

Landslide Cloud-Native Infrastructure Benchmarking

Readiness and Performance Testing

As enterprises modernize applications and migrate from virtual machines to cloud-native, containerized platforms, ensuring infrastructure readiness has become a critical priority. While Kubernetes and cloud platforms promise elastic scaling, high availability, and efficient resource utilization, teams often encounter performance variability, resource contention, and operational blind spots that impact application reliability and cost.

Validating whether infrastructure can consistently support business-critical workloads — while scaling, recovering, and reallocating resources across hybrid and multi-cloud environments — adds further complexity. Without clear performance baselines for throughput, latency, and capacity, organizations risk degraded application performance, unplanned outages, cloud overspend, and slower time-to-market.

Keysight's Landslide Cloud-Native Infrastructure Benchmarking provides purpose-built benchmarking to validate the design, deployment, and operation of enterprise cloud-native infrastructure before issues reach production.

Highlights

- **Accelerate cloud-native adoption** by validating Kubernetes infrastructure readiness before deploying business critical applications.
- **Reduce risk of outages and performance degradation** by verifying usable capacity and infrastructure behavior under real world load.
- **Improve cost efficiency** by identifying over and under-utilized compute, memory, storage, and network resources.
- **Enable data-driven infrastructure decisions** for scaling, upgrades, and hybrid or multi-cloud strategy.
- **Purpose-built benchmarking** across CPU, memory, storage, and networking to support modern enterprise workloads.
- **Seamless observability** integration with Grafana, Prometheus, and Kafka for real time insight and faster troubleshooting.

Solution Overview

Landslide Cloud-Native Infrastructure Benchmarking (Cloud Engine) benchmarks the performance, efficiency, and behavior of Kubernetes infrastructure and cloud-based resources under realistic, production-like conditions. By simulating diverse enterprise application workloads, platform and site reliability engineering (SRE) teams gain actionable insight into scalability limits, performance bottlenecks, and resource allocation efficiency.

Benchmarking spans central processing unit (CPU), memory, storage, and networking, providing infrastructure teams with the data they need to validate capacity, optimize configurations, compare environments, and plan growth with confidence. The solution integrates seamlessly with CI/CD pipelines, enabling continuous infrastructure validation as part of modern DevOps and platform engineering workflows.

