Kubernetes

Data Infrastructure Insights

NetApp
February 03, 2026

This PDF was generated from https://docs.netapp.com/us-en/data-infrastructure-insights/kubernetes_landing_page.html on February 03, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Kubernetes

Kubernetes Cluster Overview

The Data Infrastructure Insights Kubernetes Explorer is a powerful tool for displaying the overall health and usage of your Kubernetes clusters and allows you to easily drill down into areas of investigation.

Clicking on **Dashboards > Kubernetes Explorer** opens the Kubernetes Cluster list page. This overview page contains table of the Kubernetes clusters on your tenant.

Cluster list

The cluster list displays the following information for each cluster on your tenant:

- Cluster **Name**. Clicking on a cluster name will open the **detail page** for that cluster.
- **Saturation** percentages. Overall Saturation is the highest of CPU, Memory, or Storage Saturation.
- Number of **Nodes** in the cluster. Clicking this number will open the Node list page.
- Number of **Pods** in the cluster. Clicking this number will open the Pod list page.
- Number of **Namespaces** in the cluster. Clicking this number will open the Namespace list page.
- Number of **Workloads** in the cluster. Clicking this number will open the Workload list page.

Refining the Filter

When you are filtering, as you begin typing you are presented with the option to create a **wildcard filter** based on the current text. Selecting this option will return all results that match the wildcard expression. You can also create **expressions** using NOT or AND, or you can select the "None" option to filter for null values in the field.

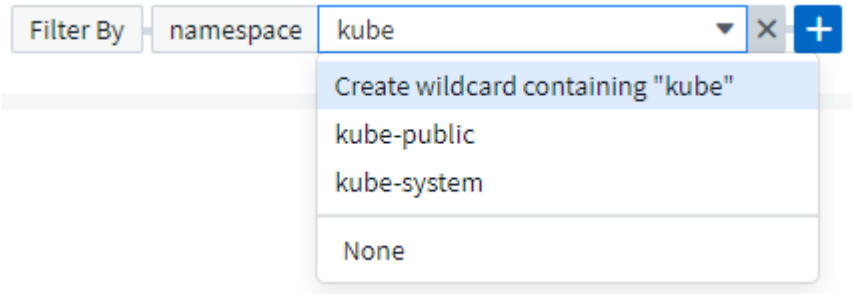

Filters based on wildcards or expressions (e.g. NOT, AND, "None", etc.) display in dark blue in the filter field. Items that you select directly from the list are displayed in light blue.

Kubernetes filters are contextual, meaning for example that if you are on a specific node page, the pod_name filter only lists pods related to that node. Moreover, if you apply a filter for a specific namespace, then the pod_name filter will list only pods on that node *and* in that namespace.

Note that Wildcard and Expression filtering works with text or lists but not with numerics, dates or booleans.

Before Installing or Upgrading the NetApp Kubernetes Monitoring Operator

Read this information before installing or upgrading the Kubernetes Monitoring Operator.

Component	Requirement
Kubernetes version	Kubernetes v1.20 and above.
Kubernetes distributions	AWS Elastic Kubernetes Service (EKS) Azure Kubernetes Service (AKS) Google Kubernetes Engine (GKE) Red Hat OpenShift Rancher Kubernetes Engine (RKE) VMware Tanzu
Linux OS	Data Infrastructure Insights does not support nodes that are running with Arm64 architecture. Network monitoring: must be running Linux kernel version 4.18.0 or above. Photon OS is not supported.
Labels	Data Infrastructure Insights supports monitoring of Kubernetes nodes that are running Linux, by specifying a Kubernetes node selector that looks for the following Kubernetes labels on these platforms: Kubernetes v1.20 and above: Kubernetes.io/os = linux Rancher + cattle.io as orchestration/Kubernetes platform: cattle.io/os = linux
Commands	The curl and kubectl commands must be available.; for best results, add these commands to the PATH.
Connectivity	kubectl cli is configured to communicate with the target K8s cluster, and have Internet connectivity to your Data Infrastructure Insights environment. If you are behind a proxy during installation, follow the instructions in the Configuring Proxy Support section of the Operator installation. For accurate audit and data reporting, synchronize the time on the Agent machine using Network Time Protocol (NTP) or Simple Network Time Protocol (SNTP).

