COLVIN RUN DEVSECOPS ENABLEMENT

Powered by Harness & Datadog

Colvin Run Networks, Inc.

Proprietary and Confidential | January 2023

Colvin Run

www.colvinrun.net

harness

DATADOG

DISTRIBUTION STATEMENT B: Distribution authorized to U.S. Government only | Proprietary Information (DFARS - SBIR Data Rights): Ded 19, 2022 | Other requests for this document shall be referred the government technical representative listed in SBIR Contract FA864923P0243, with applicable data rights.

CONFIDENTIAL & PROPRIETARY

Devsecops@colvinrun.net

harness

DATADOG

SOLUTION AREAS

Colvin Run rapidly delivers user-driven solutions that enable organizations to harness their data like never before.

01 Curated Intelligence

02 Workflow Management

03 Data Management

04 Network Management

05 IT Assurance

INDUSTRY-LEADING TECHNOLOGY PARTNERS

MicroStrategy

tableau SOFTWARE

denodo

Allied Telesis

servicenow

IBM

harness

databricks

Microsoft

aws

Digital Bazaar

Colvin Run

DEVSECOPS ENABLEMENT

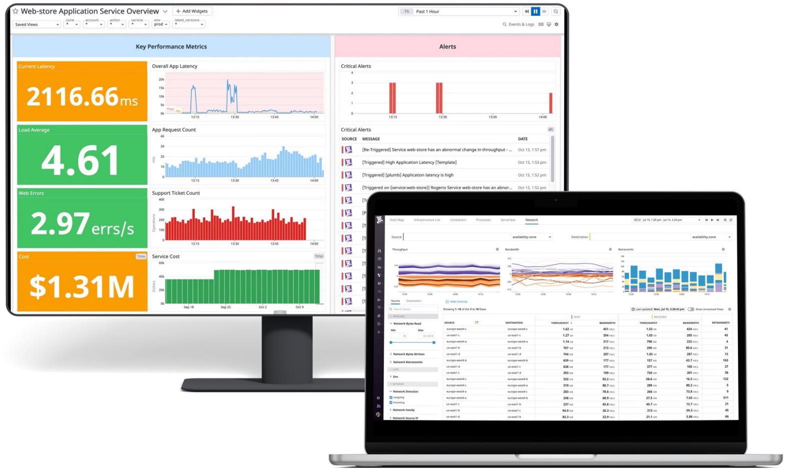

Goals

- Integrate existing Harness and Datadog solutions into an USAF DevSecOps implementation.
- Develop and demonstrate end-user enablement with Colvin Run's solution.

Approach

- Contextualize Harness modules and Datadog's software.
- Develop detailed integration steps to create compatibility between both platforms to USAF requirements.
- Demonstrate the Colvin Run capability with rapid prototypes and path to production.

Benefits

- Up to $1.2M for 12 months in AI/ML development funding from AFWERX SBIR with zero cost matching.
- Shape existing legacy third-party technologies into modernized solution.
- Fast, intuitive, visible analytics with ATO-ready software that is FedRAMP pending.

harness

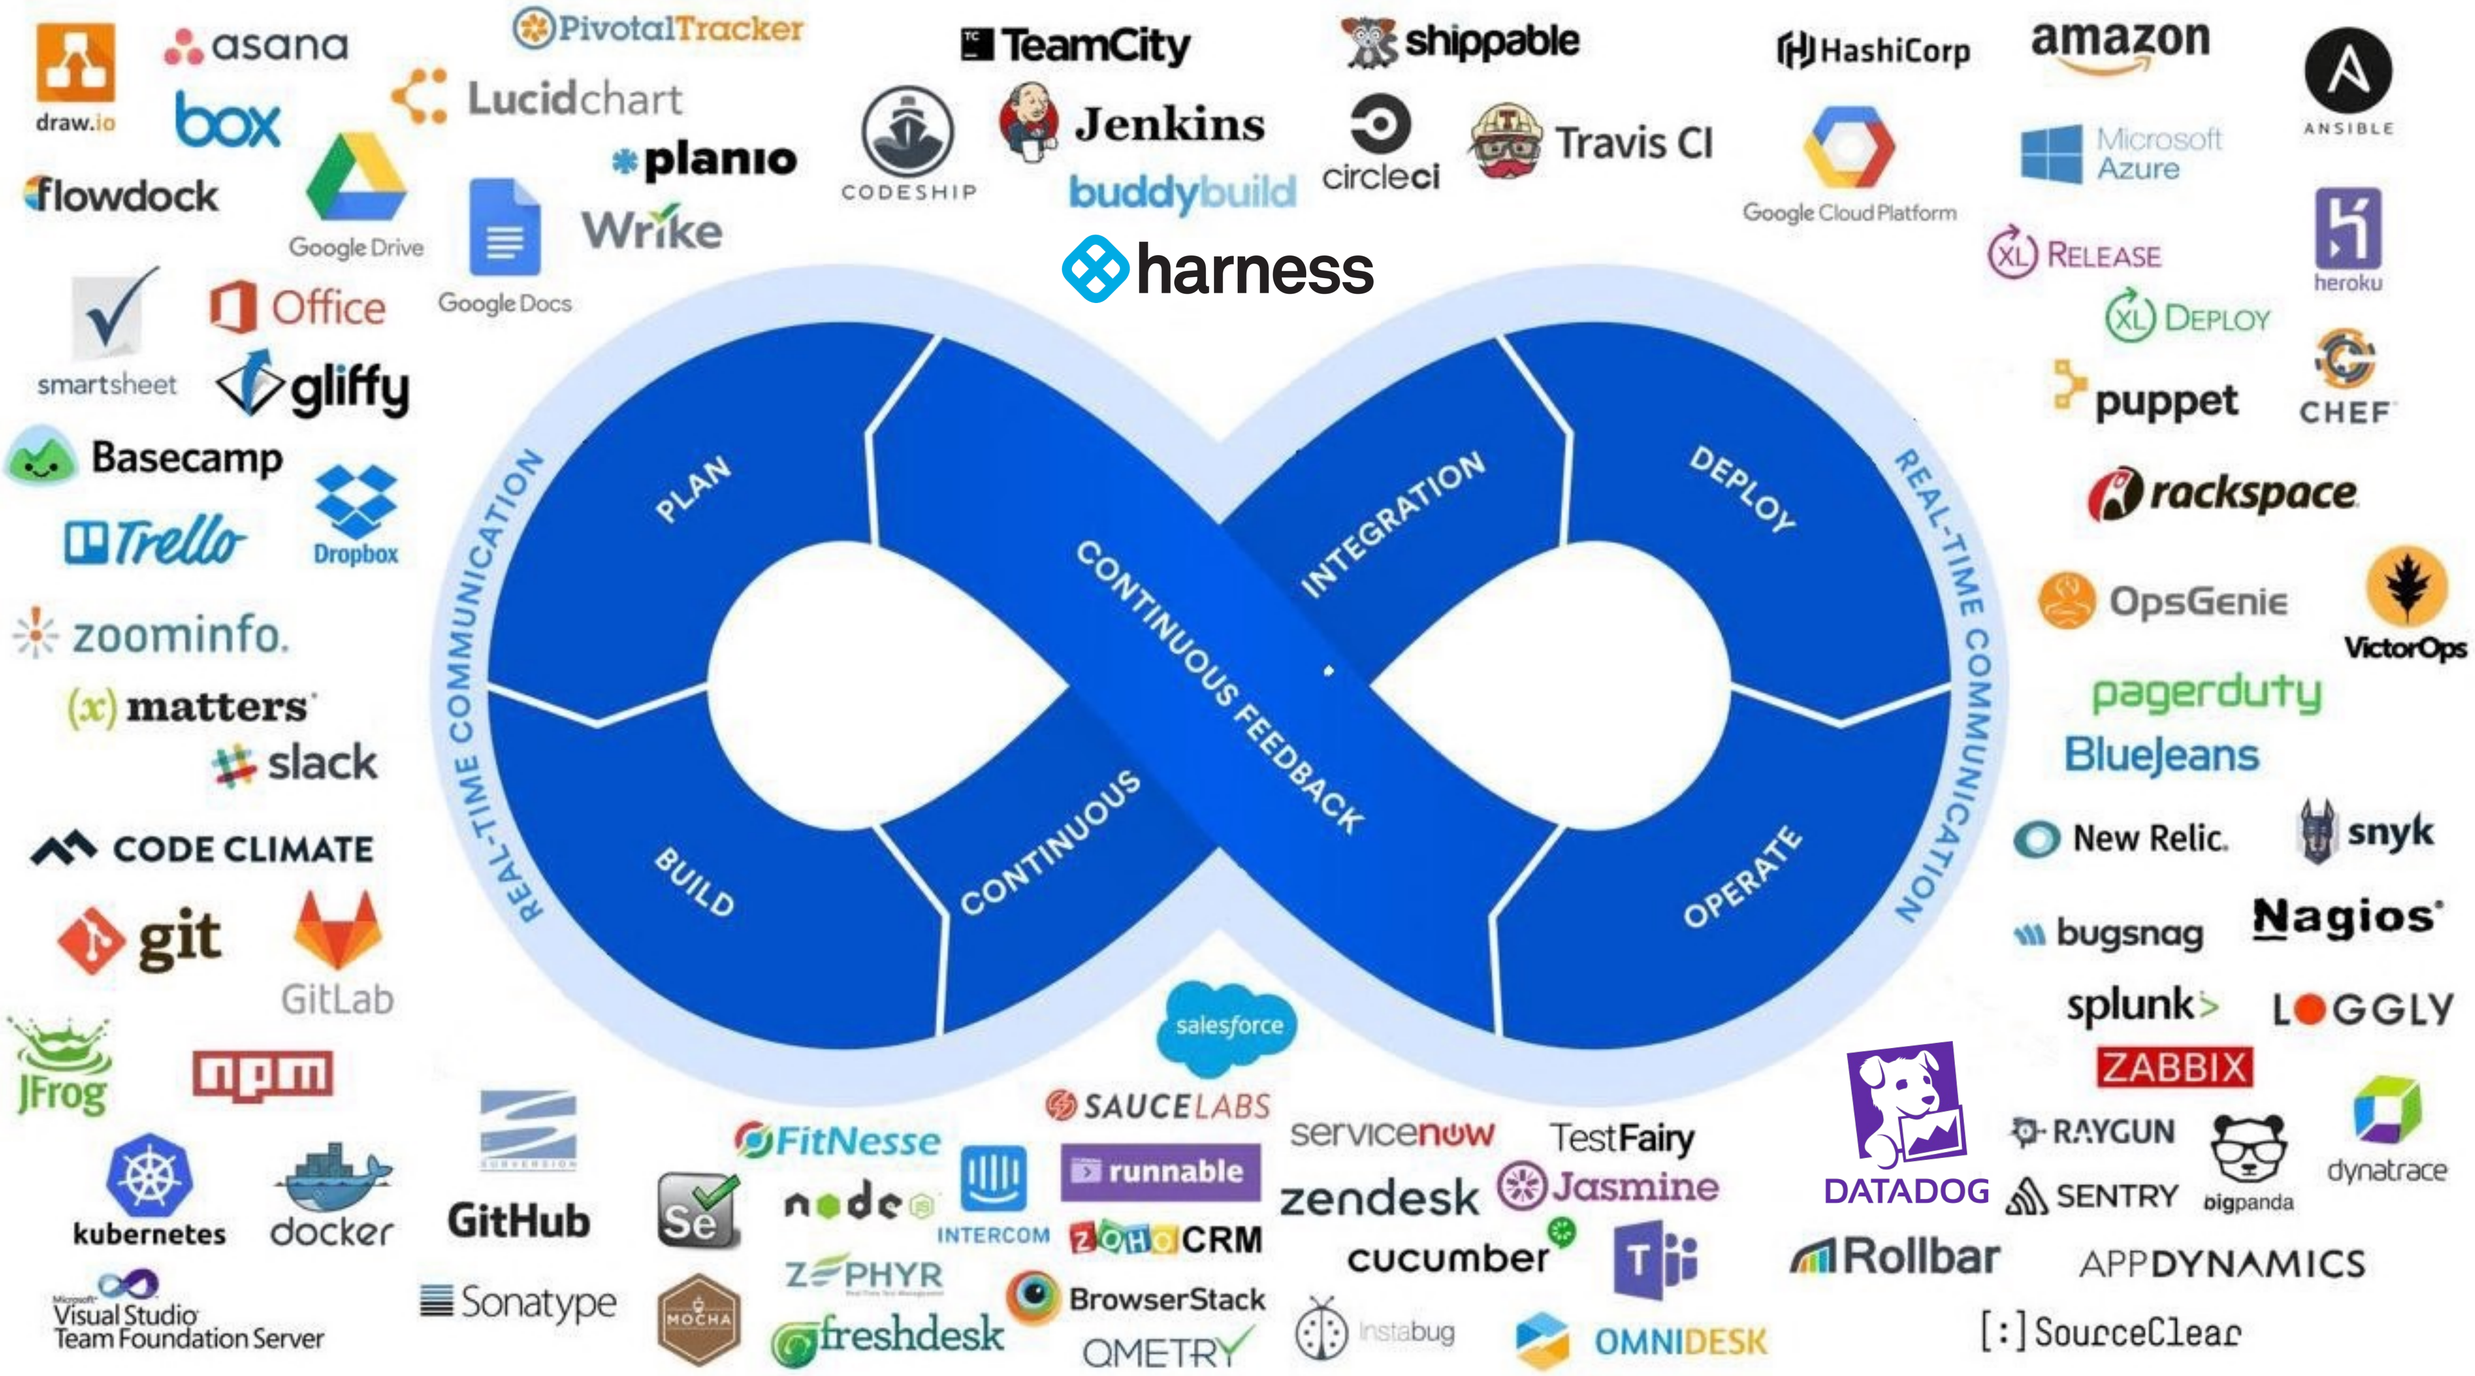
PLAN
BUILD
CONTINUOUS
INTEGRATION
DEPLOY
OPERATE
CONTINUOUS FEEDBACK
REAL-TIME COMMUNICATION
REAL-TIME COMMUNICATION
draw.io
asana
box
flowdock
Google Drive
PivotalTracker
Lucidchart
planio
Wrike
Google Docs
TeamCity
CODESHIP
Jenkins
buddybuild
shippable
circleci
Travis CI
HashiCorp
Google Cloud Platform
amazon
Microsoft Azure
ANSIBLE
XL RELEASE
heroku
XL DEPLOY
puppet
CHEF
rackspace
OpsGenie
VictorOps
pagerduty
BlueJeans
New Relic.
snyk
bugsnag
Nagios
splunk>
LOGGLY
ZABBIX
DATADOG
RAYGUN
SENTRY
bigpanda
dynatrace
Rollbar
APPDYNAMICS
[:] SourceClear
smartsheet
Office
gliffy
Basecamp
Trello
Dropbox
zoominfo.
(x) matters
slack
CODE CLIMATE
git
GitLab
JFrog
npm
kubernetes
docker
Microsoft Visual Studio Team Foundation Server
SUBVERSION
GitHub
Sonatype
Se
MOCHA
FitNesse
node
ZEPHYR
freshdesk
INTERCOM
salesforce
SAUCELABS
runnable
ZOHO CRM
BrowserStack
QMETRY
servicenow
TestFairy
zendesk
Jasmine
cucumber
Instabug
OMNIDESK

