
    - V-256970 - The Red Hat Enterprise Linux operating system must be configured to allow sending email notifications of configuration changes and adverse events to designated personnel.

## CEM for Windows

---

You can deploy CEM for Windows to help ensure that your servers and workstations on Microsoft Windows operating systems comply with CIS Benchmarks.

To get started, review the basic concepts and then follow the instructions to deploy CEM in your environment. See [Getting started](#) on page 220.

- [Release notes](#) on page 212

Review the release notes to learn about updates and resolved issues in the Compliance Enforcement Module (CEM) for Windows.

- [Getting started](#) on page 220

Learn the basic concepts associated with the Compliance Enforcement Modules and then review the steps for deploying CEM in your environment.

- [Installing CEM](#) on page 222

Before you install CEM, complete the preparation steps: review the system requirements, install and configure Puppet Enterprise (PE) or open source Puppet, and purchase CEM. To help avoid issues, install and evaluate CEM in a *test* environment before you install CEM in a production environment.

- [Upgrading CEM](#) on page 225

You can upgrade CEM for Windows to take advantage of the latest features, fixes, and improvements. To help ensure a smooth upgrade process, complete the preparation tasks first.

- [Configuring CEM](#) on page 227

Configuration of CEM is optional. If you installed CEM and assigned the `cem_windows` class to one or more node groups, the Center for Internet Security (CIS) Server Level 1 profile is enforced automatically during the next Puppet run. However, if the default values leave your infrastructure in an undesirable state, or if you want to customize compliance to meet your organization's requirements, you can configure CEM.

- [Reference: Benchmarks and controls](#) on page 233

For help with configuring CEM, review the [CEM Windows Reference](#) on Puppet Forge.

## Release notes

---

Review the release notes to learn about updates and resolved issues in the Compliance Enforcement Module (CEM) for Windows.

### v1.5.2

Released 19 March 2024

CEM for Windows v1.5.2 introduces updates to enhance protection of Windows Server systems. Default values were changed for three Center for Internet Security (CIS) controls, thus helping to ensure that the controls will be correctly enforced to protect the `winreg` registry key and internal system objects.

### Resolved issues

- For Windows Server 2016, 2019, and 2022, the implementation of CIS Controls 2.3.10.8 and 2.3.10.9 was corrected. For both controls, the default value of the `value` parameter was changed to `Machine`. By enforcing these controls, you can help to prevent attackers from accessing sensitive configuration data in the `winreg` registry key.
- For Windows Server 2016, 2019, and 2022, the implementation of CIS Control 2.3.15.2 was updated to specify the correct path for the `path` parameter. By enforcing this control, you can help to prevent unauthorized users from modifying internal system objects.
- A default value was changed to help ensure that CIS Control 18.6.4.1 can be enforced without disrupting operations on Windows Server 2022 systems. CIS Control 18.6.4.1 enforces Domain Name System resolution over HTTPS (DoH) to help protect systems against spoofing and man-in-the-middle attacks. Previously, the default setting of `Enabled: Require DoH` could prevent agent nodes from reporting to the Puppet primary server. To resolve the issue, the setting was changed to `Enabled: Allow DoH` to ensure that DoH is allowed but not required.

### v1.5.1

Released 6 October 2023

- **Changed**
  - Introduced a change that is designed to simplify CEM for Windows configuration. In previous releases, CEM for Windows was configured to ignore controls related to the renaming of Administrator and Guest accounts. This configuration was designed to avoid rare cases in which the control settings could cause Puppet run failures. As a result of this default behavior, users who wanted to enable the controls had to specify an `ignore list` that did not include the controls. Specifying the controls in an `only list` was not helpful because the `ignore list` overrode the `only list`. To resolve this issue, the default setting of the `ignore list` was changed to empty.
- **Fixed**
  - Fixed an issue that prevented some user-specified configuration options from being applied. The issue affected only some parameters on some controls.

### v1.5.0

Released 22 August 2023

- **Changed**
  - This release includes updates that are designed to enhance security on Microsoft Windows 10 Enterprise, Windows Server 2019, and Windows Server 2016 operating systems:
    - For users of the Microsoft Windows 10 Enterprise operating system, the Center for Internet Security (CIS) Benchmark was upgraded from v1.12.0 to v2.0.0. For a list of control updates, see [Control updates introduced for CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0](#) on page 233.
    - For users of the Windows Server 2019 operating system, the CIS Benchmark was upgraded from v1.3.0 to v2.0.0. For a list of control updates, see [Control updates introduced for CIS Microsoft Windows Server 2019 Benchmark v2.0.0](#) on page 234.
    - For users of the Windows Server 2016 operating system, the CIS Benchmark was upgraded from v1.4.0 to v2.0.0. For a list of control updates, see [Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v2.0.0](#) on page 235.
  - The CEM for Windows documentation now provides more detailed upgrade instructions, including preparation steps that you can take to help ensure a smooth upgrade. See [Upgrading CEM](#) on page 225.

- **Fixed**
  - Fixed an issue related to the `cem_domain_controller` fact, which was incorrectly reporting a value of `false` in all instances. Now, the `cem_domain_controller` fact correctly reports a value of `true` when CEM for Windows runs on a domain controller.

### v1.4.0

Released 27 June 2023

- **Added**
  - Enforcement of the Center for Internet Security (CIS) Microsoft Windows Server 2022 Benchmark v2.0.0.
- **Changed**
  - `cem_windows` no longer supports the use of legacy configuration as of this update. Legacy configuration refers to configurations of `cem_windows` used prior to the release of v1.1.0. `cem_windows` is no longer compatible with configurations that were used before v1.1.0. Please update any legacy configuration to the current standard of configuring `cem_windows`.

### v1.3.0

Released 15 December 2022

This release includes updates for users of the Microsoft Windows Server 2016 operating system. With this release, users can enforce Center for Internet Security (CIS) Microsoft Windows Server 2016 Benchmark v1.4.0. For a list of control updates, see [Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v1.4.0](#) on page 236.

### v1.2.3

Released 25 October 2022

- **Added**
  - Added a Puppet Puppet Bolt task, `cem_delete_securitypolicy_inf`, to use for error resolution. The Puppet Bolt task resolves a corruption error that can affect the temporary file that is used by Desired State Configuration (DSC) to manage the local security policy:
    - The error is indicated by the following message in the Puppet run log:

```
Index operation failed; the array index evaluated to null
```
    - To resolve the error, run the `cem_delete_securitypolicy_inf` task and re-run Puppet on the affected node.
- **Changed**
  - The product documentation was revised to improve usability and retrievability:
    - The change log was migrated from Puppet Forge to the central location for Puppet documentation, Puppet Docs. The change log was renamed to [Release notes](#) on page 212.
    - The readme file was transformed into a series of topics with a structure similar to other Puppet documentation. The CEM topics are now available on Puppet Docs, starting with [Introducing the Compliance Enforcement Modules](#) on page 173.
    - The [Reference](#) and [Dependencies](#) documentation, which is generated automatically, remains on Puppet Forge.
- **Fixed**
  - Fixed an error that prevented catalog retrieval from Puppet Enterprise (PE) during Continuous Delivery for Puppet Enterprise pipeline runs. This error occurred when the impact analysis tool was used to set up a temporary environment, which was then deleted. The `_FILE_` variable continued to point to the deleted

environment. As a result, the Puppet run returned an error message: Could not retrieve catalog from remote server.

## **v1.2.2**

Released 10 August 2022

### **Fixed**

- Fixed typos in Microsoft Windows firewall logging paths managed by the following controls:
  - CIS Windows 10
    - 9.1.5
    - 9.2.5
    - 9.3.7
  - CIS Windows Server 2016
    - 9.1.5
    - 9.2.5
    - 9.3.7
  - CIS Windows Server 2019
    - 9.1.5
    - 9.2.5
    - 9.3.7

- Fixed an issue that could cause the following controls to not be enforced:
  - CIS Windows 10
    - 18.9.17.2
    - 18.9.64.1
    - 18.9.65.3.10.1
    - 18.9.65.3.10.2
    - 18.9.65.3.2.1
    - 18.9.72.1
    - 18.9.75.1
    - 18.9.103.1
  - CIS Windows Server 2016
    - 18.9.45.10.1
  - CIS Windows Server 2019
    - 18.9.41.1
    - 18.9.45.1
    - 18.9.47.11.1
    - 18.9.65.3.10.1
    - 18.9.65.3.10.2
    - 18.9.65.3.2.1
    - 18.9.65.3.3.1
    - 18.9.65.3.3.3
    - 18.9.65.3.3.4
    - 18.9.67.2
    - 18.9.72.1
    - 18.9.89.1
    - 18.9.90.3
    - 18.9.102.2.2
    - 18.9.103.1
    - 18.9.47.5.1.2

### **v1.2.1**

Released 31 May 2022

#### **Fixed**

- Fixed a bug related to profile configuration on Microsoft Windows 10 nodes.

### **v1.2.0**

Released 24 May 2022

#### **Changed**

- Updated the Center for Internet Security (CIS) Windows Server 2019 Benchmark to version 1.3.0.

- **Fixed**
  - Resolved issues leading to scan failures for the following CIS controls on Windows Server 2019:
    - 9.3.7
    - 9.2.5
    - 9.1.5
    - 18.9.108.4.1
    - 18.9.65.3.9.1
    - 18.8.3.1
    - 18.8.21.5
    - 18.5.21.1
    - 18.4.x
    - 18.2.1

### v1.1.2

Released 12 May 2022

- **Changed**
  - Updated the minimum required version of the `dsc/auditpolicydsc` module to `1.4.0-0-4`. That dependency contains bug fixes and features required by `cem_windows`. Update your Puppetfile accordingly.
- **Fixed**
  - Updated the default value for the Windows Attack Surface Reduction (ASR) rules to `Audit` instead of `Block`.
    - While the value of `Audit` is not CIS-compliant, setting the ASR rules to `Block` prevented the Puppet agent from successfully configuring the node.
    - If you see Puppet run errors like `Could not evaluate: undefined method []' for nil:NilClass` when enforcing CEM, manually set the Windows ASR rules to `Audit`. To learn more about Windows ASR rules, see [Attack surface reduction rules overview](#).
  - Fixed an issue that applied more controls to a node than required by the configured profile and level.
  - Fixed an issue that caused controls that should be ignored to be applied. This issue occurred when the controls were mapped to a parameter of a resource that was not ignored.
  - Fixed several issues related to configuration backward-compatibility.

