


## Enterprise Server 3.19 release notes

---

# Enterprise Server 3.19.4

[Download GitHub Enterprise Server 3.19.4](#)

March 12, 2026

 This is not the [latest release](#) of Enterprise Server. Please use the latest release for the latest security, performance, and bug fixes.

---

### 3.19.4: Security fixes [↗](#)

- **HIGH:** An attacker with push access to a repository could execute arbitrary code on the instance by injecting malicious values into Git push options. The push options were not properly sanitized before being included in internal headers used for Git operations, allowing the attacker to override internal metadata fields and achieve remote code execution. GitHub has requested CVE ID [CVE-2026-3854](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **HIGH:** An authenticated attacker could execute arbitrary JavaScript in another user's browser session. The vulnerability was an HTML-escaping flaw in task list rendering that allowed malicious task list items in issues or comments to bypass Content Security Policy protections. GitHub has requested CVE ID [CVE-2026-2266](#) for this vulnerability, which was reported via the GitHub Bug Bounty program. GitHub has requested [CVE ID CVE-2026-2266](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **MEDIUM:** An attacker could use the REST API endpoints `/search/commits` or `/search/issues` with a personal access token (classic) that lacks the `repo` scope to retrieve results from private or internal repositories by using the `repo:OWNER/REPO` qualifier. GitHub has requested CVE ID [CVE-2026-3582](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **MEDIUM:** An attacker with read access to a repository and write access to a project could bypass repository write permissions to modify issue and pull request labels, assignees, and other metadata by adding duplicate items to the project. GitHub has requested CVE ID [CVE-2026-3306](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).

### 3.19.4: Bug fixes [↗](#)

- Users experienced delays or failures when performing Git operations over HTTP. The operations could hang indefinitely due to a deadlock in the babeld service.
- For repositories using the Python uv package manager, `uv.lock` files were not included in the dependency graph. This prevented Dependabot from detecting vulnerable dependencies or providing security updates for those files.
- On instances with a cluster configuration, the OpenTelemetry Collector configuration file contained extraneous blank lines in the blackbox exporter section, resulting in improperly formatted YAML.
- Administrators could not load the Replication page in the Management Console due to an Internal Server Error.
- Some metric exporters enabled in the OpenTelemetry observability stack would have an incorrect `instance` label.
- After an upgrade, `ghe-config-apply` could fail to remove some pre-upgrade Docker images and report `Error response from daemon: conflict: unable to delete <id>`.
- In the Management Console, when an administrator uploaded an invalid license file, the page could fail to display an error notification.
- Administrators for instances using the collectd metrics stack saw empty `git fetch caching` graphs on the Management Console monitoring page.
- After upgrading, `ghe-config-apply` failed to start services including HAProxy and Redis. Docker images were incorrectly removed during the upgrade process, preventing services from starting.
- On the dependency graph page, users saw a banner promoting automatic dependency submission despite the feature being unavailable on GitHub Enterprise Server. The banner also linked to documentation that was inaccessible.
- The Copilot Code Review branch rule was unintentionally visible.
- On instances with GitHub Actions enabled, Actions workflow runs could be silently skipped when creating many issues rapidly via the API. Previously, some "issue opened" webhooks were processed before the new issue was saved to the database, causing the event to be dropped and the workflow not to start.

- Enterprise owners experienced slow loading and timeouts when updating personal access token lifetime policies for enterprises with many organizations.
- Users experienced failures when migrating repositories with releases using GitHub Enterprise Importer. Migrations failed to import release assets that were incompletely uploaded at the time of export, as the export archive referenced assets without including the corresponding files.
- When a security configuration was applied using the repositories table on any page other than the first, the table temporarily displayed the first page of repositories.
- Users saw an option to enable "Push protection for yourself" for public repositories in their user settings. This feature is not applicable to GitHub Enterprise Server. Users who enabled the feature experienced unexpected behavior with push protection in repositories that are visible to the entire enterprise. Push protection remains available at the repository, organization, and enterprise levels.
- After upgrading to 3.19, site administrators could not delete users via the Management Console or the REST API. User deletion attempts returned a `500 Internal Server Error`.

### 3.19.4: Changes [↗](#)

- To improve performance on large instances, HAProxy automatically scales its thread count based on available CPUs and uses higher connection limits for high-traffic backend services including GitHub Actions, database connections, job queues, and package registry. Administrators can override the thread count using `ghe-config haproxy-nbthread` if needed.
- API consumers can update issues via the REST API (PATCH) using fine-grained permissions for closing and reopening issues, and for setting milestones.
- API consumers can update issues via the GraphQL `updateIssue` mutation using fine-grained permissions for closing and reopening issues, and for setting milestones.

### 3.19.4: Known issues [↗](#)

- During an upgrade of GitHub Enterprise Server, custom firewall rules are removed. If you use custom firewall rules, you must reapply them after upgrading.
- During the validation phase of a configuration run, a `No such object` error may occur for the Notebook and Viewscreen services. This error can be ignored as the services should still correctly start.

- If the root site administrator is locked out of the Management Console after failed login attempts, the account does not unlock automatically after the defined lockout time. Someone with administrative SSH access to the instance must unlock the account using the administrative shell. For more information, see [Troubleshooting access to the Management Console](#).
- In some situations, large `.adoc` files stored in a repository do not render properly in the web UI. The raw contents are still available to view as plaintext.
- Admin stats REST API endpoints may time out on appliances with many users or repositories. Retrying the request until data is returned is advised.
- When following the steps for [Replacing the primary MySQL node](#), step 14 (running `ghe-cluster-config-apply`) might fail with errors. If this occurs, re-running `ghe-cluster-config-apply` is expected to succeed.
- Running `ghe-config-apply` as part of the steps for [Replacing a node in an emergency](#) may fail with errors if the node being replaced is still reachable. If this occurs, shutdown the node and repeat the steps.
- If a hotpatch upgrade requires the `haproxy-frontend` service to be restarted, the restart will hang if there are existing long-lived connections, such as browser web sockets or Git operations. No new connections will be accepted for up to 5 minutes. Any existing unfinished connections at this time will be disconnected.
- When restoring data originally backed up from a 3.13 or greater appliance version, the Elasticsearch indices need to be reindexed before some of the data will show up. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- When enabling automatic update checks for the first time in the Management Console, the status is not dynamically reflected until the "Updates" page is reloaded.
- When restoring from a backup snapshot, a large number of `mapper_parsing_exception` errors may be displayed.
- When initializing a new cluster, nodes with the `consul-server` role should be added to the cluster before adding additional nodes. Adding all nodes simultaneously creates a race condition between nomad server registration and nomad client registration.
- Admins setting up cluster high availability (HA) may encounter a spokes error when running `ghe-cluster-repl-status` if a new organization and repositories are created before using the `ghe-cluster-repl-bootstrap` command. To avoid this issue, complete the cluster HA setup with `ghe-cluster-repl-bootstrap` before creating new organizations and repositories.

