Enterprise Server 3.18 release notes

Enterprise Server 3.18.7

Download GitHub Enterprise Server 3.18.7

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.

3.18.7: 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.
- **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.
- **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.

3.18.7: 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.
- 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>`.
- 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.
- 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.

3.18.7: 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) or the GraphQL `updateIssue` mutation using fine-grained permissions for closing and reopening issues, and for setting milestones.

3.18.7: 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.
- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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.
- Unexpected elements may appear in the UI on the repository overview page for locked repositories.
- 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.
- Custom NTP settings are removed during the upgrade process.
- 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 3.18.6

Download GitHub Enterprise Server 3.18.6

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 3.18.6 has been unpublished for operational reasons. Please use the most recent available patch release of 3.18. [Updated: 2026-3-13]

3.18.6: 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.

- **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.

- **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.

3.18.6: 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.

- 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>`.

- 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.

- 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.

3.18.6: 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) or the GraphQL `updateIssue` mutation using fine-grained permissions for closing and reopening issues, and for setting milestones.

3.18.6: 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.
- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.

- 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 `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`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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.
- Unexpected elements may appear in the UI on the repository overview page for locked repositories.
- 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.

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 3.18.5

Download GitHub Enterprise Server 3.18.5

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.

3.18.5: 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.

3.18.5: 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.

3.18.5: 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.
- 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.
- Organization creation would fail with a 500 error when the system attempted to verify CAPTCHA responses even when no CAPTCHA challenge would be presented to the user.

- 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.

3.18.5: 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.

3.18.5: 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.

- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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.
- Unexpected elements may appear in the UI on the repo overview page for locked repositories.
- 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.
- Custom NTP settings are removed during the upgrade process.
- 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 3.18.4

Download GitHub Enterprise Server 3.18.4

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.

3.18.4: 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.

3.18.4: 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.
- Deleted organizations were still accessible from users' organization dashboards for a period of time.
- 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.
- 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.
- On high-availability clusters with Elasticsearch Cross Cluster Replication (CCR) enabled, replication failed if the datacenter and consul-datacenter values didn’t match.

3.18.4: 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.

3.18.4: 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.
- 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.
- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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.

- Unexpected elements may appear in the UI on the repository page for locked repositories.
- 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.
- Custom NTP settings are removed during the upgrade process.

Enterprise Server 3.18.3

Download GitHub Enterprise Server 3.18.3

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.

3.18.3: Security fixes

- **HIGH:** An attacker could inject HTML elements with IDs that collided with server-initialized data islands due to insufficient sanitization. When a privileged user viewed crafted content in certain Project views, these injected elements could overwrite critical application state objects, resulting in unintended server-side POST requests or other unauthorized backend interactions. GitHub has requested CVE ID CVE-2025-14046 for this vulnerability, which was reported via the GitHub Bug Bounty program.

3.18.3: Bug fixes

- Due to a regression in a recent patch release, Dependabot did not respond to some commands on pull requests, such as rebases, because webhook deliveries to loopback addresses were

blocked. Webhook deliveries to the Dependabot endpoint now succeed, although deliveries to other endpoints on loopback addresses are still blocked.

3.18.3: 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."
- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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`.

- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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.
- Unexpected elements may appear in the UI on the repo overview page for locked repositories.
- 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.
- Custom NTP settings are removed during the upgrade process.

Enterprise Server 3.18.2

Download GitHub Enterprise Server 3.18.2

December 02, 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.

3.18.2: Security fixes

- **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. [Updated: 2026-01-06]
- **HIGH:** A privilege escalation vulnerability was identified in GitHub Enterprise Server that allowed an authenticated Enterprise admin to gain root SSH access to the appliance by exploiting a symlink escape in pre-receive hook environments. By crafting a malicious repository and environment, an attacker could replace system binaries during hook cleanup and execute a payload that adds their own SSH key to the root user's authorized keys—thereby granting themselves root SSH access to the server. To exploit this vulnerability, the attacker needed to have enterprise admin privileges. This vulnerability has been assigned CVE-2025-11578 and was reported through the GitHub Bug Bounty program.
- Packages have been updated to the latest security versions.