The Current State of SDLC

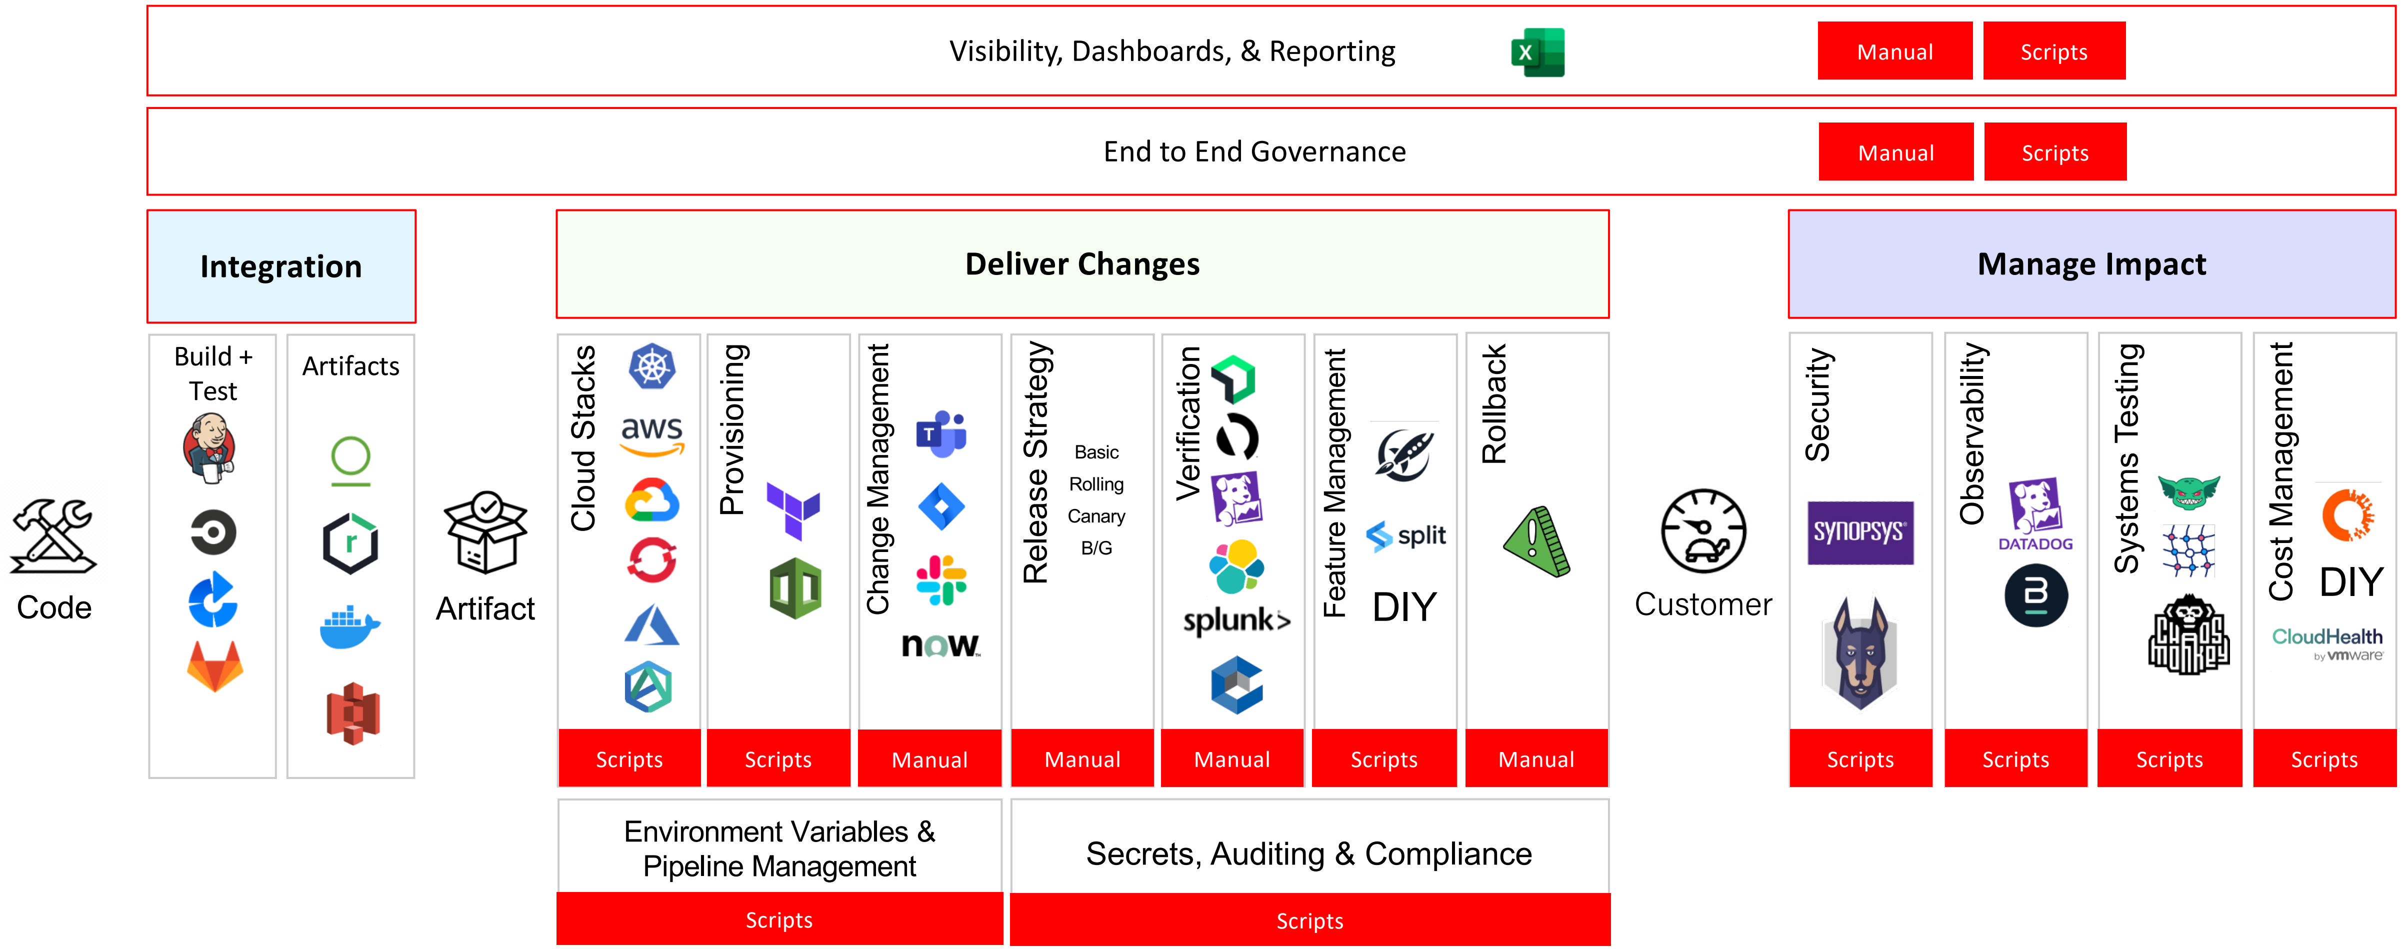

Visibility, Dashboards, & Reporting — Manual | Scripts

End to End Governance — Manual | Scripts

Code

Integration

- Build + Test
- Artifacts

Artifact

Deliver Changes

- Cloud Stacks — Scripts
- Provisioning — Scripts
- Change Management — Manual
- Release Strategy (Basic, Rolling, Canary, B/G) — Manual
- Verification — Manual
- Feature Management (DIY) — Scripts
- Rollback — Manual

