

Foglight® for Container Management 7.3.0
User and Administration Guide



© 2024 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Where next meets now are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. "Apache HTTP Server", Apache, "Apache Tomcat" and "Tomcat" are trademarks of the Apache Software Foundation. Google is a registered trademark of Google Inc. Android, Chrome, Google Play, and Nexus are trademarks of Google Inc. Red Hat, JBoss, the JBoss logo, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries. CentOS is a trademark of Red Hat, Inc. in the U.S. and other countries. Fedora and the Infinity design logo are trademarks of Red Hat, Inc. Microsoft, .NET, Active Directory, Internet Explorer, Hyper-V, Office 365, SharePoint, Silverlight, SQL Server, Visual Basic, Windows, Windows Vista and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. AIX, IBM, PowerPC, PowerVM, and WebSphere are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Java, Oracle, Oracle Solaris, PeopleSoft, Siebel, Sun, WebLogic, and ZFS are trademarks or registered trademarks of Oracle and/or its affiliates in the United States and other countries. SPARC is a registered trademark of SPARC International, Inc. in the United States and other countries. Products bearing the SPARC trademarks are based on an architecture developed by Oracle Corporation. OpenLDAP is a registered trademark of the OpenLDAP Foundation. HP is a registered trademark that belongs to Hewlett-Packard Development Company, L.P. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. MySQL is a registered trademark of MySQL AB in the United States, the European Union and other countries. Novell and eDirectory are registered trademarks of Novell, Inc., in the United States and other countries. VMware, ESX, ESXi, vSphere, vCenter, vMotion, and vCloud Director are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Sybase is a registered trademark of Sybase, Inc. The X Window System and UNIX are registered trademarks of The Open Group. Mozilla and Firefox are registered trademarks of the Mozilla Foundation. "Eclipse", "Eclipse Foundation Member", "EclipseCon", "Eclipse Summit", "Built on Eclipse", "Eclipse Ready" "Eclipse Incubation", and "Eclipse Proposals" are trademarks of Eclipse Foundation, Inc. IOS is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. Apple, iPad, iPhone, Mac OS, Safari, Swift, and Xcode are trademarks of Apple Inc., registered in the U.S. and other countries. Ubuntu is a registered trademark of Canonical Ltd. Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. OpenSUSE, SUSE, and YAST are registered trademarks of SUSE LCC in the United States and other countries. Citrix, AppFlow, NetScaler, XenApp, and XenDesktop are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. AlertSite and DéjàClick are either trademarks or registered trademarks of Boca Internet Technologies, Inc. Samsung, Galaxy S, and Galaxy Note are registered trademarks of Samsung Electronics America, Inc. and/or its related entities. MOTOROLA is a registered trademark of Motorola Trademark Holdings, LLC. The Trademark BlackBerry Bold is owned by Research In Motion Limited and is registered in the United States and may be pending or registered in other countries. Quest is not endorsed, sponsored, affiliated with or otherwise authorized by Research In Motion Limited. Ixia and the Ixia four-petal logo are registered trademarks or trademarks of Ixia. Opera, Opera Mini, and the O logo are trademarks of Opera Software ASA. Tevron, the Tevron logo, and CitraTest are registered trademarks of Tevron, LLC. PostgreSQL is a registered trademark of the PostgreSQL Global Development Group. MariaDB is a trademark or registered trademark of MariaDB Corporation Ab in the European Union and United States of America and/or other countries. Vormetric is a registered trademark of Vormetric, Inc. Intel, Itanium, Pentium, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Debian is a registered trademark of Software in the Public Interest, Inc. OpenStack is a trademark of the OpenStack Foundation. Amazon Web Services, the "Powered by Amazon Web Services" logo, and "Amazon RDS" are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Infobright, Infobright Community Edition and Infobright Enterprise Edition are trademarks of Infobright Inc. POLYCOM®, RealPresence® Collaboration Server, and RMX® are registered trademarks of Polycom, Inc. All other trademarks and registered trademarks are property of

their respective owners.

Legend

- **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

- ! **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

- ⓘ **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Foglight for Container Management User and Administration
Guide Foglight Version- 7.3.0
Software Version- 7.3.0

Contents

Understanding Foglight for Container Management	6
About Foglight for Container Management	6
Architecture	7
Sizing Your Monitored Environment	7
Foglight Management Server Requirements	8
Kubernetes Agent Requirements	8
Docker Swarm Agent Requirements	8
Getting Started	9
Prerequisite	9
Creating and Activating Agent	19
Configuring data collection interval	25
Using Foglight for Container Management	27
Kubernetes	28
Monitoring Kubernetes Clusters	28
Monitoring Kubernetes Nodes	31
Monitoring Kubernetes Workloads	35
Monitoring Kubernetes Pods	40
Monitoring Kubernetes Other Components	41
Alarms	43
Capacity Management	44
Cost	45
Optimizer	50
Administration	54
Docker Swarm	60
Monitoring Docker Containers	60
Monitoring Docker Hosts	62
Monitoring Docker Swarm Clusters	66
Monitoring Docker Swarm Services	66
Alarms	67
Analytics	68
Kubernetes analytics	68
Docker Swarm analytics	70
Domains and Object Groups	73
Domains	73
Object Groups	73
Reference	75
Metrics	75
Kubernetes metrics	75
Docker Swarm metrics	79
Rules	79
Kubernetes	80
Docker Swarm	92

Modifying alarms and threshold values	95
Customization	95
About Us	100
Technical support resources	100

Understanding Foglight for Container Management

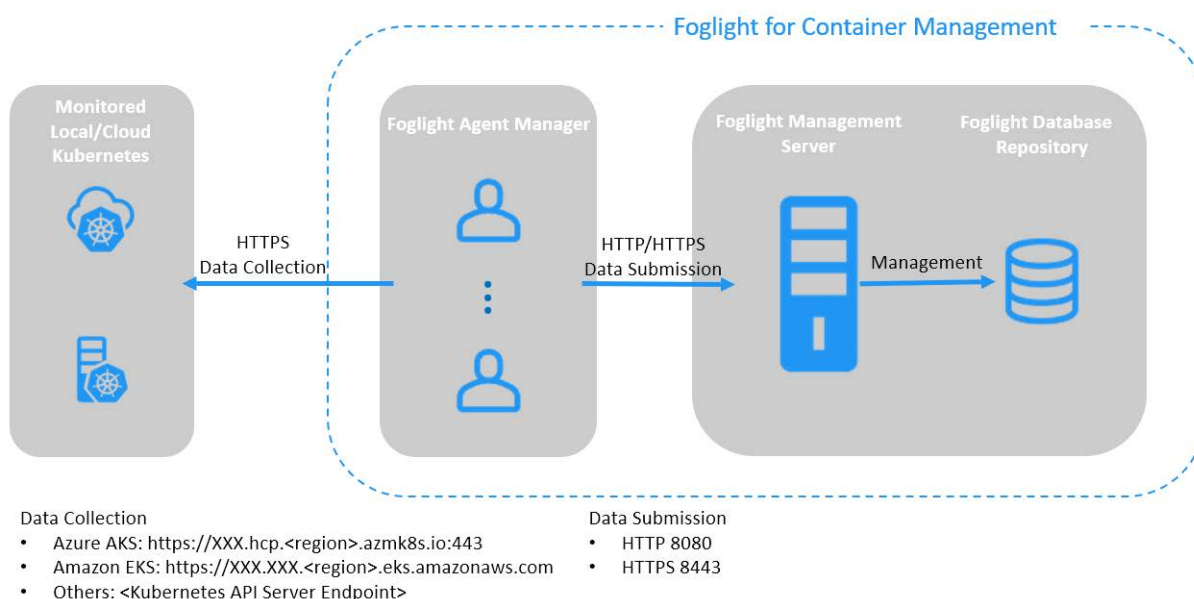
- [About Foglight for Container Management](#)
- [Architecture](#)
- [Sizing Your Monitored Environment](#)
 - [Foglight Management Server Requirements](#)
 - [Kubernetes Agent Requirements](#)
 - [Docker Swarm Agent Requirements](#)
- [Getting Started](#)
 - [Prerequisite](#)
 - [Creating and Activating Agent](#)
 - [Configuring data collection interval](#)

About Foglight for Container Management

Containers are a method of operating system virtualization that allow you to run an application and its dependencies in resource-isolated processes. Foglight® for Container Management simplifies this process by tracking each container, the resources it consumes, and the remaining compute of the container host, as well as providing you with the cluster information and pre-configured rules with notifications identifying the problem of your clusters.

Architecture

Figure 1. Components of Foglight for Container Management



Foglight for Container Management consists of three main components:

- Foglight Management Server and Foglight Database Repository — Responsible for managing, alerting, and viewing the collected data. Both components can be set to run on the same machine or reside on separate machines.
- Agent Manager — Hosts the monitoring Kubernetes agents.
- Docker Swarm clusters — Manages containerized applications in a clustered environment.
- Kubernetes clusters — Manages containerized applications in a clustered environment.