Component	Requirement
Other	If you are running on OpenShift 4.6 or higher, you must follow the OpenShift Instructions in addition to ensuring these pre-requisites are met.
API Token	If you are re-deploying the Operator (i.e. you are updating or replacing it), there is no need to create a new API token; you can re-use the previous token.

Important Things to Note Before You Start

If you are running with a proxy, have a custom repository, or are using OpenShift, read the following sections carefully.

Also read about Permissions.

Configuring Proxy Support

There are two places where you may use a proxy on your tenant in order to install the NetApp Kubernetes Monitoring Operator. These may be the same or separate proxy systems:

- Proxy needed during execution of the installation code snippet (using "curl") to connect the system where the snippet is executed to your Data Infrastructure Insights environment
- Proxy needed by the target Kubernetes cluster to communicate with your Data Infrastructure Insights environment

If you use a proxy for either or both of these, to install the NetApp Kubernetes Operating Monitor you must first ensure that your proxy is configured to allow good communication to your Data Infrastructure Insights environment. For example, from the servers/VMs from which you wish to install the Operator, you need to be able to access Data Infrastructure Insights and be able to download binaries from Data Infrastructure Insights.

For the proxy used to install the NetApp Kubernetes Operating Monitor, before installing the Operator, set the *http_proxy/https_proxy* environment variables. For some proxy environments, you may also need to set the *no_proxy environment* variable.

To set the variable(s), perform the following steps on your system **before** installing the NetApp Kubernetes Monitoring Operator:

1. Set the *https_proxy* and/or *http_proxy* environment variable(s) for the current user:
   a. If the proxy being setup does not have Authentication (username/password), run the following command:

```
export https_proxy=<proxy_server>:<proxy_port>
```

   b. If the proxy being setup does have Authentication (username/password), run this command:

```
export
http_proxy=<proxy_username>:<proxy_password>@<proxy_server>:<proxy_po
rt>
```

For the proxy used for your Kubernetes cluster to communicate with your Data Infrastructure Insights environment, install the NetApp Kubernetes Monitoring Operator after reading all of these instructions.

Configure the proxy section of AgentConfiguration in operator-config.yaml before deploying the NetApp Kubernetes Monitoring Operator.

```
agent:
  ...
  proxy:
    server: <server for proxy>
    port: <port for proxy>
    username: <username for proxy>
    password: <password for proxy>

    # In the noproxy section, enter a comma-separated list of
    # IP addresses and/or resolvable hostnames that should bypass
    # the proxy
    noproxy: <comma separated list>

    isTelegrafProxyEnabled: true
    isFluentbitProxyEnabled: <true or false> # true if Events Log enabled
    isCollectorsProxyEnabled: <true or false> # true if Network
Performance and Map enabled
    isAuProxyEnabled: <true or false> # true if AU enabled
  ...
...
```

Using a custom or private docker repository

By default, the NetApp Kubernetes Monitoring Operator will pull container images from the Data Infrastructure Insights repository. If you have a Kubernetes cluster used as the target for monitoring, and that cluster is configured to only pull container images from a custom or private Docker repository or container registry, you must configure access to the containers needed by the NetApp Kubernetes Monitoring Operator.

Run the “Image Pull Snippet” from the NetApp Monitoring Operator install tile. This command will log into the Data Infrastructure Insights repository, pull all image dependencies for the operator, and log out of the Data Infrastructure Insights repository. When prompted, enter the provided repository temporary password. This command downloads all images used by the operator, including for optional features. See below for which features these images are used for.