3.18.2: Bug fixes

- Administrators may have experienced delays with configuration runs after a reboot if `ghe-reconfigure.service` was still activating, impacting run performance and stability.
- After an administrator clicked the "Clear Dependencies" button in the dependency graph section of the site admin dashboard, dependencies were not deleted as expected.
- On instances with a "No Proxy" setting configured for GitHub Actions with MinIO or AWS remote blob providers, administrators sometimes experienced failures reading or writing Actions logs, artifacts, or caches because some traffic was incorrectly routed through the instances proxy.
- New Microsoft Teams integrations failed to set up because the required `tenant_id` field was missing from the configuration, following Microsoft's deprecation of multi-tenant bot creation.

- Administrators running the `ghe-repl-decommission` script received an error.
- Site administrators using the Management Console would see overly verbose error messages on the maintenance page. These error messages were not cleared when a new request was made, and no message was displayed when maintenance mode changes were saved successfully.
- Some instances were blocked from upgrading to GitHub Enterprise Server 3.18.1 because of an issue with concurrent execution of certain database migrations.
- Teams granted all-repository organization roles could not be requested as reviewers in pull requests.
- An "Invite member" button intended only for GitHub.com was displayed on the enterprise "People" tab.
- Administrators who had upgraded to the previous patch release may have observed a significant increase in executions of the `SecurityOverviewAnalytics::UpdateFeatureStatusSummaryJob`, causing background job queue saturation, service delays, reduced stability, and lower performance for environments using security overview analytics.
- Link previews did not appear in Slack conversations when messages were delivered through socket mode, affecting the visibility of linked GitHub content.
- Users who accessed GitHub Enterprise Server with Chromium-based browsers experienced slow logout performance.
- Audit log searches could temporarily miss recent events or show incomplete results right after new index creation at the start of a month. Administrators now experience reduced lag between the creation of monthly audit log search indexes and their availability for searches and write operations.
- When new Elasticsearch indexes were created, index routing memos could go to a read-only MySQL replica and fail, causing delays in audit log indexing after monthly rollovers. The memos are now written to the primary database rather than a read-only replica.

3.18.2: Changes

- Site administrators can monitor NUMA (Non-Uniform Memory Access) performance metrics using newly enabled collectors in `node_exporter` and dedicated Grafana dashboard panels.

- Site administrators configuring server monitoring using OpenTelemetry metrics will see improvements and adjustments to the metrics and monitoring options. The changes focus on better security, enhanced observability, and improved configuration consistency in multi-node deployments. In addition, it is easier to export Grafana dashboards and import them into your own environments.
- Administrators can add security key-backed (SK) SSH certificate authorities.
- Administrators and users experience faster and more efficient searching of GitHub Actions workflow runs, with lower compute and networking resource usage. Searches for workflow runs within a repository are now always scoped to an associated repository.
- `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]

3.18.2: 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. See "Troubleshooting access to the Management Console."
- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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 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.

- Unexpected elements may appear in the UI on the repository overview page for locked repositories.
- 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.
- Custom Network Time Protocol (NTP) settings are removed during the upgrade process.

Enterprise Server 3.18.1

Download GitHub Enterprise Server 3.18.1

November 10, 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.

3.18.1: Security fixes

- **CRITICAL:** Redis has been upgraded to version 6.2.20 to address CVE-2025-49844 (also known as RediShell). Administrators should apply this update promptly to mitigate potential security risks.
- **HIGH:** An attacker could execute arbitrary code in the context of other users' browsers by supplying a malicious `label:` value that was injected into the DOM without proper sanitization. This could be triggered when a user visits a crafted Issues search URL, enabling session hijacking, account takeover, and recovery code exfiltration. GitHub has requested CVE ID CVE-2025-11892 for this vulnerability, which was reported via the GitHub Bug Bounty program.
- **LOW:** When a user updated a classic Personal Access Token (PAT) to remove all scopes instead of revoking the PAT, the change was silently ignored and the PAT continued to grant its

previously held permissions. To mitigate this issue, GitHub updated the token management logic to correctly clear scopes when no scope is provided.

- Packages have been updated to the latest security versions.

