COMPARISON

Helix GitSwarm vs. GitHub Enterprise

Why Switch to GitSwarm Now?

Git has been a remarkable success in many ways, but users often have to augment its command line tools to make their workflows more productive. With no formal standards available, organizations have adapted a variety of tools, each with their own specialties and limitations.

GitHub is the leading solution for hosted Git repositories, especially for open source, but as an enterprise solution it has severe limitations. Delivered as a closed virtual appliance, GitHub Enterprise has few configuration options for administrators beyond the network setup. Customers have to rely on GitHub to deliver security patches and to provide performance improvements and bug fixes for any components that are part of the appliance.

GitHub itself does not address the fundamental issues with Git such as security, scalability for large projects and large binary files.

Any user of GitHub Enterprise, or anyone considering it, should evaluate their workflows and toolset needs – now and for the next several years. After all, products under development may be the most valuable intellectual property organization possesses, and its designers and developers are the most valuable personnel. Teams would be advised to not "make do" or compromise with such crucial assets.

Contents

Summary

GitHub Enterprise provides a subset of the tools needed to address some of the limitations of Git; however, it does not address important needs for modern, high-performance, high-quality development, including:

- No file-level access control
- File size limited to 100MB:
  - Large file support only in early access preview
  - Requires separate client installation for all users
- Limited deployment options:
  - Available only as a virtual machine image
  - No native backup, clustering or replication technology

Perforce Helix offers a compelling solution that is a more complete Git management toolset, including support for the largest files, security down to the individual file level, immutable audit trails, hugely scalable repositories and much more. Legendary support is available for Git users as well as those using Helix clients. Our Professional Services teams have decades of experience deploying version management solutions in the most demanding environments, and are available to ensure your switch to Perforce Helix is as smooth as possible.

Comparison Matrix

Basic Capabilities	GitHub Enterprise	GitSwarm + Helix
Repository	Git repos only with the standard limitations on file and repo size, limited security options, and rewritable history.	Git or Helix for local repositories, Helix for shared repositories with unlimited scalability and rich security.
Workflow	Clone, fork, and collaborate via merge requests.	Clone, fork, and collaborate via merge requests alongside file locking for digital assets that cannot be merged.
Visualizations	Contributor, commits and code frequency graphs.	Progress graphs, branching diagrams, unified activity feed, revision graph, and timeline view.
Deployment Options	Deploy on bare metal, in a virtual machine or in the cloud.	Helix supports classic central workflow as well as distributed workflows through its native DVCS capabilities. It is easy to sync data from Git repositories into Helix.

Advanced Capabilities	GitHub Enterprise	GitSwarm + Helix
Narrow Cloning	N/A	Clone tiny, Git-sized slices of huge, Helix-sized assets.
Synchronization	N/A	Automated mirroring from GitSwarm to Helix, whose federated architecture synchronizes around the world.
Scalability	Single instance model	Helix is highly scalable with clustering and federated server options.
Reliability	Single point of failure	Helix clusters have a high-availability option.
Traceability	Destructible history	Mutable local history. Immutable shared history. Full developer freedom with full auditing and regulatory compliance in Helix.
Deployment Options	Closed virtual machine only with no root access	Runs virtualized and on metal, full access to the server.
Threat Detection	N/A	Behavioral analytics to detect if intellectual property is at risk (add-on option).

Repository

GitSwarm

Depending on your project needs, you can use Git or Helix for your master repository. With a Helix master, you have unlimited scalability, file-level access control, immutable history and more. Automatic mirroring between Git and Helix masters ensures all contributors can access all the assets using their preferred workflows and tools.

GitHub Enterprise

Can use Git only as a repository format, which imposes limits on the practical sizes of repos, file sizes and limited security options.

Why Should I Care?

Git has practical limitations in size and scope: Repositories that exceed 1GB tend to take a long time to clone and determine the status of local files, which slows developers. Repositories grow because of the depth of history, the number of files in the project and the size of the files. A Git management system that stores only Git repositories cannot change these limitations.

Git has no built-in security, authorization and protection against destructive history rewrite. Instead, this must be provided by the Git management system.

Workflow

GitSwarm

Clone, fork and merge request workflows. Optional file locking for content that can't easily be merged.

GitHub Enterprise

Clone, fork and merge request workflows.

Why Should I Care?

Modern development often includes digital assets in form of binary files such as graphics, media or output from a continuous integration system as well as plain text such as source code. These binary files often cannot be merged and therefore require a different workflow than DVCS; a centralized approach with file locking is a better choice for these files. The result is a hybrid approach suited to contributors' needs.

A solution that can offer a hybrid approach without sacrificing a unified single repository with global visibility and security is best suited to the task.

Visualizations

GitSwarm

Graphical overviews of progress including progress graphs, branching diagrams, unified activity feed, revision graph, and timeline view.