Core Operator Functionality and Kubernetes Monitoring

- netapp-monitoring
- kube-rbac-proxy
- kube-state-metrics
- telegraf
- distroless-root-user

Events Log

- fluent-bit
- kubernetes-event-exporter

Network Performance and Map

- ci-net-observer

Push the operator docker image to your private/local/enterprise docker repository according to your corporate policies. Ensure that the image tags and directory paths to these images in your repository are consistent with those in the Data Infrastructure Insights repository.

Edit the monitoring-operator deployment in operator-deployment.yaml, and modify all image references to use your private Docker repository.

```
image: <docker repo of the enterprise/corp docker repo>/kube-rbac-
proxy:<kube-rbac-proxy version>
image: <docker repo of the enterprise/corp docker repo>/netapp-
monitoring:<version>
```

Edit the AgentConfiguration in operator-config.yaml to reflect the new docker repo location. Create a new imagePullSecret for your private repository, for more details see *https://kubernetes.io/docs/tasks/configure-pod-container/pull-image-private-registry/*

```
agent:
  ...
  # An optional docker registry where you want docker images to be pulled
from as compared to CI's docker registry
  # Please see documentation for
xref:{relative_path}task_config_telegraf_agent_k8s.html#using-a-custom-or-
private-docker-repository[using a custom or private docker repository].
  dockerRepo: your.docker.repo/long/path/to/test
  # Optional: A docker image pull secret that maybe needed for your
private docker registry
  dockerImagePullSecret: docker-secret-name
```

OpenShift Instructions

If you are running on OpenShift 4.6 or higher, you must edit the AgentConfiguration in *operator-config.yaml* to enable the *runPrivileged* setting:

```
# Set runPrivileged to true SELinux is enabled on your kubernetes nodes
runPrivileged: true
```

Openshift may implement an added level of security that may block access to some Kubernetes components.

Permissions

If the cluster you are monitoring contains Custom Resources which do not have a ClusterRole which aggregates to view, you will need to manually grant the operator access to these resources to monitor them with Event Logs.

1. Edit *operator-additional-permissions.yaml* before installing, or after installing edit the resource *ClusterRole/<namespace>-additional-permissions*
2. Create a new rule for the desired apiGroups and resources with the verbs ["get", "watch", "list"]. See https://kubernetes.io/docs/reference/access-authn-authz/rbac/
3. Apply your changes to the cluster

```
kubectl -n <NAMESPACE> delete agent netapp-ci-agent-monitoring-netapp
kubectl delete crd agents.monitoring.netapp.com
kubectl -n <NAMESPACE> delete role netapp-ci-agent-manager netapp-ci-kube-
state-metrics
kubectl delete clusterrole netapp-ci-<NAMESPACE>-additional-permissions
netapp-ci-<NAMESPACE>-agent-manager netapp-ci-<NAMESPACE>-agent-secret
netapp-ci-<NAMESPACE>-agent-view-plus netapp-ci-<NAMESPACE>-change-
observer-view-plkubectl get us netapp-ci-<NAMESPACE>-kube-state-metrics
netapp-ci-<NAMESPACE>-net-observerkubectl
kubectl delete clusterrolebinding netapp-ci-<NAMESPACE>-additional-
permissions netapp-ci-<NAMESPACE>-agent-manager netapp-ci-<NAMESPACE>-
agent-secret netapp-ci-<NAMESPACE>-agent-view netapp-ci-<NAMESPACE>-agent-
view-plus netapp-ci-<NAMESPACE>-change-observer-additional-permissions
netapp-ci-<NAMESPACE>-change-observer-secret netapp-ci-<NAMESPACE>-change-
observer-view netapp-ci-<NAMESPACE>-change-observer-view-plus netapp-ci-
<NAMESPACE>-event-exporter netapp-ci-<NAMESPACE>-kube-state-metrics
netapp-ci-<NAMESPACE>-net-observer
kubectl delete netapp-ci-<NAMESPACE>-psp-nkmo
kubectl delete ns <NAMESPACE>
```

Kube-state-metrics Counters

Refer to the Kube-state-metrics documentation for more information on the metric counters.

