

Monitoring Kubernetes clusters, applications, and infrastructure

The power of integrated tools

 **NetApp**



Contents

➔	Kubernetes adoption challenges	3
➔	Cloud Insights: Intuitive Kubernetes cluster visualization	5
➔	Cloud Insights: Node and pod monitoring	8
➔	Cloud Insights: Kubernetes, application, and infrastructure troubleshooting	10
➔	Getting started with Cloud Insights	11



Since its release in 2015, Kubernetes has dramatically changed the IT landscape. But years later, the question remains: What’s the best way to monitor Kubernetes, the infrastructure that it touches, and the applications that run within its clusters?

For someone who’s accustomed to managing traditional IT infrastructure, Kubernetes introduces another “moving part”—and significant complexity. It also promises to relieve the burden of workload placement. The task has grown from monitoring servers, virtual machines, networks, and file servers to now also dealing with new concepts of clusters, pods, containers, name-spaces, labels, persistent volumes, persistent volume claims, auto-scaling, and more. Furthermore, the new Kubernetes concepts have real-world impacts on the operation of traditional infrastructure components.

This e-book highlights some essential guidelines that will help you succeed with Kubernetes. It also explores how NetApp® Cloud Insights enables nonexperts to effectively monitor Kubernetes, the underlying infrastructure, and the applications and services that run on Kubernetes.



Kubernetes adoption challenges

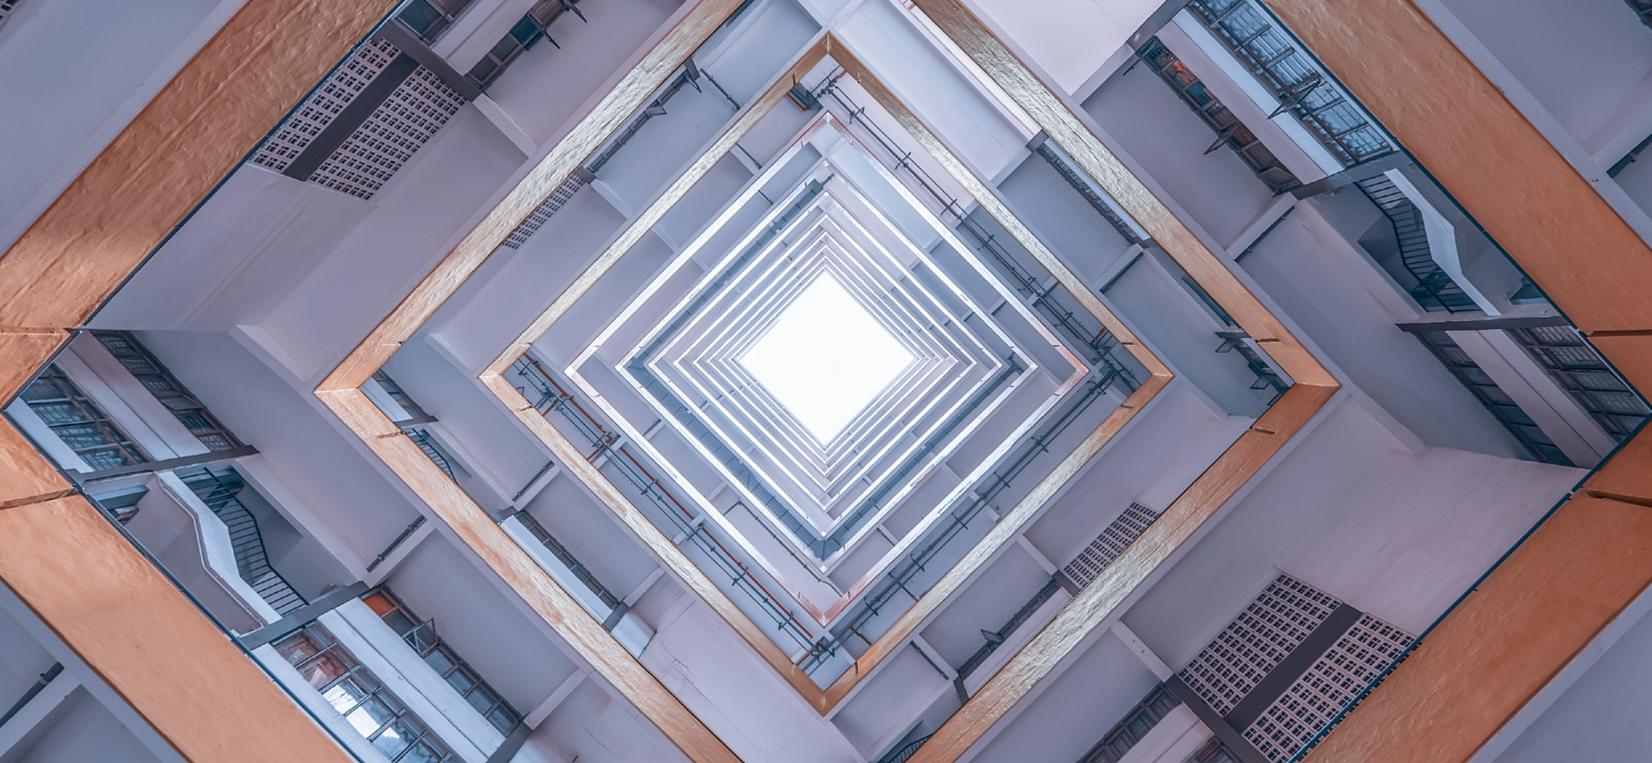
There was a time when monitoring all the metrics to keep a system healthy was a task for multiple teams. First, there was the application team, with its dashboards and custom metrics. Then there were operations teams handling multiple servers and allocating resources to each service according to its needs. Lastly, the IT teams monitored physical machines and hardware changes. This hierarchical approach was effective but not very agile.