Sizing Your Monitored Environment

Consider the possibility of a great amount of objects being collected, analyzed, and maintained by the application, several aspects of the underlying server must be taken into account. The sizing of the supporting clusters and containers depends on the complexity of the underlying environment. Sufficient processing power and CPU memory are required to support effective collection, server data handling, and analytics.

i | **NOTE:** Currently Quest validates the environment with up to 10000 containers. If your environment beyonds this scale, contact Quest Support.

Foglight Management Server Requirements

The minimum system requirements of the Foglight Management Server vary from the scale of clusters. The scale of clusters is determined by running containers.

Table 1. Foglight Management Server requirements

Operating System	Maximum Containers	Foglight		Agent Manager	
		JVM Settings	# of CPUs	JVM Settings	# of CPUs
Windows 64-bit	1000	Xms Xmx=4G	2	Xms Xmx=4G	2
	5000	Xms Xmx=8G	4	Xms Xmx=8G	4
	10000	Xms Xmx=12G	6	Xms Xmx=12G	6
Linux 64-bit	1000	Xms Xmx=4G	2	Xms Xmx=4G	2
	5000	Xms Xmx=8G	4	Xms Xmx=8G	4
	10000	Xms Xmx=12G	6	Xms Xmx=12G	6

If you are using an embedded Agent Manager, make sure to use the sum resources of both Foglight and Agent Manager.

Kubernetes Agent Requirements

Kubernetes Agent collects inventory and metrics every 5 minutes by default. Refer to [Configuring data collection interval](#) for details about how to change the collection interval.

Table 2. Kubernetes Agent requirements

Maximum Containers	Kubernetes Agent Collection Interval (minutes)	
	Inventory	Metrics
500	5	5
1000	10	10
5000	30	30
10000	60	60

Table 2 is the recommendations for local Kubernetes clusters. If you deploy Kubernetes clusters on the Cloud Provider Kubernetes Service, consider your network rate and change your configurations based on different Cloud Provider and different region/zone of your cluster.

Docker Swarm Agent Requirements

Docker Swarm Agent collects inventory and metrics every 5 minutes by default. Refer to [Configuring data collection interval](#) for details about how to change the collection interval.

Table 3. Docker Swarm Agent requirements

Maximum Containers	Docker Swarm Agent Collection Interval (minutes)	
	Inventory	Metrics
500	5	5
1000	10	10
5000	30	30

Table 3 is the recommendations for local Docker Swarm clusters. For cloud environment, consider network rate and change configurations based on different Cloud Provider and different region/zone.

Getting Started

- Prerequisite
 - Kubernetes Agent
 - [Preparing the Kubernetes credential](#)
 - [Enabling Heapster service in monitored environment](#)
 - [Enabling Prometheus service in monitored environment](#)
 - Docker Swarm Agent
 - [Preparing Docker Swarm Agent credentials](#)
 - [Enabling Docker Remote API for monitored docker host](#)
 - [Uploading Docker Swarm Agent credentials](#)
- Creating and Activating Agent
 - [Creating and Activating a Kubernetes Agent](#)
 - [Creating and Activating a Docker Swarm Agent](#)
- [Configuring data collection interval](#)

Prerequisite

Kubernetes Agent

Each Kubernetes Agent monitors the assets inside the selected Kubernetes Service Providers. To enable the data collection, complete the following prerequisites before create agent.

- [Preparing the Kubernetes credential](#)
- [Enabling Heapster service in monitored environment](#)

Preparing the Kubernetes credential

The Kubernetes configuration file named *KubeConfig* is a standard configuration of Kubernetes and is required for Kubernetes agents to access the cluster. Foglight for Container Management verifies and supports the local Kubernetes and the following Cloud Kubernetes Service Providers. Based upon your environment, select either of approaches to get your *KubeConfig* file:

i | NOTE: Data from different Kubernetes Agents with the same cluster name will be merged into one cluster.

- [Local Kubernetes](#)
- [Azure Kubernetes Service \(AKS\)](#)
- [Amazon Elastic Container Service for Kubernetes \(EKS\)](#)
- [Google Cloud Platform Container Engine \(GKE\)](#)
- [IBM Cloud Kubernetes Service](#)
- [OpenShift Origin](#)

Local Kubernetes

If you build a Kubernetes cluster locally, find this *KubeConfig* file under the `/etc/kubernetes/admin.kubeconfig` on your master node.

Azure Kubernetes Service (AKS)

Before generating the Kubernetes credentials, record the following information:

- Azure Username
- Azure Password
- Azure Subscription Number
- The name of your AKS Cluster Resource Group
- The name of your AKS cluster

Download the [Azure Command Line Interface](#) and install it in your local platform, and then follow steps below to generate your Kubernetes credential:

- 1 Run the command `az login`.

Then a browser shows up, directing you to the Azure Portal where you should enter your Azure Username and Password to complete the authentication.

- 2 Run the command: `az account set --subscription <azure subscription number>`
- 3 Run the command: `az aks get-credentials --resource-group <azure resource group name> --name <azure cluster name>`
- 4 Find the Kubernetes configuration file under `<USER_HOME>/.kube/config` on your local platform.

i | **NOTE:** The token in this Kubernetes configuration file will get expired after two years. If you don't want the credential gets expired, refer to [Foglight Container Tools](#) for detail.

Amazon Elastic Container Service for Kubernetes (EKS)

Follow the Amazon EKS official guide [Getting Started with Amazon EKS](#). Follow the guide and complete [Create a kubeconfig for Amazon EKS](#). in the end of the guide.


i | **NOTE:** If you don't want the credential gets expired, refer to [Foglight Container Tools](#) for detail.

Google Cloud Platform Container Engine (GKE)

Download the [Google Cloud Client tool](#) and install it in your local platform, and then follow steps below to generate your Kubernetes credential.

To generate the intermediate Kubernetes credential for your cluster:

- 1 Log in to your Kubernetes cluster, click **Connect** next to your cluster name.



- 2 Click to copy the command below, and then run this command.

Connect to the cluster

You can connect to your cluster via command-line or using a dashboard.

Command-line access

Configure `kubectl` command line access by running the following command:

```
$ gcloud container clusters get-credentials jane-gke-demo --zone us-central1-a --project dulcet-bucksaw-208510
```

[Run in Cloud Shell](#)

- Find the intermediate Kubernetes configuration file under `<USER_HOME>/.kube/config` on your local platform. The following is the example of this intermediate Kubernetes configuration file.

NOTE: This Kubernetes configuration file cannot be used as the agent credential because the token in this file will get expired soon and “`cmd-path`” of the token directs to your local platform.

```
kind: Config
preferences: {}
users:
- name: gke_dulcet-bucksaw-208510_us-central1-a_jane-gke-demo
  user:
    auth-provider:
      config:
        access-token: ya29.G1zuBVkzkoVcl1VUY_7yXMS0DpQ7z7ahGzFA_f2o08FhxZjDICKRFDaw5ytBc9dHBT90yYk
        cmd-args: config config-helper --format=json
        cmd-path: C:\Users\jwang7\AppData\Local\Google\Cloud SDK\google-cloud-sdk\bin\gcloud.cmd
        expiry: 2018-07-04T06:44:59Z
        expiry-key: '{.credential.token_expiry}'
        token-key: '{.credential.access_token}'
      name: gcp
```

- Open Google Cloud Client tool and run the following commands to create a Kubernetes service account that grants with the `cluster-admin` role and the access to your Google Kubernetes Engine (GKE) cluster.

- `kubectl create serviceaccount <service account name>`
- `kubectl create clusterrolebinding <cluster role binding name> --clusterrole=cluster-admin -serviceaccount=default:<service account name>`

“`default`” in the above command is the namespace name of this service account name. You can also change to other namespace names, as needed.

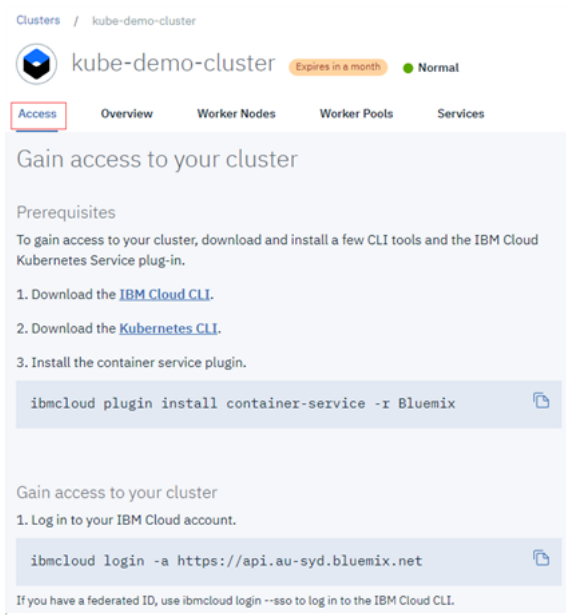
- `kubectl describe serviceaccount <service account name>`

You will get the response similar as below. Record the `<secret name>` for later use.

```
C:\>kubectl describe serviceaccount jane-gke-sa
Name: jane-gke-sa
Namespace: default
Labels: <none>
Annotations: <none>
Image pull secrets: <none>
Mountable secrets: jane-gke-sa-token-x2n6w
Tokens: jane-gke-sa-token-x2n6w
Events: <none>
```

- `kubectl describe secret <secret name>`

You will get response similar as below. Record the token value (exclude “token:”) for later use.