Use the following links to access information for these kube state metrics counters:

1. ConfigMap Metrics
2. DaemonSet Metrics
3. Deployment Metrics
4. Ingress Metrics
5. Namespace Metrics
6. Node Metrics
7. Persistent Volume Metrics

8. Persistant Volume Claim Metrics
9. Pod Metrics
10. ReplicaSet metrics
11. Secret metrics
12. Service metrics
13. StatefulSet metrics

The table below lists the possible options for the *AgentConfiguration* file:

Component	Option	Description
agent		Configuration options that are common to all components that the operator can install. These can be considered as "global" options.
	dockerRepo	A dockerRepo override to pull images from customer's private docker repositories instead of the Data Infrastructure Insights docker repository. Default is the Data Infrastructure Insights docker repository.
	dockerImagePullSecret	Optional: A secret for the customer's private repository.
	clusterName	Free text field that uniquely identifies a cluster across all customer clusters. This should be unique across a Data Infrastructure Insights tenant. Default is what the customer enters in the UI for the "Cluster Name" field.
	proxy Format: proxy: server: port: username: password: noProxy: isTelegrafProxyEnabled: isAuProxyEnabled: isFluentbitProxyEnabled: isCollectorProxyEnabled:	Optional to set proxy. This is usually the customer's corporate proxy.
telegraf		Configuration options that can customize the telegraf installation of the Operator
	collectionInterval	Metrics collection interval, in seconds (Max=60s)
	dsCpuLimit	CPU Limit for telegraf ds
	dsMemLimit	Memory limit for telegraf ds
	dsCpuRequest	CPU request for telegraf ds

Component	Option	Description
	dsMemRequest	Memory request for telegraf ds
	rsCpuLimit	CPU Limit for telegraf rs
	rsMemLimit	Memory limit for telegraf rs
	rsCpuRequest	CPU request for telegraf rs
	rsMemRequest	Memory request for telegraf rs
	runPrivileged	Run the telegraf DaemonSet's *telegraf-mountstats-poller* container in privileged mode. Set this to true if SELinux is enabled on your Kubernetes nodes.
	runDsPrivileged	Set runDsPrivileged to true to run the telegraf DaemonSet's telegraf container in privileged mode.
	batchSize	See Telegraf configuration documentation
	bufferLimit	See Telegraf configuration documentation
	roundInterval	See Telegraf configuration documentation
	collectionJitter	See Telegraf configuration documentation
	precision	See Telegraf configuration documentation
	flushInterval	See Telegraf configuration documentation
	flushJitter	See Telegraf configuration documentation
	outputTimeout	See Telegraf configuration documentation
	dsTolerations	telegraf-ds additional tolerations.
	rsTolerations	telegraf-rs additional tolerations.
	skipProcessorsAfterAggregators	See Telegraf configuration documentation
	unprotected	See this known Telegraf issue. Setting *unprotected* will instruct the Kubernetes Monitoring Operator to run Telegraf with the `--unprotected` flag.
	insecureK8sSkipVerify	If telegraf is unable to verify certificate due to lack of IP SANs, try enabling verification skip
kube-state-metrics		Configuration options that can customize the kube state metrics installation of the Operator
	cpuLimit	CPU limit for kube-state-metrics deployment
	memLimit	Mem limit for kube-state-metrics deployment
	cpuRequest	CPU request for kube state metrics deployment
	memRequest	Mem request for kube state metrics deployment