Adding Kubernetes to your IT landscape brings with it the need for new skills to navigate and to tame a new control plane. Although hiring an expert or learning the appropriate skills yourself can be a challenge, Kubernetes has a logical foundation that's easy to grasp:

- Application components are arranged as containers.
- Related or dependent containers are defined within pods (along with their storage needs).
- Configuration is externally applied to pods.
- Everything is submitted to the Kubernetes cluster, which distributes the execution of the pods across the available nodes and keeps them running.

Even with Kubernetes mastery, you still must monitor and maintain infrastructure relationships and interactions. Each node of a Kubernetes cluster is a virtual machine or a bare-metal computer. Each pod might mount storage from file servers. It's important to know that the operating system and components of the server system that are acting as Kubernetes nodes are functional. The related storage systems must be healthy, available, and running smoothly with appropriate quality-of-service classes and sufficient capacity available.

And there are the applications to consider, too. Kubernetes provides namespaces and labels to organize and to manage application components. Namespaces act as isolators for the naming of Kubernetes components and are convenient modules for application management. You can use labels to define applications (or their subcomponents, such as microservices) within or across namespaces. Monitored metrics at the application level map to Kubernetes pods that are running on nodes; those nodes, in turn, run on virtual machines or servers and may mount external storage. There are connectivity chains from applications to the infrastructure, with Kubernetes placed squarely in the middle.



Therefore, the additional adoption challenge of Kubernetes is to integrate its monitoring with the applications and the infrastructure. Problems can surface at any layer in the overall system, and time to resolution is improved by access to tools that stitch everything together. But understanding Kubernetes is not the end of the journey.

We at NetApp can't help you choose a Kubernetes distribution and version, and we can't help with the initial configuration, node selection, or distribution of your workloads. But we can help you monitor Kubernetes and the associated infrastructure with NetApp Cloud Insights, a cloud-hosted, software-as-a-service solution for on-premises and cloud-based systems.

Cloud storage for Kubernetes

Kubernetes accelerates the development of new cloud-native applications and enables the modernization of existing ones. Although the containers that these applications execute are ephemeral, their data is not.

With NetApp cloud storage services, you get a persistent, fully managed, cost-effective, and secure storage environment for Kubernetes. By delivering elasticity, scalable performance, and shared file access across multiple cloud providers and regions, these services promote consume storage according to application demands.

NetApp Trident is a fully supported open-source project that automates the provisioning of Kubernetes persistent storage for stateful applications. This automation provides for dynamic increases or decreases of performance based on workload.

Find out more:

- [NetApp Cloud Volumes Service for AWS](#)
- [NetApp Cloud Volumes ONTAP® for Kubernetes](#)
- [NetApp Trident](#)



Cloud Insights: Intuitive Kubernetes cluster visualization

Knowing the importance of monitoring Kubernetes and applications and infrastructure, let's first address Kubernetes.

Any Kubernetes cluster consists of nodes that run on virtual machines or servers. Across those nodes, the Kubernetes software coordinates the execution of pods that define the container-packaged software to be run. Each pod exists within one of many namespaces and may be labeled with additional metadata that can be used for specialized management. Overall resource utilization across the cluster and any warnings or alerts that correspond to nodes and pods should surface in a clusterwide visualization. At this highest level of concern, these things are what you need to know about your Kubernetes cluster.