Environment Variables & Pipeline Management — Scripts

Secrets, Auditing & Compliance — Scripts

Customer

Manage Impact

- Security — Scripts
- Observability — Scripts
- Systems Testing — Scripts
- Cost Management (DIY) — Scripts

☒ Slow ☒ Brittle ☒ Unreliable ☒ Costly

COLVIN RUN DEVSECOPS ENABLEMENT

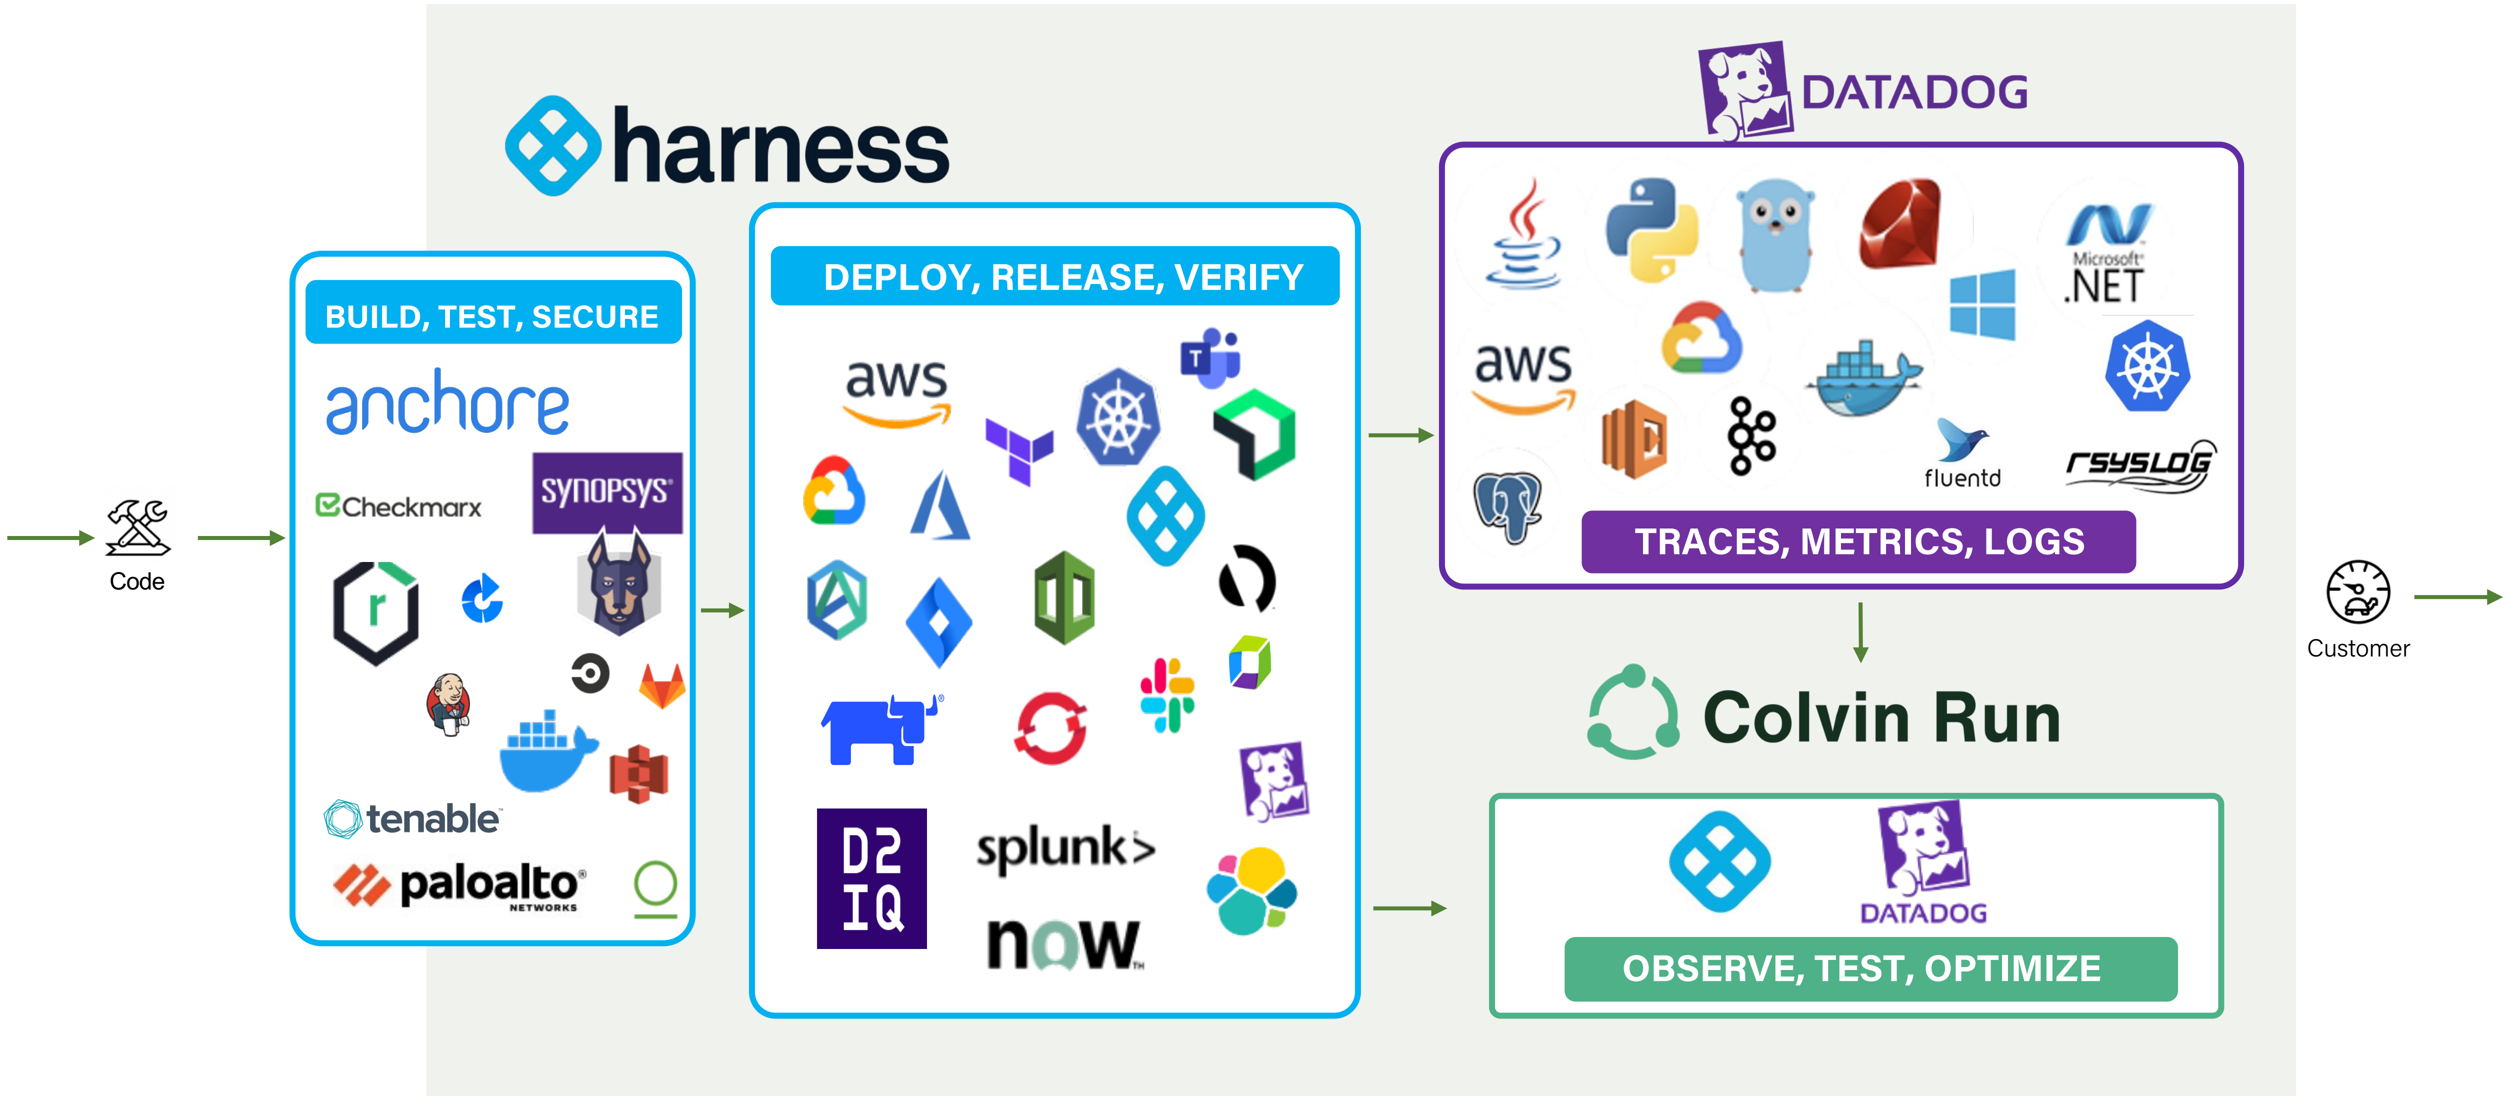

Code

harness

BUILD, TEST, SECURE

DEPLOY, RELEASE, VERIFY

DATADOG

TRACES, METRICS, LOGS

Colvin Run

OBSERVE, TEST, OPTIMIZE

Customer

COLVIN RUN DEVSECOPS ENABLEMENT

Unleash DevSecOps with our Integrated Enablement Platform

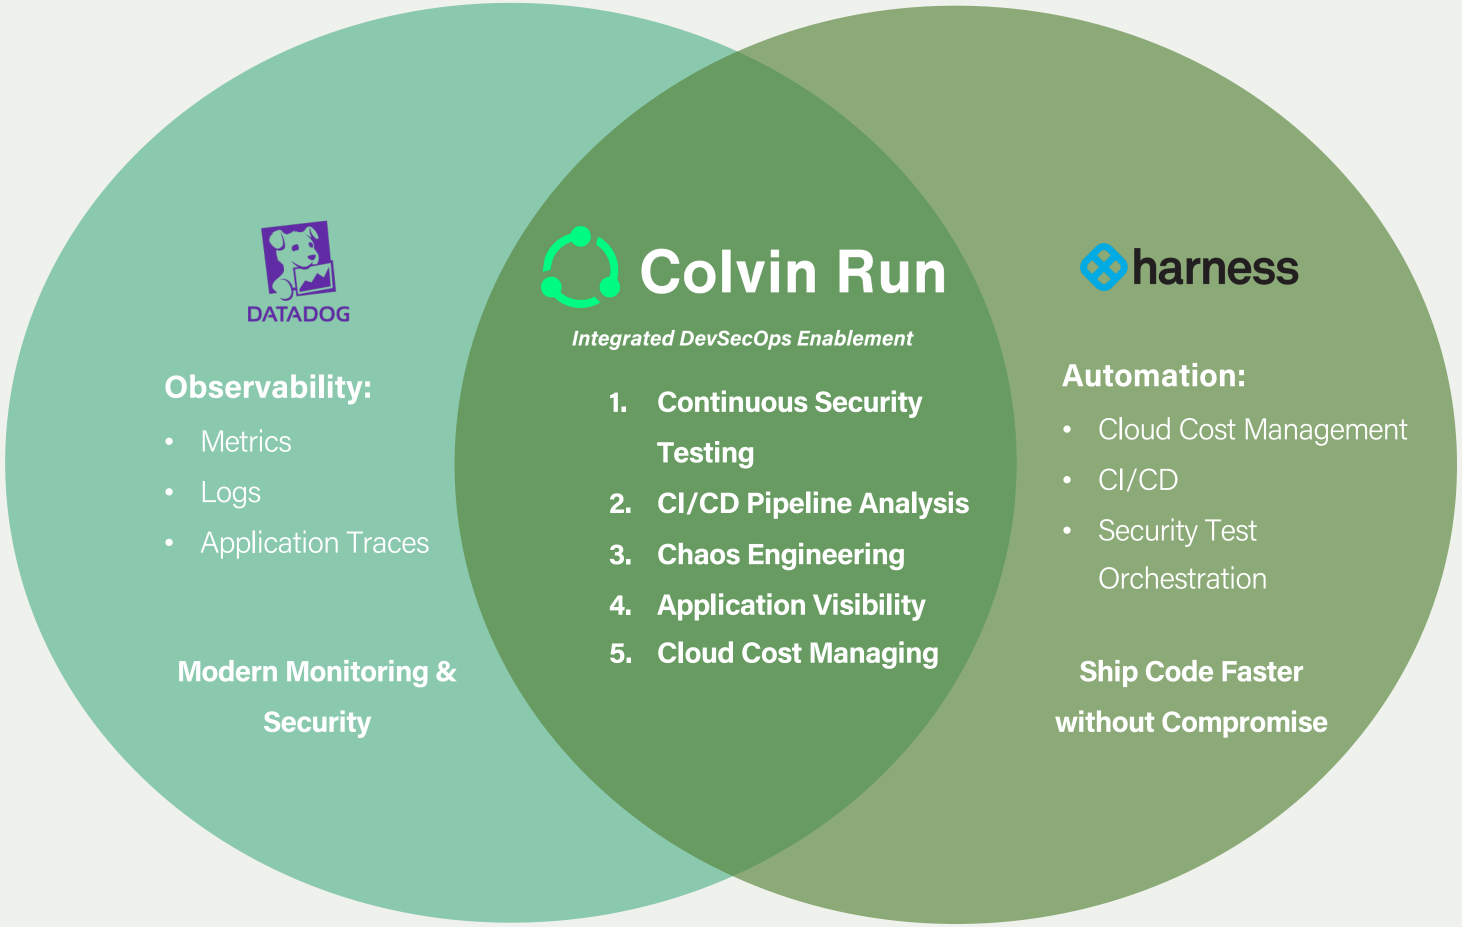

DATADOG

Observability:

- Metrics
- Logs
- Application Traces

Modern Monitoring & Security

Colvin Run

Integrated DevSecOps Enablement

1. **Continuous Security Testing**
2. **CI/CD Pipeline Analysis**
3. **Chaos Engineering**
4. **Application Visibility**
5. **Cloud Cost Managing**

harness

Automation:

- Cloud Cost Management
- CI/CD
- Security Test Orchestration

Ship Code Faster without Compromise

Continuous Security Deployment—
Automatically run security scanners at the right stages of the pipeline while monitoring integrations.

Pipeline Speed & Analysis—
Build & deploy mission-critical applications faster while continuously testing.

Chaos Engineering—
Run sets of chaos experiments on target system and observe results in order to make reliability improvements.

Transform DevOps Observability—
Visibility into every environment and every service.

Cloud Cost Transparency—
Understand granular spending and cost report while having transparency across infrastructures.

FAST. INTUITIVE. VISIBLE.

harness

DATADOG

CONTINUOUS SECURITY TESTING

When to Engage

Integrated with CI/CD

- Spending time maintaining tools and manually scripting