Component	Option	Description
	resources	a comma separated list of resources to capture. example: cronjobs,daemonsets,deployments,ingresses,jobs,namespaces,nodes,persistentvolumeclaims, persistentvolumes,pods,replicasets,resourcequotas,services,statefulsets
	tolerations	kube-state-metrics additional tolerations.
	labels	a comma separated list of resources for which kube-state-metrics should capture labels example: cronjobs=[*],daemonsets=[*],deployments=[*],ingresses=[*],jobs=[*],namespaces=[*],nodes=[*], persistentvolumeclaims=[*],persistentvolumes=[*],pods=[*],replicasets=[*],resourcequotas=[*],services=[*],statefulsets=[*]
logs		Configuration options that can customize logs collection and installation of the Operator
	readFromHead	true/false, should fluent bit read the log from head
	timeout	timeout, in secs
	dnsMode	TCP/UDP, mode for DNS
	fluent-bit-tolerations	fluent-bit-ds additional tolerations.
	event-exporter-tolerations	event-exporter additional tolerations.
	event-exporter-maxEventAgeSeconds	event-exporter max event age. See https://github.com/jkroepke/resmoio-kubernetes-event-exporter
	fluent-bit-containerLogPath	By default, the Fluentbit DaemonSet will mount the /var/log and /var/lib/docker/containers host paths to access/read the Kubernetes container logs. If Kubernetes has been configured to place container logs in a non-default location, use this option to modify the Fluentbit DaemonSet to mount the non-default path.
workload-map		Configuration options that can customize the workload map collection and installation of the Operator.
	cpuLimit	CPU limit for net observer ds
	memLimit	mem limit for net observer ds
	cpuRequest	CPU request for net observer ds
	memRequest	mem request for net observer ds
	metricAggregationInterval	metric aggregation interval, in seconds
	bpfPollInterval	BPF poll interval, in seconds

Component	Option	Description
	enableDNSLookup	true/false, enable DNS lookup
	l4-tolerations	net-observer-l4-ds additional tolerations.
	runPrivileged	true/false - Set runPrivileged to true if SELinux is enabled on your Kubernetes nodes.
change-management		Configuration options for Kubernetes Change Management and Analysis
	cpuLimit	CPU limit for change-observer-watch-rs
	memLimit	Mem limit for change-observer-watch-rs
	cpuRequest	CPU request for change-observer-watch-rs
	memRequest	mem request for change-observer-watch-rs
	workloadFailureDeclarationIntervalSeconds	Interval after which a non-successful deployment of a workload will be marked as failed, in seconds
	workloadDeployAggrIntervalSeconds	Frequency at which workload deployments are combined and sent, in seconds
	nonWorkloadDeployAggrIntervalSeconds	Frequency at which non-workload deployments are combined and sent, in seconds
	termsToRedact	A set of regular expressions used in env names and data maps whose value will be redacted Example terms:"pwd", "password", "token", "apikey", "api-key", "jwt"
	additionalKindsToWatch	A comma separated list of additional kinds to watch from the default set of kinds watched by the collector
	kindsToIgnoreFromWatch	A comma separated list of kinds to ignore from watching from the default set of kinds watched by the collector
	logRecordAggrIntervalSeconds	Frequency with which log records are sent to CI from the collector
	watch-tolerations	change-observer-watch-ds additional tolerations. Abbreviated single line format only. Example: '{key: taint1, operator: Exists, effect: NoSchedule},{key: taint2, operator: Exists, effect: NoExecute}'

Detail Section

Each page of the Kubernetes Explorer displays a set of usage graphs that may include CPU, Memory, and Storage. Below these graphs are summaries and lists of the top objects in each category, with links to underlying details. For example, *Node* shows pods and containers, *Pod* shows PVCs and related objects and containers, etc. The following illustration shows an example of the detailed information on a *Node* page:

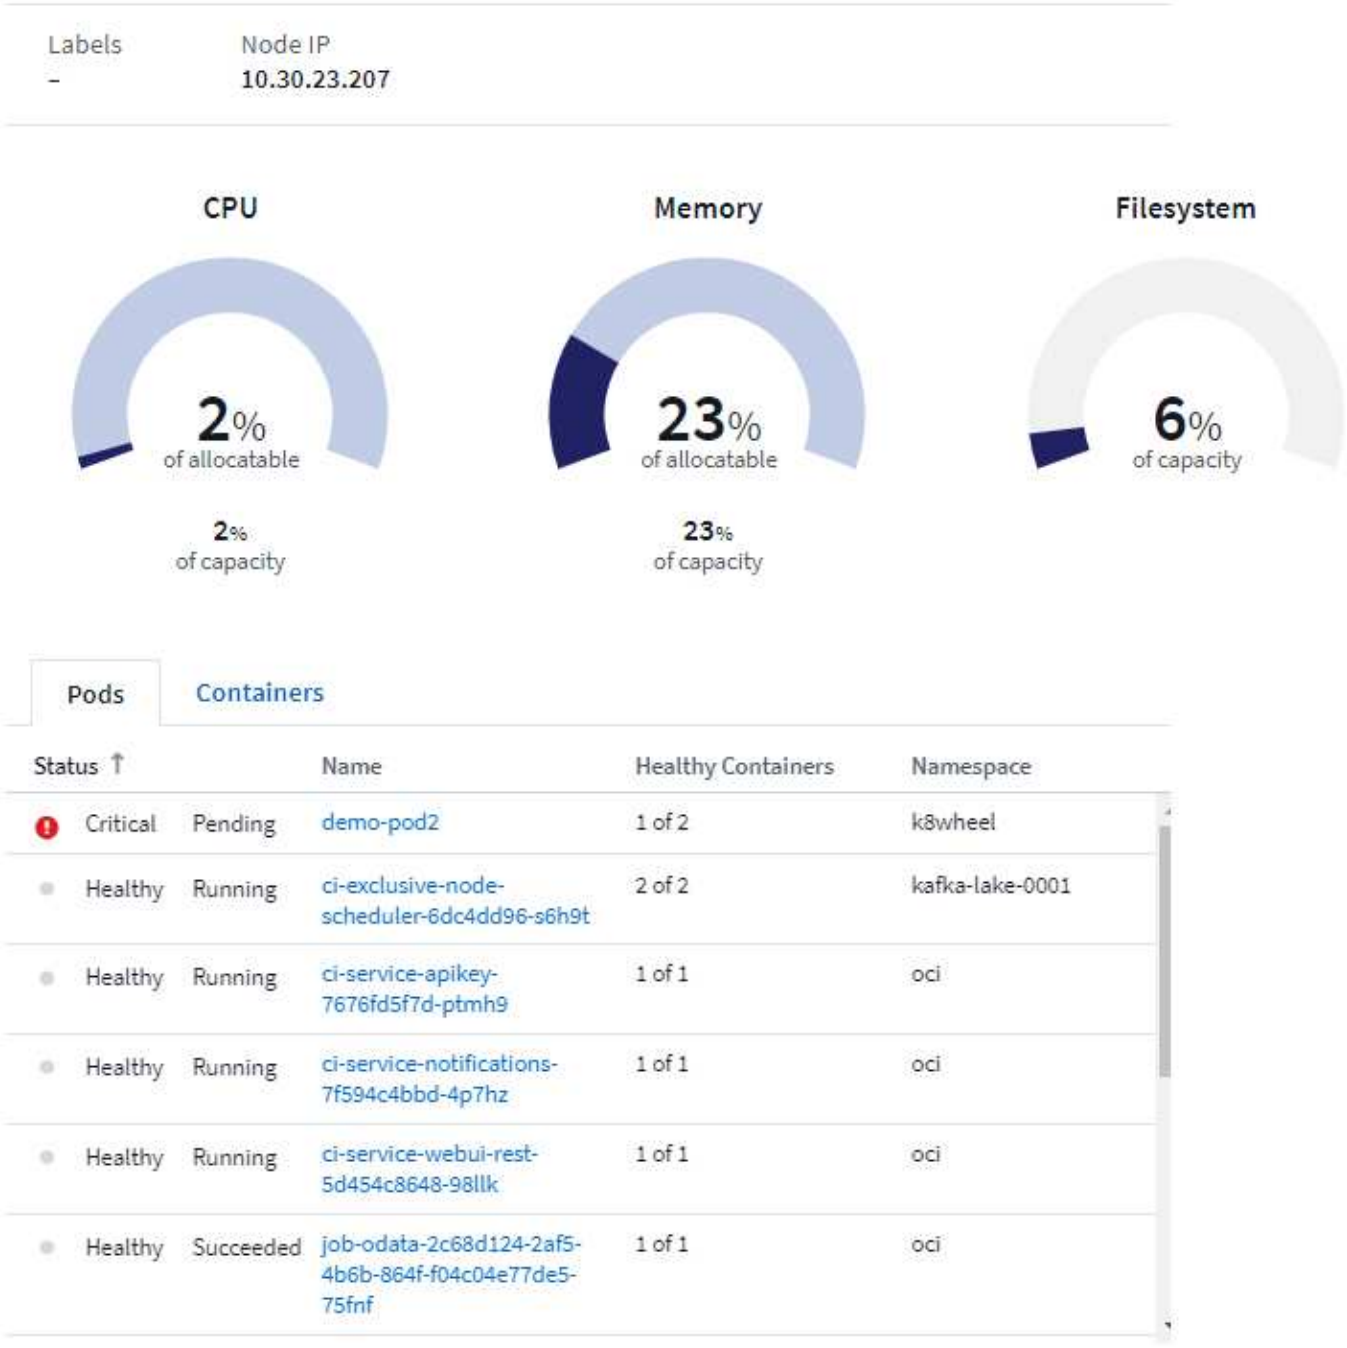

Resources experiencing alerts will show at the top of the lists. Click on the affected resource to drill into it for more detail.

Keep the following in mind when reading the gauges.

The dark blue band shows the amount used.

- When viewed against the *light blue band*, the dark blue shows used as the % of allocatable