NetApp Cloud Insights gives you a sortable view of all your Kubernetes clusters, containing node, pod, and namespace information. Furthermore, any warnings and alerts within a cluster, that are attached to the nodes or the pods, grab your attention so that you can quickly begin troubleshooting, triage, and resolution. You can dive down into cluster details from this comprehensive view.

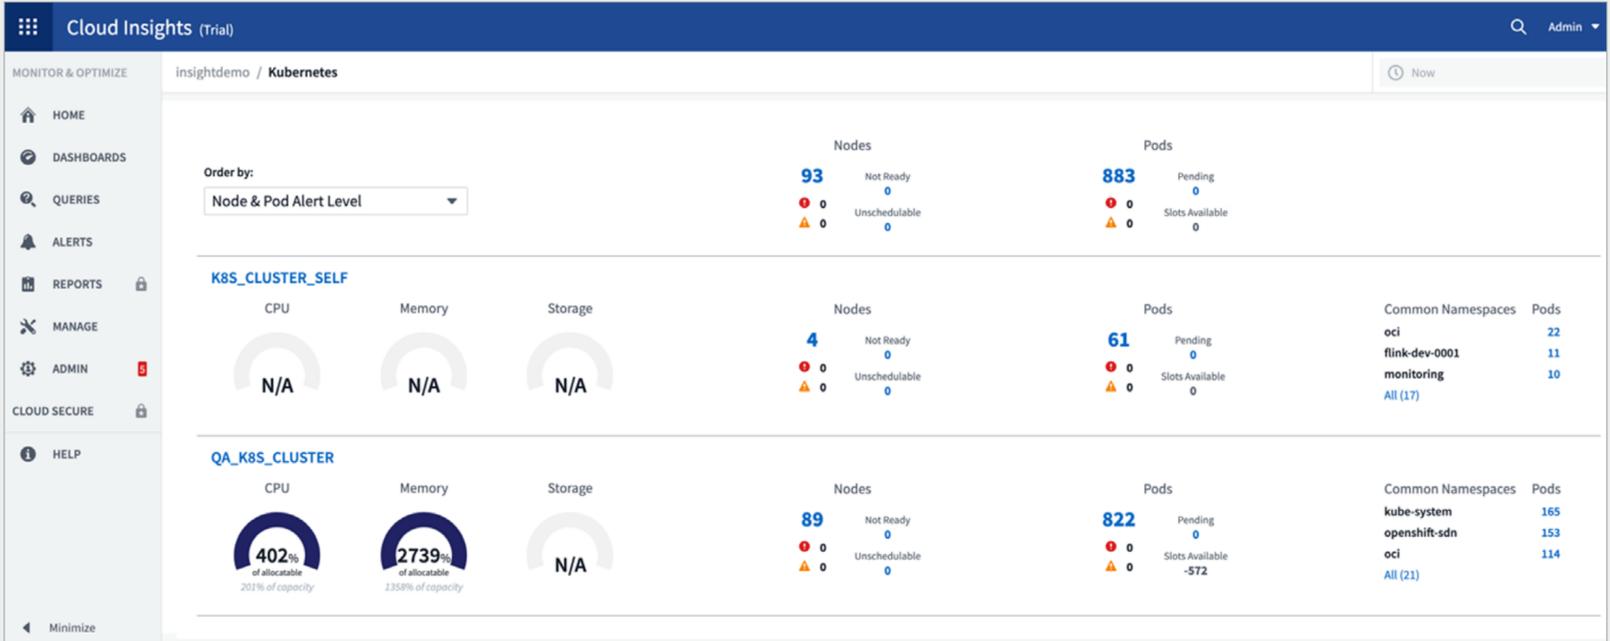


Figure 1) Overview of monitored Kubernetes clusters—Node & Pod Alert Level.

Cloud Insights topology mapping for Kubernetes clusters enables nonexperts to identify issues and dependencies down to the container level. It centralizes cluster, node, and application metrics alongside metrics from resources and services.

When you dive down to view a single Kubernetes cluster, it's important that you can quickly understand whether the cluster is healthy. And if it's not healthy, you must know where to direct your attention to quickly identify and triage issues. In a properly organized display like with Cloud Insights, you don't have to be a Kubernetes expert to understand the overall state of the cluster and where to go to find more detail.

In a cluster-centric view, a display of aggregated information is critical. The cumulative resources of a Kubernetes cluster are summations across the nodes of the pods (showing node states) and aggregate the utilization of compute resources. By viewing resource utilization information, you can better understand how much additional load a cluster can handle. The excess resources provide overhead capabilities for handling automatically scaled-out pods, handling additional workloads, or rebalancing pods if a subset of nodes experience trouble.