- In a cluster, the host running restore requires access to the storage nodes via their private IPs.
- On an instance hosted on Azure, commenting on an issue via email means the comment is not added to the issue.
- After a restore, existing outside collaborators cannot be added to repositories in a new organization. This issue can be resolved by running `/usr/local/share/enterprise/ghe-es-search-repair` on the appliance.
- After a geo-replica is promoted to be a primary by running `ghe-repl-promote`, the Actions workflow of a repository does not have any suggested workflows.
- When publishing npm packages in a workflow after restoring from a backup to GitHub Enterprise Server 3.13.5.gm4 or 3.14.2.gm3, you may encounter a `401 Unauthorized` error from the GitHub Packages service. This can happen if the restore is from an N-1 or N-2 version and the workflow targets the npm endpoint on the backup instance. To avoid this issue, ensure the access token is valid and includes the correct scopes for publishing to GitHub Packages.
- The setting to define private registries at the organization level for code scanning is only available if Dependabot is also enabled for the instance.
- Upgrading or hotpatching to 3.19.1 may fail on nodes that have been continuously upgraded from versions older than 2021 (for example GHES version 2.17). If this issue occurs, you will see log entries prefixed with `invalid secret` in `ghe-config.log`. If you are running nodes from these older versions, it is recommended not to upgrade to 3.19.1.
- When applying an enterprise security configuration to all repositories (for example, enabling secret scanning or code scanning across all repositories), the system immediately enqueues enablement jobs for every organization in the enterprise simultaneously. For enterprises with a large number of repositories, this can result in significant system load and potential performance degradation. If you manage a large enterprise with many organizations and repositories, we recommend applying security configurations at the organization level rather than at the enterprise level in the UI. This allows you to enable security features incrementally and monitor system performance as you roll out changes.

---

## Enterprise Server 3.19.3

[Download GitHub Enterprise Server 3.19.3](#)

March 10, 2026

🚩 This is not the [latest patch release](#) of this release series, and this is not the [latest release](#) of Enterprise Server. Please use the latest release for the latest security, performance, and bug fixes.

**Warning:** GitHub Enterprise Server 3.19.3 has been unpublished for operational reasons. Please use the most recent available patch release of 3.19. [Updated: 2026-03-13]

### 3.19.3: Security fixes [↗](#)

- **HIGH:** An attacker with push access to a repository could execute arbitrary code on the instance by injecting malicious values into Git push options. The push options were not properly sanitized before being included in internal headers used for Git operations, allowing the attacker to override internal metadata fields and achieve remote code execution. GitHub has requested CVE ID [CVE-2026-3854](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **HIGH:** An authenticated attacker could execute arbitrary JavaScript in another users browser session by crafting a malicious task list item in an issue or comment that exploited an HTML-escaping flaw in task list rendering, bypassing Content Security Policy protections. GitHub has requested [CVE ID CVE-2026-2266](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **MEDIUM:** An attacker could use the REST API endpoints `/search/commits` or `/search/issues` with a personal access token (classic) that lacks the `repo` scope to retrieve results from private or internal repositories by using the `repo:OWNER/REPO` qualifier. GitHub has requested CVE ID [CVE-2026-3582](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **MEDIUM:** An attacker with read access to a repository and write access to a project could bypass repository write permissions to modify issue and pull request labels, assignees, and other metadata by adding duplicate items to the project. GitHub has requested CVE ID [CVE-2026-3306](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).

### 3.19.3: Bug fixes [↗](#)

- Users experienced delays or failures when performing Git operations over HTTP. The operations could hang indefinitely due to a deadlock in the babeld service.

- For repositories using the Python uv package manager, `uv.lock` files were not included in the dependency graph. This prevented Dependabot from detecting vulnerable dependencies or providing security updates for those files.
- On instances with a cluster configuration, the OpenTelemetry Collector configuration file contained extraneous blank lines in the blackbox exporter section, resulting in improperly formatted YAML.
- Administrators could not load the Replication page in the Management Console due to an Internal Server Error.
- Some metric exporters enabled in the OpenTelemetry observability stack would have an incorrect `instance` label.
- After an upgrade, `ghe-config-apply` could fail to remove some pre-upgrade Docker images and report `Error response from daemon: conflict: unable to delete <id>`.
- In the Management Console, when an administrator uploaded an invalid license file, the page could fail to display an error notification.
- Administrators for instances using the collectd metrics stack saw empty `git fetch caching` graphs on the Management Console monitoring page.
- After upgrading, `ghe-config-apply` failed to start services including HAProxy and Redis. Docker images were incorrectly removed during the upgrade process, preventing services from starting.
- On the dependency graph page, users saw a banner promoting automatic dependency submission despite the feature being unavailable on GitHub Enterprise Server. The banner also linked to documentation that was inaccessible.
- The Copilot Code Review branch rule was unintentionally visible.
- On instances with GitHub Actions enabled, Actions workflow runs could be silently skipped when creating many issues rapidly via the API. Previously, some "issue opened" webhooks were processed before the new issue was saved to the database, causing the event to be dropped and the workflow to not start. After this fix, workflow runs start reliably for all rapid issue creations, regardless of timing.
- Enterprise owners experienced slow loading and timeouts when updating personal access token lifetime policies for enterprises with many organizations.
- Users experienced failures when migrating repositories with releases using GitHub Enterprise Importer. Migrations failed to import release assets that were incompletely uploaded at the time

of export, as the export archive referenced assets without including the corresponding files.

- When a security configuration was applied using the repositories table on any page other than the first, the table temporarily displayed the first page of repositories.
- Users saw an option to enable "Push protection for yourself" for public repositories in their user settings. This feature is not applicable to GitHub Enterprise Server. Users who enabled the feature experienced unexpected behavior with push protection in repositories that are visible to the entire enterprise. Push protection remains available at the repository, organization, and enterprise levels.
- After upgrading to 3.19, site administrators could not delete users via the Management Console or the REST API. User deletion attempts returned a `500 Internal Server Error`.

### 3.19.3: Changes [↗](#)

- To improve performance on large instances, HAProxy automatically scales its thread count based on available CPUs and uses higher connection limits for high-traffic backend services including GitHub Actions, database connections, job queues, and package registry. Administrators can override the thread count using `ghe-config haproxy-nbthread` if needed.
- API consumers can update issues via the REST API (PATCH) using fine-grained permissions for closing and reopening issues, and for setting milestones.
- API consumers can update issues via the GraphQL `updateIssue` mutation using fine-grained permissions for closing and reopening issues, and for setting milestones.