By default IBM Cloud Kubernetes Service uses certificate authority file and token/refresh token. However, certificate authority data and service account should be used in the Kubernetes Agent credential. After you successfully test your connection through “`kubectl get nodes`”, follow the steps below to generate the Kubernetes Agent credential.

- 1 Run the command `kubectl config view --minify=true --flatten -o json`. You will get an output similar as below, then record the `<certificate authority data>` for later use.

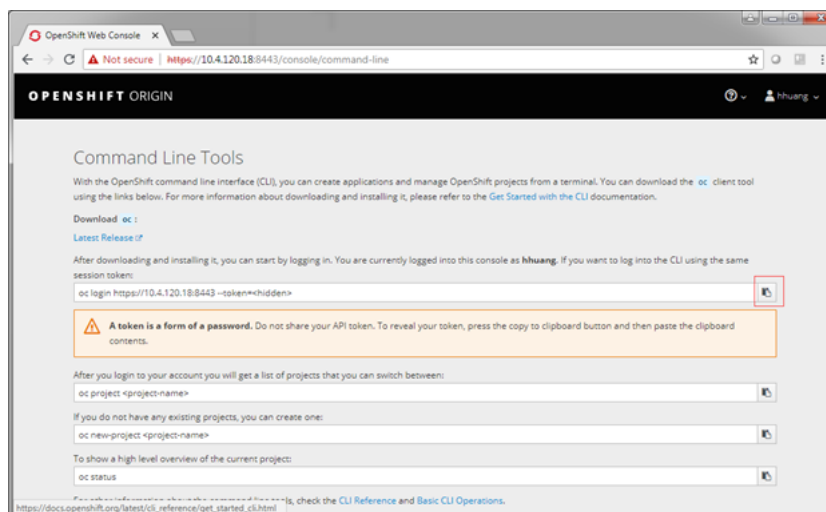
```
C:\Users\juang7\k.kube>kubectl config view --minify=true --flatten -o json
{
  "kind": "Config",
  "apiVersion": "v1",
  "preferences": {},
  "clusters": [
    {
      "name": "kube-demo-cluster",
      "cluster": {
        "server": "https://130.198.66.34:30244",
        "certificate-authority-data": "LS0tLS1CRUdJTTBDRUJUSUJzJG0FURSB0tL
S0tCk1JSUZYUENDQXk5ZGF3S0tJk1Uz3d0eKeYURURd05UUNWAGNOTkR0eE10S0hNFRF03T1RRM1dqQ
TUN0GN3T1ZZRFZBUUREQ2oUcK1G6STJNU0ZtURSaE1GazBOUE5oUUpGH5tCmpNeK0zTURNEfLR0dNe
F0ZF4R0Ct-NWakR1Z5TF0aE10S0Mk3Uptk0ns3Foz21H0x0cQKRRU2B0U9YD0c4Q010S0U0Z0tDq
NdFQ111dKxL1pcc89CckJH1H1PT0wenBpR0e5M3REYU0veGtYm1ueJmZc1ZU0U5ay1R6a21xqJUMU
zUBRkFKtU0FQNEVYRGJmZnpHcKk3M20zZkDL2huMh0CjNRQ7Q0U254Z1RRU1E7R0R2M0k3UkQcD1p1t
zZy2JUanErcKfU13g1NKFamRrncprUGoXN1UzTjz2SzeVudUsKY1puS0dQ0U5MSExOHFaZ2NkQ3Jja
FNTdEw0R2N3MkQ0NEZyq3dRb1N1TJ0n0XBYWFRERkxTd65yeXpX0UkZR0p9m2FrahQ308s5Y21U1B6F
0B0q2PEZ40cJ0q0q5eBJDcmF0NG1DTHkxREZ6GUMUMH3TPhNERPhyz20E0xckR3Uzkb0MHBT0960
GpS0ZU2T1u0U3EUPBDe0pVY1PM091Q110b25n0n1EK1BZcEckRzdT20N3VTJAx0c1ZzQkA110dU04S
zR0ckPSNlu2ctR1d6Q0pPc194cjk2dy94Tn14Uk11UkMPC2RzTUg1N3hPS3pgN1J0b2E3h0PZ0oYL
3ZnhL200Qkx1U0PU9QU0EhybitJS09Z21Q5csh2c0xentQ11dz0Fp1RD1qMnk4UURNeEd2L2Zyb1hRT
pplCkduYkxuekAMZUCR0c3eG1t0XZS0UBo0CtGymFxbatGU1Z5d2Nk01FgSnZ63T1uZEp6ZnQ42k91Z
0RMejBteFgR08eah1dS011PUZZU3ANTUe2RC0M011S0e0u11h0y0a0e00u14R3V4U3N1Muzc03U0
J0bhdBaD1oRC9U0p0p0NUU3X11NkcU12V3U0FHC3NDMTZB2J0Z1Z0Nk9o0U5YXp0UHTUWR6ZndRR
0dxcjB4NHMMU1U23021dFC1BZSEUMR0JM0y4M5KcEMzEFU0XDf0U0FHT1FNRT03SFFZRFZME9CQ
11FRkcxUEyazQxZ21TcndfX1tRkMkDzdsb0t1cE5NqjhhQTFUZE13U01NQnFBRkcxUEyazQxZ21Tc
ndfX1tRkM3N2xta2U0tk1Bd0dBHUURRkDRRgnQU1CQVYd0RRMUplB1p0aHZjtkFRRUx0UUFEE2dJQ
kF0c29J0M0y0j0e0c4T1Nk0c05andU02M042LUB1JTC0M011R0anR0M0c0E00A99EU1Umb0P0SJR0a
0g5M09nR0k0Z0vz03ZP0hIN0F014T316U0GcxT24v0y9kbbU5QW0kaj0BHUU0U0EhCd0HUWd0tS0y1tck1T0
21Z0UJj030vU10e0U0pZ1dn0Xh1U3p2MzR0eTE00tnazE5VWZSYkRQ0v01U0npU0011c0pCTUdZ0NzT0
U1ie1M0M0d01N49p0jLZR1R0d11FU1U0Z089U0GZnaJNq0XA2Q21z0Mz1re11w0U05CjZPcjd1b0U0L0Y0M
0U1dnR0RG0c0J2Q1dnad0E0d0R0d2tn0n00T0pp0cJ0D0B0U0dE0nR0T1y0c0nEFOU0U0N0G0k0MU05T1EzT
E110W1k0tU0c11h03p0808c0U0C0ng4eTUZ200Mk0ZEvR1Boe1N25n0M0E0F042R080u0n0c0p5U
0Z1M1JBR0z1aK3H0Q2U0Uk10YjFkY2x6akt0F0Q0k3JKX25v149F00Z1Z1B1Vmd1U09U0R3d3Vt0t6SjZ0a
HRI0m0JUTU0zAZS00p0Q0dmUj1XTE00cFznTIG1YSDdrZkM4QjFaQ1N0cE2Z40Z0eXN0Yq1k2RD11d2tm5
Gp0Z1J2e0U0K0UxSRFZ0eH060G0v0WT0k02hd0n0v0ZnF0y0W1EL3U0HTUNZRjFTU2N2T1B1U0v0UGN2cUNF0
LIZ293REZ0cJ01Hq02YVh0L31xcnplZUJNMjgrZ1UWRER1azJXW0c2U1ZqQnU3b3Y4QUN0bFh0Y2NQ0
0Z0N010R1d1e0d0U0j0hZUW0C0R0d0E0HTk1h01Zn0N1e0U0h0b25T003Z25R0T112R0Vz0G0v0z2Y21LR
1kxZUM0S0D0v0Za13p0M0t0U0E1xak0K0U0BoNFZT811xa2Zc0t0L80cR0U5E1ENFU1RJk1DQURFLS0tL
S0k"
      }
    }
  ]
}
```

- 2 Run the command `kubectl create serviceaccount <service account>`.
- 3 Run the command `kubectl describe serviceaccount <service account>`. You will get a response similar as below, then record `<service account secret>` (in this sample, it is `jane-sa-token-xkqrk`) for later use.

```
C:\Users\juang7\k.kube>kubectl describe serviceaccount jane-sa
Name: jane-sa
Namespace: default
Labels: <none>
Annotations: <none>
Image pull secrets: <none>
Mountable secrets: jane-sa-token-xkqrk
Tokens: jane-sa-token-xkqrk
Events: <none>
```


Download the [OpenShift Command Line Interface](#) and install it in your local platform, and then follow steps below to generate your Kubernetes credential:

- 1 Log in to OpenShift and generate an intermediate Kubernetes configuration file.
 - a After logging in to OpenShift, click **Command Line Tools** on the upper right.
 - b Click the button next to the *Session token* field, copy the command, and then paste it in your local Command Line Tool. Make sure to find the intermediate Kubernetes configuration file under `<USER_HOME>/.kube/config` on your local platform.