3.18.1: Bug fixes

- Initializing a cluster configuration for the first time could fail with `Error: Validation preflight-check`.
- Administrators running the `ghe-repl-start-all` command may have encountered replicas remaining in an enabled state after a failed operation, causing subsequent configuration updates to execute on unintended nodes. Replicas now revert to a disabled state if the command fails.
- Setting up MySQL replication on secondary replica nodes was inefficient and consumed unnecessary root disk space.
- Administrators and users who accessed dashboard panels experienced issues with the CPU panel, navigation between dashboards, and a missing home dashboard.
- Administrators could not generate support bundles on stateless high availability nodes because the `ghe-support-bundle` command failed when attempting to query Elasticsearch on nodes without the `elasticsearch-server` role.
- After an upgrade, administrators found that Elasticsearch allocation remained set to "none," causing subsequent upgrades to fail. Enterprise upgrades now correctly set allocation to "all" after configuration is applied, preventing upgrade blocks.
- When running the `system-requirements` check as part of the `ghe-cluster-config-check` command prior to the initialization of a new cluster, the check request would fail because it exceeded the overall request timeout.
- Creating an organization would fail with a 500 or validation error if a maximum lifetime policy for personal access tokens was set to less than 366 days in the enterprise settings.
- Announcements scheduled using the `expires_at` timestamp in ISO 8601 format were not parsing the specified time correctly, resulting in the time component always being ignored.
- On pull requests in organization-owned repositories, users could not request reviews from teams with the "All-repository read" organization role.
- Administrators experienced 500 errors when attempting to run Dependabot from the Security tab, to scan repositories for dependency vulnerabilities.

- On instances with thousands of organizations and roles, opening the security overview page for an organization or any other organization-level pages accessible via the Security tab triggered inefficient database queries that could degrade performance for other users.
- Administrators who had upgraded to the previous patch release may have observed a significant increase in executions of the `SecurityOverviewAnalytics::UpdateFeatureStatusSummaryJob`, causing background job queue saturation, service delays, reduced stability, and lower performance for environments using security overview analytics.
- On instances where GitHub Actions workflows require approval to run on pull requests from forked repositories, workflows remained queued indefinitely after users clicked "Approve and run".
- The GitHub system user was not always properly set on startup, occasionally surfacing in authentication errors or failed secret scanning jobs in logs.

3.18.1: Changes

- Elasticsearch deprecation warnings, which are logged to index files in new versions of Elasticsearch, have been disabled. These warnings provided no value to administrators, and in some cases could block upgrades of instances in high-availability or cluster configurations.
- Logging of configuration runs is improved with streamlined logging for different configuration phases. Phase-specific logs are written to both the main log file (`ghe-config.log`) and the console for better visibility.
- Users can no longer view Git objects, such as commits and tags, that exceed the maximum size limit of 10 MB.

3.18.1: 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.

- On an instance with the HTTP `X-Forwarded-For` header configured for use behind a load balancer, all client IP addresses in the instance's audit log erroneously appear as 127.0.0.1.
- 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, shut down 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. You can also trigger the reindexing by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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 more 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 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.
- Unexpected elements may appear in the UI on the repository overview page for locked repositories.
- 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.

Enterprise Server 3.18.0

Download GitHub Enterprise Server 3.18.0

October 14, 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.

3.18.0: Features

- **Instance services**
  - Operators use OpenTelemetry metrics to monitor the appliance. This feature is currently in public preview and should only be used in preproduction environments. You can also export Prometheus metrics to third-party observability systems. See About OpenTelemetry metrics.
  - Site administrators can enable a larger limit of **50,000 items** on GitHub Projects, rather than the default limit of 1,200.

    a. After upgrading your instance to 3.18, wait for the `memex-project-items` index to be migrated and the `memex-project-items` index repair job to complete.

    If you run the following commands before the index repair is completed, project data will appear to be missing from any partially repaired projects. This problem will resolve itself once the repair completes.

    b. Access the administrative shell and run the following command.

    ```shell
    ghe-config app.github.projects-increased-limits-enabled true
    ```

    c. Apply the configuration with `ghe-config-apply`. This will restart services and may cause a brief interruption for users.

- **APIs**
  - For push webhook events, the `html_url` and `url` fields return different values. The `html_url` field returns the repository URL (e.g., https://github.com/), while the `url` field provides the API URL (e.g., https://api.github.com/repos/). Previously, both fields returned the same link, unlike other webhook events like `pull_request`.