### 3.19.3: Known issues [↗](#)

- During an upgrade of GitHub Enterprise Server, custom firewall rules are removed. If you use custom firewall rules, you must reapply them after upgrading.
- During the validation phase of a configuration run, a `No such object` error may occur for the Notebook and Viewscreen services. This error can be ignored as the services should still correctly start.
- If the root site administrator is locked out of the Management Console after failed login attempts, the account does not unlock automatically after the defined lockout time. Someone with administrative SSH access to the instance must unlock the account using the administrative shell. For more information, see [Troubleshooting access to the Management Console](#).

- In some situations, large `.adoc` files stored in a repository do not render properly in the web UI. The raw contents are still available to view as plaintext.
- Admin stats REST API endpoints may time out on appliances with many users or repositories. Retrying the request until data is returned is advised.
- When following the steps for [Replacing the primary MySQL node](#), step 14 (running `ghe-cluster-config-apply`) might fail with errors. If this occurs, re-running `ghe-cluster-config-apply` is expected to succeed.
- Running `ghe-config-apply` as part of the steps for [Replacing a node in an emergency](#) may fail with errors if the node being replaced is still reachable. If this occurs, shutdown the node and repeat the steps.
- If a hotpatch upgrade requires the `haproxy-frontend` service to be restarted, the restart will hang if there are existing long-lived connections, such as browser web sockets or Git operations. No new connections will be accepted for up to 5 minutes. Any existing unfinished connections at this time will be disconnected.
- When restoring data originally backed up from a 3.13 or greater appliance version, the Elasticsearch indices need to be reindexed before some of the data will show up. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- When enabling automatic update checks for the first time in the Management Console, the status is not dynamically reflected until the "Updates" page is reloaded.
- When restoring from a backup snapshot, a large number of `mapper_parsing_exception` errors may be displayed.
- When initializing a new cluster, nodes with the `consul-server` role should be added to the cluster before adding additional nodes. Adding all nodes simultaneously creates a race condition between nomad server registration and nomad client registration.
- Admins setting up cluster high availability (HA) may encounter a spokes error when running `ghe-cluster-repl-status` if a new organization and repositories are created before using the `ghe-cluster-repl-bootstrap` command. To avoid this issue, complete the cluster HA setup with `ghe-cluster-repl-bootstrap` before creating new organizations and repositories.
- In a cluster, the host running restore requires access the storage nodes via their private IPs.
- On an instance hosted on Azure, commenting on an issue via email means the comment is not added to the issue.

- After a restore, existing outside collaborators cannot be added to repositories in a new organization. This issue can be resolved by running `/usr/local/share/enterprise/ghe-es-search-repair` on the appliance.
- After a geo-replica is promoted to be a primary by running `ghe-repl-promote`, the Actions workflow of a repository does not have any suggested workflows.
- When publishing npm packages in a workflow after restoring from a backup to GitHub Enterprise Server 3.13.5.gm4 or 3.14.2.gm3, you may encounter a `401 Unauthorized` error from the GitHub Packages service. This can happen if the restore is from an N-1 or N-2 version and the workflow targets the npm endpoint on the backup instance. To avoid this issue, ensure the access token is valid and includes the correct scopes for publishing to GitHub Packages.



GitHub Docs

Version: Enterprise Server 3.19 ▾



☰ Enterprise administrators / Release notes


from versions older than 2021 (for example GHEES version 2.17). If this issue occurs, you will see log entries prefixed with `invalid secret` in `ghe-config.log`. If you are running nodes from these older versions, it is recommended not to upgrade to 3.19.1.

- When applying an enterprise security configuration to all repositories (for example, enabling secret scanning or code scanning across all repositories), the system immediately enqueues enablement jobs for every organization in the enterprise simultaneously. For enterprises with a large number of repositories, this can result in significant system load and potential performance degradation. If you manage a large enterprise with many organizations and repositories, we recommend applying security configurations at the organization level rather than at the enterprise level in the UI. This allows you to enable security features incrementally and monitor system performance as you roll out changes.

## Enterprise Server 3.19.2

[Download GitHub Enterprise Server 3.19.2](#)

February 10, 2026

 This is not the [latest patch release](#) of this release series, and this is not the [latest release](#) of Enterprise Server. Please use the latest release for the latest security, performance, and bug fixes.

### 3.19.2: Features

- Administrators can configure advanced SMTP settings for improved email delivery performance and reliability. These settings map to Postfix configuration parameters as documented in the Postfix documentation. New options include:
  - IPv4-only relay: Route email to addresses at a specific email domain through an IPv4-only relay host. Setting `smtp.ipv4-only` to `true` configures Postfix to route all email to the domain specified in `smtp.relay-domain` through `smtp.relay-host` on port `smtp.relay-port` using IPv4 only.
  - Connection caching: Control connection reuse and caching ( `smtp.connection-cache-time-limit` , `smtp.connection-reuse-count-limit` , `smtp.connection-cache-on-demand` ).
  - Delivery concurrency: Tune parallel email delivery limits ( `smtp.destination-concurrency-limit` , `smtp.initial-destination-concurrency` , `smtp.destination-concurrency-positive-feedback` ).
  - Queue management: Configure retry timing and queue processing ( `smtp.maximal-backoff-time` , `smtp.queue-run-delay` )
  - Connection limits: Set maximum inbound SMTP connections ( `smtp.client-connection-count-limit` ).
- For administrators using geo-replication or high availability (HA), `ghe-repl` tooling supports cross-cluster replication (CCR) for Elasticsearch, improving search index replication between instances.

### 3.19.2: Security fixes

- **MEDIUM:** By supplying the migration identifier, an attacker could upload unauthorized content to another user's repository migration export due to a missing authorization check. This could cause victims to download attacker-controlled migration archives, potentially impacting the integrity of downstream repository imports. GitHub has requested a CVE ID [CVE-2026-1355](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **LOW:** GitHub Enterprise Server included React versions 19.0, 19.1, and 19.2 in its package, which contain vulnerabilities in the React Server Components protocol ([CVE-2025-55182](#), [CVE-2025-66478](#)). GitHub Enterprise Server does not use React Server Components and was not vulnerable to exploitation. React has been updated to version 19.2.1 to address findings from security scanning tools.