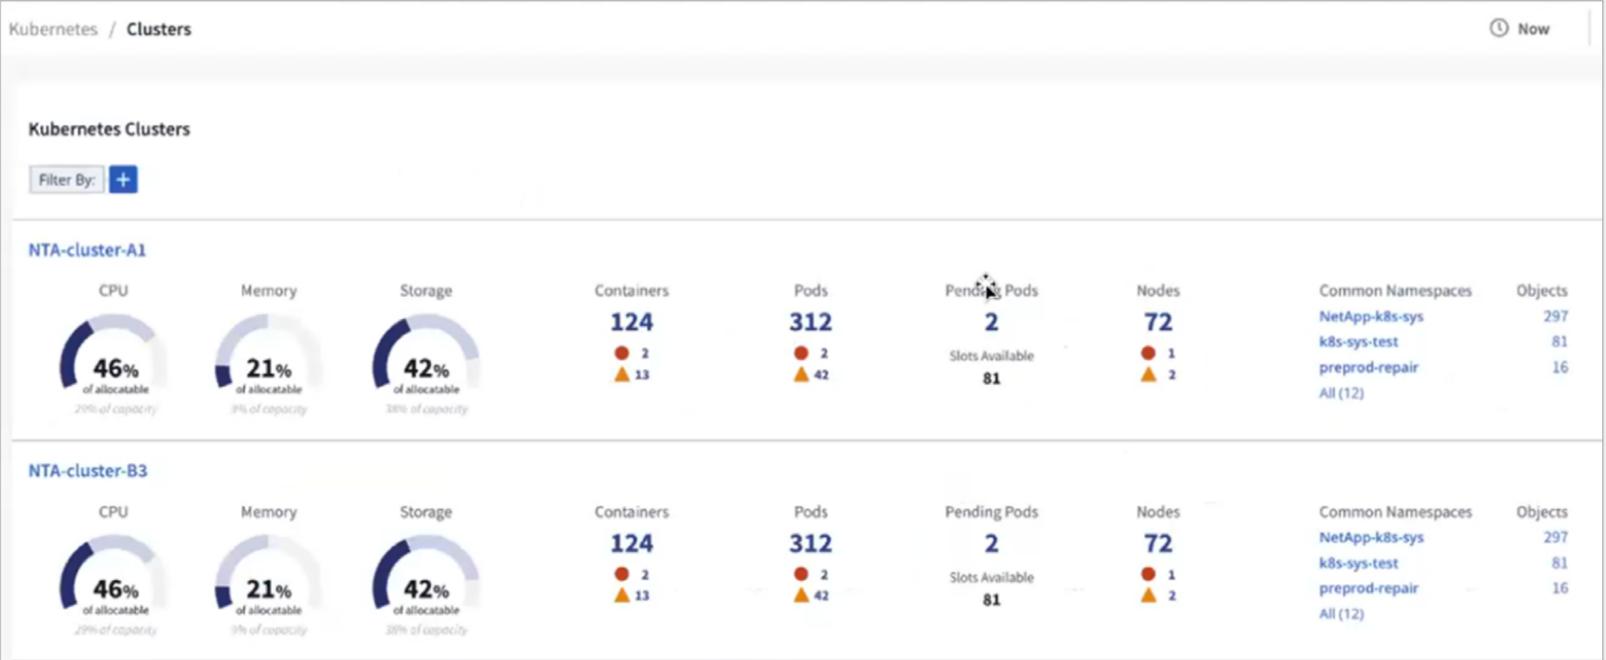


Figure 2) Overview of monitored Kubernetes clusters.

An ability to identify pods with active alerts, and the nodes that host troubled pods, is the first step in taking corrective action. NetApp Cloud Insights displays the pertinent cluster measures of node count; pods and their states of healthy, alerting, or pending; and aggregate resource utilization and clusterwide top consumers. The “Kubernetes wheel” in the lower-left corner of Figure 3 shows each cluster node in an arc, with the number of pods that are deployed in the node (with status information color-coded). You can click through for additional details across the display.