- Manual rendering of pages across browsers and tools

Intelligent Scanner Results

- Many different tools with disparate output
- Processing massive data lake from a scanner

Automated Pipeline Security

- Spending time coding end-to-end tests that break
- Minimal test coverage & stability for the application stack

Break Down Team Silos

- Struggles with verifying fixes were implemented
- Minimal collaboration across stages and teams

Interested in Continuous Security Testing? Click or Scan the QR Code to Contact Us!

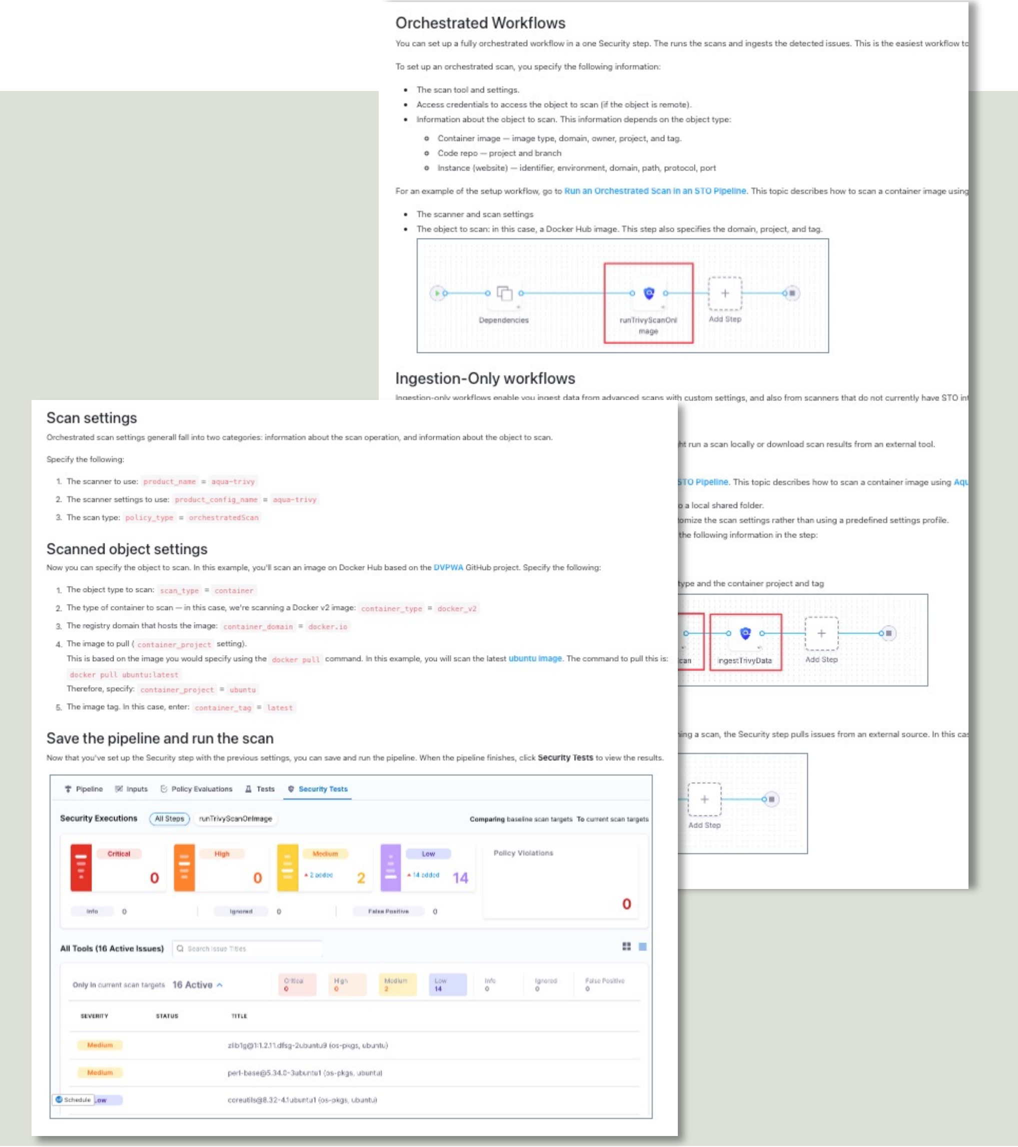

Orchestrated Workflows

You can set up a fully orchestrated workflow in a one Security step. The runs the scans and ingests the detected issues. This is the easiest workflow to

To set up an orchestrated scan, you specify the following information:

- The scan tool and settings.
- Access credentials to access the object to scan (if the object is remote).
- Information about the object to scan. This information depends on the object type:
  - Container image — image type, domain, owner, project, and tag.
  - Code repo — project and branch
  - Instance (website) — identifier, environment, domain, path, protocol, port

For an example of the setup workflow, go to Run an Orchestrated Scan in an STO Pipeline. This topic describes how to scan a container image using

- The scanner and scan settings
- The object to scan: in this case, a Docker Hub image. This step also specifies the domain, project, and tag.

Ingestion-Only workflows

Ingestion-only workflows enable you ingest data from advanced scans with custom settings, and also from scanners that do not currently have STO in

ht run a scan locally or download scan results from an external tool.