- **HIGH:** An attacker could merge their own pull request into a repository that allowed forks and for which they didn't have write access, by exploiting an incorrect authorization check in the `enable_auto_merge` mutation for pull requests in specific scenarios. Exploitation required a clean pull request status and only applied to branches without branch protection rules enabled. GitHub has requested CVE ID [CVE-2026-1999](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).
- **HIGH:** An authenticated attacker could exploit a URL redirection vulnerability in GitHub Enterprise Server to leak privileged authorization tokens by redirecting requests to an attacker-controlled domain. This could allow exfiltration of the `Actions.ManageOrgs` JWT and potential remote code execution. This vulnerability was reported via the [GitHub Bug Bounty program](#).

### 3.19.2: Bug fixes

- The Elasticsearch panel in the Operational Health dashboard of the Management Console did not correctly represent the clusters health. As a result, administrators may have seen inaccurate status indicators for Elasticsearch availability and performance.
- Alambic failed to start after reboot or upgrade if legacy multi-disk for alambic was set up.
- GitHub Enterprise Server Backup Service (preview) was disabled after upgrading.
- Running `ghe-config-apply` could fail if Redis experienced transient connectivity issues during the configuration process.
- Resolved an issue in Enterprise Manage where the Backups (Preview) tab failed to open and returned an Internal Server Error. This tab now load as expected.
- On instances initially configured more than five years ago, administrators were unable to access the Management Console after upgrading due to an outdated session secret that was below the required 64-byte length.
- Administrators were unable to access the "Updates" tab in the Management Console due to a template rendering error that displayed an Internal Server Error.
- Administrators could not use the "Clear dependencies" tool in the site admin dashboard because the required `RESET_MANIFESTS_CONSUMER_GROUP` environment variable was missing.
- When administrators configured password authentication, the Prometheus endpoint for OpenTelemetry metrics failed to expose metrics due to health check failures.

- When administrators would apply configuration changes via the management console, the state shown would occasionally briefly flicker to a failure before being marked as successful causing confusion as to whether the configuration had succeeded.
- On an instance configured behind a load balancer, users received unexpected secondary rate limit warnings during authentication when the `X-Forwarded-For` header included port numbers. This occurred because the system incorrectly ignored the header values containing ports, preventing proper client IP address identification.
- Users with read access to a repository were unable to close issues even when granted the "Close issue" fine-grained permission through custom repository roles. Permission checks were relying solely on the triager role when evaluating a users ability to close issues.
- Push rejections due to custom pre-receive hooks were not visible in the audit log.
- Users could only view webhook deliveries from the previous three days.

### 3.19.2: Changes [↗](#)

- Administrators can configure database connection pool limits for the authentication and authorization services to improve performance on instances experiencing high concurrent request volumes. The limits can be adjusted using `ghe-config` keys: `app.authnd.mysql-max-open-conns` , `app.authnd.mysql-max-idle-conns` , `app.authzd.db-resolver-max-open-conns` , and `app.authzd.db-resolver-max-idle-conns` . The default values remain unchanged (authnd: 100 max open and 100 max idle connections; authzd: 100 max open and 15 max idle connections). These settings should only be adjusted with guidance from GitHub Support.
- The `spokesctl status` command displays the current priority of repository issues based on the most recent check. Previously, the command displayed the highest priority the issue had reached since it was first detected, which could be misleading if the issue had been partially resolved.

### 3.19.2: Known issues [↗](#)

- During an upgrade of GitHub Enterprise Server, custom firewall rules are removed. If you use custom firewall rules, you must reapply them after upgrading.
- During the validation phase of a configuration run, a `No such object` error may occur for the Notebook and Viewscreen services. This error can be ignored as the services should still correctly start.

- If the root site administrator is locked out of the Management Console after failed login attempts, the account does not unlock automatically after the defined lockout time. Someone with administrative SSH access to the instance must unlock the account using the administrative shell. For more information, see [Troubleshooting access to the Management Console](#).
- In some situations, large `.adoc` files stored in a repository do not render properly in the web UI. The raw contents are still available to view as plaintext.
- Admin stats REST API endpoints may timeout on appliances with many users or repositories. Retrying the request until data is returned is advised.
- When following the steps for [Replacing the primary MySQL node](#), step 14 (running `ghe-cluster-config-apply`) might fail with errors. If this occurs, re-running `ghe-cluster-config-apply` is expected to succeed.
- Running a config apply as part of the steps for [Replacing a node in an emergency](#) may fail with errors if the node being replaced is still reachable. If this occurs, shutdown the node and repeat the steps.
- If a hotpatch upgrade requires the `haproxy-frontend` service to be restarted, the restart will hang if there are existing long-lived connections, such as browser web sockets or Git operations. No new connections will be accepted for up to 5 minutes. Any existing unfinished connections at this time will be disconnected.
- When restoring data originally backed up from a 3.13 or greater appliance version, the Elasticsearch indices need to be reindexed before some of the data will show up. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- When enabling automatic update checks for the first time in the Management Console, the status is not dynamically reflected until the "Updates" page is reloaded.
- When restoring from a backup snapshot, a large number of `mapper_parsing_exception` errors may be displayed.
- When initializing a new GHES cluster, nodes with the `consul-server` role should be added to the cluster before adding additional nodes. Adding all nodes simultaneously creates a race condition between nomad server registration and nomad client registration.
- Admins setting up cluster high availability (HA) may encounter a spokes error when running `ghe-cluster-repl-status` if a new organization and repositories are created before using the `ghe-cluster-repl-bootstrap` command. To avoid this issue, complete the cluster HA setup with `ghe-cluster-repl-bootstrap` before creating new organizations and repositories.

- In a cluster, the host running restore requires access the storage nodes via their private IPs.
- On an instance hosted on Azure, commenting on an issue via email meant the comment was not added to the issue.
- After a restore, existing outside collaborators cannot be added to repositories in a new organization. This issue can be resolved by running `/usr/local/share/enterprise/ghe-es-search-repair` on the appliance.
- After a geo-replica is promoted to be a primary by running `ghe-repl-promote`, the actions workflow of a repository does not have any suggested workflows.
- When publishing npm packages in a workflow after restoring from a backup to GitHub Enterprise Server 3.13.5.gm4 or 3.14.2.gm3, you may encounter a `401 Unauthorized` error from the GitHub Packages service. This can happen if the restore is from an N-1 or N-2 version and the workflow targets the npm endpoint on the backup instance. To avoid this issue, ensure the access token is valid and includes the correct scopes for publishing to GitHub Packages.
- The setting to define private registries at the organization level for code scanning is only available if dependabot is also enabled for the instance.
- Upgrading or hotpatching to 3.19.1 may fail on nodes that have been continuously upgraded from versions older than 2021 (i.e. 2.17). If this issue occurs, you will see log entries prefixed with `invalid secret` in `ghe-config.log`. If you are running nodes from these older versions, it is recommended not to upgrade to 3.19.1.
- When applying an enterprise security configuration to all repositories (for example, enabling Secret Scanning or Code Scanning across all repositories), the system immediately enqueues enablement jobs for every organization in the enterprise simultaneously. For enterprises with a large number of repositories, this can result in significant system load and potential performance degradation. If you manage a large enterprise with many organizations and repositories, we recommend applying security configurations at the organization level rather than at the enterprise level in the UI. This allows you to enable security features incrementally and monitor system performance as you roll out changes.
- When viewing the status of an ongoing backup on the "Backups" page of the Management Console, the backup may initially report as "incomplete" instead of "in progress". You can ignore the initial "incomplete" status because the backup is still running and will report the correct status once it has progressed further. In some configurations, such as cluster topologies, this may take up to 5 minutes.