- **Policies**
  - Enterprise administrators can create enterprise-level rulesets, and set pull request merge methods using rules. These features provide greater control and consistency across repositories within the enterprise.

  - Developers can request exceptions to push rules through a delegated bypass process, ensuring each request is reviewed, audited, and approved for transparency. Email notifications keep developers updated on approval status.

- **Secret Protection (part of Advanced Security)**
  - Secret scanning supports additional default patterns for secret protection, expanding coverage for more token formats and credential types. This enhancement helps administrators and users better prevent accidental exposure of sensitive information.
  - Organization and security admins can run a free secret risk assessment to scan their organization for aggregate insights on public leaks, private exposures, and token types. The assessment provides a dashboard with actionable data to help organizations understand and address secret leak risks. See Find secrets exposed in your organization with the secret risk assessment on the GitHub Blog.
  - Administrators and developers can use the Secret Scanning Alerts API to hide the values of detected secret literals within secret scanning alerts. This helps prevent accidental exposure of sensitive information when viewing or processing alert data. See Secret scanning alerts API now supports hiding secret literals on the GitHub Blog.

- **Code Security (part of Advanced Security)**
  - Administrators and security teams can view improved metrics for CodeQL pull request alerts on the security overview dashboard. These updates provide more precise insight into alert identification and resolution to help organizations strengthen their security posture. Dashboard data is scoped to pull requests against the default branch; future updates will expand coverage to other branches. Historical dashboard data is not retroactively updated. See Viewing metrics for pull request alerts.
  - Organization administrators with Code Security can grant Dependabot access to repositories at scale from the organization level. Options allow you to enable Dependabot access permanently for all current and future internal repositories. New API endpoints support programmatic management of repository access permissions. See It's now easier to grant Dependabot access to repositories from the organization level on the GitHub Blog.
  - Users can track the progress of code scanning alert resolution with the new "Development" section. This section highlights when an alert is introduced, addressed, or reintroduced, helping users understand the lifecycle of each alert and supporting better code security management. See Track progress on code scanning alerts with the new development section on the GitHub Blog.

- This release comes installed with version 2.21.4 of the CodeQL CLI, used in the CodeQL action for code scanning. Significant updates since the default version installed on GitHub Enterprise Server 3.17 include:

  - General availability of support for analyzing GitHub Actions workflows. See GitHub Actions workflow security analysis with CodeQL is now generally available on the GitHub Blog.
  - The GitHub Actions `actions/missing-workflow-permissions` query provides better alert messages and fix suggestions.
  - Improved Java analysis. The `java/spring-boot-exposed-actuators` query is included in the default code scanning query stack to help identify publicly exposed Spring Boot actuators.
  - Support for Swift 6.1.1, ensuring you can analyze projects built with this version.
  - The Python extractor analyzes files in hidden directories by default.
  - C/C++ improvements, including added support for more Windows APIs including file read functions, command-line and environment variable APIs, and flow models for SQLite and OpenSSL libraries.
  - Javascript and TypeScript enhancements, including:
    - Support for TypeScript 5.8, enabling analysis of the latest Typescript language features.
    - Expanded JavaScript analysis to cover Apollo Server, React Relay, SAP packages, and TanStack libraries for broader security scanning.
    - Enhanced path injection detection for several additional methods.
    - A fix for an issue where `tsconfig.json` files containing array literals and trailing commas were not correctly extracted.
    - Improved modeling of the `fastify` framework and the `shelljs` and `async-shelljs` libraries, which could result in improved analysis results for apps using them.
    - New detections of sources and sinks in Next.js and DOM element references, improving the detection of XSS issues.
  - Ruby enhancements, including:
    - Improved the `rb/useless-assignment-to-local` query, so you'll see fewer false positives and will get helpful documentation for alerts.
    - The `rb/uninitialized-local-variable` query now only generates an alert when a variable is used as a method call receiver. This should reduce noise. In addition, new help content is available for this query.
    - Calls to `super` without explicit arguments now have their implicit arguments generated, resulting in more accurate analysis.