STO Pipeline. This topic describes how to scan a container image using Aqu

o a local shared folder.

tomize the scan settings rather than using a predefined settings profile.

the following information in the step:

type and the container project and tag

ing a scan, the Security step pulls issues from an external source. In this cas

Scan settings

Orchestrated scan settings generall fall into two categories: information about the scan operation, and information about the object to scan.

Specify the following:

1. The scanner to use: `product_name` = `aqua-trivy`
2. The scanner settings to use: `product_config_name` = `aqua-trivy`
3. The scan type: `policy_type` = `orchestratedScan`

Scanned object settings

Now you can specify the object to scan. In this example, you'll scan an image on Docker Hub based on the DVPWA GitHub project. Specify the following:

1. The object type to scan: `scan_type` = `container`
2. The type of container to scan — in this case, we're scanning a Docker v2 image: `container_type` = `docker_v2`
3. The registry domain that hosts the image: `container_domain` = `docker.io`
4. The image to pull (`container_project` setting).
   This is based on the image you would specify using the `docker pull` command. In this example, you will scan the latest ubuntu image. The command to pull this is:
   `docker pull ubuntu:latest`
   Therefore, specify: `container_project` = `ubuntu`
5. The image tag. In this case, enter: `container_tag` = `latest`

Save the pipeline and run the scan

Now that you've set up the Security step with the previous settings, you can save and run the pipeline. When the pipeline finishes, click **Security Tests** to view the results.

Security Executions — All Steps — runTrivyScanOnImage — **Comparing** baseline scan targets **To** current scan targets

Critical 0
High 0
Medium ▲2 added 2
Low ▲14 added 14
Policy Violations 0
Info 0 | Ignored 0 | False Positive 0

All Tools (16 Active Issues) — Search Issue Titles

Only in current scan targets 16 Active ^ — Critical 0 — High 0 — Medium 2 — Low 14 — Info 0 — Ignored 0 — False Positive 0

SEVERITY	STATUS	TITLE
Medium		zlib1g@1:1.2.11.dfsg-2ubuntu9 (os-pkgs, ubuntu)
Medium		perl-base@5.34.0-3ubuntu1 (os-pkgs, ubuntu)
Low		coreutils@8.32-4.1ubuntu1 (os-pkgs, ubuntu)

Colvin Run

INTEGRATED CI/CD PIPELINE ANALYSIS

When to Engage

Compliance at Scale

- Can't maximize velocity
- We are re-platforming our applications to run on a more modern architecture

CI/CD Visibility

- Unable to quickly find reason for test or pipeline failure
- Lack of visibility on effect a given commit has on a pipeline

Intelligent Test Runner

- Unable to separate running tests on a given commit versus all

Pipeline health

- No metrics available to show build frequency, failure rate, average duration and 95th percentile duration

Interested in CI/CD Pipeline Analysis? Click or Scan the QR Code to Contact Us!

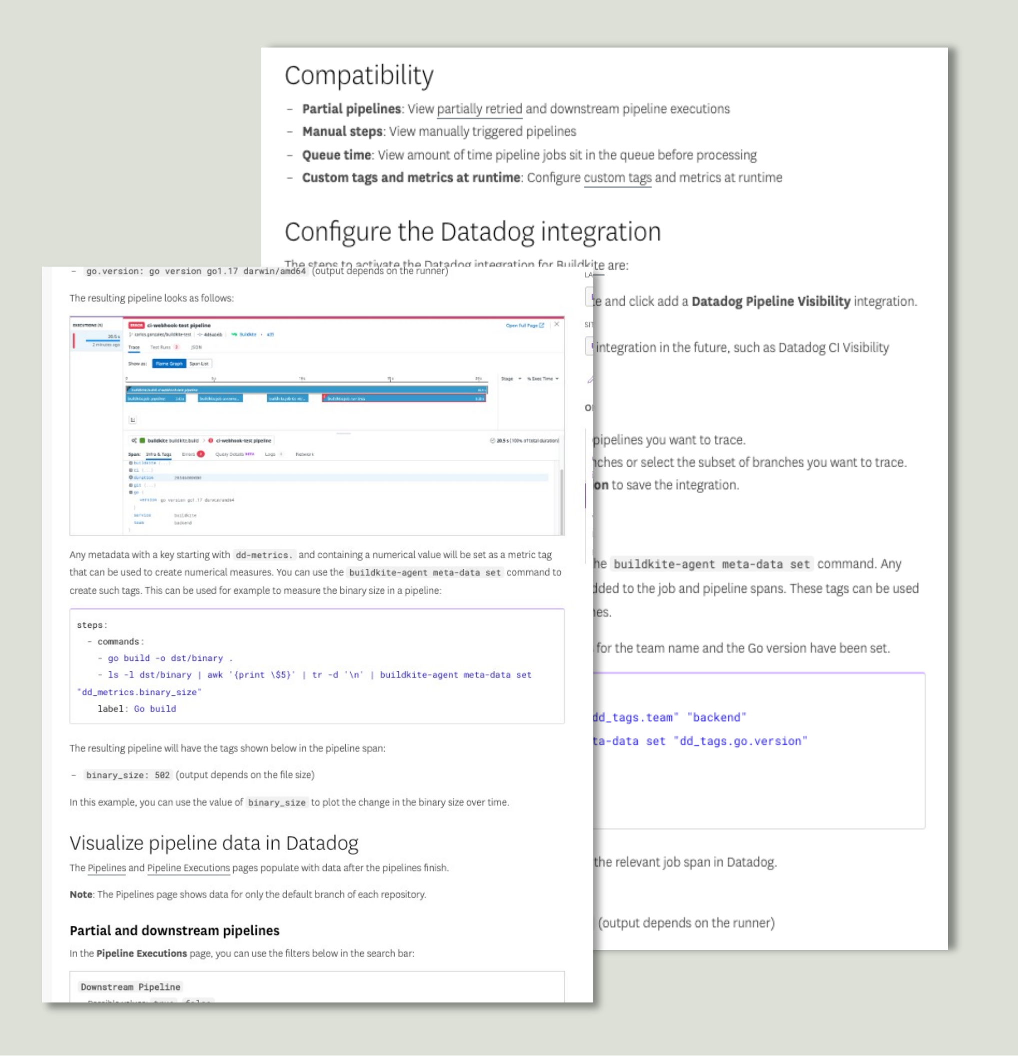

Compatibility

- **Partial pipelines**: View partially retried and downstream pipeline executions
- **Manual steps**: View manually triggered pipelines
- **Queue time**: View amount of time pipeline jobs sit in the queue before processing
- **Custom tags and metrics at runtime**: Configure custom tags and metrics at runtime

Configure the Datadog integration

The steps to activate the Datadog integration for Buildkite are:

e and click add a **Datadog Pipeline Visibility** integration.

integration in the future, such as Datadog CI Visibility

pipelines you want to trace.

ches or select the subset of branches you want to trace.

on to save the integration.

he `buildkite-agent meta-data set` command. Any

dded to the job and pipeline spans. These tags can be used

es.

for the team name and the Go version have been set.

```
dd_tags.team" "backend"
ta-data set "dd_tags.go.version"
```

the relevant job span in Datadog.

(output depends on the runner)

- `go.version: go version go1.17 darwin/amd64` (output depends on the runner)

The resulting pipeline looks as follows:

Any metadata with a key starting with `dd-metrics.` and containing a numerical value will be set as a metric tag that can be used to create numerical measures. You can use the `buildkite-agent meta-data set` command to create such tags. This can be used for example to measure the binary size in a pipeline:

```
steps:
  - commands:
    - go build -o dst/binary .
    - ls -l dst/binary | awk '{print \$5}' | tr -d '\n' | buildkite-agent meta-data set
"dd_metrics.binary_size"
    label: Go build
```

The resulting pipeline will have the tags shown below in the pipeline span:

- `binary_size: 502` (output depends on the file size)

In this example, you can use the value of `binary_size` to plot the change in the binary size over time.

Visualize pipeline data in Datadog

The Pipelines and Pipeline Executions pages populate with data after the pipelines finish.

Note: The Pipelines page shows data for only the default branch of each repository.

Partial and downstream pipelines

In the **Pipeline Executions** page, you can use the filters below in the search bar:

```
Downstream Pipeline
```

CHAOS ENGINEERING

When to Engage

Integrated into CI/CD Systems

- Application failures with excessive logging to debug, too many retries, and service timeouts

Robust Experiments

- Poor user experience
- Capacity issues and monitoring dashboards not available

Enable Observability

- Continuous infrastructure failures such as: device failures, network failures and regions not being available

Collaborative

- No centralized control plane

Interested in Chaos Engineering? Click or Scan the QR Code to Contact Us!

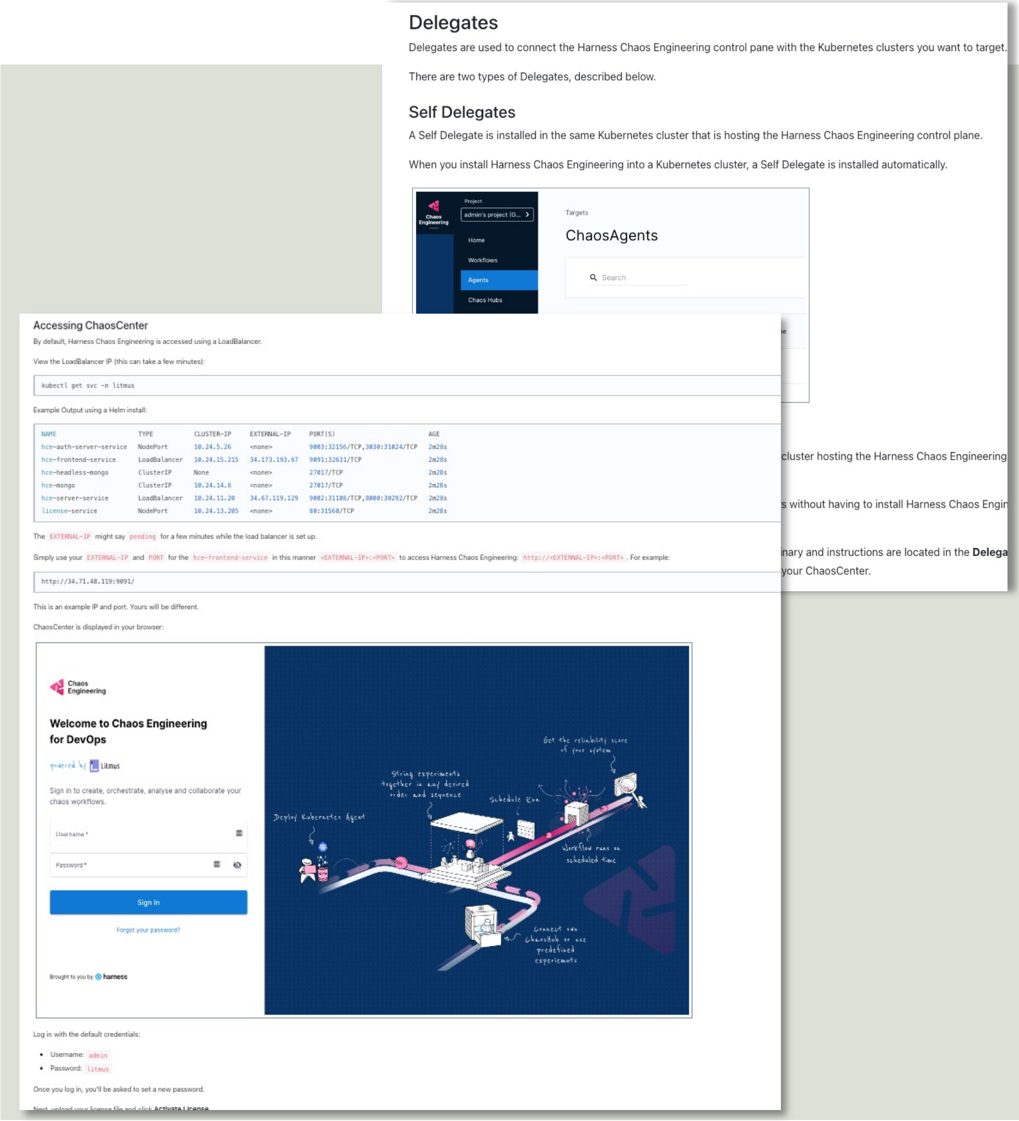

Delegates

Delegates are used to connect the Harness Chaos Engineering control pane with the Kubernetes clusters you want to target.