# Enterprise Server 3.19.1

[Download GitHub Enterprise Server 3.19.1](#)

January 06, 2026

 This is not the [latest patch release](#) of this release series, and this is not the [latest release](#) of Enterprise Server. Please use the latest release for the latest security, performance, and bug fixes.

## 3.19.1: Security fixes

- **HIGH:** An authenticated attacker with permission to configure webhooks could perform SSRF to access internal-only services on the instance, potentially disrupting background job processing. Exploitation required webhook configuration privileges and the ability to craft valid service requests. GitHub has requested CVE ID [CVE-2026-1999](#) for this vulnerability, which was reported via the [GitHub Bug Bounty](#) program.
- **HIGH:** An attacker could execute code within a victim's browser, potentially accessing sensitive information, by causing malicious HTML to be injected into the DOM when content is rendered by the Filter component found across GitHub. GitHub has requested CVE ID [CVE-2025-13744](#) for this vulnerability, which was reported via the [GitHub Bug Bounty program](#).

## 3.19.1: Bug fixes

- On instances with GitHub Actions enabled, when administrators deleted a self-hosted runner from the service, the runner process continued running on the host and did not exit automatically.
- In the "Password and authentication policies" section of the Management Console, administrators could specify invalid values for the "Login attempt limit for all users" and "Lockout time for Management Console users" settings, because inputs were not correctly validated.
- The highlighted section on the sidebar of the Management Console settings page did not always accurately reflect the content currently scrolled into view.
- After selecting "local storage" for migration storage in the Management Console, administrators found that the setting appeared to be cleared when the settings page refreshed.
- Administrators saw errors about missing or incomplete Actions cluster configuration in the Management Console, even on instances where GitHub Actions was not enabled.
- Custom network time protocol (NTP) settings could be removed after upgrades.

- Site administrators could not easily identify when a configuration run for their instance failed in the Management Console. Failed runs were indicated only by a header and steps could remain in a "pending" state.
- Site administrators could not generate a CSV list of SCIM-provisioned users with the `ghe-scim-identities-csv` command because its wrapper script was missing from `/usr/local/bin`.
- Administrators encountered inaccurate free disk space calculations when setting Elasticsearch watermarks, as incorrect methods were used for determining root and data disk sizes.
- Upgrading an instance from 3.17.x or 3.18.x to 3.19.x would reset existing observability metrics settings.
- Administrators who set the `ELASTOMER_INDEX_LOCK_BACKOFF_ATTEMPTS` environment variable to configure Elasticsearch index lock backoff attempts saw no effect, as the instance required the `ENTERPRISE_` prefix for this variable.
- Commit authors who ignored notifications from a repository did not receive secret scanning alert emails when their credentials were detected in that repository.
- The site admin bar displayed debugging information used by GitHub.
- Users could not access pull requests authored by a GitHub App when the pull request targeted a branch protected by a rule.
- On high-availability clusters with Elasticsearch Cross Cluster Replication (CCR) enabled, replication failed if the `datacenter` and `consul-datacenter` values didn't match.
- When administrators enabled GitHub Advanced Security features in bulk, enablement progress was not always tracked accurately. As a result, subsequent bulk scans for GitHub Secret Protection could be triggered or grouped incorrectly.

### 3.19.1: Changes

- Administrators can capture distributed tracing data for Nomad job allocations using the `usr/local/share/enterprise/ghe-capture-trace-data` command to help diagnose performance issues. This feature is available only on standalone instances and should be run with guidance from GitHub Support.
- Developers can see code scanning annotations listed with errors first, followed by warnings and notes, in newly generated annotation lists. Previously, annotation order was random, which

could make critical issues less visible, especially when some annotations were omitted due to high alert volume. This improves the clarity and prioritization of code scanning results.

- To help large instances run more efficiently, enterprise administrators can more easily opt out of the behavior where GitHub generates a rebase commit every time we check whether a pull request can be merged. This change consolidates prior handling of multiple repository rule variables and backend feature flags.

Now, if an administrator sets the instance's

`skip_rebase_commit_generation_from_rebase_merge_settings` configuration variable to `true`, the "Allow rebase merging" option in a repository's pull request settings becomes the source of truth for whether rebase commits are generated when mergeability is checked.

- You can configure multiple data disks to host MySQL and repository data. This capability is currently in public preview and is applicable only for standalone and high availability topologies. It does not apply to cluster topologies. For more information, see [Configuring multiple data disks](#). [Updated: 2026-01-19]

### 3.19.1: Known issues [↗](#)

- When applying an enterprise security configuration to all repositories (for example, enabling Secret Scanning or Code Scanning across all repositories), the system immediately enqueues enablement jobs for every organization in the enterprise simultaneously. For enterprises with a large number of repositories, this can result in significant system load and potential performance degradation. If you manage a large enterprise with many organizations and repositories, we recommend applying security configurations at the organization level rather than at the enterprise level in the UI. This allows you to enable security features incrementally and monitor system performance as you roll out changes.
- Upgrading or hotpatching to 3.19.1 may fail on very old nodes that have been continuously upgraded from versions older than 2021 versions (i.e. 2.17). If this issue occurs, you will see log entries prefixed with `invalid secret` in `ghe-config.log`. If you are running nodes this old, it is recommended not to upgrade to 3.19.1. If you must hotpatch to 3.19.1, first run `ghe-config 'secrets.session-manage' | tr -d '\n' | wc -c`. If the output is less than 64, run `ghe-config --unset 'secrets.session-manage'` and `ghe-config-apply` before you start the hotpatch. You can also run these same commands after the hotpatch to recover from the failure. [Updated: 2026-01-12]
- During an upgrade of GitHub Enterprise Server, custom firewall rules are removed. If you use custom firewall rules, you must reapply them after upgrading.