Infrastructure Benchmarking Capabilities

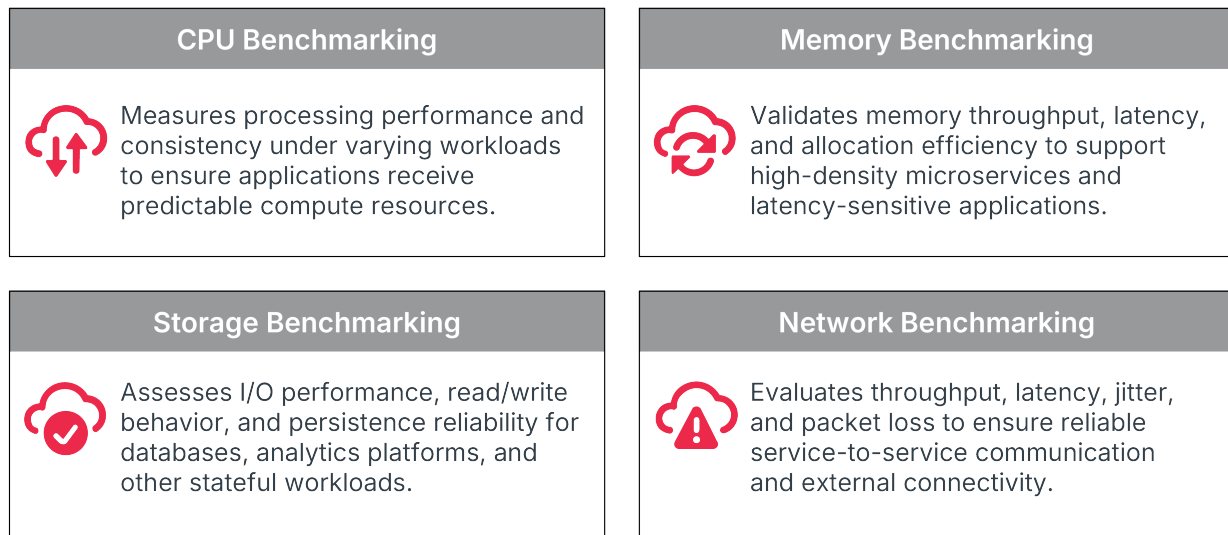


Figure 1. Chart displaying infrastructure benchmarking capabilities

Validating Infrastructure Readiness

Cloud-native infrastructure benchmarking establishes performance and resiliency baselines across Kubernetes environments. Enterprises can run configurable workload scenarios to validate how clusters behave under stress, confirm Kubernetes scheduling, failover, and recovery mechanisms, and compare configurations using A/B or shadow testing.

Comprehensive metrics collection and real-time visualization simplify analysis and accelerate informed decision-making for platform engineering, SRE, and cloud operations teams.

Key Features and Benefits

- **Infrastructure Benchmarking:** Benchmark CPU, memory, storage, and networking performance across major Kubernetes distributions and enterprise cloud platforms to validate readiness before application deployment.
- **Purpose-Built Resource Testing:** Granular measurement of compute, memory, storage I/O, and network behavior under real-world enterprise workloads.
- **Comprehensive Metrics Collection:** Deep Kubernetes-layer visibility into utilization, latency, and allocation efficiency to identify bottlenecks and optimization opportunities.
- **Visualization & Monitoring:** Real-time dashboards powered by Grafana, Prometheus, and Kafka to accelerate analysis and troubleshooting.
- **Configurable Scenarios:** User-defined workloads and parameters aligned to enterprise application profiles and service expectations.
- **Embedded Workflow Engine:** Drag-and-drop workflows combining CPU, memory, storage, and networking actions into repeatable, automated benchmark scenarios.
- **CI/CD Integration:** Native integration with DevOps pipelines to support continuous infrastructure validation as environments evolve.
- **Dynamic Allocation Validation:** Confirms Kubernetes can dynamically allocate, migrate, and recover workloads across hybrid and multi-cloud environments.

CPU Benchmarking in Action

This CPU benchmarking example shows how Landslide automates performance testing in a Kubernetes cluster. Using the Cloud Engine, operators can drive workloads across worker nodes, measure CPU utilization, and collect detailed Kubernetes-layer metrics.

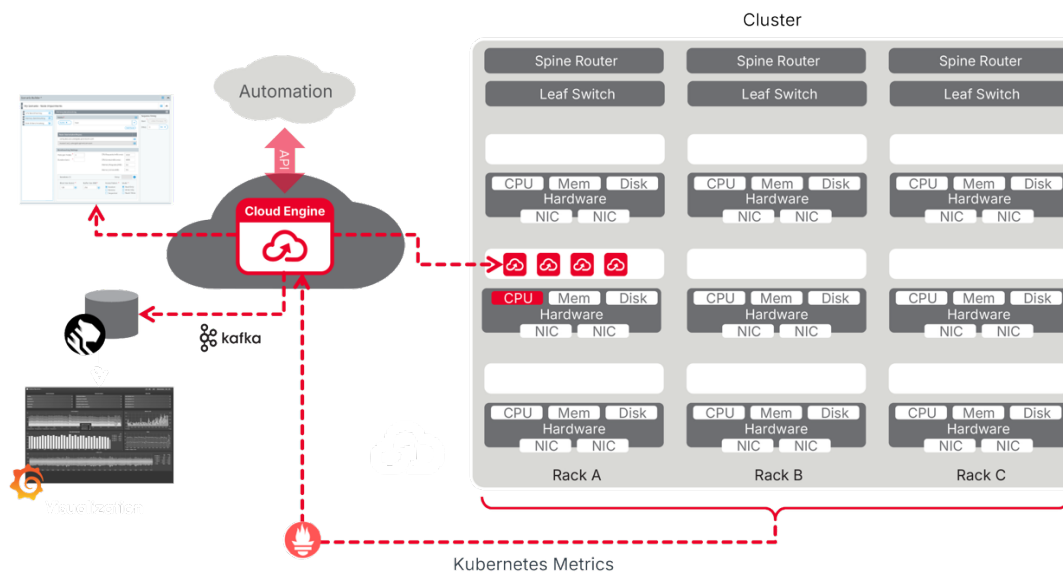


Figure 2. CPU benchmarking chart displaying performance automation

The same approach can benchmark memory, storage, and networking by generating data flows between nodes and storage systems. Metrics are streamed through Kafka and Prometheus into Grafana dashboards, delivering real-time visibility of infrastructure performance.