There are two types of Delegates, described below.

Self Delegates

A Self Delegate is installed in the same Kubernetes cluster that is hosting the Harness Chaos Engineering control plane.

When you install Harness Chaos Engineering into a Kubernetes cluster, a Self Delegate is installed automatically.

cluster hosting the Harness Chaos Engineering

s without having to install Harness Chaos Engin

inary and instructions are located in the **Delega**
your ChaosCenter.

Accessing ChaosCenter

By default, Harness Chaos Engineering is accessed using a LoadBalancer.

View the LoadBalancer IP (this can take a few minutes):

```
kubectl get svc -n litmus
```

Example Output using a Helm install:

```
NAME                     TYPE          CLUSTER-IP     EXTERNAL-IP     PORT(S)                        AGE
hce-auth-server-service  NodePort      10.24.5.26     <none>          9003:32156/TCP,3030:31024/TCP  2m28s
hce-frontend-service     LoadBalancer  10.24.15.215   34.173.193.67   9091:32631/TCP                 2m28s
hce-headless-mongo       ClusterIP     None           <none>          27017/TCP                      2m28s
hce-mongo                ClusterIP     10.24.14.6     <none>          27017/TCP                      2m28s
hce-server-service       LoadBalancer  10.24.11.20    34.67.119.129   9002:31108/TCP,8000:30292/TCP  2m28s
license-service          NodePort      10.24.13.205   <none>          80:31560/TCP                   2m28s
```

The `EXTERNAL-IP` might say `pending` for a few minutes while the load balancer is set up.

Simply use your `EXTERNAL-IP` and `PORT` for the `hce-frontend-service` in this manner `<EXTERNAL-IP>:<PORT>` to access Harness Chaos Engineering: `http://<EXTERNAL-IP>:<PORT>` . For example:

```
http://34.71.48.119:9091/
```

This is an example IP and port. Yours will be different.

ChaosCenter is displayed in your browser:

Welcome to Chaos Engineering for DevOps

Sign in to create, orchestrate, analyse and collaborate your chaos workflows.

Chaos Engineering
powered by Litmus
Username *
Password *
Sign In
Forgot your password?
Brought to you by harness
Deploy Kubernetes Agent
String experiments together in any desired order and sequence
Schedule Run
Get the reliability score of your system
Workflow runs on scheduled time
Connect own ChaosHub or use predefined experiments

Log in with the default credentials:

- Username: `admin`
- Password: `litmus`

Once you log in, you'll be asked to set a new password.

Next, upload your license file and click **Activate License**.

harness

DATADOG

APPLICATION VISIBILITY

When to Engage

Organization-wide visibility

- Lack of communication between teams and stages
- Engineers are unaware of upstream and downstream systems intertwined with their own

Collaboration

- Lack of visibility for users across the organization to evaluate data, collaborate, or share information

DevOps Automation

- Struggles with development workflows, error-prone manual processes, and over-burdening technical staff

Measure Everything

- No culture of measurement or observability

Interested in Application Visibility? Click or Scan the QR Code to Contact Us!

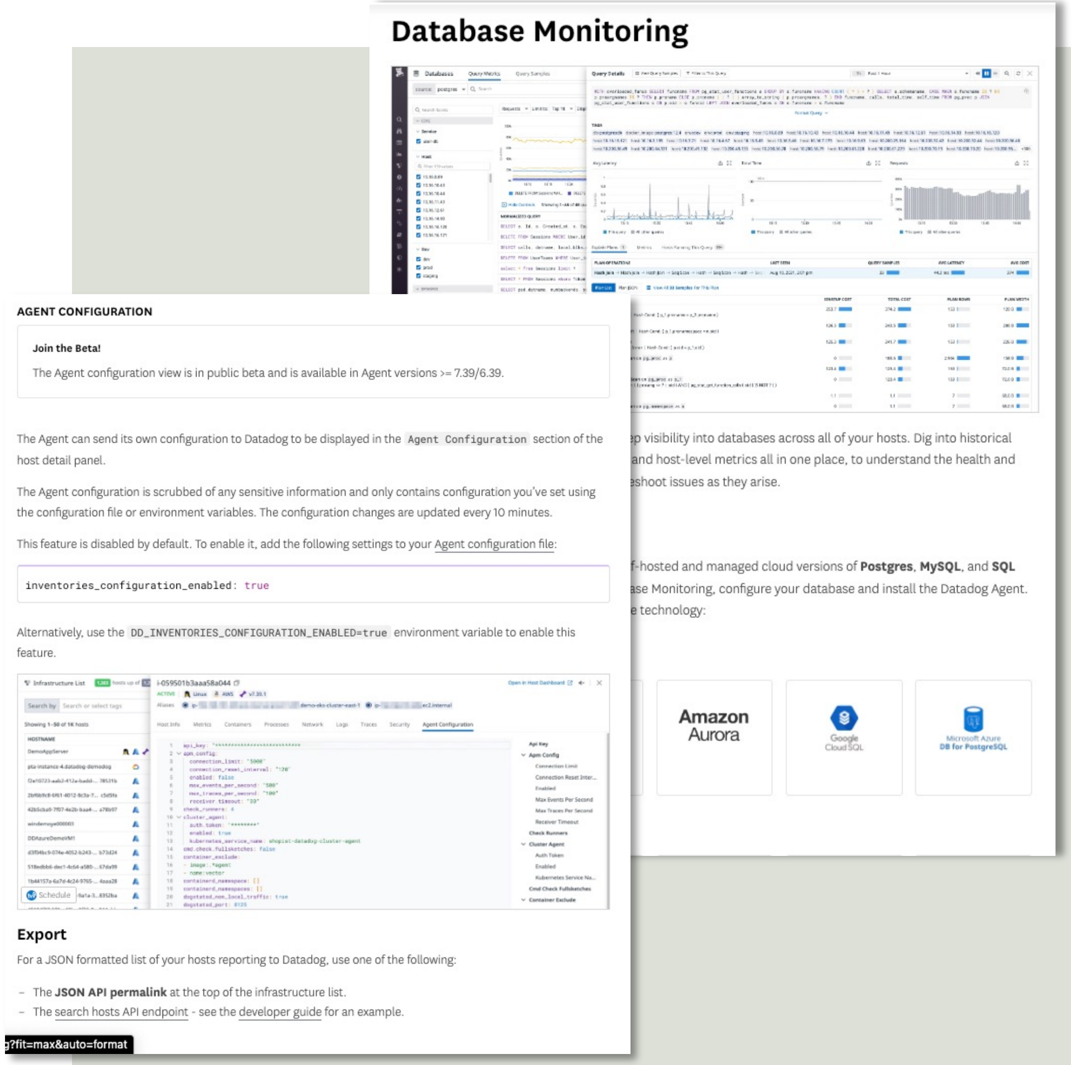

Database Monitoring

p visibility into databases across all of your hosts. Dig into historical and host-level metrics all in one place, to understand the health and eshoot issues as they arise.

f-hosted and managed cloud versions of **Postgres**, **MySQL**, and **SQL** ase Monitoring, configure your database and install the Datadog Agent. e technology:

AGENT CONFIGURATION

Join the Beta!

The Agent configuration view is in public beta and is available in Agent versions >= 7.39/6.39.

The Agent can send its own configuration to Datadog to be displayed in the `Agent Configuration` section of the host detail panel.

The Agent configuration is scrubbed of any sensitive information and only contains configuration you've set using the configuration file or environment variables. The configuration changes are updated every 10 minutes.

This feature is disabled by default. To enable it, add the following settings to your Agent configuration file:

```
inventories_configuration_enabled: true
```

Alternatively, use the `DD_INVENTORIES_CONFIGURATION_ENABLED=true` environment variable to enable this feature.

```
api_key: "****************************"
apm_config:
  connection_limit: "5000"
  connection_reset_interval: "120"
  enabled: false
  max_events_per_second: "500"
  max_traces_per_second: "100"
  receiver_timeout: "30"
check_runners: 4
cluster_agent:
  auth_token: "********"
  enabled: true
  kubernetes_service_name: shopist-datadog-cluster-agent
cmd.check.fullsketches: false
container_exclude:
- image:.*agent
- name:vector
containerd_namespace: []
containerd_namespaces: []
dogstatsd_non_local_traffic: true
dogstatsd_port: 8125
```

Export

For a JSON formatted list of your hosts reporting to Datadog, use one of the following:

- The **JSON API permalink** at the top of the infrastructure list.
- The search hosts API endpoint - see the developer guide for an example.

g?fit=max&auto=format

Colvin Run

CLOUD COST MANAGEMENT

When to Engage

Cost Transparency

- We're finding it hard to create cost clarity between teams, projects, and products

Cost Control

- We have disparate logging APM, synthetics, infra tooling today…it's not efficient and too expensive
- Costs are spiraling out of control

Cost Optimization

- We are missing context and intelligence for cost optimization

Cost Governance

- Users must manually turn on/off environments, often without clear understanding of impacts
- DevSecOps/IT are unable to respond to cost spikes quickly

Interested in Cloud Cost Management? Click or Scan the QR Code to Contact Us!

