Listicle

6 Key Steps for Kubernetes Troubleshooting

splunk>
a CISCO company

Imagine navigating a dense forest with only a flickering flashlight — the batteries are low, the sky is darkening, and you don't have a navigation system. You sense that you're off course, but without a reliable map or compass, you're left wandering aimlessly. That's what troubleshooting inside Kubernetes environments feels like without observability — disorienting, reactive, and uncertain.

Kubernetes has progressed beyond a simple container orchestrator into a comprehensive system for managing dynamic, distributed, and multicloud workloads. While this evolution has unlocked new efficiencies and scalability, it has also introduced unique troubleshooting challenges.

To effectively troubleshoot your Kubernetes environment, your teams need a structured approach that addresses the entire technology stack, from nodes to applications. You also need a service inventory so you know what exists and what to monitor. With a **comprehensive observability** solution that integrates infrastructure monitoring, APM, and logging capabilities, you will be able to continuously analyze KPIs and detect failure trends — *before* they impact production.

Read on to discover six essential steps for Kubernetes troubleshooting.

1. Identify the problem scope

(node, pod, service layer, network) to narrow down root causes quickly. Failures can manifest anywhere inside a Kubernetes environment, and each layer presents distinct challenges, making root cause analysis more complex than traditional infrastructure. Observability will help you nail down where the problem is coming from, speeding up resolution.

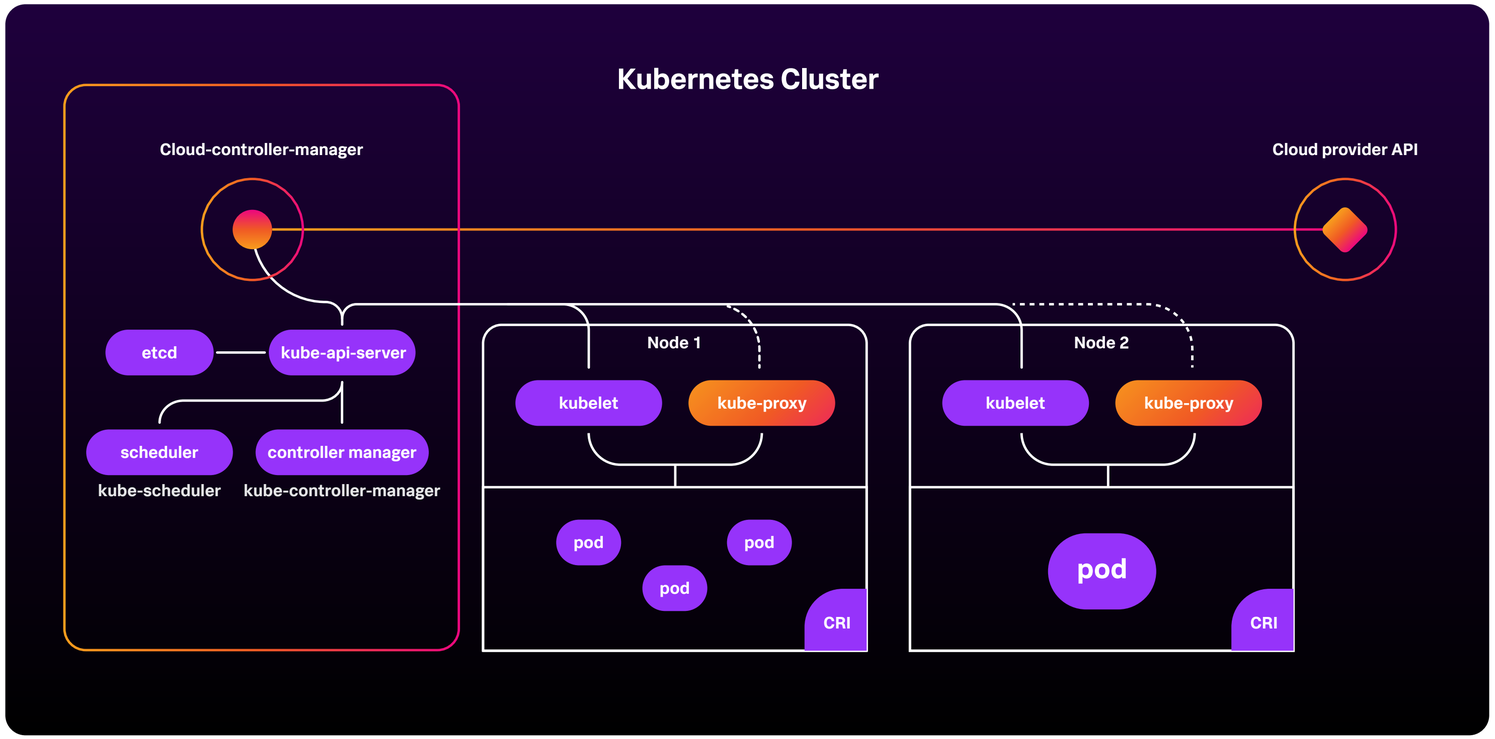

2. Conduct proactive resource allocation

using historical resource utilization trends to prevent resource exhaustion, optimize scaling cost, and mitigate performance bottlenecks. With an observability solution that offers data for post-incident review and historical insights, your team can investigate past incidents and analyze resource usage patterns to optimize allocations and prevent recurring issues.

3. Surface and analyze anomalies in key metrics

to detect early warning signs of instability, performance degradation, or resource contention before they impact service reliability. Proactive detection is key.

4. Use AI-driven insights and machine learning-powered anomaly

detection to minimize reliance on tribal knowledge, surface hidden patterns, and accelerate root cause analysis. This capability will enable users of all skill levels to operate with expert insight.

5. Check for misconfigurations or issues with downstream services

that may be contributing to performance degradation.

A comprehensive observability solution will help you ensure that dependencies, such as downstream services, data stores, and 3rd-party services are functioning as expected. It will also help you audit and validate configurations regularly, discovering when new dependencies have been added, to faster identify issues and streamline troubleshooting.

6. Implement vendor-neutral instrumentation with OpenTelemetry

to unify metrics, logs, and traces across clusters and clouds, ensuring consistent visibility and preventing vendor lock-in. Cloud-native monitoring tools are often vendor-specific and siloed, making it difficult to maintain a unified view — whether for a single application spanning clouds or separate clusters operating in different environments. Leveraging OpenTelemetry will help with this by permitting relevant data to be sent anywhere and analyzed together.

The key to troubleshooting Kubernetes environments faster and more accurately is to train your team on observability tools. Just as modern navigation systems go beyond simple maps, observability equips teams with real-time, actionable insights to improve visibility, address issues with precision, and keep Kubernetes environments running at peak performance.

To get the full Kubernetes troubleshooting roadmap, check out our in-depth ebook, **Troubleshooting Kubernetes Environments: Shifting Gears on your Journey to Digital Resilience**.

splunk> a CISCO company

25_CMP_listicle_6-tips-for-kubernetes-troubleshooting_v10