- During the validation phase of a configuration run, a `No such object` error may occur for the Notebook and Viewscreen services. This error can be ignored as the services should still correctly start.
- If the root site administrator is locked out of the Management Console after failed login attempts, the account does not unlock automatically after the defined lockout time. Someone with administrative SSH access to the instance must unlock the account using the administrative shell. For more information, see [Troubleshooting access to the Management Console](#).
- In some situations, large `.adoc` files stored in a repository do not render properly in the web UI. The raw contents are still available to view as plaintext.
- Admin stats REST API endpoints may time out on appliances with many users or repositories. Retrying the request until data is returned is advised.
- When following the steps for [Replacing the primary MySQL node](#), step 14 (running `ghe-cluster-config-apply`) might fail with errors. If this occurs, re-running `ghe-cluster-config-apply` is expected to succeed.
- Running a config apply as part of the steps for [Replacing a node in an emergency](#) may fail with errors if the node being replaced is still reachable. If this occurs, shutdown the node and repeat the steps.
- If a hotpatch upgrade requires the `haproxy-frontend` service to be restarted, the restart will hang if there are existing long-lived connections, such as browser web sockets or Git operations. No new connections will be accepted for up to 5 minutes. Any existing unfinished connections at this time will be disconnected.
- When restoring data originally backed up from a 3.13 or greater appliance version, the Elasticsearch indices need to be reindexed before some of the data will show up. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- When enabling automatic update checks for the first time in the Management Console, the status is not dynamically reflected until the "Updates" page is reloaded.
- When restoring from a backup snapshot, a large number of `mapper_parsing_exception` errors may be displayed.
- When initializing a new cluster, nodes with the `consul-server` role should be added to the cluster before adding more nodes. Adding all nodes simultaneously creates a race condition between nomad server registration and nomad client registration.

- Administrators setting up cluster high availability (HA) may encounter a spokes error when running `ghe-cluster-repl-status` if a new organization and repositories are created before using the `ghe-cluster-repl-bootstrap` command. To avoid this issue, complete the cluster HA setup with `ghe-cluster-repl-bootstrap` before creating new organizations and repositories.
- In a cluster, the host running restore requires access to the storage nodes via their private IPs.
- On an instance hosted on Azure, commenting on an issue via email meant the comment was not added to the issue.
- After a restore, existing outside collaborators cannot be added to repositories in a new organization. This issue can be resolved by running `/usr/local/share/enterprise/ghe-es-search-repair` on the appliance.
- After a geo-replica is promoted to be a primary by running `ghe-repl-promote`, the actions workflow of a repository does not have any suggested workflows.
- When publishing npm packages in a workflow after restoring from a backup to GitHub Enterprise Server 3.13.5.gm4 or 3.14.2.gm3, you may encounter a `401 Unauthorized` error from the GitHub Packages service. This can happen if the restore is from an N-1 or N-2 version and the workflow targets the npm endpoint on the backup instance. To avoid this issue, ensure the access token is valid and includes the correct scopes for publishing to GitHub Packages.
- The setting to define private registries at the organization level for code scanning is only available if Dependabot is enabled for the instance.
- In patch 3.19.1, we identified an issue in the Management Console where the Backups (Preview) and Updates tabs may fail to open and instead return an Internal Server Error. We recommend using the command line interface (CLI) for backups and updates until an updated patch is released. [Updated: 2026-01-13]
- When viewing the status of an ongoing backup on the "Backups" page of the Management Console, the backup may initially report as "incomplete" instead of "in progress." You can ignore this initial "incomplete" status; the backup is still running and will report the correct status once it has progressed further. In some configurations, such as cluster topologies, this may take up to 5 minutes.

# Enterprise Server 3.19.0

[Download GitHub Enterprise Server 3.19.0](#)

December 09, 2025

 This is not the [latest patch release](#) of this release series, and this is not the [latest release](#) of Enterprise Server. Please use the latest release for the latest security, performance, and bug fixes.

## 3.19.0: Features

### • Instance services

- You can configure which SSH and TLS ciphers are used on your instance. You can view the default ciphers and select preferred ones, providing you flexibility and ability to exclude weak ciphers. For more information, see [Configuring TLS and SSH ciphers](#).
- Starting 3.19, new installations of GHES will have OpenTelemetry metrics enabled and Collectd metrics disabled by default. You have the option to toggle between the two. Upgraded instances will retain their current settings. In about two to three releases, OpenTelemetry metrics will become the only supported metrics. To learn about OTel metrics, see [OpenTelemetry metrics](#).
- `ghe-repl-start` can now be executed without requiring a maintenance window when setting up a new replica, as long as `ghe-repl-setup` is immediately followed by `ghe-config-apply`. [Updated: 2025-12-17]
- You can configure multiple data disks to host MySQL and repository data. This capability is currently in public preview and is applicable only for standalone and high availability topologies. It does not apply to cluster topologies. For more information, see [Configuring multiple data disks](#). [Updated: 2026-01-19]

### • Migrations

- Administrators must update network allowlists with the new IP address ranges for GitHub Enterprise Importer migrations. Without this configuration, migration operations will fail due to blocked connectivity between environments.

### • APIs

- You can install GitHub Apps on the enterprise account and use them to manage your enterprise. Enterprise-installed GitHub Apps have access to a new set of permissions:
  - Managing GitHub App installations across the enterprise
  - SCIM provisioning and SSO management
  - Custom repository properties
  - Custom organization roles owned by the enterprise
  - Enterprise people management

Managing GitHub Apps across the enterprise allows you to programmatically audit, install, and uninstall GitHub Apps for all of the organizations in your enterprise using a single token. This high-powered permission enables better organization management at scale.

- Users can be made application managers of GitHub Apps owned by the enterprise. App Managers can update the application registration but do not have the ability to manage application installations.

The app manager feature has also been updated to use the roles platform, which means that organization teams can be made app managers of individual organization-owned apps, and a new Organization App Manager role can be assigned to teams and users to give them access to *all* of the apps owned by an organization. For more information, see [About GitHub App managers](#).

- **GitHub Advanced Security**

- Administrators can delegate code scanning alert dismissal to repository users. This enables responsible users to manage security findings and streamline remediation directly from the repository. The delegated alert dismissal feature is now generally available. For more information, see [the changelog](#)
- Administrators and security teams can now choose between default and advanced CodeQL setups for code scanning. The advanced setup allows for custom queries and more granular configuration, while the default setup offers a simplified workflow for standard security analysis. For more information, see [the changelog](#)
- The REST API for secret scanning now returns `first_location_detected` and `has_more_locations` fields in its responses.
- Administrators can specify which secret scanning patterns are included in push protection to enhance control over exposure prevention workflows. This update allows finer-tuning of push protected secrets.