**Upgrade requirement:** To ensure that the updates in this release take effect, you might have to restart the `pe-puppetserver` service on your Puppet primary server after Code Manager deploys the new code.

### v1.1.1

Released 7 April 2022

- **Changed**
  - Improved the display of controls in the [CEM Windows Reference](#).
- **Fixed**
  - Fixed several instances in which configurations from versions previous to v1.1.0 were not recognized. The v1.1.1 configuration is backward compatible with versions prior to v1.1.0.
  - Fixed an issue that required the `cem_windows` module to exist in the same environment as the Puppet primary server. You can now deploy the module to a different environment than your primary server. The module will be operational.
  - Fixed incorrect Puppet Strings in `init.pp` file.

**v1.1.0**

Released 24 March 2022

- **Added**
  - The documentation was updated to list the controls that will be reported as failed or unknown in Comply after `cem_windows` is applied.

**Tip:** A failed or unknown status is reported because the CIS-CAT Pro Assessor looks for registry keys that are configured by Microsoft Group Policy Objects rather than keys that are set locally by the `cem_windows` user. The CIS Windows benchmarks are designed to work only for domain-joined systems. At the time of the v1.1.0 release, CIS was working on Windows benchmarks for a standalone system to resolve the issue.

- **Changed**
  - Updated the CIS Windows 10 Benchmark to v1.12.0 to match the latest benchmark version released with Comply 2.4.0.
  - The `cem_windows` module was updated to implement a new architecture. The new architecture, applied in the background, provides more flexibility for system configuration. For details, see the [readme file](#).

**v1.0.7**

Released 16 December 2021

- **Removed**
  - Removed unnecessary resource defaults in two Windows Server 2016 control classes.

**v1.0.6**

Released 16 December 2021

- **Removed**
  - Removed unnecessary resource defaults in Windows Server 2016 control classes.

**v1.0.5**

Released 8 December 2021

- **Fixed**
  - Fixed non-idempotent Desired State Configuration (DSC) resources.
  - Fixed the registry key for Windows 10 CIS control 1.1.6. Now, this control will be properly configured.

**v1.0.4**

Released 7 December 2021

- **Added**
  - In the [readme file](#), added a link to premium content installation instructions. To use CEM, you must be a premium content subscriber.
- **Fixed**
  - Fixed an issue that caused values for the `dsc_accountpolicy` parameter to be set incorrectly.

**v1.0.3**

Released 13 October 2021

- **Fixed**
  - Fixed the default value for CIS control 2.3.1.1 to align with the expected value provided by CIS.
  - Fixed the `cem_windows::allow_local_account_rdp` parameter so that it works as intended.

### v1.0.2

Released 11 October 2021

- **Fixed**
  - Fixed firewall profiles to align with the CIS specification.

### v1.0.1

Released 30 September 2021

- **Fixed**
  - Fixed the Windows 10 Hiera name to ensure that Windows 10 can be used. For more information about Hiera, see [Configure settings with Hiera](#).

## Known issues and limitations

The current release includes known issues and limitations. In most cases, workarounds are provided.

- **On Windows Server 2022 systems, communication between agent nodes and the Puppet primary server can fail.** This issue can occur in CEM for Windows v1.5.1 and earlier when the following control is enforced: CIS Windows Server 2022 Benchmark (2.0.0) Control 18.6.4.1. In these circumstances, nodes might be prevented from sending reports to the Puppet primary server. The issue occurs because the control's default setting, `Enabled: Require DoH`, enforces Domain Name System resolution over HTTPS (DoH). The issue is resolved in CEM for Windows v1.5.2, in which the default setting was changed to `Enabled: Allow DoH`.
- **An incorrect top-level key is shown in Hiera configuration examples.** On Puppet Forge, the "Reference" section incorrectly shows `"puppetlabs-cem_windows::config:"` as the top-level key in the Hiera configuration examples. The correct top-level key is `"cem_windows::config:"`.
- **A registry key override can occur when duplicate normalized names are used to specify CIS controls.** The issue occurs because the normalized control names for two authentication settings related to Windows Remote Management (WinRM) are identical. The normalized control names are the same for both the client (18.9.102.1.1) and service (18.9.102.2.1) controls. The workaround is to configure the controls with the control numbers (18.9.102.1.1 and 18.9.102.2.1) or the normalized control numbers (`c18_9_102_1_1` and `c18_9_102_2_1`). This issue occurs only in Windows Enterprise 10 environments.
- **In a Windows Server 2016 or 2019 environment, a scan failure is reported for CIS Control 2.3.10.12.** The failure affects the following control: 2.3.10.12, (L1) Ensure 'Network access: Shares that can be accessed anonymously' is set to 'None.' This control is enforced correctly but fails Comply scans. The scans detect the backing registry value, type `REG_MULTI_SZ`. The expected value is a blank item as the first line in a multiline string. However, the Puppet module that manages registry settings does not permit the use of blank values. As a workaround, no value is set for the backing registry. A blank value and no value are functionally equivalent, resulting in the same configuration. For this reason, you can ignore the reported scan failure.
- **The Center for Internet Security (CIS) Control 18.2.1, 'Ensure LAPS AdmPwd GPO Extension / CSE is installed,' is not enforced.** Control 18.2.1 requires downloading and installing the Local Administrator Password Solution (LAPS) client from the Microsoft website. Because CEM for Windows does not support third-party Windows package managers, this software cannot be installed. In addition, the CIS scanner scans only for the presence of the LAPS client .dll file but does not confirm that LAPS is configured or functional at the domain level.
- **After an upgrade, you might have to restart Puppet Server or the `pe-puppetserver` service.** Starting with v1.1.0, CEM for Windows implements a new architecture. If you upgrade CEM from v1.0.7 or earlier to v1.1.0 or later, and you encounter errors, try restarting the `pe-puppetserver` service or restarting or reloading Puppet Server. For instructions, see [Restarting Puppet Server](#).

- **You might have to manually set Windows Attack Surface Reduction (ASR) rules to Audit.** In `cem_windows` releases prior to v1.1.2, a default value of `Block` was set in the module to comply with CIS guidelines. However, the `Block` value prevented the Puppet agent from successfully configuring the node. For this reason, the default value was changed to `Audit`, which is not CIS compliant. If you see Puppet run errors like `Could not evaluate: undefined method []' for nil:NilClass` when enforcing CEM, manually set the Windows ASR rules to `Audit`. To learn more about Windows ASR rules, see [Attack surface reduction rules overview](#).
- **Some controls can fail scans.** During a Comply scan, you might see error messages about CIS recommended guidelines that are not enforced. These error messages are triggered by bugs in the CIS-CAT Pro Assessor that is bundled with Comply. CEM correctly enforces these settings. The following controls are affected:
  - 1.1.5 - Windows Server 2016 and Windows Server 2019
  - 1.1.6 - Windows Server 2016 and Windows Server 2019
  - 2.3.10.7 - Windows Server 2016
  - 18.2.1 - Windows Server 2019
  - 18.4.1 - Windows Server 2016 and Windows Server 2019
  - 18.4.8 - Windows Server 2016
  - 18.4.9 - Windows Server 2016 and Windows Server 2019
  - 18.4.12 - Windows Server 2016
  - 18.8.21.5 - Windows Server 2016
  - 18.9.47.5.1.2 - Windows Server 2019
  - 18.9.62.3.9.1 - Windows Server 2016
- **Puppet runs are not idempotent.** If you see Desired State Configuration (DSC) resources showing corrective changes in a Puppet run, for example, `Unknown feature "custom_isync"`, you are running an incompatible version of Puppet. CEM for Windows requires that Puppet agents at the version 6 level must be v6.23.0 or later, and agents at the version 7 level must be v7.8.0 or later.
- **If the Puppet agent fails to upgrade when you use the `puppetlabs/puppet_agent` module, restart the computer or virtual machine where the Puppet agent is running to help ensure that updates are applied.**
- **If you use remote desktop protocol (RDP) to access nodes, users who are members of the groups `Guests` and local accounts will not be able to log in by default.** To provide access to these groups, set the `cem_windows::allow_local_account_rdp` parameter to `true`.
- **If non-admin users cannot log in to nodes, the issue might be related to event logs.** By default, Windows Event Log does not clear events. When the event log of a node is full, only administrators can log in. To clear the event logs manually, find the specific recommendation in your compliance framework and configure the setting. In the Windows registry, locate the following key:
 

```
HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\EventLog
\Application:Retention
```

Then, set the `Retention` value to 0.
- **You cannot disable Windows Remote Management (WinRM).** The WinRM service is required for the DSC modules and cannot be disabled.

## Getting started

---

Learn the basic concepts associated with the Compliance Enforcement Modules and then review the steps for deploying CEM in your environment.

- [Basic concepts](#) on page 221

The basic concepts include an overview of CEM and the security standards that it enforces. You can also learn about `Hiera`, a key-value store that is used to configure CEM.

- [Next steps](#) on page 222

After you are familiar with the basic concepts, you can take the next steps.

## Basic concepts

The basic concepts include an overview of CEM and the security standards that it enforces. You can also learn about [Hiera](#), a key-value store that is used to configure CEM.

### Compliance Enforcement Modules

CEM is software that automatically enforces security standards on IT infrastructures. After CEM is installed and configured, you can run Puppet Enterprise (PE) or open source Puppet on your specified nodes, and CEM automatically enforces security controls.

### Hiera

To configure CEM, you can use the *Hiera* key-value store. Hiera stores configuration data in a hierarchical structure in key-value pairs. For an introduction, see [About Hiera](#).