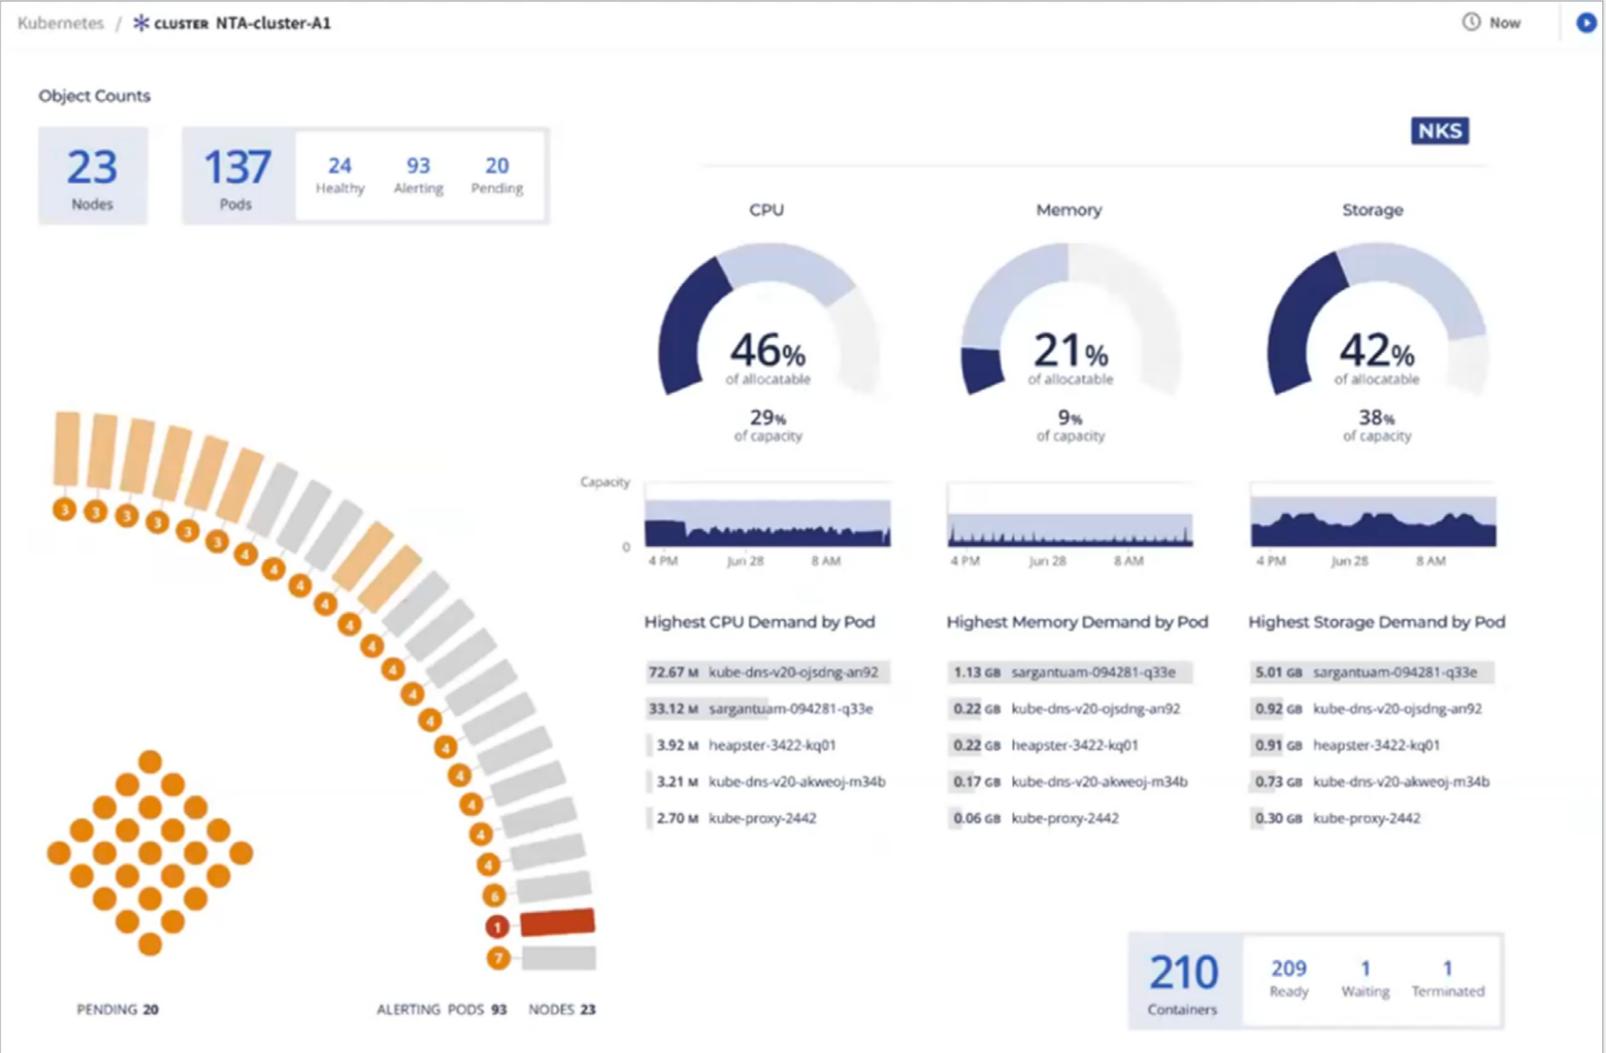


Figure 3) Exploded view of Kubernetes cluster information.