- Organization and security admins can now run a free scan to understand how their repositories are affected by secret leaks and exposures. These secret risk assessments can be run at the organization level from the `Security` tab.
- When uploading analysis results for code scanning using SARIF files, each run in a multi-run SARIF file is now processed as a separate scan. Previously, multiple runs in one SARIF file were combined into a single scan, which could cause confusion in results and reporting. For more information, see [the changelog](#).
- GitHub secret scanning now detects and alerts you on secrets found in GitHub wikis, in addition to previously supported locations, including GitHub issues, pull requests, and discussions.

Secrets, like API keys, passwords, and tokens, can hide in many places. If these leaks aren't managed correctly, each one of them could pose a substantial risk. To help protect you from leaked secrets, anywhere within your GitHub perimeter, GitHub provides visibility across all major surfaces for hundreds of supported token formats.

- This release comes installed with version 2.22.4 of the CodeQL CLI, used in the CodeQL action for code scanning. Significant updates since the default version installed on GitHub Enterprise Server 3.18 include:
  - Users can analyze Go codebases more comprehensively, as CodeQL 2.22.0 improves coverage for Go. The release extends support for Go's generics and enhances the precision of dataflow analysis, enabling identification of vulnerabilities and defects in a wider variety of Go code patterns.
  - Users working with Swift can analyze projects using Swift 6.1.2, with CodeQL now supporting this version. This enhancement enables security and quality analyses for organizations adopting the latest Swift updates.
  - Users can now analyze Rust projects using CodeQL, with Rust support available in public preview. Organizations developing in Rust can begin early adoption of vulnerability detection and quality analyses in this language. Rust support is subject to change as feedback is gathered during the preview period.
  - Users analyzing Go codebases can scan projects built with Go 1.25, as CodeQL adds support for this new Go release.
  - View more in the changelogs for versions [CodeQL 2.22.0](#), [CodeQL 2.22.1](#), and [CodeQL 2.22.4](#).

- **Dependabot**

- Administrators and security teams can prioritize security fixes using the new Dependabot metrics page. The page provides insights on open vulnerable dependencies and other metrics to inform vulnerability management. This feature is now generally available for GitHub Advanced Security customers.
- Administrators and security teams can use the new Dependabot metrics page to prioritize remediation efforts. The page displays summary metrics and detailed insights to help track code security status over time.
- Dependabot now supports Gradle lockfiles in GHES, enabling users to keep dependencies up to date and improve supply chain security by automatically creating pull requests when newer versions are detected. This helps maintainers ensure project stability and security when managing Gradle projects.
- Administrators can optionally configure Dependabot to wait for a package to reach a specified minimum age before updating dependencies in their `dependabot.yml` files.
- Administrators can configure Dependabot in the `dependabot.yml` file to create a single pull request that updates dependencies across multiple package ecosystems within a repository.
- Administrators can centrally manage configurations for private registries used by Dependabot. This allows for streamlined setup and maintenance of registry credentials, improving the workflow for managing dependencies securely across the organization.
- Users can keep `vcpkg` dependencies up to date with Dependabot version updates. For more information, see the [changelog](#).
- Administrators and users can automate version updates for Rust toolchain dependencies using Dependabot. This enhancement streamlines the process of keeping Rust environments up to date and secure, reducing manual overhead for dependency management. For details, see the [changelog](#).
- Administrators and repository maintainers can now configure Dependabot to exclude automatic pull requests for dependency manifests located in selected subdirectories. This update helps users manage updates more flexibly and avoid unnecessary PRs for specific project paths. For more information, see the [changelog](#).
- You can now choose a "Not set" option for GitHub Code Security features in your organization's security configurations. Previously, you could only enable or disable features like code scanning and Dependabot at the organization level. With the new "Not set" option, you can enforce some security settings (such as secret scanning) while letting

repository administrators decide whether to enable GitHub Code Security features on their repositories.

This update gives organizations more flexibility in managing security requirements and helps repository administrators tailor their security setup to their specific needs.

To learn more about configuring security settings at the organization level, see [Creating a custom security configuration](#).

- Administrators can configure expanded cooldown windows for Dependabot alerts, allowing more flexible alert suppression during periods of high activity. Additionally, Dependabot now supports additional package managers, simplifying workflows for enterprises using diverse ecosystems. For the full list, see [Dependabot supported ecosystems and repositories](#).
- Administrators and repository owners can manage Dependabot alerts using batched updates for dependencies. This feature reduces alert noise by grouping related alerts and allowing simultaneous remediation, streamlining workflow and improving oversight for security and maintenance.
- Dependabot can now update private Go modules hosted on enterprise registries and behind GOPROXY-compatible private proxies, as well as public modules, within the same workflow. This enables automated version and security updates for internal Go libraries.

## • GitHub Actions

- For self-hosted GitHub Actions runners on this GitHub Enterprise Server release, the minimum required version of the GitHub Actions Runner application is 2.328.0. See the release notes for this version in the [actions/runner repository](#). If your instance uses ephemeral self-hosted runners and you've disabled automatic updates, you must upgrade your runners to this version of the Runner application before upgrading your instance to this GitHub Enterprise Server release.
- Enterprise administrators can assign fine-grained permissions for GitHub Actions through custom repository roles. This update enables precise control over workflow access, improving security and flexibility for automation management in repositories.
- Administrators can enforce policies to block specific actions and require SHA-based pinning when workflows use actions from public repositories. These policies help improve security for workflows by ensuring only approved actions are used and referenced by immutable SHAs.

## • Community experience

- Users can view a repository's contributing guidelines directly from both the repository's main tab and the sidebar. This feature makes it easier for contributors to find and follow project-specific contribution instructions, supporting a more accessible and collaborative workflow.

## • Organizations

- Enterprise administrators can create custom organization roles that are available in every organization in the enterprise, setting a standard set of roles for your organization owners to assign. These roles cannot be edited by organization owners.

As part of this update, the number of custom roles that can be created in enterprises and organizations has been raised to 20 per role type and owner. This means that an organization owner can have up to 40 custom roles to pick from.

## • Repositories

- Enterprise administrators can manage rules more efficiently with the general availability of ruleset history, import, and export. Ruleset history allows tracking and rolling back changes, while import and export simplify sharing and reusing rulesets, including [GitHub's ruleset-recipes](#).
- Users can more easily explore contributors and code frequency insights with improved navigation, interactive chart legends for hiding data series, and options to view or download the data as a CSV or PNG. See [Repositories - Updated insight views \(General Availability\)](#) on the GitHub Blog.

