


AI/ML in the Harness Platform


The Industry's First Software Delivery Platform to use AI/ML to Drive DevOps Efficiency and Improve Developer Experience

At Harness, we're using AI/ML and advanced automation to eliminate the worst parts of developers' jobs, so they can focus on writing code to create new products and features. AI/ML are at the heart of the Harness Platform, helping to verify deployments, identify test optimizations, recommend cloud cost optimizations, roll back to a previous state, facilitate complex deployment patterns, detect cloud cost anomalies, and more.

Here's a look at how AI/ML and intelligent automation are used across the Harness Platform.



Build & Test More Efficiently




Build and test your software faster using Harness **Continuous Integration** (CI) Test Intelligence™, which uses ML to reduce build cycles by up to 90 percent by only running the necessary tests that are relevant to code changes. Once these tests are identified, Harness CI speeds up test cycles even more by splitting and running tests concurrently.

Harness CI uses advanced automation to dramatically simplify caching and reduce pipeline execution time. This proprietary **Cache Intelligence** automatically caches well-known directories for common build tools to avoid overloading compute processing power, significantly accelerating build times (up to 4x faster than other CI tools).

Harness **Security Test Orchestration** (STO) uses ML to analyze the mountains of data collected by security scanners, and provide prioritized vulnerability lists and remediation recommendations. This information is conveyed directly to developers in the build phase, rather than at the end of the SDLC, effectively shifting security testing left to save valuable time from late-stage rework. The result is increased application security and high delivery velocity, with less manual labor and many hours saved.


Deploy Changes with Confidence

In Harness **Continuous Delivery & GitOps**, **Continuous Verification** automatically detects bad deployments by applying ML to data and logs from observability solutions. It even provides the option to **roll back** to the previous version automatically, saving hours of time and bringing the system back to the last known good state. This helps reduce risk and provides a safety net when code is deployed to production.



Harness **Feature Flags** uses automation to release features via pipelines and for progressive delivery, enabling teams to incrementally roll out new features to a predetermined set of users. If there's ever a problem with a new feature release, only a subset of users will be affected, and the feature can simply be shut off to address the issue. Feature Flags also automates governance via pipelines and the Open Policy Agent (OPA). This ensures that whenever a change is made — either UI, code, or GitOps — it always undergoes global change governance processes, so nothing slips through the cracks to meet compliance requirements.

Manage Business Impact



Harness **Service Reliability Management** (SRM) uses AI/ML to help perform change impact analysis, cutting through the noise, so developers can rapidly identify root cause and accelerate time to restore service.


Through integration with SRM, Harness Feature Flags leverages this AI/ML functionality to quickly identify and resolve reliability issues that are caused by a specific feature. Teams are alerted as soon as an issue is detected with the exact feature or change causing it; they can then just turn the feature off, resolving the issue in seconds.

Cloud AutoStopping™ in Harness **Cloud Cost Management** (CCM) uses AI/ML to actively manage cloud resource idle time effectively, reducing cloud spend by up to 80% for non-production workloads. Idle resources are shut down when not in use and dynamically run on spot instances (as per user-defined rules) with no impact to end users.

Harness CCM uses ML to continuously monitor cloud spend, comparing it to historical usage trends to detect cost anomalies for Kubernetes clusters and cloud accounts. When a cost anomaly is detected, resource owners are notified immediately, so they can take action on cost overruns before they spiral out of control.

Optimize Processes — Without Manual Labor

Harness **Software Engineering Insights** gives engineering leaders the insight they need to drive continuous improvement with visibility into the entire development process across DevOps tooling. An automated workflow engine triggers notifications in response to specific events, like CI triaging. By pinpointing blockers throughout the SDLC, engineering leaders can immediately take corrective action that helps drive efficiency and deliver value faster.



Learn More

When budgets are tight and teams need to deliver more with less, engineering organizations need innovative solutions that improve efficiency. Learn more about how Harness' AI/ML and advanced automation can help you deliver more software with less — all while delivering a great developer experience.



Request a personalized demo today