- Support for analyzing Kotlin applications up to version 2.2.0x, and dropped support for the 1.5.x series of Kotlin. The minimum supported Kotlin version is now 1.6.0.
- C# enhancements, including:
  - Enhancements to the `cs/missed-readonly-modifier` query, reducing false positives.
  - The `cs/gethashcode-is-not-defined` and `cs/uncontrolled-format-string` queries detect more potential issues, helping administrators identify risks more effectively.
  - The false positive rate for the query `cs/web/missing-function-level-access-control` has been reduced by improving the detection of authorization checks.
  - The true positive rate for the `cs/invalid-string-formatting` query has been increased by accounting for methods and additional overloads of existing format-like methods.
- Removed hardcoded credential queries from all query suites across multiple languages (C#, Go, Java/Kotlin, JavaScript/TypeScript, Python, Ruby, and Swift) to reduce noise and duplication of alerts from GitHub Secret Protection. See CodeQL no longer detects hardcoded secrets on the GitHub blog.

- **Dependabot**
  - Users can schedule custom update frequencies for Dependabot version updates by using cron expressions in `schedule.interval` in the Dependabot configuration file. This enhances the predefined intervals of daily, weekly, and monthly to provide more flexible scheduling options that meet specific needs.
  - Users can use Dependabot version updates to automatically keep Helm dependencies up to date. For projects that use Helm as a package manager, Dependabot can ensure dependencies stay current with the latest releases. See Dependabot version updates now support Helm on the GitHub Blog.
  - Users can use an improved checkbox UI to grant point-in-time access across their repository portfolio. New API endpoints support programmatic management of repository access permissions. See It's now easier to grant Dependabot access to repositories from the organization level on the GitHub Blog.
  - Users can use the `has:patch` filter with the Dependabot REST API to quickly identify dependencies that have available patches. This streamlines the process of addressing vulnerabilities and staying up-to-date with dependency maintenance. See Dependabot API now contains has:patch in general availability on the GitHub Blog.

  - Dependabot is generally available for execution on self-hosted GitHub Actions runners managed within Kubernetes clusters using Actions Runner Controller (ARC), providing auto-scaling, workload isolation, and improved resource management. Additionally, Dependabot support for running within a virtual network (vNet) in self-hosted runner environments is now generally available, enabling secure, isolated dependency updates with network-level governance. See Dependabot support for virtual network (vNet) and Actions Runner Controller (ARC) is generally available.

- **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 2.324.0. 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.

  - Repository users can pin specific workflows to the top of the workflows list on the Actions workflow page, making frequently used workflows easier to access and manage across the repository.

  - Users can use CodeQL code scanning to detect security vulnerabilities in GitHub Actions workflows. CodeQL automatically analyzes workflows to detect common vulnerabilities such as missing required permissions or inputs without proper validation. See GitHub Actions workflow security analysis with CodeQL is now generally available on the GitHub Blog.

  - Administrators using the Actions runner controller can configure metrics collection to address performance issues caused by high cardinality. This change allows customers to tailor metric granularity to better meet their reporting and observability needs.

  - Administrators can configure custom annotations and resource settings for the Actions Runner Controller (ARC), enabling integration with deployment tools like ArgoCD and Helm. This flexibility allows alignment with preferred DevOps workflows and supports advanced deployment strategies.

- **Community experience**

  - Users who view an organization's activity feed experience improved performance as the feed runs on a newer infrastructure. Push events are grouped into a single card, showing recent activity in chronological order, instead of individual lines for each event.

- **Organizations**
  - Users can use regex to ensure custom properties match data structures like email addresses or patterns relevant to your organization.
  - Organization members experience faster load times and improved responsiveness in the organizational feed. These performance improvements help users more efficiently review updates and activities within their organizations.

- **Repositories**
  - Enterprise owners can enrich repositories with consistent metadata across the entire enterprise using enterprise custom properties. Existing organization-level custom properties can also be promoted to the enterprise level.

- **Issues**
  - Repository administrators can control whether merged pull requests automatically close linked issues with a new repository setting. This change addresses feedback from teams who prefer to keep issues open for additional QA or process steps after merging a pull request.
  - Users can perform advanced issue searches using the AND and OR keywords and nested searches using both the REST and GraphQL APIs. This enhancement enables more precise queries to find exactly the set of issues needed for tracking and reporting.
  - Users can manage issue types in GitHub Issues and Projects via the REST API, enabling automation of issue type creation, updates, deletions, and assignments to issues.
  - Users can close issues as duplicates of others, improving issue management clarity. In addition, the REST API supports viewing, adding, removing, and reprioritizing sub-issues, enabling automation of issue hierarchies. See Close issue as a duplicate, REST API for sub-issues, and more on the GitHub blog.
  - Organization administrators can standardize issue management by creating issue types across repositories. See Managing issue types in an organization.
  - Users can access an improved Issues dashboard page at HOSTNAME.com/issues featuring saved views to create and save custom queries across repositories and organizations, and a new "Recent activity" view to find relevant work.