Technical Specifications

Landslide Cloud Engine is built as a tiered, microservices-based application designed to validate the readiness of Kubernetes environments.

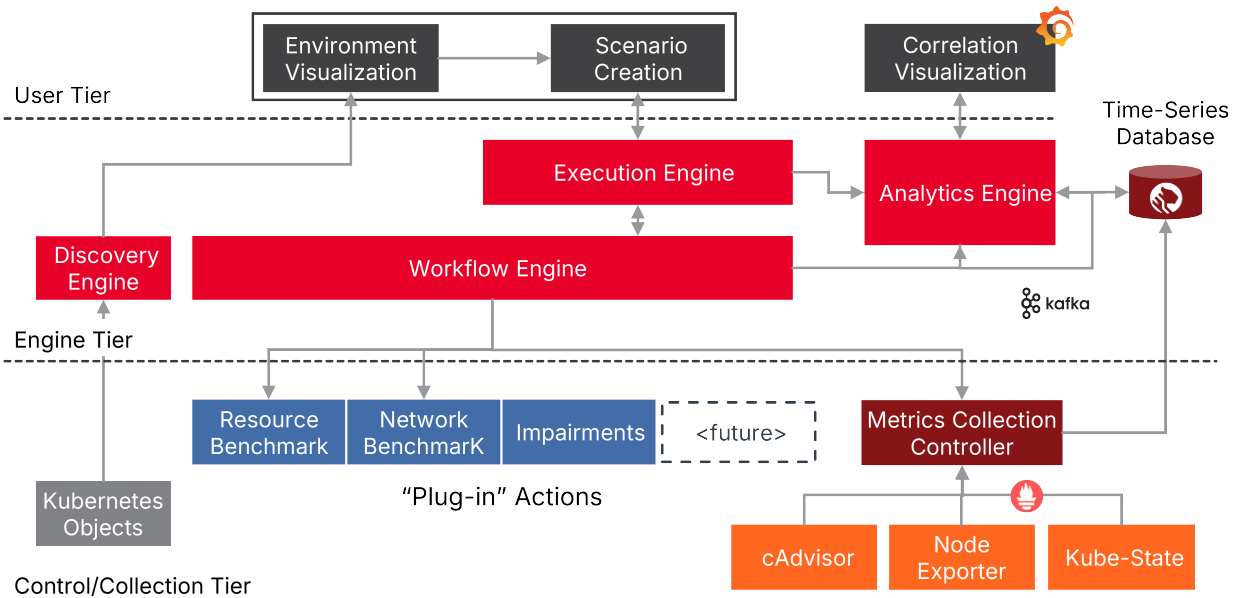


Figure 3. Flow chart displaying Kubernetes environments

Enterprise-Ready Architecture

Landslide Cloud Engine is delivered as a modular, microservices-based application designed to validate Kubernetes environments at scale. At the user tier, teams interact through intuitive interfaces for environment visualization, scenario creation, and correlation dashboards.

At the engine and control tiers, workflow and execution engines orchestrate benchmarking scenarios using plug-in actions for resource testing and telemetry collection. Performance data is gathered via Kubernetes exporters such as cAdvisor, Node Exporter, and kube-state-metrics (KSM), streamed through Kafka into an analytics engine, stored in a time-series database, and visualized through Grafana dashboards.

This flexible architecture supports extensibility, automation, and long-term investment protection for enterprise cloud platforms.

Ordering Information

Part Number	Description
L-CNFV-BASE	Validation Platform/Application (3 Scenarios)
L-CNIB-CPUMEM	Cloud-Native Infrastructure Benchmarking CPU/Memory
L-CNIB-CPUMEM-ANLT	Cloud-Native Infrastructure CPU/Mem Analytics
L-CNIB-STOR	Cloud-Native Infrastructure Benchmarking Storage
L-CNIB-STOR-ANLT	Cloud-Native Infrastructure Storage Analytics

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.



This information is subject to change without notice. © Keysight Technologies, 2026, Published in USA, June 1, 2026, 3126-1234.EN