### Center for Internet Security (CIS)

The *Center for Internet Security, Inc.*, is a nonprofit organization that strives to protect IT infrastructures through collaboration and innovation. Contributors to the organization include security experts from government, business, and academia who develop and maintain internationally recognized security standards. For more information, see [Center for Internet Security](#).

### CIS Benchmarks

CIS develops and maintains *CIS Benchmarks*, which are configuration recommendations for product families. For example, if your nodes run on the Microsoft Windows 10 Enterprise operating system, you can enforce the CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0, Level 1. For an overview, see [CIS Benchmarks](#). For a list of supported benchmarks, see [Prepare to install the module](#) on page 222.

### CIS profiles and levels

Each CIS Benchmark has a *profile*, which consists of a level and an applicability.

The *level* refers to the degree of protection:

- *Level 1* is intended to be practical and prudent, providing a clear security benefit without inhibiting the use of the technology.
- *Level 2* extends the Level 1 profile to provide additional protection for systems in which security is paramount. Level 2 can affect a system's performance and usability while promoting enhanced security.

The *applicability* refers to the affected system component. For example, if a benchmark has a profile of `Level 1 - Member Server`, the benchmark provides Level 1 (basic) security protections for servers.

### CIS controls

Each CIS Benchmark consists of *controls*, which are also called *recommendations* or *rules*. Each control is a security safeguard. For example, a control might disable the use of Bluetooth communication technologies on the protected system because Bluetooth transmissions can be intercepted. Or a control might specify that passwords must consist of at least 14 characters.

To learn about the CIS controls that are enforced by CEM, go to the [CEM Windows Reference](#) on Puppet Forge and click a benchmark to see the list of enforced controls. The control description starts with its config ID and name, for example:

```
1.1.1 - (L1) Ensure 'Enforce password history' is set to '24 or more password(s)'
```

The anatomy of a CIS control is as follows:

- **Parameters:** Configuration options, along with data types and default values.

- **Supported Levels:** The supported levels for the control, for example, Level 1.
- **Supported Profiles:** The applicability of the control, for example, `member_server`.
- **Hiera Configuration Example:** Snippet of code that can be used to configure the control in Hiera.
- **Alternate Config IDs:** The alternate config IDs for a control. If you configure the control in Hiera, you can use any of the listed config IDs. However, you cannot mix and match types within a configuration; you must use a single type of config ID.
- **Resource:** The name of the Puppet resource that enforces the control.

To enforce a CIS control on a node, you add Hiera data to your control repository. A control repository is the location where CEM configuration data is stored. For example, to enforce a CIS Benchmark with a profile of `Level 1 - Member Server` and specify only one control (for file system integrity), you would enter the following values:

```
control-repo/data/nodes/<node name>.yaml
cem_windows::benchmark: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 only:
 - 'ensure_filesystem_integrity_is_regularly_checked'
```

## Next steps

After you are familiar with the basic concepts, you can take the next steps.

Follow the instructions:

1. To prepare for the installation and install CEM, follow the instructions in [Installing CEM](#) on page 222.
2. To configure CEM, follow the instructions in [Configuring CEM](#) on page 227.

## Installing CEM

Before you install CEM, complete the preparation steps: review the system requirements, install and configure Puppet Enterprise (PE) or open source Puppet, and purchase CEM. To help avoid issues, install and evaluate CEM in a *test* environment before you install CEM in a production environment.

- [Prepare to install the module](#) on page 222

To help ensure the successful deployment of CEM, complete the preparation steps.

- [Install and evaluate the module in a test environment](#) on page 223

In some cases, compliance controls can negatively impact services that run on nodes. To help avoid possible issues, install and evaluate CEM in a *test* environment before running CEM in a production environment.

- [Install the module in the production environment](#) on page 224

To deploy CEM, you must install the module in the production environment and then classify the nodes on which you want to enforce Center for Internet Security (CIS) compliance.

- [Uninstall the module](#) on page 225

To stop using CEM, you can uninstall the `cem_windows` module. Alternatively, to stop using CEM on one or more nodes, you can declassify the nodes to remove their association with the `cem_windows` class.

## Prepare to install the module

To help ensure the successful deployment of CEM, complete the preparation steps.

1. Review the following table to ensure that CEM can meet your organization's requirements. `cem_windows` supports the following operating systems and Center for Internet Security (CIS) Benchmarks:

| Operating system    | Framework             | Profile                |
|---------------------|-----------------------|------------------------|
| Windows Server 2022 | CIS Benchmarks v2.0.0 | Member Server, Level 1 |
| Windows Server 2019 | CIS Benchmarks v2.0.0 | Member Server, Level 1 |

| Operating system      | Framework             | Profile                       |
|-----------------------|-----------------------|-------------------------------|
| Windows Server 2016   | CIS Benchmarks v2.0.0 | Member Server, Level 1        |
| Windows 10 Enterprise | CIS Benchmarks v2.0.0 | Corporate Enterprise, Level 1 |

**Restriction:** The `domain_controller` profile is not supported for any CIS controls.

**Tip:** CEM uses Desired State Configuration (DSC) modules and the `validation_mode` parameter to ensure that resources do not remain in a "flapping" state. For more information, see [securitypolicydsc](#).

- To manage nodes with Puppet Enterprise (PE), install PE. You can install any version in the 2021.7 or 2023 release stream. For instructions, see [Installing](#).
- To manage nodes with open source Puppet, install Puppet 7 or 8. For instructions, see [Install Puppet](#).
- Review the dependencies to ensure that your infrastructure will meet the requirements. Go to Puppet Forge and review the [Dependencies](#) tab.
- If you installed PE, follow the instructions in [Configuring Puppet Enterprise](#).
- If you installed open source Puppet, follow the instructions in [Configure Puppet settings](#).
- If you are using Puppet 7, verify that the agent and server are at v7.8.0 or later. You can also use any level of Puppet 8.
- To purchase CEM software, contact a Puppet by Perforce sales representative. For more information, see [Contact Us](#).
- Optionally, to help avoid issues during deployment to a production environment, you can install and evaluate CEM in a test environment. For instructions, see [Install and evaluate the module in a test environment](#) on page 223.

## Install and evaluate the module in a test environment

In some cases, compliance controls can negatively impact services that run on nodes. To help avoid possible issues, install and evaluate CEM in a *test* environment before running CEM in a production environment.

- Learn about the CIS Benchmark that you plan to enforce:
  - For a list of supported CIS Benchmarks, see [Prepare to install the module](#) on page 222.
  - For details about CIS Benchmarks and associated controls, see the [CEM Windows Reference](#) on Puppet Forge. You can also download benchmark information from the [CIS Benchmarks List](#). If you are using CEM with Puppet Comply, you can view details about the benchmarks in Comply.
- Make a list of any CIS controls that you plan to enable, disable, or configure to meet your organization's requirements.

For example, if a control specifies that a password must be changed every 60 days, but your organization requires a password change every 30 days, you can change the expected value for the associated control.

**Tip:** For the sake of simplicity, some users review the controls and enable only a *limited subset* to meet their organization's requirements.

- Identify a test environment. Many users follow the instructions in [Environments](#). You can also use any alternative method that works for you:
  - For Puppet Enterprise (PE), create a test node group and then assign the `cem_linux` class to that node group.
  - For open source Puppet, follow the instructions in [Classifying nodes](#). Ensure that the CEM module is included on the test nodes.
- Download CEM from Puppet Forge. CEM is available as a subscription. For more information, see the [Premium content page](#).
- If the host server is connected to the internet, install the module by following the instructions in [Installing modules from the Forge by using an internet connection](#).
- If the host server is not connected to the internet, install the module by following the instructions in [Installing modules from the Forge in an air-gapped environment](#).

7. Verify that the CEM module is successfully installed in the test environment.

**Tip:** If the installation was successful, you can find `cem_windows` in the following directory:

```
/etc/puppetlabs/code/environments/<environment_name>/modules/cem_windows
```

8. Implement any other configuration updates that you identified in Step 2. Take the following actions:

- a. Specify the updates as described in [Configuring CEM](#) on page 227.

**Tip:** You can simplify configuration by using the Hiera key-value store as described in [Getting started with Hiera](#). For examples, see [Basic configuration examples](#) on page 230 and [Advanced configuration example](#) on page 231.

- b. Ensure that the updates are deployed to the test environment. For example, if you are using Hiera and Code Manager, you must update the Hiera YAML files, commit the changes to the appropriate branch of your control repository, and trigger [Code Manager](#). If you are using open source Puppet, you would follow the same procedure but trigger `r10k`.

9. To detect and resolve any errors, take the following actions:

- a. Look for errors in Puppet runs in your test environment.
- b. If you detect errors, review and update your configuration. For help with configuration options, see the [CEM Windows Reference](#).
- c. If the configuration is correct but errors persist, enable debug logging on the Puppet primary server and review the `puppetserver.log` file. For more information, see [Puppet Server logging](#).

**Tip:** In the log list, CEM errors are prefixed with `CEM` or `cem`.

- d. Optionally, for additional insight into errors, enable tracing and debugging when you run the Puppet agent.
- e. If you are unable to resolve the errors, take one of the following actions:
  - PE users can post a question in the [#compliance](#) Slack channel in the [Puppet Community](#) or open a ticket with [Puppet Support](#).
  - Open source Puppet users can post a question in the [#compliance](#) Slack channel in the [Puppet Community](#), open a ticket on the [cem\\_issues webpage in GitHub](#), or open a ticket with [Puppet Support](#). As an open source Puppet user, your options vary depending on the support package that you purchased with CEM.

## Install the module in the production environment

To deploy CEM, you must install the module in the production environment and then classify the nodes on which you want to enforce Center for Internet Security (CIS) compliance.

### Before you begin

If you have not done so, install and evaluate CEM in a *test* environment before running CEM in a production environment. In this way, you can help to detect and avoid possible issues. For instructions, see [Install and evaluate the module in a test environment](#) on page 223.