## • Issues

- Users can duplicate issues to any repository with a Duplicate issue action in the sidebar. The new form prepopulates title, description, assignees, labels, type, projects, and milestone, helping reuse formats, split large tasks, and create variants across repositories. Edit details before creation to tailor scope.
- Users can attach a wider range of code, data, document, image, audio, and log files in issues, pull requests, discussions, and comments: .py .yaml .yml .css .xml .html .htm .js .sql .java .c .cpp .sh .php .ts .tsx .cs .ipynb .pdb .xlsm .tsv .drawio .bin .rtf .doc .debug .msg .eml .copilotmd .bmp .tif .tiff .mp3 and .wav.

## • Commits

- Users benefit from a refreshed commit details page that enhances code review and navigation. The improved experience displays comment counts directly in the file tree, enables seamless switching between unified and split views, and introduces settings for line height and minimizing comments shown in diffs.

- **Pull requests**

- The improved "Files changed" experience for pull requests introduces a streamlined interface with enhanced navigation and filtering options, making it easier to review and manage changes. This feature is in public preview and subject to change.
- Pull request search in the web interface and via GraphQL and REST APIs now uses Elasticsearch as its dedicated backend, matching the existing issues search infrastructure. This update improves reliability and helps prevent timeouts when searching for pull requests in large repositories.

- **Accessibility**

- Improved accessibility for pull request reviewer status indicators. Users with assistive technologies can more easily identify reviewer status, supporting a more inclusive code review experience across pull requests. For more information, see [About pull request reviews](#).

### 3.19.0: Changes

- The code viewer and editor consistently respect each user's defined tab width preference across files and sessions. Previously, tab width settings could be inconsistently applied, causing code to display with unexpected indentation. This update ensures a uniform code viewing experience.
- The default tab size for code rendering is now set to 4 spaces instead of 8. This change provides a more consistent and readable display for code across the platform, aligning with common coding standards and improving the experience for developers who view or review code.
- Email notifications for issues and pull requests include additional headers to improve filtering and organization in email clients. These new custom headers give users and administrators more options for managing and sorting notification emails.
- Enterprises using IP allowlists should verify and update their network settings to include the newly required IP ranges for importer migrations. Failure to allow these addresses prevents

successful migrations.

- Projects now support up to 50,000 active items and 10,000 archived items. The previous limit was 1,200 items total. There is no option to opt out of this increased limit.
- YJIT, Ruby's Just-In-Time (JIT) compiler, is enabled by default. Users may experience faster application performance and improved resource efficiency across their instance. [Updated: 2025-12-17]

### 3.19.0: Known issues [↗](#)

- Custom firewall rules are removed during the upgrade process.
- During the validation phase of a configuration run, a No such object error may occur for the Notebook and Viewscreen services. This error can be ignored as the services should still correctly start.
- If the root site administrator is locked out of the Management Console after failed login attempts, the account does not unlock automatically after the defined lockout time. Someone with administrative SSH access to the instance must unlock the account using the administrative shell. For more information, see "[Troubleshooting access to the Management Console](#)."
- In some situations, large `.adoc` files stored in a repository do not render properly in the web UI. The raw contents are still available to view as plaintext.
- Admin stats REST API endpoints may timeout on appliances with many users or repositories. Retrying the request until data is returned is advised.
- When following the steps for [Replacing the primary MySQL node](#), step 14 (running `ghe-cluster-config-apply`) might fail with errors. If this occurs, re-running `ghe-cluster-config-apply` is expected to succeed.
- Running a config apply as part of the steps for [Replacing a node in an emergency](#) may fail with errors if the node being replaced is still reachable. If this occurs, shutdown the node and repeat the steps.
- When restoring data originally backed up from a 3.13 or greater appliance version, the elasticsearch indices need to be reindexed before some of the data will show up. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.

- When initializing a new GHES cluster, nodes with the `consul-server` role should be added to the cluster before adding additional nodes. Adding all nodes simultaneously creates a race condition between nomad server registration and nomad client registration.
- Admins setting up cluster high availability (HA) may encounter a spokes error when running `ghe-cluster-repl-status` if a new organization and repositories are created before using the `ghe-cluster-repl-bootstrap` command. To avoid this issue, complete the cluster HA setup with `ghe-cluster-repl-bootstrap` before creating new organizations and repositories.
- In a cluster, the host running restore requires access the storage nodes via their private IPs.
- On an instance hosted on Azure, commenting on an issue via email meant the comment was not added to the issue.
- After a restore, existing outside collaborators are unable to be added to repositories in a new organization. This issue can be resolved by running `/usr/local/share/enterprise/ghe-es-search-repair` on the appliance.
- After a geo-replica is promoted to be a primary by running `ghe-repl-promote`, the actions workflow of a repository does not have any suggested workflows.
- When publishing npm packages in a workflow after restoring from a backup to GitHub Enterprise Server 3.13.5.gm4 or 3.14.2.gm3, you may encounter a `401 Unauthorized` error from the GitHub Packages service. This can happen if the restore is from an N-1 or N-2 version and the workflow targets the npm endpoint on the backup instance. To avoid this issue, ensure the access token is valid and includes the correct scopes for publishing to GitHub Packages.
- Users may see a mismatch between repository-level Dependabot alerts and the overall Security Risk dashboard metrics. This can be resolved by reloading the page.
- The setting to define private registries at the organization level for code scanning is only available if dependabot is also enabled for the instance.
- When viewing the status of an ongoing backup on the "Backups" page of the Management Console, the backup may initially report as "incomplete" instead of "in progress". This initial status can be ignored; the backup is still running and will report the correct status once it has progressed further. In some configurations, such as cluster topologies, this may take up to 5 minutes.

### 3.19.0: Closing down

- As announced in [this previous blog post](#), GitHub will stop supporting basic authentication to APIs using a username and password in the coming versions of GHES. Instead of using

password authentication, [create a personal access token](#) in limited situations like testing. You should authenticate apps in production by using the web applications flow. For more information, see [Authorizing OAuth apps](#)

- The "reviewers" configuration option for Dependabot pull requests is retired. Reviewers are now determined by repository CODEOWNERS files. If your workflow depended on the "reviewers" option, update your automation to use CODEOWNERS for assigning pull request reviewers.
- Starting 3.21, networking-related syscalls will be disabled by default in the pre-receive hook environment. For enhanced security, hook environments will be placed in dedicated network namespaces. You will be able to override the default setting by setting pre-receive-hook-networking to enabled. As an alternative to many pre-receive hooks, see [About rulesets](#).
- In 3.20, we will be retiring `Telegraf`. For context, this was a dark-shipped service running in the background and not part of any customer workflows. If you have discovered it and notice it is missing in a future release, we want to you to know we have intentionally removed it.

## Legal

© 2026 GitHub, Inc. [Terms](#) [Privacy](#) [Status](#) [Pricing](#) [Expert services](#) [Blog](#)