- c On your local platform, browse to open this configuration file. You may see the context similar to the following. Record `<config-cluster-name>` for later use.

```
1 apiVersion: v1
2 clusters:
3   - cluster:
4       insecure-skip-tls-verify: true
5       server: <cluster-ip-port>
6       name: <config-cluster-name>
7 contexts:
8   - context:
9       cluster: <config-cluster-name>
10      namespace: <project-name>
11      user: <config-user-name>
12      name: <config-context-name>
13 current-context: <config-context-name>
14 kind: Config
15 preferences: {}
16 users:
17   - name: <config-user-name>
18     user:
19       token: <access-token>
```

- 2 The token generated in step 1 will be expired after 4 hours, however Foglight for Container Management needs a permanent Kubernetes credential. So you need to create a service account with **"cluster-admin"** role, and then get the authorization code (not expired) of this service account to generate our permanent Kubernetes credential.

- a Run the command `oc project <project-name>`.
 - b Run the command `oc create serviceaccount <service-account-name>`.
You can check if your service account has been created successfully using the command:
`kubectl get serviceaccounts`
 - c Run the command `oc serviceaccounts get-token <service-account-name>`. Then you will get a token `<service-account-token>` like below. Record this token for later use.

```
"eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCJ9.eyJpc3MiOiJrdWJ1cm5ldGVzL3N1cnZpY2VhY2NvdW50Iiwia3ViZXRVZXRlc3VybnkiOiJpYWN1YWNjb3VudC9uYW1lc3BhY2UiOiJkZWZhZDw0Iiwia3ViZXRVZXRlc3VybnkiOiJpYWN1YWNjb3VudC9zZWN5ZXQubmFtZSI6Im9zLWZkbyJ9"
```

```
WluLXRva2VuLWY0a2ZsIiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9zZXJ2aWNlLW
FjY291bnQubmFtZSI6Im9zLWFkbWluIiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNjb3VudC9
zZXJ2aWNlLWFjY291bnQudWlkIjoiODMzNGU0NTQtNzQ1Yy0xMWU4LWFmNmEtMDA1MDU2YjY3
NDFhIiwic3ViIjoic3lzdGVtOnNlcnZpY2VhY2NvdW50OmRlZmF1bHQ6b3MtYWRTaW4iOiJpPqDSWkUAN9F
H_AoXy2U1elkHN_Bs9IR1xo0zNCJlwcY0h3zuQnrkOFi8gVpX1I77uhApp7oIjPqDSWkUAN9F
6mP_tNdGwJsqrMhYEMOtCLnnIM61BYxIcABvwr66a0Z3Gn0D7EM5M_7XgKDC16ON3W5NaH0D8
DpVTYqkxQ49u3qt4gqrcjVCaSSDNWlgGxY4KOIDrUbKkdgaRKzeD9o4Bv9VbYICqyxwoUebku
JAcHiXGIcSe-ozS_zroPiltT5HW-RY0Pn3Fp3zBnydiokna0-mXot5lqoYc-
R6E1U9YSrAOhWm9Q8ipiut6OczXbmLPM4DYve6dmHi_j5FquCqhod-Q1A7aPw"
```

- d Run the following command to grant your service account with the “cluster-admin” privilege:

```
kubectl create clusterrolebinding <cluster-role-binding-name> --
clusterrole=cluster-admin --serviceaccount=default:<service-account-
name>.
```
- 3 Generate a permanent Kubernetes configuration file and save it under `<USER_HOME>/<config file/credential>`.
 - a Open and edit the intermediate configuration file.
 - b Use `kubectl` to add user credentials, create new context, in the end change the existing contexts to the ones that you added in step 2.

For example,

```
kubectl config set-credentials <credential-name> --token=<service-
account-token>

kubectl config set-context <new-context-name> --cluster=<config-cluster-
name> --user=<credential-name> --namespace=<project-name>

kubectl config use-context <new-context-name>
```

- c Save the current Kubernetes configuration file.

Enabling Heapster service in monitored environment

There are various approaches to enable Heapster on your Kubernetes cluster. Visit [Heapster official website](#) to determine the approach that you are going to deploy your Heapster service, or you can follow instructions in <https://github.com/foglight/container> to deploy your service.

Some of the cloud platform Kubernetes service has enabled Heapster service for the cluster. If you have connected to the cluster, run the following command to check: `kubectl cluster-info`

Enabling Prometheus service in monitored environment

There are various approaches to enable Prometheus on your Kubernetes cluster. Visit [Prometheus Official Website](#) to determine the approach that you are going to deploy your Prometheus service. Or you can refer to configuration management systems, such as helm (<https://github.com/helm/charts/tree/master/stable/prometheus>) or ansible (<https://github.com/cloudalchemy/ansible-prometheus>) to simplify your installation process.

In the Foglight Container **Administration** dashboard, a simple template is provided for you to deploy Prometheus service onto your cluster. For more information, see the Deploy/Migrate section in [Metrics Collector](#) on page 56.

Docker Swarm Agent

Each Docker Swarm Agent monitors the assets in one docker host. Docker Remote API needs to be enabled for the Docker Swarm Agent collecting data from the docker host. If TLS is enabled to secure the Docker Remote API, credential for Docker Swarm Agent needs to be prepared. Complete the following prerequisites before create agent.

- [Preparing Docker Swarm Agent credentials](#)
- [Enabling Docker Remote API for monitored docker host](#)
- [Uploading Docker Swarm Agent credentials](#)

Preparing Docker Swarm Agent credentials

If TLS enabled to secure Docker Remote API, then complete the following guide to get the credentials for Docker Swarm Agent for the docker host. Otherwise, continue with [Enabling Docker Remote API for monitored docker host](#) on page 17

Refer to the [official guide](#) to generate the keys. Be aware that, during generating the keys, the Foglight Agent Manager host address should be in the allow access list.

Docker Swarm Agent needs following credentials, you can get them when you finish the [official guide](#).

- CA Public Key (ca.pem in [official guide](#))
- Client Public Key (cert.pem in [official guide](#))
- Client Private Key (key.pem in [official guide](#))

Enabling Docker Remote API for monitored docker host

Change `ExecStart` in docker service startup script as below.

Non-TLS secured

```
ExecStart=/usr/bin/dockerd -H tcp://0.0.0.0:2375 -H  
unix:///var/run/docker.sock
```

i | **NOTE:** Access should be allowed to the TCP port 2375

TLS secured

If TLS enabled, complete [Preparing Docker Swarm Agent credentials](#) on page 17 first, then you will get the ca.pem, server-cert.pem and server-key.pem mentioned in the [official guide](#).

```
ExecStart=/usr/bin/dockerd --tlsverify --tlscacert=ca.pem --tlscert=server-  
cert.pem --tlskey=server-key.pem -H tcp://0.0.0.0:2375 -H  
unix:///var/run/docker.sock
```

i | **NOTE:** Access should be allowed to the TCP port 2375

Then restart docker service.

Uploading Docker Swarm Agent credentials

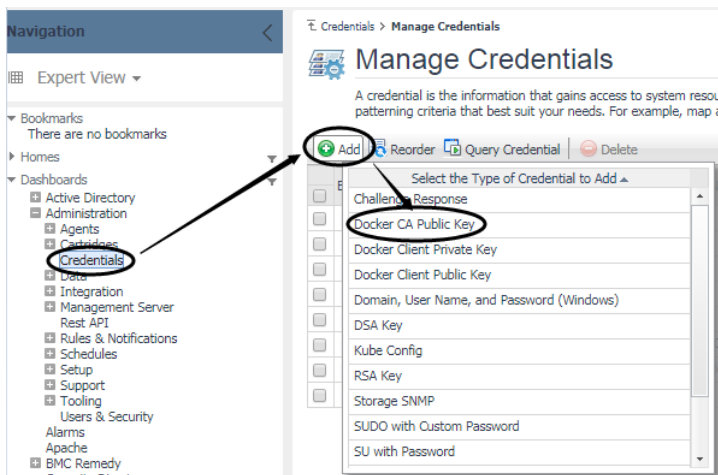
If TLS is enabled to secure Docker Remote API, go through this section to upload the credential for Docker Swarm Agent. Otherwise, skip this section.

When complete [Preparing Docker Swarm Agent credentials](#) on page 17, following credentials should be generated.

- CA Public Key
- Client Public Key
- Client Private Key

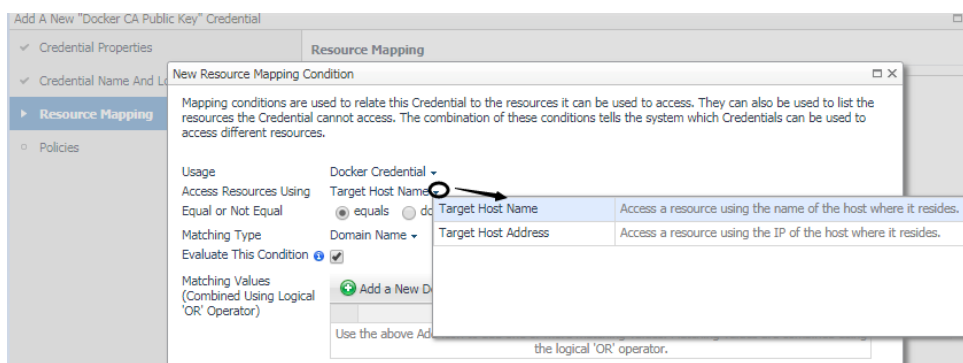
These are the credentials for Docker Swarm Agent, complete the following steps to upload the credentials.