1. If you have not done so, download CEM from Puppet Forge. The module is available as a premium content subscription. For more information, see the [Premium content page](#).
2. If the host server is connected to the internet, install the module by following the instructions in [Installing modules from the Forge by using an internet connection](#).
3. If the host server is not connected to the internet, install the module by following the instructions in [Installing modules from the Forge in an air-gapped environment](#).
4. If you installed Puppet Enterprise (PE), specify the nodes on which you want PE to run and enforce compliance. To specify the nodes, open the PE console and assign the `cem_windows` class to one or more node groups. For instructions about creating and classifying node groups, see [Grouping and classifying nodes](#).
5. If you installed open source Puppet, specify the nodes on which you want Puppet to run and enforce compliance. Follow the instructions in [Classifying nodes](#).

## Uninstall the module

To stop using CEM, you can uninstall the `cem_windows` module. Alternatively, to stop using CEM on one or more nodes, you can declassify the nodes to remove their association with the `cem_windows` class.



**CAUTION:** If you uninstall CEM or declassify one or more nodes, Puppet Enterprise (PE) or open source Puppet no longer runs on the affected nodes to enforce Center for Internet Security (CIS) compliance. When you uninstall CEM or declassify a node, the node does not revert to its pre-CEM state.

Take one of the following actions:

- To uninstall the CEM module, follow the instructions in [Uninstalling modules](#).
- On PE, to declassify nodes, can remove the nodes from the node group that is associated with the `cem_windows` class. For instructions, see [Remove nodes from a node group](#). Alternatively, to stop using CEM on an entire node group, you can remove the `cem_windows` class from the node group. For instructions, see [Remove classes from a node group](#).
- On open source Puppet, ensure that `cem_windows` is no longer included in the manifest that applies the `cem_windows` class.

## Upgrading CEM

You can upgrade CEM for Windows to take advantage of the latest features, fixes, and improvements. To help ensure a smooth upgrade process, complete the preparation tasks first.

- [Prepare to upgrade the module](#) on page 225

Before you upgrade the module to a new version, learn about the new version and familiarize yourself with any updates that were implemented to comply with Center for Internet Security (CIS) Benchmarks. Then, test the upgrade in a non-production environment and troubleshoot any issues.

- [Upgrade the module](#) on page 227

After you complete the preparation steps, you are ready to upgrade CEM for Windows.

### Prepare to upgrade the module

Before you upgrade the module to a new version, learn about the new version and familiarize yourself with any updates that were implemented to comply with Center for Internet Security (CIS) Benchmarks. Then, test the upgrade in a non-production environment and troubleshoot any issues.

1. To learn about a new CEM release, review the [Release notes](#) on page 212, which provide information about major updates, such as the introduction of additional operating systems and new or changed CIS Benchmarks. You can also learn about defect fixes, minor changes, and known issues.

**Tip:** For an example of the changes associated with a benchmark update, see the CEM for Windows v1.5.0 updates in the [Release notes](#) on page 212.

2. Learn about the CIS Benchmark that you plan to enforce in the new release. Review the associated controls and configuration options:
  - For a list of supported benchmarks in CEM, see [Prepare to install the module](#) on page 222.
  - For details about the supported benchmarks and associated controls, see the [CEM Windows Reference](#) on Puppet Forge. You can also download benchmark information from the [CIS Benchmarks List](#). If you are using CEM with Puppet Comply, you can view details about the benchmarks in Comply.
3. Identify a test environment. Many users follow the instructions in [Environments](#). You can also use any alternative method that works for you.

Testing is important because CEM can make hundreds of changes to a system, and many of those changes are critical to components. By testing and troubleshooting issues in advance, you can help to prevent issues later.

4. In the test environment, upgrade CEM. To help simplify the process, you can use a [Puppetfile](#):
  - a. In the Puppetfile, update the version in the CEM module declaration. For example, to upgrade `cem_windows` to version 1.5.1, specify the CEM declaration as shown:

```
mod 'puppetlabs/cem_windows', '1.5.1'
```

- b. Commit the change to the appropriate branch or test environment.
  - c. Deploy the change by using Code Manager or r10k. For instructions about using Code Manager and r10k, see [Managing code with Code Manager](#).
5. Verify that the CEM module is successfully upgraded in the test environment.
6. Determine whether the CEM configuration must be updated. If so, make a list of the required updates.

In many cases, you can upgrade CEM *without* additional configuration. However, if the relevant CIS Benchmark was updated in the release, the CEM configuration typically requires updates because controls might have been added or removed, or their numbers or titles might have changed. For example, if your configuration uses a "normalized number" control ID ("c1\_1\_1" or similar), pay close attention to any controls whose number changed because you must update the corresponding control ID in your configuration. If the release notes indicate that the number for control "1.1.1" changed to "1.1.2," you must replace "c1\_1\_1" in your configuration with "c1\_1\_2."

Control updates are documented in [Reference: Benchmarks and controls](#) on page 233. For example, if you are upgrading CEM to v1.5.0, you would go to the reference section and then locate the information for your operating system:

- [Control updates introduced for CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0](#) on page 233
- [Control updates introduced for CIS Microsoft Windows Server 2019 Benchmark v2.0.0](#) on page 234
- [Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v2.0.0](#) on page 235

7. Implement any required configuration updates. This step is crucial to help prevent CEM errors caused by misconfiguration.

You can simplify configuration by using the Hieradata key-value store as described in [Getting started with Hieradata](#). Take the following actions:

- a. Update the configuration of the test environment. For instructions, see [Find and set configuration options](#) on page 228.
  - b. Ensure that the updates are deployed to the test environment. For example, if you are using Hieradata and Code Manager, you must update the Hieradata YAML files, commit the changes to the appropriate branch of your control repository, and trigger Code Manager.
8. To detect and resolve any errors, take the following actions:
  - a. Look for errors in Puppet runs in your test environment.
  - b. If you detect errors, review and update your configuration. For help with configuration options, see the [CEM Windows Reference](#).
  - c. If the configuration is correct but errors persist, enable debug logging on the Puppet primary server and review the `puppetserver.log` file. For more information, see [Puppet Server logging](#).

**Tip:** In the log list, CEM errors are prefixed with `CEM` or `cem`.

- d. Optionally, for additional insight into errors, enable tracing and debugging when you run the Puppet agent.
  - e. If you are unable to resolve the errors, post a question in the [#compliance](#) Slack channel in the [Puppet Community](#) or open a ticket with [Puppet Support](#).

9. To plan the upgrade in the *production* environment, consider the following questions:
- Should the upgrade occur in stages or on all nodes simultaneously?
  - What is the risk of the upgrade, and is further testing required to mitigate the risk?
  - In the rare event that the upgrade is not successful, what is the plan to roll back the changes, given the fact that CEM cannot revert changes automatically?

**Tip:** For assistance with these questions, you can contact [Puppet Support](#). Support engineers are prepared to assist you before, during, and after the upgrade.

## Upgrade the module

After you complete the preparation steps, you are ready to upgrade CEM for Windows.

1. Update the CEM declaration in the Puppetfile. Specify the CEM version number to which you are upgrading. For example, to upgrade `cem_windows` to version 1.5.1, specify the CEM declaration as shown:

```
mod 'puppetlabs/cem_windows', '1.5.1'
```

For instructions about modifying the Puppetfile, see [Declare Forge modules in the Puppetfile](#).

2. Commit the change to the appropriate branch.
3. Deploy the change by using Code Manager or r10k.
4. If you determined during the preparation process that configuration updates are required in CEM, implement the configuration updates. For instructions, see [Configuring CEM](#) on page 227.

**Tip:** Starting with v1.1.0, CEM for Windows implements a new architecture. If you upgrade CEM from v1.0.7 or earlier to v1.1.0 or later, and you encounter errors, try restarting the `pe-puppetserver` service or restarting or reloading Puppet Server. For instructions, see [Restarting Puppet Server](#).

## Configuring CEM

Configuration of CEM is optional. If you installed CEM and assigned the `cem_windows` class to one or more node groups, the Center for Internet Security (CIS) Server Level 1 profile is enforced automatically during the next Puppet run. However, if the default values leave your infrastructure in an undesirable state, or if you want to customize compliance to meet your organization's requirements, you can configure CEM.

For example, if a CIS control sets the maximum password age at 365 days, but your organization requires a password change every 90 days, you can configure CEM accordingly.



### CAUTION:

Incorrect configuration of CEM can cause operational or security issues. CEM can make hundreds of changes to a system, and many of those changes are critical to components. For this reason, before you update the configuration, test the planned updates on one or two nodes. Evaluate the results and resolve any issues before implementing your configuration updates in a production environment. For instructions, see [Install and evaluate the module in a test environment](#) on page 223.

As an alternative, you can implement the CEM default settings, which are fully CIS compliant. However, depending on the complexities of your system environment, some default settings might not be appropriate. For this reason, to help ensure a secure configuration, review any controls and settings that you plan to implement and validate the controls in limited testing before implementing them in production.

You configure CEM by using the Hiera key-value store in your control repository. For more information, see [About Hiera](#) and [Getting started with Hiera](#).

For general information about configuration options, see [Overview of configuration options](#) on page 228.

For detailed information about configuration options, see the [CEM Windows Reference](#).

For configuration examples, see [How to configure the module: Examples and guidelines](#) on page 229.

- [Overview of configuration options](#) on page 228

Configuration options include top-level options, framework options, and Center for Internet Security (CIS)-specific options.

- [How to configure the module: Examples and guidelines](#) on page 229

The following examples demonstrate the use of CEM in a production environment.

## Overview of configuration options

Configuration options include top-level options, framework options, and Center for Internet Security (CIS)-specific options.

If you installed CEM and assigned the `cem_windows` class to a node group, the default profile is enforced. However, to customize CEM to meet your organization's requirements, you can configure benchmarks by using Hieradata. For more information, see [Hiera](#).