Cloud Insights: Node and pod monitoring

Kubernetes monitoring can be further segmented into node and pod monitoring.

Node monitoring means evaluating the health of the entire set of nodes that make up a Kubernetes cluster. Some key monitoring metrics are:

- **Node resources.** Each Kubernetes node has its own resource utilization to track, such as CPU usage, memory page faults, and available disk space. By tracking this information, you can see any potential pressure on resources and see the overhead within the cluster (on a per-node basis) that's available for autoscaling and additional workloads.
- **Node alerts.** Any alerts that concern nodes and their ability to host pods affect the overall functioning of your cluster. To keep everything running as smoothly as possible, attend to these alerts right away.
- **Pods deployed.** The number of pods running on each node exposes cluster vulnerabilities and enables you to see how Kubernetes has handled each deployment. For example, a key system should have more than one pod running in different nodes and should be balanced by a load balancer.

If no services are running, it doesn't matter whether a cluster is available and healthy. Pod-specific metrics can be divided into two categories: pod and container metrics.

- **Pod metrics.** A pod's Kubernetes metrics are the metrics that the orchestrator extracts. With these metrics, your administrator can track ongoing deployments or see how many instances are running. If the number of running pods is lower than expected, it may indicate that the cluster is out of resources.
- **Container metrics.** Much like node resources, containers, too, have their own resource utilization metrics to track. And these resources can also be constrained. Knowing how allocated resources are being used at the container level is crucial for you to identify potential problems and bottlenecks.