  - The GitHub Issues interface is faster and easier to use, with a filter bar featuring autocomplete and syntax highlighting, a "create more" option for quick issue creation, alphabetical sorting of issue forms and templates, a copy link button for sharing issues, and improved loading for long issues.
  - Users can find issues more efficiently using advanced search with AND, OR, and parentheses for nested searches. See Filtering and searching issues and pull requests.
  - Users can organize large tasks by breaking issues into sub-issues. Sub-issues create a nested structure, making it easier to track progress and manage work within a project.

- **Pull requests**
  - Repository and organization administrators can use the new merge method rule for rulesets to control which merge methods—merge commit, squash, or rebase—are allowed on targeted branches when merging pull requests via the UI or APIs. This ensures consistency and simplifies workflows across branches.

3.18.0: 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. 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 an appliance with version 3.13 or greater, the Elasticsearch indices must be reindexed before the data will display. This happens via a nightly scheduled job. It can also be forced by running `/usr/local/share/enterprise/ghe-es-search-repair`.
- An organization-level code scanning configuration page is displayed on instances that do not use GitHub Advanced Security or code scanning.
- 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 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.

- The entry for Private Registries in the organization settings menu is not visible unless Dependabot is enabled.
- Upgrading to this version from GHES 3.16.10 and higher or 3.17.6 and higher will cause the upgrade to fail due to this version containing an older version of MySQL. To avoid this issue please upgrade to GHES 3.18.1 or higher.

  [Updated: 2025-11-24]

3.18.0: Closing down

- In GitHub Enterprise Server 3.20, GitHub will retire the security manager API in favor of the organization roles API. See Notice of breaking changes: Security manager REST API will be retired and replaced with the organization roles REST API on the GitHub blog
- Microsoft Exchange Online is retiring SMTP basic authentication in September 2025. If your GitHub Enterprise Server instance uses this method to send email, delivery may fail after the retirement date. Microsoft recommends switching to a supported alternative. As another option, you may consider using an SMTP OAuth proxy such as email-oauth2-proxy, though this is not officially supported. For details and configuration guidance, see the Microsoft announcement and the proxy's documentation.

3.18.0: Retired

- The /explore functionality, including the Activity and Trending pages, is no longer available. Users can no longer access these pages to discover trending repositories or recent activity.
- The ability to bulk convert issues to discussions using labels is closing down. Users can continue to convert individual issues to discussions manually using the "Convert to discussion" option. See Moderating discussions.
- GitHub Actions users should update workflows that modify check run statuses via the REST API. GitHub will restrict the ability to change check run status for runs created by Actions to prevent inconsistent state changes. Review your workflows to ensure compatibility with this update.
- Deployment permissions in GitHub Actions workflows have changed. Workflows using the deployment protection rule or required reviewers must now explicitly grant write or admin permissions to the GITHUB_TOKEN for successful deployment. Update workflows to avoid disruptions.

- The announcement banner GraphQL fields have been replaced. Users can now manage instance-wide announcements through updated GraphQL fields, improving consistency and control for administrators. The existing individual fields following the `announcementX` pattern have been removed, and the new fields are within the `announcementBanner` object.

- Automatic watching of repositories and teams is closing down. Users will no longer be auto-subscribed when joining organizations or teams, reducing notification noise and confusion. Existing auto-watching subscriptions remain unchanged; users stay subscribed to previously watched repositories or teams. See Configuring notifications.

- The issue template using a single `ISSUE_TEMPLATE.md` file is no longer supported. Users should migrate to using multiple issue templates stored in the `.github/ISSUE_TEMPLATE/` directory to provide a better issue reporting experience. See About issue and pull request templates. [Updated: 2025-12-04]

Legal

© 2026 GitHub, Inc. Terms Privacy Status Pricing Expert services Blog