On the **Administration > Credentials > Manage Credentials** dashboard, click **Add**, and then select Docker CA Public Key or Docker Client Public Key or Docker Client Private Key to upload related credentials. Take Docker CA Public Key as an example.

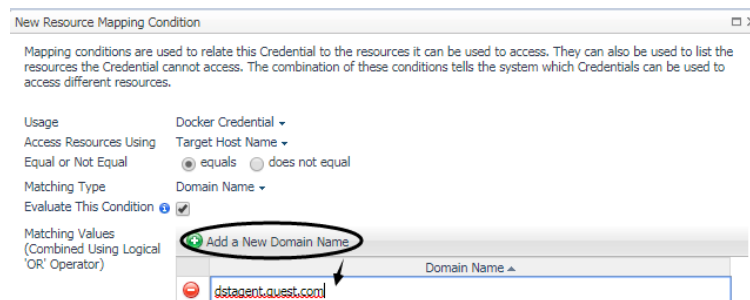


In the **Add a New “Docker CA Public Key” Credential** dialog box, specify the following values:


- Credential Properties: Click **Load from file** to import Docker CA Public Key, and then click **Next**.
- Credential Name And Lockbox: Specify a unique name for this credential, and then click **Next**.
- Resource Mapping: Click **Add**. In the **New Resource Mapping Condition** dialog box, choose Target Host Name or Target Host Address for the monitored docker host.



If choose **Target Host Name**, then enter the host name of the monitored docker host.



If choose **Target Host Address**, then enter the IP address of the monitored docker host.



Click **Add** to finish editing **New Resource Mapping Condition** and back to **Resource Mapping**. And then click **Finish**.

The **Docker CA Public Key** has been uploaded and mapped to the docker host. To monitor this docker host, upload **Docker Client Public Key** and **Docker Client Private Key** by following the above steps.

Creating and Activating Agent

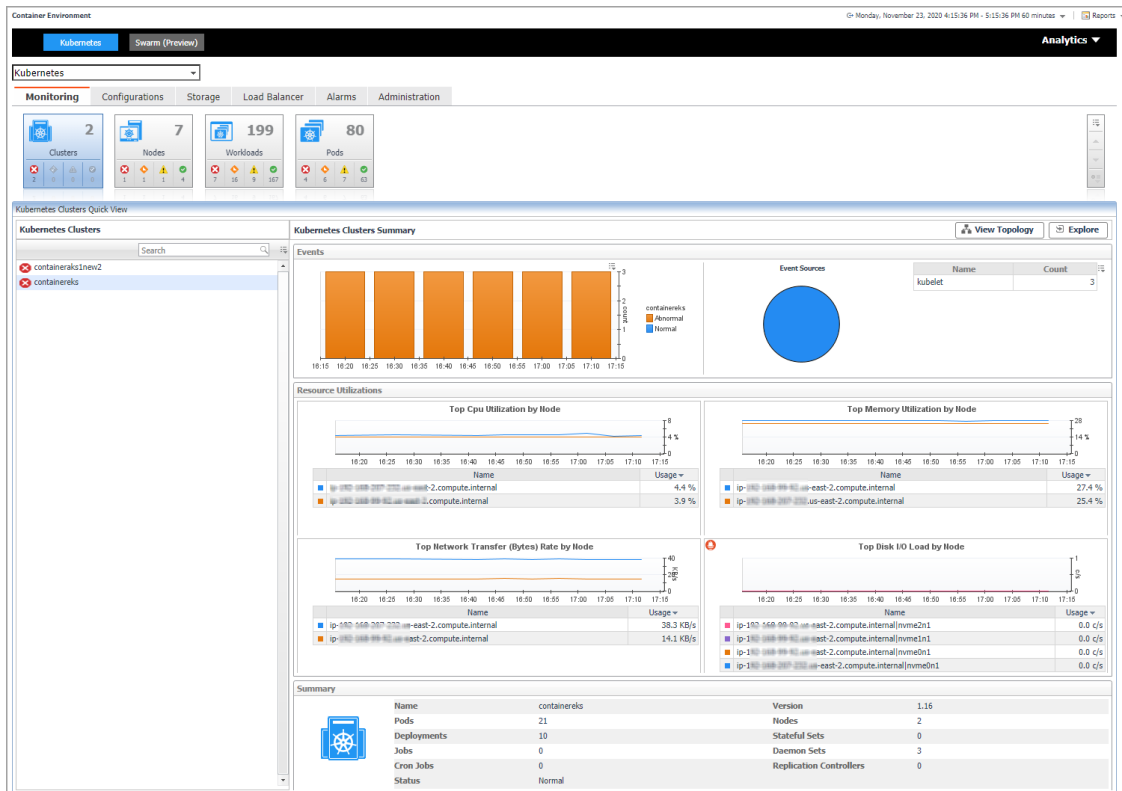
Foglight for Container Management supports Kubernetes Agent and Docker Swarm Agent.

- [Creating and Activating a Kubernetes Agent](#)
- [Creating and Activating a Docker Swarm Agent](#)

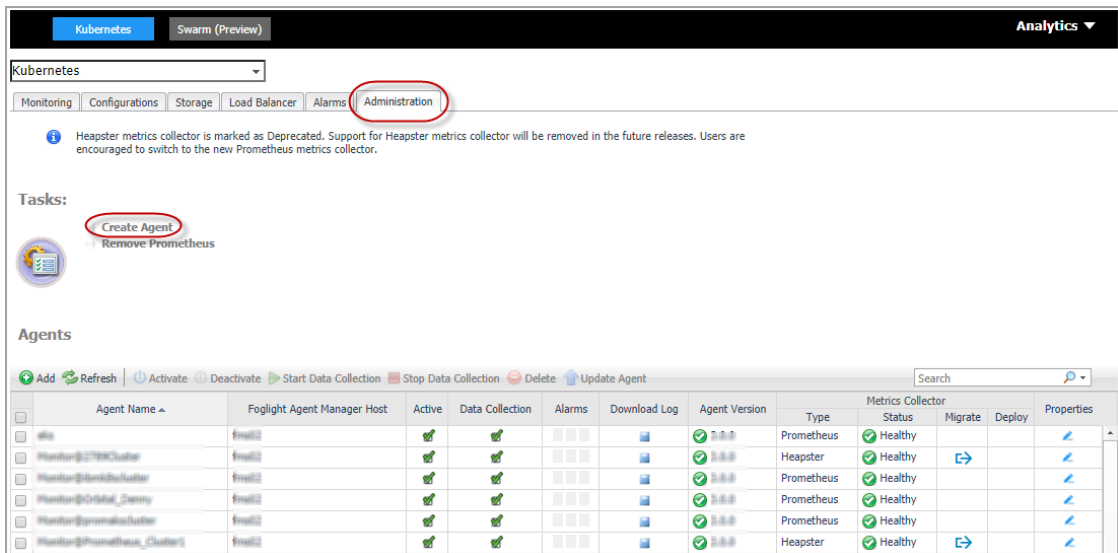
Creating and Activating a Kubernetes Agent

To create a Kubernetes agent on a monitored host:

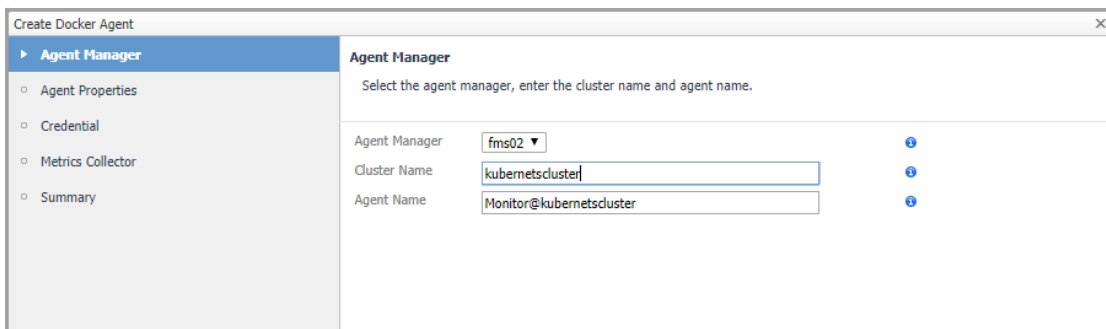
- 1 Log in to the Foglight browser interface and make sure the left Navigation panel is open.
- 2 On the navigation panel, from **Standard View** click **Container Environment** or from **Expert View** click **Dashboards > Container**. Then the Container dashboard will display as below.



- 3 In the Container dashboard, click **Administration** tab, and then click **Create Agent**. The **Create Docker Agent** wizard opens.