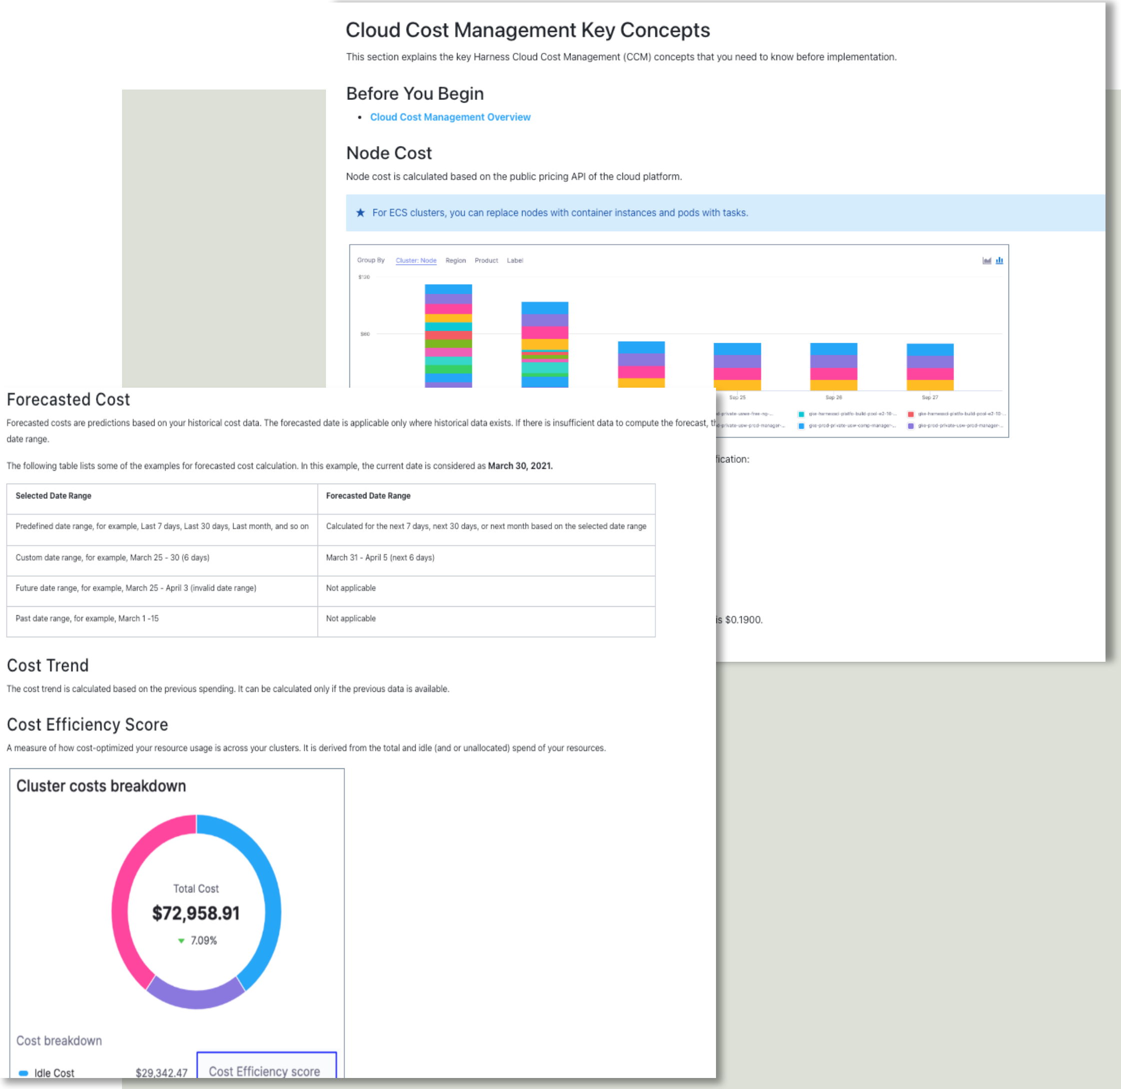

Cloud Cost Management Key Concepts

This section explains the key Harness Cloud Cost Management (CCM) concepts that you need to know before implementation.

Before You Begin

- Cloud Cost Management Overview

Node Cost

Node cost is calculated based on the public pricing API of the cloud platform.

★ For ECS clusters, you can replace nodes with container instances and pods with tasks.

Group By
Cluster: Node
Region
Product
Label
$120
$80
Sep 25
Sep 26
Sep 27
gke-harnessci-platfo-build-pool-e2-10-...
gke-prod-private-usw-comp-manager-...
gke-harnessci-platfo-build-pool-e2-10-...
gke-prod-private-usw-prod-manager-...

fication:

is $0.1900.

Forecasted Cost

Forecasted costs are predictions based on your historical cost data. The forecasted date is applicable only where historical data exists. If there is insufficient data to compute the forecast, th date range.

The following table lists some of the examples for forecasted cost calculation. In this example, the current date is considered as **March 30, 2021.**

Selected Date Range	Forecasted Date Range
Predefined date range, for example, Last 7 days, Last 30 days, Last month, and so on	Calculated for the next 7 days, next 30 days, or next month based on the selected date range
Custom date range, for example, March 25 - 30 (6 days)	March 31 - April 5 (next 6 days)
Future date range, for example, March 25 - April 3 (invalid date range)	Not applicable
Past date range, for example, March 1 -15	Not applicable

Cost Trend

The cost trend is calculated based on the previous spending. It can be calculated only if the previous data is available.

Cost Efficiency Score

A measure of how cost-optimized your resource usage is across your clusters. It is derived from the total and idle (and or unallocated) spend of your resources.

harness

DATADOG

COLVIN RUN ADVANTAGES

IMPLEMENTATION

- Colvin Run works directly with you to get DevSecOps up and running with our Harness + Datadog powered solutions
- Evaluate current processes & infrastructure to find best ways to implement without disruption
- Create product roadmap for migration

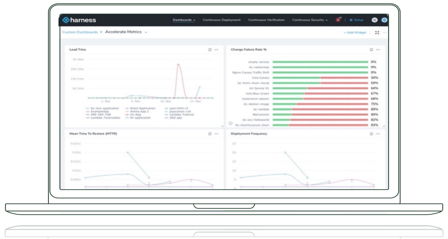

TRAINING

- Developed variety of training options
- Detailed documentation for customer-specific configurations
- Assist DevOps in adopting new products
- Security orchestration best practices

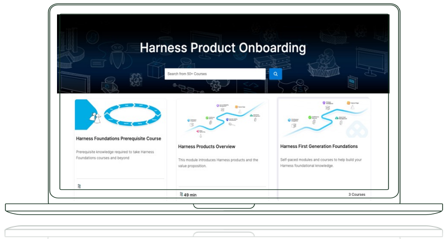
Harness Product Onboarding

Search from 50+ Courses

Harness Foundations Prerequisite Course

Prerequisite knowledge required to take Harness Foundations courses and beyond

Harness Products Overview

This module introduces Harness products and the value proposition.

49 min

Harness First Generation Foundations

Self-paced modules and courses to help build your Harness foundational knowledge.

3 Courses

ONGOING SUPPORT

- Always on hand to handle any questions, training requests, or post-implementation issues
- Ticket-generated support for proactive assistance
- Accelerate initiatives like cloud migration and onboard new services

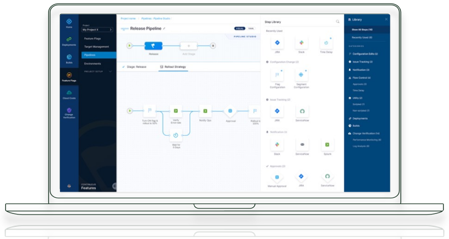

Colvin Run

SBIR PHASE III (1/2)

Colvin Run's Small Business Innovation Research (SBIR) Phase I and II contracts provide Federal customers with a statutory justification & approval (J&A) for sole-source contracting.

- SBIR is a statutory program (15 U.S.C § 638), not a vehicle
- Federal Acquisition Regulation (FAR) 6.302-5 states that "Full and open competition need not be provided when…a statue expressly authorizes or requires that the acquisition be made through another agency or from a specified source."
- For purposes of a J&A, the Federal agency simply states that the SBIR Phase III award is derived from, extends, or completes efforts made under prior SBIR funding agreements and is authorized pursuant to 15 U.S.C 638(r)(4)
- SBIR Phase III contracts receive small business credit, cannot be protested, have no subcontracting limit, and have no award value limit.

EXAMPLE COLVIN RUN SBIR PHASE III CONTRACTING ACTION

In Federal FY 2021, a key decision maker within the Small Business Administration's Office of Policy, Planning and Liaison (SBA-OPPL) was seeking cutting-edge data analyses for Federal procurement data.

Working with SBA-OPPL's cognizant Contracting Office, the decision maker was able to expediently issue an SBIR Phase III contract directly to Colvin Run in under 20 calendar days.

The SBA Contracting Officer referenced one of Colvin Run's prior SBIR Phase II contracts (US Navy Contract No. N68335-21-C-0001). While Colvin Run's original SBIR Phase II Contract was a different use case aimed at building a platform to analyze large volumes of P-8A Poseidon aircraft data, the SBA Contracting Officer was able to see that the SBIR Phase II Contract was rooted in data analytics, which could be extended to the SBA use case.

Competition conducted under Colvin Run's prior SBIR Phase II award satisfied FAR competition requirements.

The pre-competed data analytics SBIR was all that was required for justification to work directly with Colvin Run.

SBIR PHASE III (2/2)

Colvin Run can receive a sole-source SBIR Phase III contract from any Federal organization for requirements that derive from, extend, or complete work performed under our prior SBIR Phase I and II contracts.

DIRECT	SBIR Phase III contracts can be issued on a FAR-compliant sole-source basis, having already satisfied competition requirements in Phases I and II
RAPID	SBIR Phase III contracts can be awarded within a matter of weeks
EFFICIENT	Unlike other Small Business Set-Aside Vehicles, SBIR Phase III contracts have no 51% work-share requirement
	SBIR Phase III contracts can be structured as an IDIQ to address various, emerging requirements
	The GSA Office of Assisted Acquisition Services (AAS) is an expert at establishing SBIR Phase III vehicles
FLEXIBLE	No limits on the dollar size of an SBIR Phase III procurement
	Acquire products, production, services, research, research and development, or any combination
	Use any color of government funds (e.g. O&M, FMS, RDT&E) from any Agency
SECURE	Contracts or Task Orders can be issued with requirements for UNCLAS, SECRET, TS/SCI

MAIDEN
Maritime Agile Intelligent Data Exploitation Network
DOD SBIR TECHNOLOGY • P-8A MARITIME BIG DATA ANALYTICS

SHOPMAN
Supply Hub for Operational Predictive Maintenance & Analytics
POWERED BY MICROSTRATEGY

Big Data Analytics
- Analytics Environment
- SECRET Clearance
- Cross Domain Engineering
- ISR ML Applications
- Containerized Applications

Contract N6833519C0437

Microelectronics Blockchain
- Quantitative Assurance
- Trusted Supply Chain
- Standards Implementation
- HW/SW Security
- Blockchain Solutions

Contract HQ072720C0001

Maintenance Modernization
- User-centered CBM+ Apps
- Authoritative Data Environment
- Management Decision Aid
- Predictive Asset Management
- Logistics Applications

Contract FA864920C0113

Data Integration & Fusion
- Analytics Environment
- Applied Data Science
- Workflow Automation
- Cloud Compute Optimization
- Process Modernization

Contract FA875020P1626

Cybersecurity & Training
- Hybrid IT Infrastructure
- Cyber Operations
- Cloud Migration Support
- Hardware Enclave
- IT Process Automation

Contract FA864920P0673

Harness Software Delivery Platform

harness

Build & Test

Continuous Integration — Automate build/test, and packaging of code to artifacts

Security Testing Orchestration — Consolidate/deduplicate all app security testing cross-enterprise while natively integrating into CI/CD

Chaos Engineering — Leverage ChaosNative—the company behind the LitmusChaos CNCF Project—to experiment on a production system and promote software resilience.