GitHub Enterprise

Contributor, commits and code frequency graphs.

Why Should I Care?

Visual tools help project managers and team leads to determine the state of the project, and developers to track changes to individual contributions. This can greatly speed up bug fixes and resource management, and will help to determine the readiness of a project towards the next milestone.

Deployment Options

GitSwarm

GitSwarm is provided as a simple package. This allows administrators to deploy GitSwarm on the Linux distribution of their choice.

GitHub Enterprise

Available only as a closed virtual appliance (based on a customized Linux Ubuntu 12.04 LTS) with no root access.

Why Should I Care?

Enterprises usually have stringent rules about which Linux distributions are used. They require full access to ensure maintainability and security instead of relying on the provider to supply security patches and essential updates.

Narrow Cloning

GitSwarm

Clone tiny, Git-sized slices of huge, Helix-sized assets. Let contributors concentrate on just the files they need.

GitHub Enterprise

Not available.

Why Should I Care?

Large Git repositories (especially so-called "monorepos") are very successful for allowing easy sharing and code reuse. However, they become unmanageable with the limits of Git repositories and hinder productivity through large cloning and status-update wait times. In the majority of cases users do not actually need access to the whole repository – only a subset to accomplish their work. The ability to clone narrowly is not available in native Git, which makes it impossible to work on a small subset of the total repository without losing the ability to see global changes and assess their impact on the whole project.

Synchronization

GitSwarm

Automated mirroring from GitSwarm to Helix, whose federated architecture synchronizes around the world.

GitHub Enterprise

Not available.

Why Should I Care?

Federating Helix servers allows for intelligent distribution of content. For example, a remote team that mostly performs "pulls" from the master to perform QA tasks would benefit from a different configuration to a remote development team. Combining the power of Perforce Helix with the flexibility and ubiquity of Git gives developers the best of both worlds.

Scalability

GitSwarm

Helix is highly scalable with clustering and federated server options built into the server platform.

GitHub Enterprise

Single instance model.

Why Should I Care?

In addition to needing the flexibility to manage how assets are distributed across a global WAN, scalability within a data center is used to ensure high performance as projects and teams become larger and more complex. If this is provided only as an add-on option cost and complexity increase.

Reliability

GitSwarm

Helix clusters have a free high-availability option enabling automated failover if a server should suffer an outage.

GitHub Enterprise

Single instance model.

Why Should I Care?

The central repository contains the intellectual property of the entire enterprise. If that repository cannot be accessed because of a failure, developers cannot push and share their changes anymore, and Continuous Integration engines cease to work, delaying projects.

Traceability

GitSwarm

Immutable history, full auditing and regulatory compliance are held within the repository.

GitHub Enterprise

Destructible history.

Why Should I Care?

For most enterprises, auditing and regulatory burdens require that the file history be immutable and all access be logged. Git is inherently open and fragile, prone to overwriting through pushed changes. Some Git management solutions offer optional protection for production branches, but are very limited with respect to access control, logging and auditing.

Threat Detection

GitSwarm

Behavioral analytics learn normal behavior and detect if intellectual property is at risk from insiders that have access to the repository but may be causing risks deliberately or accidentally. (Add-on option)

GitHub Enterprise

Not available.

Why Should I Care?

IP theft is a growing problem in the knowledge-based industry. The greatest danger often comes not from the outside, where access is restricted by firewalls and other security measures, but from the inside. A careless or even malicious user who has access to sensitive content can cause tremendous harm to the enterprise, which is often not even aware that intellectual property has left the premises or has no way of proving it.

Conclusion

Helix GitSwarm provides a compelling alternative to GitHub Enterprise. Even if projects appear to be small or simple today, over time they will no doubt grow and become more complicated. At that point, it's hard to switch tools, and thus tools with limitations force artificial decisions about repository organization and/or workflow. It is recommended you take a close look at the tools you have in place and ensure they are fit for your purposes today and for the next five years. Perforce Helix provides just such an environment.

Perforce Helix, including GitSwarm, is available free of charge for personal or professional use by small teams, for open source projects and for educational organizations. For more information and downloads, visit perforce.com.

North America
Perforce Software Inc.
2320 Blanding Ave
Alameda, CA 94501
USA
Phone: +1 510.864.7400
info@perforce.com

Europe
Perforce Software UK Ltd.
West Forest Gate
Wellington Road
Wokingham
Berkshire RG40 2AT
UK
Phone: +44 (0) 1189 771020
uk@perforce.com

Australia
Perforce Software Pty. Ltd.
Suite 3, Level 10
221 Miller Street
North Sydney
NSW 2060
AUSTRALIA
Phone: +61 (0)2 8912-4600
au@perforce.com

PERFORCE

perforce.com