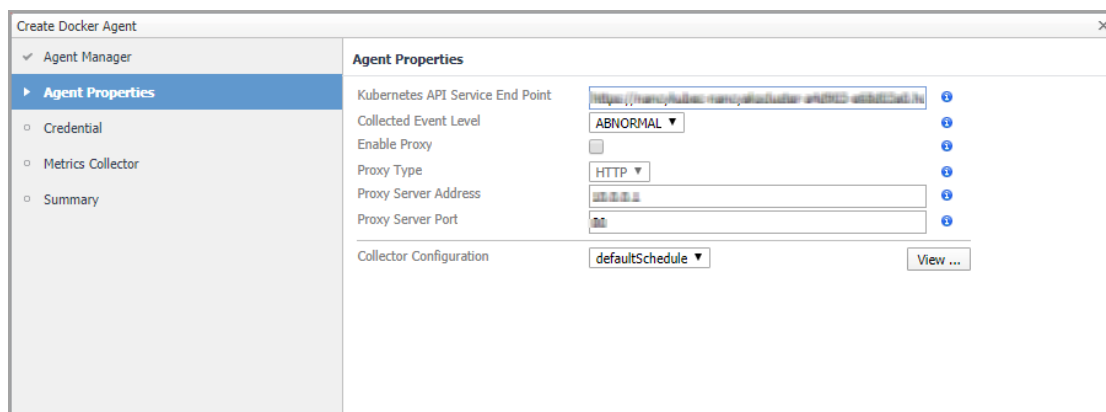


4 *Agent Manager*: specify the following values, and then click **Next**.



- *Agent Manager*: The Kubernetes Agent will be create in the selected Agent Manager.
- *Cluster Name*: Customized cluster name which identifies a Kubernetes cluster.
- *Agent Name*: Auto-generated agent name. You can change the name according to your requirement. It should be a unique name.

5 *Agent Properties*



- *Kubernetes API Service End Point*: The format is: `http(s)://<url:port>`. If you have a Kubeconfig file, retrieve this endpoint from the file (path: clusters -> cluster -> server). If there are multiple clusters, find the current context related cluster server URL.
- *Collected Event Level*: Set the collected event level, including *ABNORMAL* and *ALL*. *ALL* will collect both abnormal and normal events while *Abnormal* only collects abnormal events.

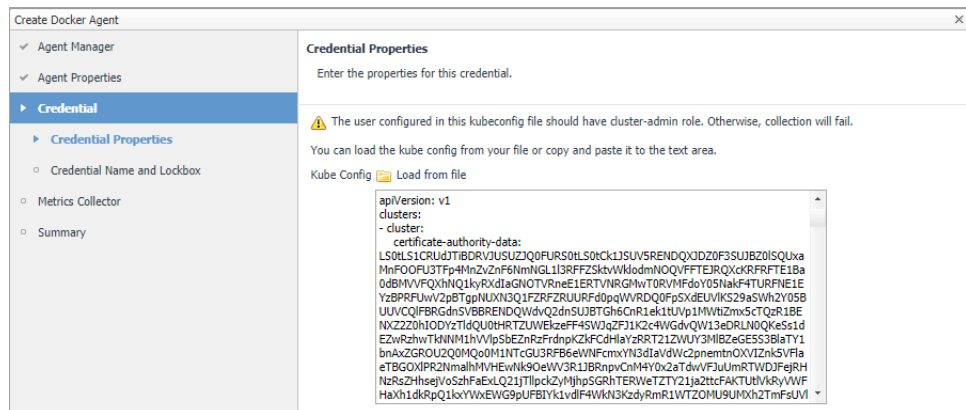
- **Enable Proxy:** To enable the proxy, select the checkbox. Enter the *Proxy Server Address* and *Proxy Server Port* information.
- **Collector Configuration:** Used to configure collection interval for inventory and metrics. You can change the collector intervals of defaultSchedule, however, this will affect all the Kubernetes Agents. Or you can create a new scheduler, configure the collector intervals, and then assign this scheduler to this agent.

Collector intervals to configure:

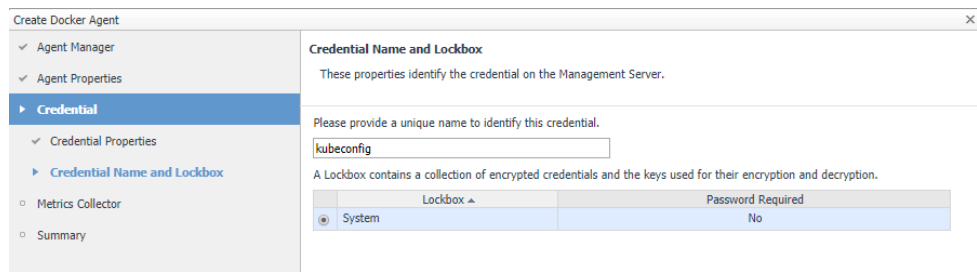
- inventory interval (inventory data collection interval).
- metrics interval (performance metrics data collection interval).

6 Credential

- If no credentials were found for the provided resource, configure credentials:
 - **Credential Properties:** Click **Load from file** to upload the credential and click **Next**.



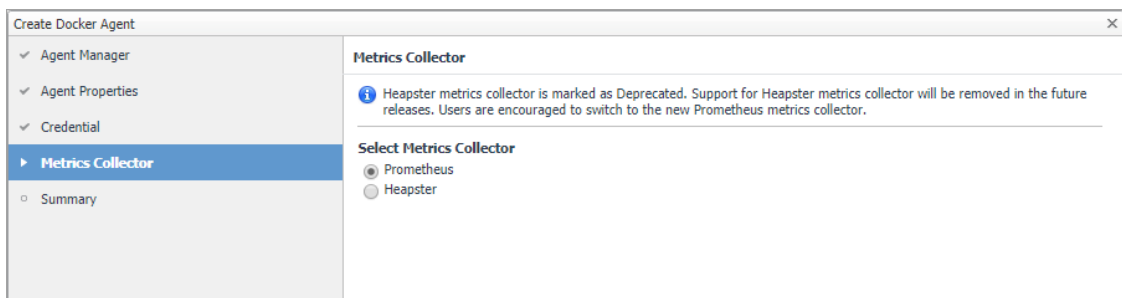
- **Credential Name and Lockbox:** give a name for the credential, choose a lockbox, then click **Next**.



- If an existing credential is detected, go to **Metrics Collector** directly.

7 Metrics Collector

Currently both **Prometheus metrics collector** and **Heapster metric collector** are supported. However, users are encouraged to use the Prometheus metrics collector.



Prometheus metrics collector

- Have existing Prometheus in your cluster.

i | **NOTE:** Ensure that you have Prometheus in your cluster before this step. We will check the service existence and health status after you click Next.

- Select Prometheus Type.

- Default type: allows Agent to use KubeConfig credentials to access the Prometheus service. Enter the existing Prometheus service namespace and name to configure the metrics collector.

- OpenShift Prometheus type: you need to provide OpenShift Route Hostname, Username, and Password for Agent to access the target OpenShift Prometheus.

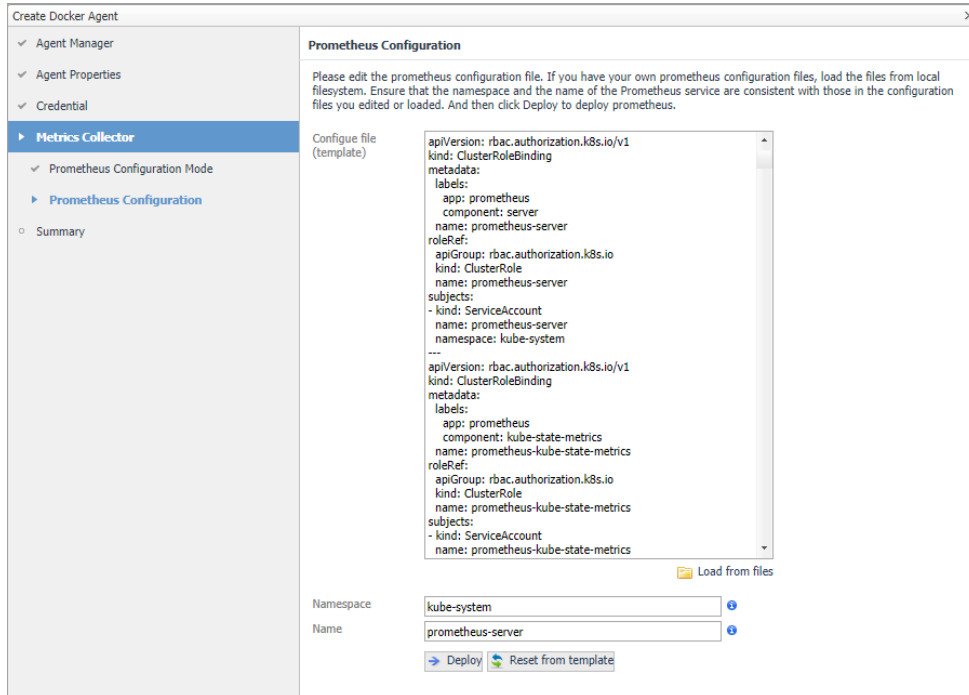
i | **NOTE:** To find the Route Hostname, go to the **OPENSHIFT ORIGIN** dashboard, **Application > Routes**.