amount. This matches the "% of allocatable" value shown (81 in the example below).
- When viewed against the *white background*, the dark blue shows used as the % of total capacity. This matches the "% of capacity" value shown (63 in this example).

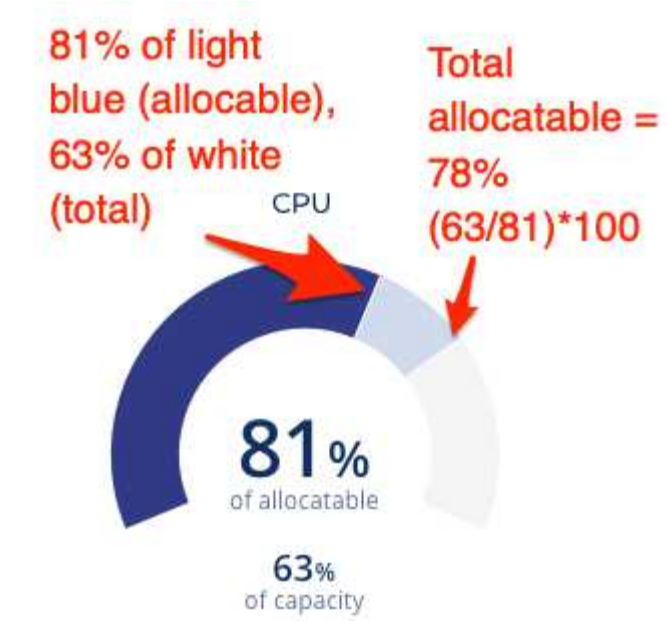

Details

Hovering over a circle displays a summary of information for that service.