**Tip:** When you use Hieradata to specify CEM configuration options, the configuration looks different from the configuration for other Puppet products. The reason is that the `config` variable receives a hash, and the key values in the hash control the CEM variables in the configuration.



**CAUTION:** CEM's default settings are fully CIS compliant. Too much customization can cause your configurations to be noncompliant.

## Find and set configuration options

You can find configuration options in the [CEM Windows Reference](#) on Puppet Forge. A control description starts with a config ID and name, for example:

```
1.1.1 - (L1) Ensure 'Enforce password history' is set to '24 or more
password(s)'
```

The anatomy of a CIS control is as follows:

- **Parameters:** Configuration options for a control, along with the data type and default value.
- **Supported Levels:** The supported levels, for example, Level 1.
- **Supported Profiles:** The applicability of the control. For example, a control with a profile of `member_server` is applicable to server components.
- **Hiera Configuration Example:** Snippet of Hieradata that can be used to configure a control.
- **Alternate Config IDs:** The alternate config IDs for a control. Any of these config IDs, along with the full control name, can be used as a key in the `control_config` hash.
- **Resource:** The name of the Puppet resource that enforces the control.

## Guidelines for specifying CIS config IDs

You can specify controls in the `control_config` hash by referencing the full control name, the control number, the normalized control name, or the normalized control number. You *cannot* mix and match these forms and must pick a single config ID form to use for your config. Full control names and control numbers are copied verbatim from the benchmarks and are case-sensitive. Normalized control names have lowercase letters and contain only alphanumeric characters and underscores. Normalized control numbers are always prefixed with a `c` and contain only numeric characters separated by underscores.

Example of alternative config IDs:

- Full control name: `(L1) Ensure 'Enforce password history' is set to '24 or more password(s)'`
- Control number: `1.1.1`

- Normalized control name:  
ensure\_enforce\_password\_history\_is\_set\_to\_24\_or\_more\_passwords
- Normalized number: c1\_1\_1

### Top-level configuration options

These options are configured at the top level of the module.

In Hiera, these options are prefixed with `cem_windows`:

- `framework` - Enum[ 'cis' ] - the compliance framework to use. CEM supports only `cis`. Default: `cis`.
- `config` - Optional[Hash] - the location for all non-top-level configuration options. Default: `undef`.
- `allow_on_domain_controller` - Boolean - If `cem_windows` detects that it is running on a domain controller, CEM does not enforce controls and logs a warning to inform the user. In this way, CEM helps to prevent the enforcement of compliance settings on domain controllers that could negatively impact an entire domain. Default: `false`.
- `enable_long_paths` - Boolean - Enables support for long path names in the Windows registry. Setting this option to `false` can cause issues with some Desired State Configuration (DSC) modules used in `cem_windows`. Default: `true`.
- `privileged_user` - Optional[String] - If the Puppet agent does not run under a user with local administrator privileges, you must supply the name of a user with local administrator privileges. This is used by DSC to enforce a state on a machine. Default: `undef`.
- `privileged_password` - Sensitive[Any] - If you specified a privileged user, use this option to specify a password for that user account. Default: `undef`.
- `allow_local_account_rdp` - Boolean - By default, `cem_windows` disables remote desktop protocol (RDP) access for non-domain accounts. If you set this option to `true`, local accounts on the node can make RDP connections to the node. Default: `false`.

### Framework configuration options

The framework configuration options are available as key-value pairs within the `cem_windows::config`: hash.

- `control_configs` - Optional[Hash] — location for all rule-specific configurations. Default: `undef`.
- `only` - Optional[Array[String]] — takes an array of control class names (`manifests/benchmarks/<benchmark>/controls/*.pp`). The classes specified here are included in the catalog. Takes precedence over the `ignore`: option. Default: `undef`.
- `ignore` - Optional[Array[String]] — takes an array of control class names (`manifests/benchmarks/<benchmark>/controls/*.pp`). The classes specified here are not included in the catalog. If `only`: is specified, this option has no effect. Default: `undef`.

### CIS-specific configuration options

The CIS-specific configuration options are available as key-value pairs within the `cem_windows::config`: hash. These options are applicable only to the CIS compliance framework.

- `profile` - Optional[Enum[ 'member\_server', 'corporateenterprise' ]] — the name of the benchmark profile. `corporateenterprise` is supported only on Windows 10 Enterprise operating systems. Default for Windows Server operating systems: `member_server`. Default for Windows 10 Enterprise operating systems: `corporate_enterprise`.
- `level` - Optional[Enum[ '1', '2' ]] — the name of the profile level. The only value supported by CEM is 1. Default: 1.

For more details about configuration options, see [Reference: Benchmarks and controls](#) on page 233.

## How to configure the module: Examples and guidelines

The following examples demonstrate the use of CEM in a production environment.

- [Basic configuration examples](#) on page 230

When you specify a compliance framework, CEM is configured to provide rule enforcement and configuration for that framework. For example, to enforce the Center for Internet Security (CIS) Server Level 1 benchmark for a node,

you must classify the node with the `cem_windows` class, set the `benchmark` parameter to `cis`, and run Puppet. To learn more about CEM configuration, see the following examples.

- [Advanced configuration example](#) on page 231

Building on the basic configuration examples, the `control_configs` section specifies advanced options for controls.

- [Run Desired State Configuration resources as a specific user](#) on page 231

Desired State Configuration (DSC) requires local administrator privileges to modify Windows resources. Typically, the Puppet agent runs under a user account with these permissions. However, if the Puppet agent on a node does not have local administrator permissions, you can use Hiera to configure a user account that does have the required permissions.

- [Allow local accounts to access nodes](#) on page 231

To allow a local user account to access a node with remote desktop protocol (RDP), set the top-level option `allow_local_account_rdp` to `true`.

- [Enforce specific rules](#) on page 232

To configure CEM to enforce only specific rules, use the `only` key.

- [Ignore specific rules](#) on page 232

To configure CEM to ignore specific rules, use the `ignore` key.

- [Customize rules](#) on page 232

You can customize most rules by using the `control_configs` key and supplying the key with a hash value.

- [Rename the Administrator and Guest accounts](#) on page 232

To help protect your infrastructure, rename the `Administrator` and `Guest` accounts.

## Basic configuration examples

When you specify a compliance framework, CEM is configured to provide rule enforcement and configuration for that framework. For example, to enforce the Center for Internet Security (CIS) Server Level 1 benchmark for a node, you must classify the node with the `cem_windows` class, set the `benchmark` parameter to `cis`, and run Puppet. To learn more about CEM configuration, see the following examples.

### Example 1

In the following example, CEM applies only the following controls on a Windows 10 node: `'c1_1_1'` and `'c2_3_1_1'`.

1. Add the following Hiera data to your control repository, `control_repo`:

```
control-repo/data/nodes/<node name>.yaml
cem_windows::benchmark: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 only:
 - 'c1_1_1'
 - 'c2_3_1_1'
```

2. Classify the node with the `cem_windows` class.
3. Run Puppet.

### Example 2

In the following example, CEM applies all controls in the configured benchmark and profile on a Windows 10 node except for the following controls: `'c1_1_1'` and `'c2_3_1_1'`.

1. Add the following Hiera data to your control repository, `control_repo`:

```
control-repo/data/nodes/<node name>.yaml
cem_windows::benchmark: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
```

```
ignore:
 - 'c1_1_1'
 - 'c2_3_1_1'
```

2. Classify the node with the `cem_windows` class.
3. Run Puppet.

### Advanced configuration example

Building on the basic configuration examples, the `control_configs` section specifies advanced options for controls.

1. Add the following code to the node's Hiera file:

```
cem_windows::benchmark: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 only:
 - 'c1_1_1'
 - 'c2_3_1_1'
 - 'c2_3_7_2'
 - 'c2_3_1_3'
 - 'c2_3_1_4'
 - 'c2_3_1_5'
 - 'c2_3_1_6'
 - 'c2_3_10_1'
 - 'c18_9_47_5_1_2'
 control_configs:
 'c18_9_47_5_1_2':
 default_for_all_rules: 'Audit'
 'c2_3_1_6':
 dsc_accounts_rename_administrator_account:
 - 'adminaccountname'
```

2. Classify the node with the `cem_windows` class.
3. Run Puppet.

### Run Desired State Configuration resources as a specific user

Desired State Configuration (DSC) requires local administrator privileges to modify Windows resources. Typically, the Puppet agent runs under a user account with these permissions. However, if the Puppet agent on a node does not have local administrator permissions, you can use Hiera to configure a user account that does have the required permissions.

Use a configuration based on the following structure:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::privileged_user: <user name>
cem_windows::privileged_pass: <user password>
```

### Allow local accounts to access nodes

To allow a local user account to access a node with remote desktop protocol (RDP), set the top-level option `allow_local_account_rdp` to `true`.

Use a configuration based on the following structure:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::allow_local_account_rdp: true
```

## Enforce specific rules

To configure CEM to enforce only specific rules, use the `only` key.

Use a configuration that is similar to the following example:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::framework: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 only:
 - 'c18_9_97_1_1'
 - 'c18_9_97_1_2'
```

## Ignore specific rules

To configure CEM to ignore specific rules, use the `ignore` key.

Use a configuration that is similar to the following example:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::framework: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 ignore:
 - 'c18_9_97_1_1'
 - 'c18_9_97_1_2'
```

**Restriction:** The `only` key and the `ignore` key are mutually exclusive, with `only` taking precedence. If you specify both keys, CEM does not use the value of the `ignore` key and enforces only the rules specified with `only`.

## Customize rules

You can customize most rules by using the `control_configs` key and supplying the key with a hash value.

To customize rules, use a configuration based on the following structure:

```
<recommendation name>:
 <recommendation param>: <value>
```

For example, to configure the rules used in the previous examples, the configuration would look like this:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::framework: 'cis'
cem_windows::config:
 profile: 'member_server'
 level: '1'
 control_configs:
 c18_9_97_1_1:
 allowbasic: '0'
 c18_9_97_1_2:
 allowunencryptedtraffic: '0'
```

## Rename the Administrator and Guest accounts

To help protect your infrastructure, rename the Administrator and Guest accounts.

1. To rename the local Administrator account to `user_1`, use the following configuration:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::config:
 control_configs:
 c2_3_1_5:
 value: 'user_1'
```

**Restriction:** If you do not specify a name for the Administrator account, the account is renamed to `magic` by default.