Deploy and Manage

Continuous Delivery — Build deployment pipelines in minutes, and safely deploy artifacts to production

Feature Flags — Release and manage features at scale

Create Reliability & Efficiency

Cloud Cost Management — Active cloud cost insights, recommendations, and automated optimizations

Service Reliability Management — Deliver on SLOs and SLAs, and ensure reliability with error budgets

Workloads

- Cloud-Native Apps
- Traditional Apps
- GitOps
- Mobile Apps
- Database Schemas
- IoT Code
- Big Data Pipelines
- ML Models
- Packaged SaaS Apps

The Three Pillars of Observability

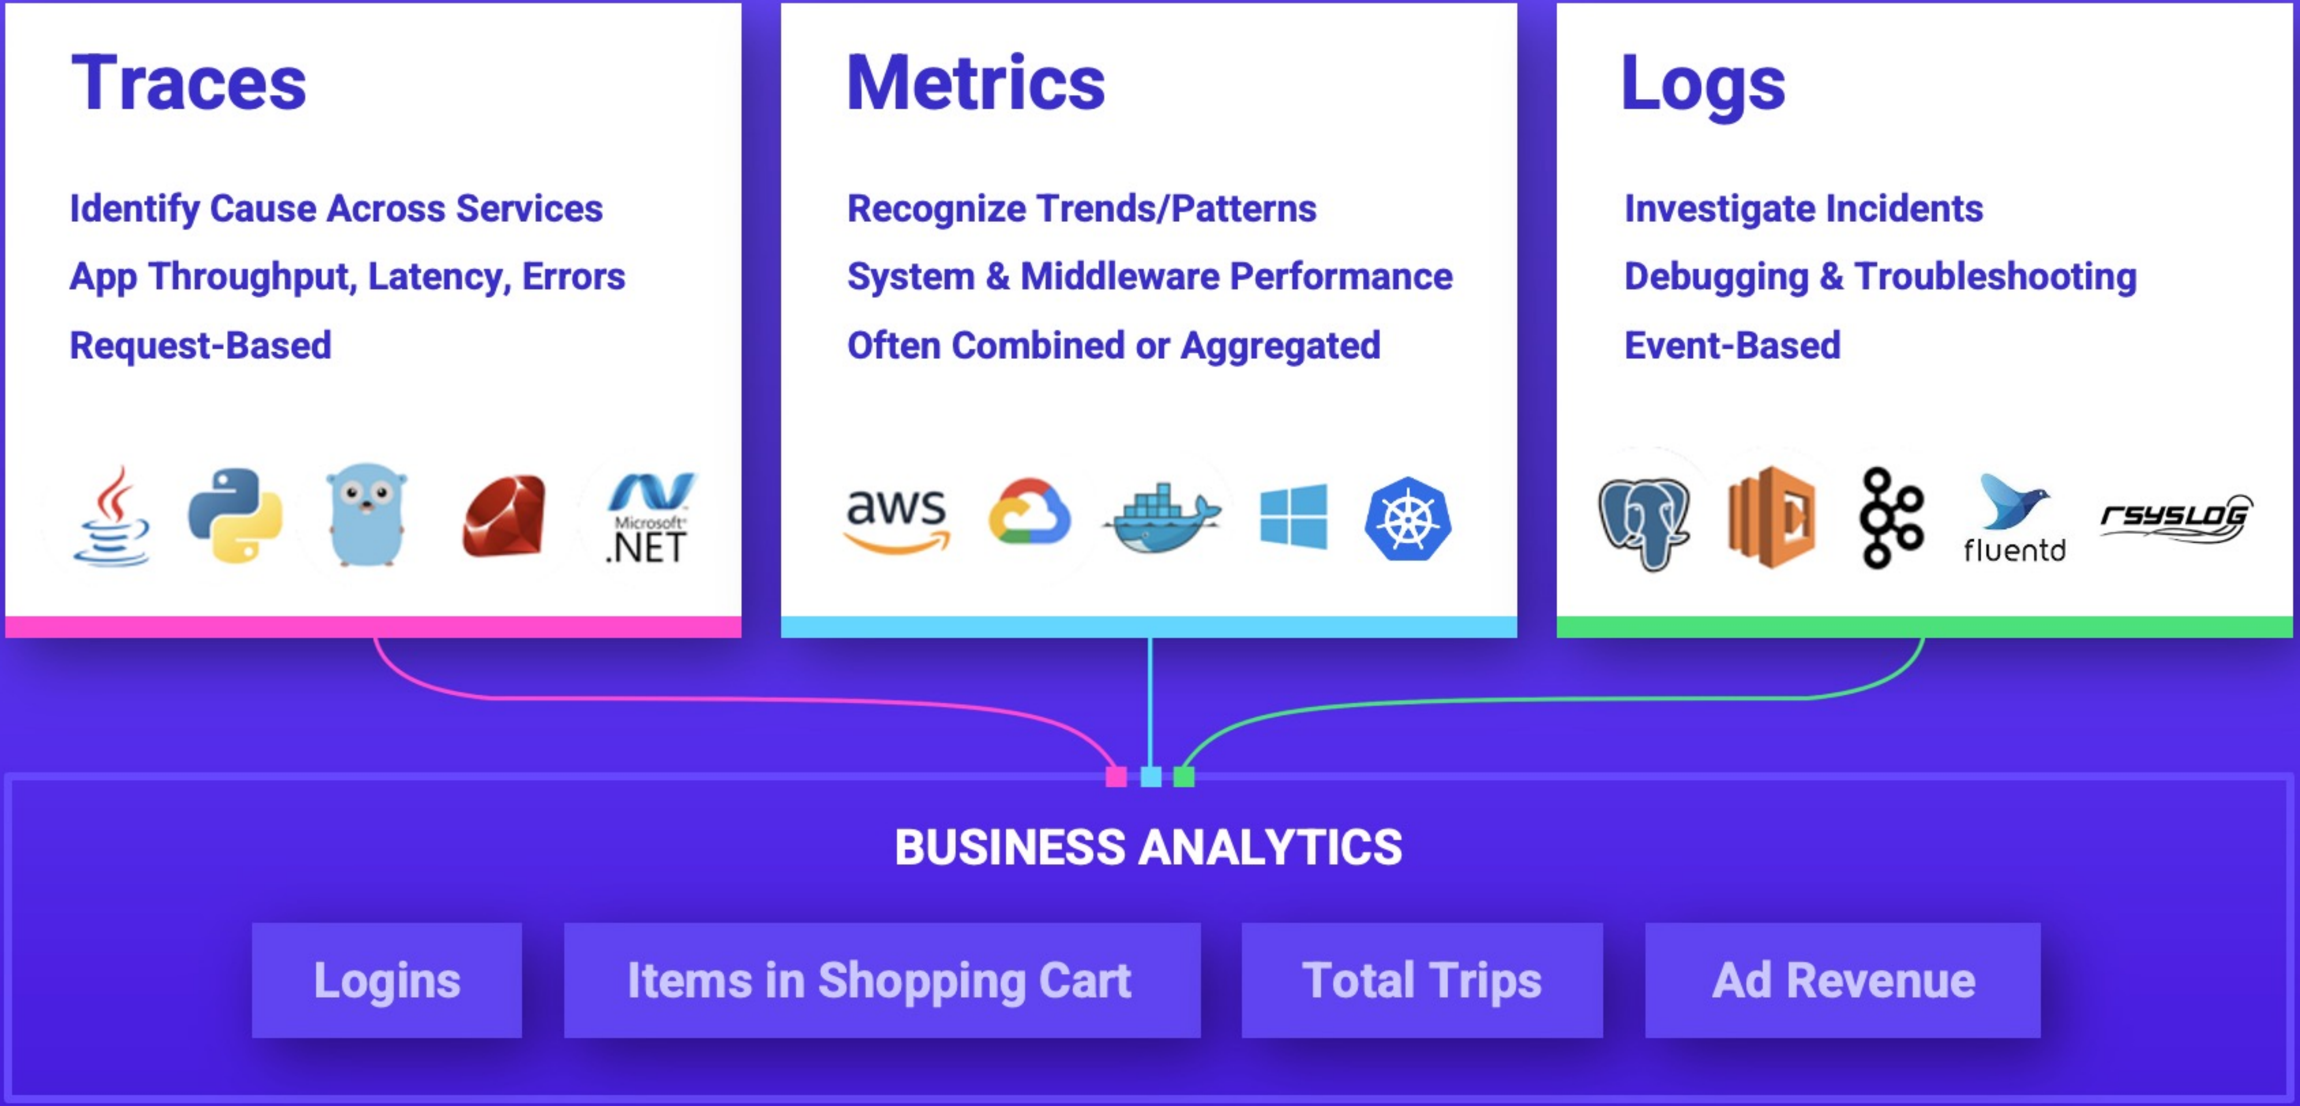

Traces

Identify Cause Across Services

App Throughput, Latency, Errors

Request-Based

Metrics

Recognize Trends/Patterns

System & Middleware Performance

Often Combined or Aggregated

Logs

Investigate Incidents

Debugging & Troubleshooting

Event-Based

BUSINESS ANALYTICS

Logins

Items in Shopping Cart

Total Trips

Ad Revenue

ADMINISTRATIVE & CONTACT INFO

SBIR CONTRACT SELECTIONS

Phase I	10
Phase II	9
Phase III	2

NAICS CODES

541330	Engineering Services
541511	Custom Computer Programming
541512	Computer Systems Design
541519	Other Computer Related Services
541611	General Management Consulting
541715	Research and Development

SBIR AWARDED TOPIC AREAS

- Automation
- Blockchain
- Business Intelligence
- Cloud
- Data Analytics
- Machine Learning
- Maintenance
- Modernization
- Networking

Nikhil Shenoy, CEO | 703-967-1967 | nikhil@colvinrun.net
Gabby Cabrera, Analyst | 980-328-9095 | gabby@colvinrun.net

Web	www.colvinrun.net
Phone	571.207.7487
Email	info@colvinrun.net
UEI	Z475KM9G68V8
CAGE	800J6