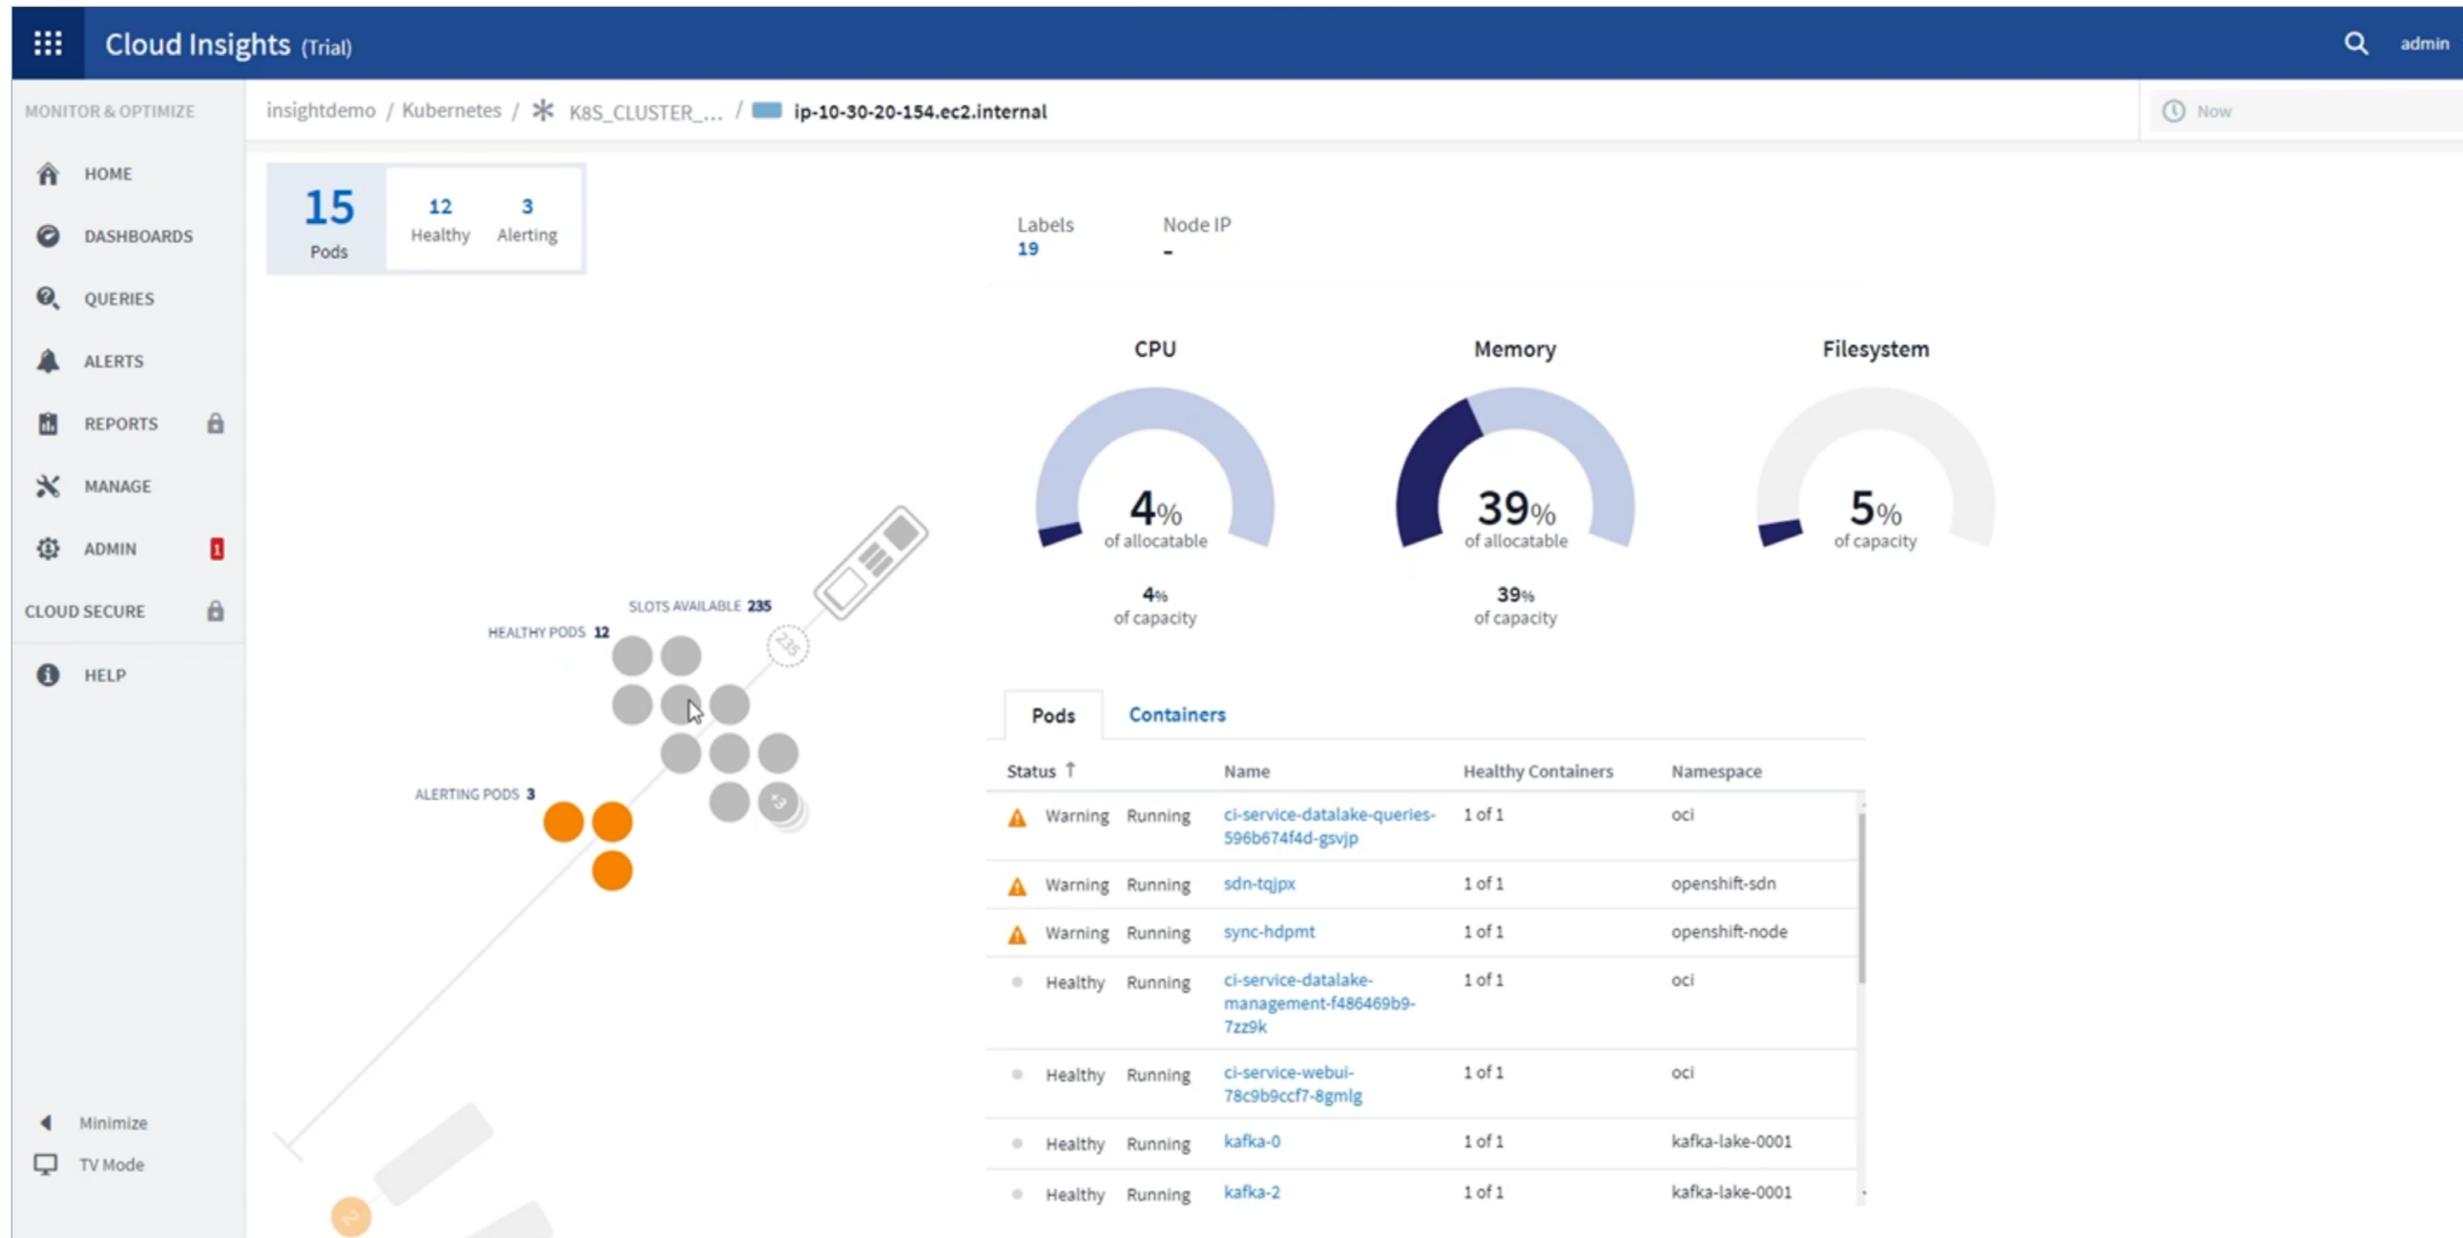


Figure 4) Cluster view

Cloud Insights: Kubernetes, application, and infrastructure troubleshooting

A NetApp Cloud Insights dashboard can incorporate multiple views, sharing information across the entire IT stack, from the application, through Kubernetes, and down to the physical resources that are used. For example, an application owner can create a dashboard in Cloud Insights that shows key performance indicators (KPIs) for a web application. An infrastructure administrator or other team member might not understand the top-level KPIs, but other KPIs in the same dashboard enable them to monitor and to troubleshoot the application.

This connect-the-dots troubleshooting approach enables your teams to monitor and to manage everything that's associated with your Kubernetes environment. Your teams can drill down to see an intuitive view of any single cluster, application, or infrastructure item that's reporting problems and get actionable information.

Because your administrators can visualize and alert on the entire application stack in one place, Cloud Insights is an invaluable resource for your operations team. A systemwide issue might not be due to a single problem; the cause might be a combination of problems. For example, with Cloud Insights, you can correlate front-end response times with back-end server loads and with the number of instances of a container so that you can quickly identify and resolve an issue.

If you're moving to hybrid cloud operations that span both private and public clouds or if you're running applications on Kubernetes in more than one public cloud, Cloud Insights makes monitoring and management easy. You get visibility and actionable information across all your environments, enabling your teams to monitor your cloud operations more effectively.

Getting started with Cloud Insights

NetApp Cloud Insights incorporates specific Kubernetes tools and views so that your teams can effectively monitor and troubleshoot your Kubernetes environments on premises and in the public cloud. The Kubernetes cluster explorer provides an overview of all the Kubernetes clusters in your environment, and topology mapping gives you a detailed visual overview of each Kubernetes cluster.



To see Cloud Insights in action for yourself, register for a 30-day free trial.

Go to [NetApp Cloud Central](#) to learn more about NetApp Cloud Insights and to start your free trial.

And if you're ready to dive even deeper, NetApp has several Kubernetes resources that are targeted to different needs.

More Cloud Insights and Kubernetes resources

- [NetApp Cloud Insights on Cloud Central](#)
- [NetApp Kubernetes hub](#)
- [Cloud Insights app modernization e-book](#)
- [NetApp Trident access](#)
- [NetApp Astra™ information](#)
- [Ocean from Spot by NetApp](#)



About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services, and applications to the right people— anytime, anywhere.