2. To rename the local Guest account to `user_2`, use the following configuration:

```
control-repo/data/nodes/winserv2019.contoso.com.yaml

cem_windows::config:
 control_configs:
 c2_3_1_6:
 value: 'user_2'
```

**Restriction:** If you do not specify a name for the Guest account, the account is renamed to `pumpkin` by default.

## Reference: Benchmarks and controls

---

For help with configuring CEM, review the [CEM Windows Reference](#) on Puppet Forge.

In the [CEM Windows Reference](#), locate the relevant CIS Benchmark and then the control.

For information about control updates that were implemented between CEM releases, see the following topics:

- [Control updates introduced for CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0](#) on page 233  
The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Enterprise 10 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.12.0, to the new benchmark resulted in module updates.
- [Control updates introduced for CIS Microsoft Windows Server 2019 Benchmark v2.0.0](#) on page 234  
The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2019 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.3.0, to the new benchmark resulted in module updates.
- [Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v2.0.0](#) on page 235  
The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2016 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.4.0, to the new benchmark resulted in module updates.
- [Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v1.4.0](#) on page 236  
The Compliance Enforcement Module (CEM) for Windows v1.3.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2016 Benchmark v1.4.0. The transition from the previous CIS Benchmark, v1.3.0, to the new benchmark resulted in module updates.

### Control updates introduced for CIS Microsoft Windows 10 Enterprise Benchmark v2.0.0

The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Enterprise 10 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.12.0, to the new benchmark resulted in module updates.

- **Added**

- The following CIS controls are added in this release:
  - 1.2.3 (L1) Ensure 'Allow Administrator account lockout' is set to 'Enabled'
  - 18.4.2 (L1) Ensure 'Configure RPC packet level privacy setting for incoming connections' is set to 'Enabled'
  - 18.4.6 (L1) Ensure 'LSA Protection' is set to 'Enabled'
  - 18.6.4.1 (L1) Ensure 'Configure NetBIOS settings' is set to 'Enabled: Disable NetBIOS name resolution on public networks'
  - 18.7.2 (L1) Ensure 'Configure Redirection Guard' is set to 'Enabled: Redirection Guard Enabled'
  - 18.7.3 (L1) Ensure 'Configure RPC connection settings: Protocol to use for outgoing RPC connections' is set to 'Enabled: RPC over TCP'
  - 18.7.4 (L1) Ensure 'Configure RPC connection settings: Use authentication for outgoing RPC connections' is set to 'Enabled: Default'
  - 18.7.5 (L1) Ensure 'Configure RPC listener settings: Protocols to allow for incoming RPC connections' is set to 'Enabled: RPC over TCP'
  - 18.7.6 (L1) Ensure 'Configure RPC listener settings: Authentication protocol to use for incoming RPC connections' is set to 'Enabled: Negotiate' or higher
  - 18.7.7 (L1) Ensure 'Configure RPC over TCP port' is set to 'Enabled: 0'
  - 18.7.9 (L1) Ensure 'Manage processing of Queuespecific files' is set to 'Enabled: Limit Queue-specific files to Color profiles'
  - 18.9.25.1 (L1) Ensure 'Allow Custom SSPs and APs to be loaded into LSASS' is set to 'Disabled'
  - 18.10.17.1 (L1) Ensure 'Enable App Installer' is set to 'Disabled'
  - 18.10.17.2 (L1) Ensure 'Enable App Installer Experimental Features' is set to 'Disabled'
  - 18.10.17.3 (L1) Ensure 'Enable App Installer Hash Override' is set to 'Disabled'
  - 18.10.17.4 (L1) Ensure 'Enable App Installer msappinstaller protocol' is set to 'Disabled'
  - 18.10.82.1 (L1) Ensure 'Enable MPR notifications for the system' is set to 'Disabled'

- **Changed**

- The following CIS controls were updated:
  - 18.9.89 'Allow Windows Ink Workspace' now has expected values of 'Enabled: On, but disallow access above lock' or 'Enabled: Disabled'.
  - 18.10.87 (L1) 'Turn on PowerShell Transcription' was set to 'Disabled' but now has an expected value of 'Enabled'.
  - 18.3.5 (L1) Ensure 'Limits print driver installation to Administrators' is set to 'Enabled' has a new number: 18.7.8.

- **Removed**

- The following CIS controls were removed:
  - 2.3.1 (L1) Ensure 'Accounts: Administrator account status' is set to 'Disabled'
  - 18.5.4 (L1) Ensure 'Configure DNS over HTTPS (DoH) name resolution' is set to 'Enabled: Allow DoH' or higher

## Control updates introduced for CIS Microsoft Windows Server 2019 Benchmark v2.0.0

The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2019 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.3.0, to the new benchmark resulted in module updates.

- **Added**

- The following CIS controls are added in this release:
  - 1.2.3 (L1) Ensure 'Allow Administrator account lockout' is set to 'Enabled'
  - 18.4.2 (L1) Ensure 'Configure RPC packet level privacy setting for incoming connections' is set to 'Enabled'
  - 18.4.6 (L1) Ensure 'LSA Protection' is set to 'Enabled'
  - 18.6.4.1 (L1) Ensure 'Configure NetBIOS settings' is set to 'Enabled: Disable NetBIOS name resolution on public networks'
  - 18.7.4 (L1) Ensure 'Configure RPC connection settings: Use authentication for outgoing RPC connections' is set to 'Enabled: Default'
  - 18.7.5 (L1) Ensure 'Configure RPC listener settings: Protocols to allow for incoming RPC connections' is set to 'Enabled: RPC over TCP'
  - 18.7.7 (L1) Ensure 'Configure RPC over TCP port' is set to 'Enabled: 0'
  - 18.7.9 (L1) Ensure 'Manage processing of Queuespecific files' is set to 'Enabled: Limit Queue-specific files to Color profiles'
  - 18.10.17.1 (L1) Ensure 'Enable App Installer' is set to 'Disabled'
  - 18.10.17.2 (L1) Ensure 'Enable App Installer Experimental Features' is set to 'Disabled'
  - 18.10.17.3 (L1) Ensure 'Enable App Installer Hash Override' is set to 'Disabled'
  - 18.10.17.4 (L1) Ensure 'Enable App Installer msappinstaller protocol' is set to 'Disabled'

- **Changed**

- The following CIS controls were updated:
  - 18.3.5 (L1) Ensure 'Limits print driver installation to Administrators' is set to 'Enabled' has a new number: 18.7.8.
  - 18.9.89 'Allow Windows Ink Workspace' now has expected values of 'Enabled: On, but disallow access above lock' or 'Enabled: Disabled'.
  - 18.10.87 (L1) 'Turn on PowerShell Transcription' was set to 'Disabled' but now has an expected value of 'Enabled'.

- **Removed**

- The following CIS controls were removed:
  - 2.3.1 (L1) Ensure 'Accounts: Administrator account status' is set to 'Disabled'
  - 18.5.4 (L1) Ensure 'Configure DNS over HTTPS (DoH) name resolution' is set to 'Enabled: Allow DoH' or higher

## Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v2.0.0

The Compliance Enforcement Module (CEM) for Windows v1.5.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2016 Benchmark v2.0.0. The transition from the previous CIS Benchmark, v1.4.0, to the new benchmark resulted in module updates.

- **Added**
  - The following CIS controls are added in this release:
    - 1.2.3 (L1) Ensure 'Allow Administrator account lockout' is set to 'Enabled'
    - 18.4.2 (L1) Ensure 'Configure RPC packet level privacy setting for incoming connections' is set to 'Enabled'
    - 18.4.6 (L1) Ensure 'LSA Protection' is set to 'Enabled'
    - 18.6.4.1 (L1) Ensure 'Configure NetBIOS settings' is set to 'Enabled: Disable NetBIOS name resolution on public networks'
    - 18.7.2 (L1) Ensure 'Configure Redirection Guard' is set to 'Enabled: Redirection Guard Enabled'
    - 18.7.3 (L1) Ensure 'Configure RPC connection settings: Protocol to use for outgoing RPC connections' is set to 'Enabled: RPC over TCP'
    - 18.7.4 (L1) Ensure 'Configure RPC connection settings: Use authentication for outgoing RPC connections' is set to 'Enabled: Default'
    - 18.7.5 (L1) Ensure 'Configure RPC listener settings: Protocols to allow for incoming RPC connections' is set to 'Enabled: RPC over TCP'
    - 18.7.6 (L1) Ensure 'Configure RPC listener settings: Authentication protocol to use for incoming RPC connections' is set to 'Enabled: Negotiate' or higher
    - 18.7.7 (L1) Ensure 'Configure RPC over TCP port' is set to 'Enabled: 0'
    - 18.7.9 (L1) Ensure 'Manage processing of Queuespecific files' is set to 'Enabled: Limit Queue-specific files to Color profiles'
    - 18.9.47.5.1 'Configure Attack Surface Reduction rules' is set to 'Enable'
    - 18.10.59.4 (L2) Ensure 'Allow search highlights' is set to 'Disabled'
- **Changed**
  - The following CIS controls were updated:
    - 18.3.5 (L1) Ensure 'Limits print driver installation to Administrators' is set to 'Enabled' has a new number: 18.7.8.
    - 18.9.89 'Allow Windows Ink Workspace' has new expected values: 'Enabled: On, but disallow access above lock' or 'Enabled: Disabled'.
- **Removed**
  - The following CIS controls were removed:
    - 2.3.1 (L1) Ensure 'Accounts: Administrator account status' is set to 'Disabled'
    - 18.5.4 (L1) Ensure 'Configure DNS over HTTPS (DoH) name resolution' is set to 'Enabled: Allow DoH' or higher

## Control updates introduced for CIS Microsoft Windows Server 2016 Benchmark v1.4.0

The Compliance Enforcement Module (CEM) for Windows v1.3.0 introduces enforcement for Center for Internet Security (CIS) Microsoft Windows Server 2016 Benchmark v1.4.0. The transition from the previous CIS Benchmark, v1.3.0, to the new benchmark resulted in module updates.

- **Added**

- The following CIS controls are added in this release:
  - Control 5.2, Ensure 'Print Spooler (Spooler)' is set to 'Disabled' (MS only). This control setting disables the print spooler service by default and thus helps to prevent security vulnerabilities.
  - Control 18.3.5, (L1) Ensure 'Limits print driver installation to Administrators' is set to 'Enabled.' This control setting helps to ensure that only administrators can install printer drivers and thus reduces security risks.
  - Control 18.5.4.1, (L1) Ensure 'Configure DNS over HTTPS (DoH) name resolution' is set to 'Enabled: Allow DoH' or higher. This control helps to protect against Domain Name System (DNS) spoofing and thus can help to prevent man-in-the-middle (MITM) attacks.
  - Control 18.6.2, (L1) Ensure 'Point and Print Restrictions: When installing drivers for a new connection' is set to 'Enabled: Show warning and elevation prompt.' This policy setting ensures that a warning is displayed when users create a printer connection by using point-and-print functionality.
  - Control 18.6.3, Ensure 'Point and Print Restrictions: When updating drivers for an existing connection' is set to 'Enabled: Show warning and elevation prompt.' This policy setting controls whether a warning is displayed when users update a printer driver for a connection that uses point-and-print functionality. Warnings help to guard against security vulnerabilities.
  - Control 18.8.7.2, (L1) Ensure 'Prevent device metadata retrieval from the Internet' is set to 'Enabled.' The impact of this control is that users without administrator privileges cannot install third-party software for peripheral devices. Instead, authorized system administrators install approved software.
  - Control 18.9.14.1, (L1) Ensure 'Turn off cloud consumer account state content' is set to 'Enabled.' This control is designed to prevent data leakage by ensuring that state information related to cloud consumer accounts is not available in an enterprise-managed environment.
  - Control 18.9.17.1, (L1) Ensure 'Allow Diagnostic Data' is set to 'Enabled: Diagnostic data off (not recommended)' or 'Enabled: Send required diagnostic data.' This policy setting controls the amount of diagnostic data reported to Microsoft. The default setting ensures that only minimal data is reported to help keep Microsoft Windows current, secure, and operational.
  - Control 18.9.17.3, (L1) Ensure 'Disable OneSettings Downloads' is set to 'Enabled.' This policy setting controls whether Windows attempts to connect with the OneSettings service to download configuration settings. Because transmission of data to a third-party vendor can present a security risk, the control disables these downloads.
  - Control 18.9.17.5, (L1) Ensure 'Enable OneSettings Auditing' is set to 'Enabled.' This policy setting helps to ensure that the Windows operating system keeps a log of attempts to connect with the OneSettings service. The logs can be useful for troubleshooting and to help prevent unauthorized access to the system.
  - Control 18.9.17.6, (L1) Ensure 'Limit Diagnostic Log Collection' is set to 'Enabled.' This policy setting helps to protect security by ensuring that additional diagnostic logs and information about crash dumps are not sent to Microsoft.
  - Control 18.9.17.7, (L1) Ensure 'Limit Dump Collection' is set to 'Enabled.' This policy setting helps to reduce the risk of sending sensitive information to Microsoft.
  - Control 18.9.47.9.4, (L1) Ensure 'Turn on script scanning' is set to 'Enabled.' This policy setting helps to ensure that scripts are scanned before they are run on the system.
  - Control 18.9.108.4.1, (L1) Ensure 'Manage preview builds' is set to 'Disabled.' This policy setting helps to prevent the installation of preview builds, which are more likely to introduce defects and security vulnerabilities.
  - Control 18.9.108.4.2, (L1) Ensure 'Select when Preview Builds and Feature Updates are received' is set to 'Enabled: 180 or more days.' This policy setting helps to ensure that new preview builds and feature updates are received 180 or more days after their release by Microsoft. The purpose of the delay is to ensure that software defects have been detected and fixed.
  - Control 19.7.8.5, (L1) Ensure 'Turn off Spotlight collection on Desktop' is set to 'Enabled.' This policy setting helps to ensure that users cannot select 'Spotlight collection' as a personalization option. As a result, users cannot display and download daily images from Microsoft to the desktop.

- **Changed**

- The following CIS controls were updated with new expected values:
  - The expected value is changed for CIS Control 18.8.3.1, Ensure 'Include command line in process creation events' is set to 'Disabled.' The expected value is now Enabled. The control affects security audit events. When this setting is enabled, any user who has read access to security audit events can read the command-line arguments for any successfully created process.
  - A control was reintroduced with the following new number and new expected value: 18.9.100.1, (L1) Ensure 'Turn on PowerShell Script Block Logging' is set to 'Enabled.' In previous releases, the same control had a different number (18.9.96.1) and the expected value was Disabled. The new policy setting helps to ensure that PowerShell script logs are available and can be used to troubleshoot attack incidents.
  - CIS Control 18.9.16.1, (L1) Ensure 'Allow Telemetry' is set to 'Enabled: 0 - Security [Enterprise Only]' or 'Enabled: 1 – Basic,' was replaced by a control with the same number. The new control is (L1) Ensure 'Do not

display the password reveal button' is set to 'Enabled.' The change means that a password reveal button will not be available. When a user enters a password, the password will be hidden.

- For some CIS Controls, only the numbers changed. Previous control numbers are listed first:
  - 18.3.5, (L1) Ensure 'NetBT NodeType configuration' is set to 'Enabled: P-node (recommended).' The new number is 18.3.6.
  - 18.3.6, Ensure 'WDigest Authentication' is set to 'Disabled.' The new number is 18.3.7.
  - 18.5.4.1, (L1) Ensure 'Turn off multicast name resolution' is set to 'Enabled.' The new number is 18.5.4.2.
  - 18.8.47.5.1, Ensure 'Microsoft Support Diagnostic Tool: Turn on MSDT interactive communication with support provider' is set to 'Disabled.' The new number is 18.8.48.5.1.
  - 18.8.47.11.1, (L2) Ensure 'Enable/Disable PerfTrack' is set to 'Disabled.' The new number is 18.8.48.11.1.
  - 18.8.49.1, (L2) Ensure 'Turn off the advertising ID' is set to 'Enabled.' The new number is 18.8.50.1.
  - 18.8.52.1.1, (L2) Ensure 'Enable Windows NTP Client' is set to 'Enabled.' The new number is 18.8.53.1.1.
  - 18.8.52.1.2, (L2) Ensure 'Enable Windows NTP Server' is set to 'Disabled' (MS only). The new number is 18.8.53.1.2.
  - 18.9.13.1, (L1) Ensure 'Turn off Microsoft consumer experiences' is set to 'Enabled.' The new number is 18.9.14.2.
  - 18.9.14.1, (L1), Ensure 'Require pin for pairing' is set to 'Enabled: First Time' OR 'Enabled: Always.' The new number is 18.9.15.1.
  - 18.9.15.2, (L1) Ensure 'Enumerate administrator accounts on elevation' is set to 'Disabled.' The new number is 18.9.16.2.
  - 18.9.16.2, (L2) Ensure 'Configure Authenticated Proxy usage for the Connected User Experience and Telemetry service' is set to 'Enabled: Disable Authenticated Proxy usage.' The new number is 18.9.17.2.
  - 18.9.16.3, (L1) Ensure 'Do not show feedback notifications' is set to 'Enabled.' The new number is 18.9.17.4.
  - 18.9.16.4, (L1) Ensure 'Toggle user control over Insider builds' is set to 'Disabled.' The new number is 18.9.17.8.
  - 18.9.26.1.1, (L1) Ensure 'Application: Control Event Log behavior when the log file reaches its maximum size' is set to 'Disabled.' The new number is 18.9.27.1.1.
  - 18.9.26.1.2, (L1) Ensure 'Application: Specify the maximum log file size (KB)' is set to 'Enabled: 32,768 or greater.' The new number is 18.9.27.1.2.
  - 18.9.26.2.1, (L1) Ensure 'Security: Control Event Log behavior when the log file reaches its maximum size' is set to 'Disabled.' The new number is 18.9.27.2.1.
  - 18.9.26.2.2, (L1) Ensure 'Security: Specify the maximum log file size (KB)' is set to 'Enabled: 196,608 or greater.' The new number is 18.9.27.2.2.
  - 18.9.26.3.1, (L1) Ensure 'Setup: Control Event Log behavior when the log file reaches its maximum size' is set to 'Disabled.' The new number is 18.9.27.3.1.
  - 18.9.26.3.2, (L1) Ensure 'Setup: Specify the maximum log file size (KB)' is set to 'Enabled: 32,768 or greater.' The new number is 18.9.27.3.2.
  - 18.9.26.4.1, (L1) Ensure 'System: Control Event Log behavior when the log file reaches its maximum size' is set to 'Disabled.' The new number is 18.9.27.4.1.
  - 18.9.26.4.2, (L1) Ensure 'System: Specify the maximum log file size (KB)' is set to 'Enabled: 32,768 or greater.' The new number is 18.9.27.4.2.
  - 18.9.30.2, (L1) Ensure 'Turn off Data Execution Prevention for Explorer' is set to 'Disabled.' The new number is 18.9.31.2.
  - 18.9.30.3, (L1) Ensure 'Turn off heap termination on corruption' is set to 'Disabled.' The new number is 18.9.31.3.
  - 18.9.30.4, (L1) Ensure 'Turn off shell protocol protected mode' is set to 'Disabled.' The new number is 18.9.31.4.
  - 18.9.39.1, (L2) Ensure 'Turn off location' is set to 'Enabled.' The new number is 18.9.41.1.
  - 18.9.43.1, (L2) Ensure 'Allow Message Service Cloud Sync' is set to 'Disabled.' The new number is 18.9.45.1.

- 18.9.44.1, (L1) Ensure 'Block all consumer Microsoft account user authentication' is set to 'Enabled.' The new number is 18.9.46.1.
- 18.9.45.3.1, (L1) Ensure 'Configure local setting override for reporting to Microsoft MAPS' is set to 'Disabled.' The new number is 18.9.47.4.1.
- 18.9.45.3.2, (L2) Ensure 'Join Microsoft MAPS' is set to 'Disabled.' The new number is 18.9.47.4.2.
- 18.9.45.4.3.1, (L1) Ensure 'Prevent users and apps from accessing dangerous websites' is set to 'Enabled: Block.' The new number is 18.9.47.5.3.1.
- 18.9.45.5.1, (L2) Ensure 'Enable file hash computation feature' is set to 'Enabled.' The new number is 18.9.47.6.1.
- 18.9.45.8.1, (L1) Ensure 'Scan all downloaded files and attachments' is set to 'Enabled.' The new number is 18.9.47.9.1.
- 18.9.45.8.2, (L1) Ensure 'Turn off real-time protection' is set to 'Disabled.' The new number is 18.9.47.9.2.
- 18.9.45.8.3, (L1) Ensure 'Turn on behavior monitoring' is set to 'Enabled.' The new number is 18.9.47.9.3.
- 18.9.45.10.1, (L2) Ensure 'Configure Watson events' is set to 'Disabled.' The new number is 18.9.47.11.1.
- 18.9.45.11.1, (L1) Ensure 'Scan removable drives' is set to 'Enabled.' The new number is 18.9.47.12.1.
- 18.9.45.11.2, (L1) Ensure 'Turn on e-mail scanning' is set to 'Enabled.' The new number is 18.9.47.12.2.
- 18.9.45.14, (L1) Ensure 'Configure detection for potentially unwanted applications' is set to 'Enabled: Block.' The new number is 18.9.47.15.
- 18.9.45.15, (L1) Ensure 'Turn off Microsoft Defender AntiVirus' is set to 'Disabled.' The new number is 18.9.47.16.
- 18.9.56.1, (L1) Ensure 'Prevent the usage of OneDrive for file storage' is set to 'Enabled.' The new number is 18.9.58.1.
- 18.9.62.1, (L2) Ensure 'Turn off Push To Install service' is set to 'Enabled.' The new number is 18.9.64.1.
- 18.9.63.2.2, (L1) Ensure 'Do not allow passwords to be saved' is set to 'Enabled.' The new number is 18.9.65.2.2.
- 18.9.63.3.2.1, (L2) Ensure 'Restrict Remote Desktop Services users to a single Remote Desktop Services session' is set to 'Enabled.' The new number is 18.9.65.3.2.1.
- 18.9.63.3.3.1, (L2) Ensure 'Do not allow COM port redirection' is set to 'Enabled.' The new number is 18.9.65.3.3.1.
- 18.9.63.3.3.2, (L1) Ensure 'Do not allow drive redirection' is set to 'Enabled.' The new number is 18.9.65.3.3.2.
- 18.9.63.3.3.3, (L2) Ensure 'Do not allow LPT port redirection' is set to 'Enabled.' The new number is 18.9.65.3.3.3.
- 18.9.63.3.3.4, (L2) Ensure 'Do not allow supported Plug and Play device redirection' is set to 'Enabled.' The new number is 18.9.65.3.3.4.
- 18.9.63.3.9.1, (L1) Ensure 'Always prompt for password upon connection' is set to 'Enabled.' The new number is 18.9.65.3.9.1.
- 18.9.63.3.9.2, (L1) Ensure 'Require secure RPC communication' is set to 'Enabled.' The new number is 18.9.65.3.9.2.
- 18.9.63.3.9.3, (L1) Ensure 'Require use of specific security layer for remote (RDP) connections' is set to 'Enabled: SSL.' The new number is 18.9.65.3.9.3.
- 18.9.63.3.9.4, (L1) Ensure 'Require user authentication for remote connections by using Network Level Authentication' is set to 'Enabled.' The new number is 18.9.65.3.9.4.
- 18.9.63.3.9.5, (L1) Ensure 'Set client connection encryption level' is set to 'Enabled: High Level.' The new number is 18.9.65.3.9.5.
- 18.9.63.3.10.1, (L2) Ensure 'Set time limit for active but idle Remote Desktop Services sessions' is set to 'Enabled: 15 minutes or less, but not Never (0).' The new number is 18.9.65.3.10.1.
- 18.9.63.3.10.2, (L2) Ensure 'Set time limit for disconnected sessions' is set to 'Enabled: 1 minute.' The new number is 18.9.65.3.10.2.
- 18.9.63.3.11.1, (L1) Ensure 'Do not delete temp folders upon exit' is set to 'Disabled.' The new number is 18.9.65.3.11.1.

- 18.9.63.3.11.2, (L1) Ensure 'Do not use temporary folders per session' is set to 'Disabled.' The new number is 18.9.65.3.11.2.
- 18.9.64.1, (L1) Ensure 'Prevent downloading of enclosures' is set to 'Enabled.' The new number is 18.9.66.1.
- 18.9.65.2, (L2) Ensure 'Allow Cloud Search' is set to 'Enabled: Disable Cloud Search.' The new number is 18.9.67.2.
- 18.9.65.3, (L1) Ensure 'Allow indexing of encrypted files' is set to 'Disabled.' The new number is 18.9.67.3.
- 18.9.70.1, (L2) Ensure 'Turn off KMS Client Online AVS Validation' is set to 'Enabled.' The new number is 18.9.72.1.
- 18.9.81.1.1, (L1) Ensure 'Configure Windows Defender SmartScreen' is set to 'Enabled: Warn and prevent bypass.' The new number is 18.9.85.1.1.
- 18.9.85.1, (L2) Ensure 'Allow suggested apps in Windows Ink Workspace' is set to 'Disabled.' The new number is 18.9.89.1.
- 18.9.85.2, (L1) Ensure 'Allow Windows Ink Workspace' is set to 'Enabled: On, but disallow access above lock' OR 'Disabled' but not 'Enabled: On.' The new number is 18.9.89.2.
- 18.9.86.1, (L1) Ensure 'Allow user control over installs' is set to 'Disabled.' The new number is 18.9.90.1.
- 18.9.86.2, (L1) Ensure 'Always install with elevated privileges' is set to 'Disabled.' The new number is 19.7.43.1.
- 18.9.86.3, (L2) Ensure 'Prevent Internet Explorer security prompt for Windows Installer scripts' is set to 'Disabled.' The new number is 18.9.90.3.
- 18.9.87.1, (L1) Ensure 'Sign-in and lock last interactive user automatically after a restart' is set to 'Disabled.' The new number is 18.9.91.1.
- 18.9.96.2, (L1) Ensure 'Turn on PowerShell Transcription' is set to 'Disabled.' The new number is 18.9.100.2.
- 18.9.98.1.1, (L1) Ensure 'Allow Basic authentication' is set to 'Disabled.' The new number is 18.9.102.2.1.
- 18.9.98.1.2, (L1) Ensure 'Allow unencrypted traffic' is set to 'Disabled.' The new number is 18.9.102.2.3.
- 18.9.98.1.3, (L1) Ensure 'Disallow Digest authentication' is set to 'Enabled.' The new number is 18.9.102.1.3.
- 18.9.98.2.1, (L1) Ensure 'Allow Basic authentication' is set to 'Disabled.' The new number is 18.9.102.2.1.
- 18.9.98.2.2, (L2) Ensure 'Allow remote server management through WinRM' is set to 'Disabled.' The new number is 18.9.102.2.2.
- 18.9.98.2.3, (L1) Ensure 'Allow unencrypted traffic' is set to 'Disabled.' The new number is 18.9.102.2.3.
- 18.9.98.2.4, (L1) Ensure 'Disallow WinRM from storing RunAs credentials' is set to 'Enabled.' The new number is 18.9.102.2.4.
- 18.9.99.1, (L2) Ensure 'Allow Remote Shell Access' is set to 'Disabled.' The new number is 18.9.103.1.
- 18.9.100.2.1, (L1) Ensure 'Prevent users from modifying settings' is set to 'Enabled.' The new number is 18.9.105.2.1.
- 18.9.103.1.3, (L1) Ensure 'Select when Quality Updates are received' is set to 'Enabled: 0 days.' The new number is 18.9.108.4.3.
- 18.9.103.2, (L1) Ensure 'Configure Automatic Updates' is set to 'Enabled.' The new number is 18.9.108.2.1.
- 18.9.103.3, (L1) Ensure 'Configure Automatic Updates: Scheduled install day' is set to '0 - Every day.' The new number is 18.9.108.2.2.
- 18.9.103.4, (L1) Ensure 'No auto-restart with logged on users for scheduled automatic updates installations' is set to 'Disabled.' The new number is 18.9.108.1.1.
- 19.7.43.1, (L1) Ensure 'Always install with elevated privileges' is set to 'Disabled.' The new number is 18.9.90.2.

- **Removed**
  - The following CIS controls are no longer available:
    - Control 18.9.96.1 was removed and replaced by 18.9.100.1, (L1) Ensure 'Turn on PowerShell Script Block Logging' is set to 'Enabled.'
    - Control 18.9.103.1.1 was deprecated and is now removed.
    - Control 18.9.103.1.2 was deprecated and is now removed.

## Copyright and trademark notices

---

© 2024 Puppet, Inc., a Perforce company. All rights reserved.

Puppet and other identified trademarks are the property of Puppet, Inc., Perforce Software, Inc., or an affiliate. Such trademarks are claimed and/or registered in the U.S. and other countries and regions. All third-party trademarks are the property of their respective holders. References to third-party trademarks do not imply endorsement or sponsorship of any products or services by the trademark holder. Contact Puppet, Inc., for further details.