Your OpenShift Login User should have the privilege to access to the Prometheus.

- Do not have existing Prometheus in your cluster.

A *Prometheus Configuration* wizard will appear. To configure the metrics collector, either use the default Prometheus template or upload your own Prometheus deployment .yaml files through *Load from files*.

- i** **NOTE:** Whenever you change the content in Configure file (template), ensure that the *Namespace* and *Name* fields are consistent with the Prometheus service configurations in your Configure file (template).
If you are using the default template, there is no need to change anything.
For a full version of the default template, see <https://github.com/Foglight/Container/tree/master/prometheus>



Reset from template: Helps you to reset the Configure file (template) content to the default template.

Deploy: Creates Prometheus components to your cluster with the Configure file. After clicking *Deploy*, a progress message will be displayed.

- If deployed successfully, a succeeded message will be displayed. Close the *Validation Result* page and then click **Next** to finish your agent creation process.
- If failed to deploy the Prometheus configuration, a *Validation Result* page will be displayed with possible solutions.

Heapster metric collector

Enter the existing Heapster service namespace and name to configure the metrics collector.

- i** **IMPORTANT:** Deploy the Heapster service to your cluster manually before creating an agent with Heapster as metric collector. Otherwise, connection test to your Heapster service will fail, and you cannot proceed to the next step.

8 Summary: click **Finish**.

Summary	
Agent Manager	192.168.1.3
Cluster	openshiftcluster
Agent Name	Monitor@openshiftcluster
Kubernetes API Service End Point	http://openshiftcluster.default.svc:8083
Prometheus Type	OpenShift Prometheus
Route Hostname	https://prometheus-kube-system.router.default.svc.cluster.local
OpenShift Username	openshift
OpenShift Password	*****
Collected Event Level	ABNORMAL
Enable Proxy	false
Proxy Type	HTTP
Proxy Server Address	192.168.1.1
Proxy Server Port	8080
Collector Configuration	defaultSchedule
Credential	192.168.1.1
Lockbox	System
New or Existing Credential	New
Lockbox Password	The Lockbox is not a password secured Lockbox.

9 Then, the agent will be created and activated automatically.

Creating and Activating a Docker Swarm Agent

Each Docker Swarm Agent monitored one docker host. If the docker host belongs to a Docker Swarm cluster, it will be considered as a manager/worker node. Otherwise, it will be considered to be a standalone docker host.

i **NOTE:** For a Docker Swarm cluster, you should create one Docker Swarm Agent for one host in the cluster, and if you want to monitor the whole cluster environment, you need to create all the Docker Swarm Agents for all the hosts in the cluster.

To create a Docker Swarm agent on a monitored host:

- 1 Login in to the Foglight browser interface and make sure the left navigation panel is open.
- 2 On the navigation panel, under **Dashboards**, click **Administration > Agents > Agent Status**.
The **Agent Status** dashboard opens.
- 3 In the **Agent Status** dashboard, click **Create Agent**.
The **Create Agent** wizard opens.
- 4 **Host Selector:** Select the monitored host that you want to monitor with the Docker Swarm agent instance that you are about to create, and then click **Next**.
i **NOTE:** In order to select the host, the Foglight Agent Manager must be installed and running on the monitored host.
- 5 **Agent Type and Instance Name:** Specify the following values, and then click **Next**.
 - **Agent Type:** Select DockerSwarmAgent from the agent type list.

- **Agent Name:** Specify the name of the agent instance that you are about to create using either of the following approaches:
 - **Generic Name:** This option is selected by default. A generic name is a combination of the host name and the agent type and uses the following syntax: `agent_type@host_name`.
 - **Specify Name:** Type that name in the *Name* field. For example, `MyAgent`.
- 6 On the **Summary** page, review the choices you have made, and then click **Finish**.
The *Agents* table refreshes automatically, showing the new Docker Swarm Agent.
 - 7 On the *Agents* table, select the Docker Swarm Agent that you create, click **Edit Properties**, and then click **Modify the private properties for this agent**.
 - 8 In the *Agents* properties view, check if the following values have been configured based upon your environment:

- **Name:** give a name to the monitored docker host, it should be unique.
 - **Host Name:** IP address or host name of the monitored docker host.
 - **Docker Remote API End Point:** Docker Remote API endpoint of the monitored docker host. For more information, see [Enabling Docker Remote API for monitored docker host](#) on page 17.
 - **Swarm Name:** specify the swarm cluster name for display. If the swarm name is kept as “default”, then the cluster name will be displayed as “default (cluster ID)” on the dashboard. If a customized name is input here, then the customized cluster name will be displayed on the dashboard.
- i | NOTE:** Ensure that the docker host inside the same cluster has the same configuration for Swarm Name.
- 9 Return back to the *Agents* table, select the above property changed Docker Swarm Agent, and then click **Activate**.

The new Docker Swarm Agent is created and data will be shown on the **Monitoring** tab after a few minutes.

Configuring data collection interval

The default data collection interval of agents is set to 5 minutes by default. Foglight for Container Management enables you to change this collection interval as needed.

- i | NOTE:** Changing the data collection interval will take effect for all Kubernetes agents and Docker Swarm agents.

To configure the data collection interval:

- 1 On the navigation panel, under **Dashboards**, select **Administration > Agents > Agent Status**.

- 2 On the *Agent Status* dashboard, select the Kubernetes agent that you use to monitoring the container environment, and then click **Edit Properties**.
- 3 In the *Edit Properties* dashboard, click **Edit** next to the *Collector Config* field.
- 4 In the KubernetesAgent or DockerSwarmAgent Collector Config dialog box, change the following values, as needed:
 - *Inventory Collector*: Specifies the interval for collecting components.
 - *Metrics Collector*: Specifies the interval for collecting metrics.
- 5 Click **Save**.

Using Foglight for Container Management

- **Kubernetes**
 - Monitoring Kubernetes Clusters
 - Monitoring Kubernetes Nodes
 - Monitoring Kubernetes Workloads
 - Monitoring Kubernetes Pods
 - Monitoring Kubernetes Other Components
 - Alarms
 - Capacity Management
 - Cost
 - Optimizer
 - Administration
- **Docker Swarm**
 - Monitoring Docker Containers
 - Monitoring Docker Hosts
 - Monitoring Docker Swarm Clusters
 - Monitoring Docker Swarm Services
 - Alarms
- **Analytics**
 - Kubernetes analytics
 - Heatmap analytics
 - Scatter Plot analytics
 - Docker Swarm analytics
 - Heatmap analytics
 - Scatter Plot analytics
- **Domains and Object Groups**
 - Domains
 - Object Groups

Kubernetes

Monitoring Kubernetes Clusters

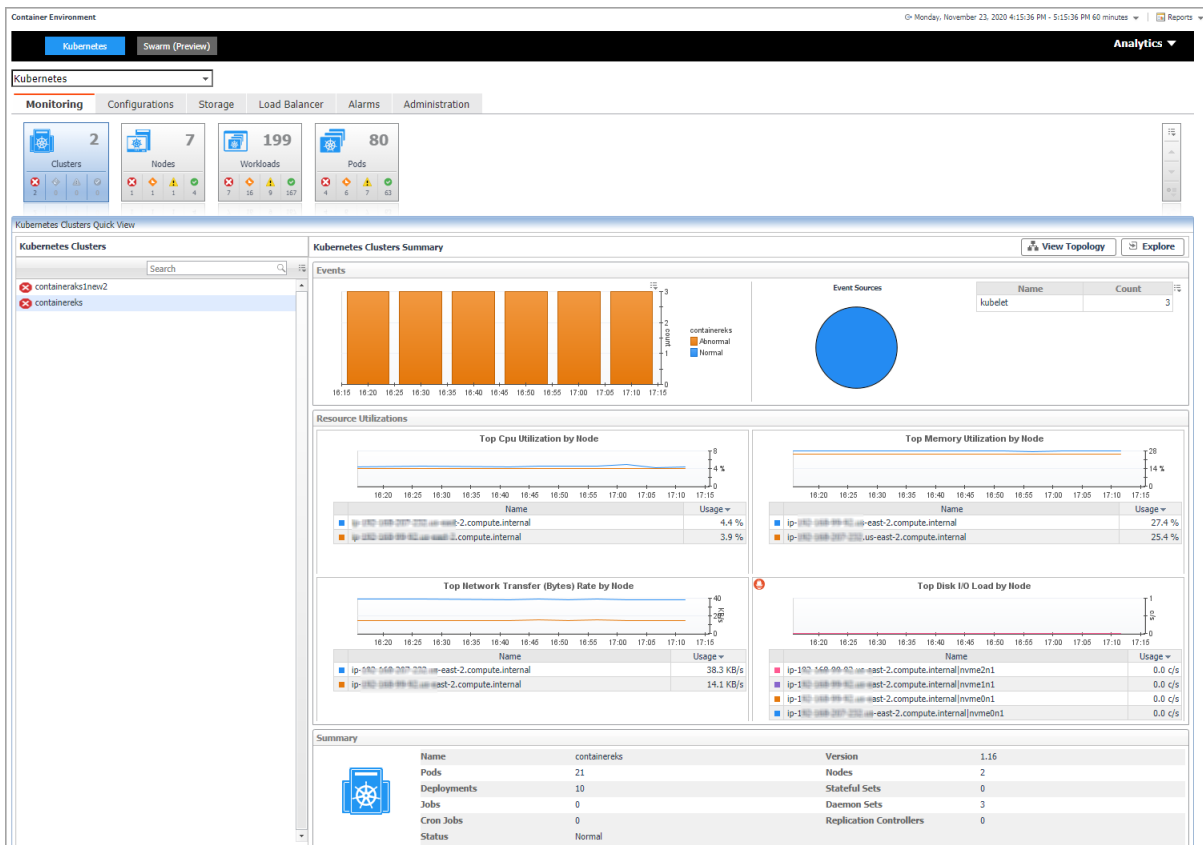
Kubernetes cluster is a group of Kubernetes resources. There are two kinds of nodes inside a cluster, Kubernetes master and Kubernetes nodes. Kubernetes master is responsible for maintaining the desired state of your cluster which Kubernetes node is responsible to run your application and cloud workflows. The *Kubernetes Cluster Quick View*, which appears after clicking **Monitoring > Clusters**, shows the data collected about the selected clusters and namespaces. This view consists of the following two panes:

- The **Kubernetes Clusters** tree view, which appears on the left of *Kubernetes Clusters Quick View*, lists the clusters existing in the monitored Kubernetes environment.
- The **Kubernetes Clusters Summary view**, which appears on the right after you select an individual cluster in the **Kubernetes Clusters** tree view.

Kubernetes Clusters Summary view

The **Kubernetes Clusters Summary** view appears on the right when you select a node in the **Kubernetes Clusters** tree view.

Figure 2. Kubernetes Clusters Summary view



The **Kubernetes Clusters Summary** view displays the following data:

- **Events:** The events occur on the selected Kubernetes cluster over a selected period of time, which includes:

- The column chart on the left: Shows the timeline of the occurred events, which indicates at what time and how many events have occurred.
 - The pie chart on the right- Event Sources: Shows the events distribution for different event source.
- **Resource Utilizations:** The resource utilization for the selected Kubernetes cluster over a selected period of time, which includes the following:
 - **Top CPU Utilization by Node:** shows the nodes that CPU Utilization are top N highest.
 - **Top Memory Utilization by Node:** shows the nodes that Memory Utilization are top N highest.
 - **Top Network Transfer (Bytes) Rate by Node:** shows the nodes that Network Transfer (Bytes) Rate are top N highest.
 - **Top Disk I/O Load by Node:** shows the nodes that Disk I/O Load are top N highest.
 - 🚨: The icon indicates this metrics is collected by Prometheus metrics collector.
- **Summary:** Displays the detailed information about the selected Kubernetes cluster, including *Name*, *Version*, *Pods*, *Nodes*, *Deployments*, *Stateful Sets*, *Jobs*, and *Replica Sets*.

Click **Explore** on the upper right of the **Kubernetes Clusters Summary** view to open the **Cluster Explorer view**, which shows more detailed information about this Kubernetes cluster.

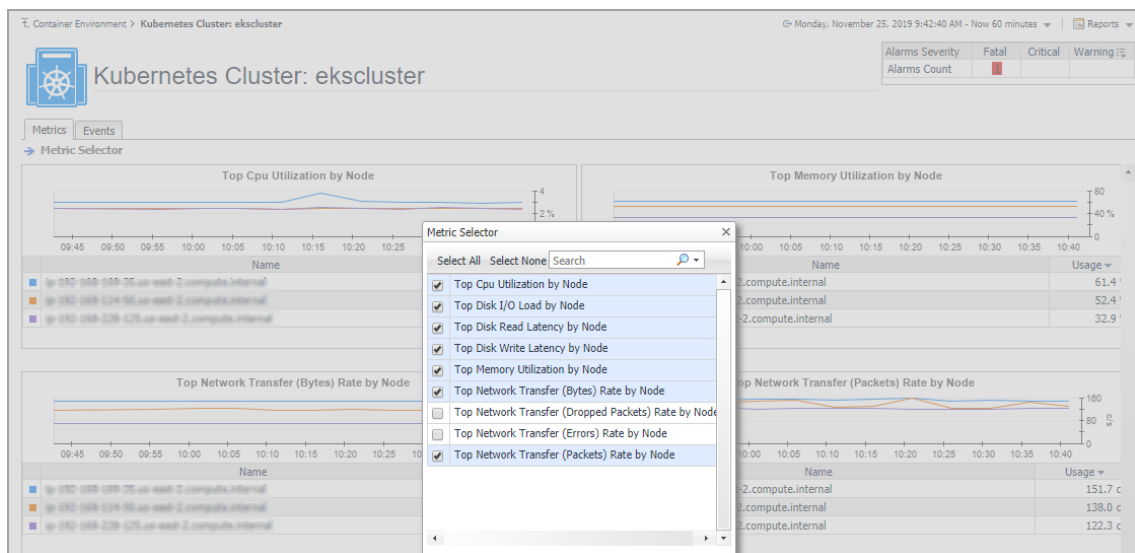
Click **View Topology** on the upper right of the **Kubernetes Clusters Summary** view to open the **Cluster Topology view**, which shows the topology graph from the application accessible aspect.

Cluster Explorer view

The *Cluster Explorer* view opens when you click **Explore** in the **Kubernetes Clusters Summary** view, which includes the following tabs:

- **Metrics tab:** The *Metrics* tab displays a *Metric Selector* allowing you to choose the metrics to be plotted on this dashboard. Charts of *CPU Usage* and *Memory Usage* are presented by default.

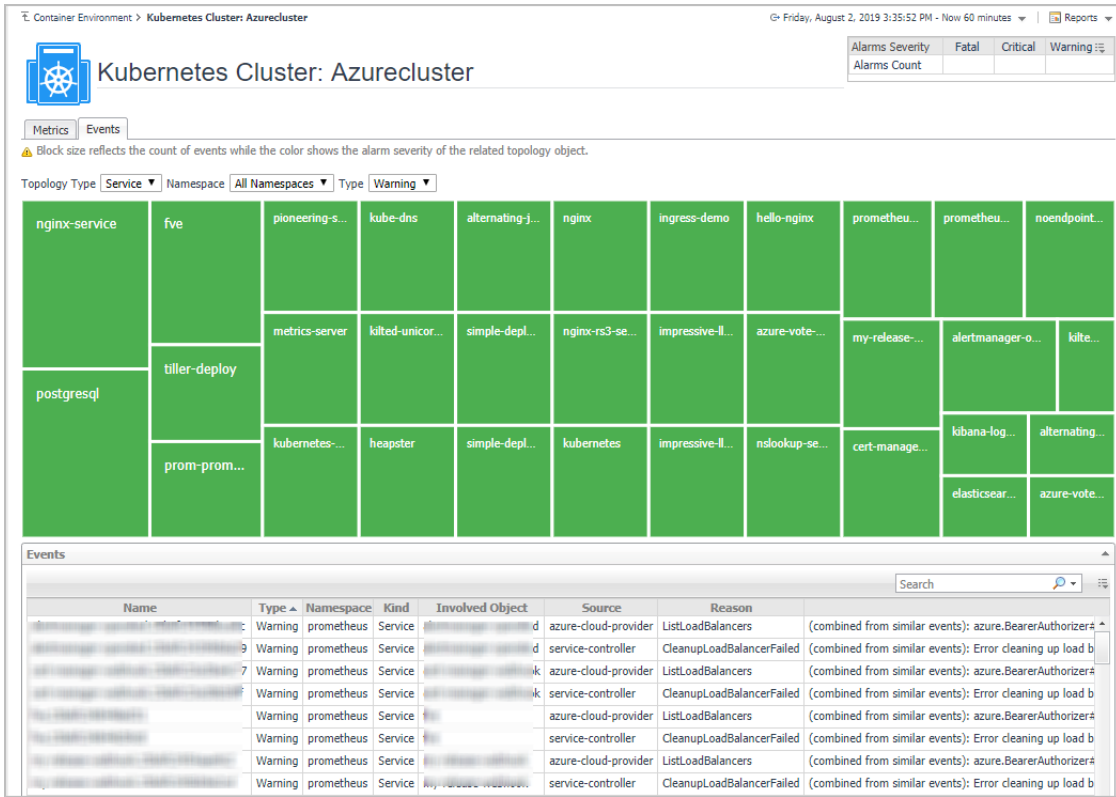
Figure 3. Kubernetes Clusters Explorer view Metrics tab



- **Events tab:** The *Events* tab shows a Heat Map of the events occur in this cluster. Heat maps will be refreshed automatically when you change either of the following fields:
 - **Topology Type:** Indicates the Kubernetes components on which the event occurs, including Pod, Node, and Service.
 - **Namespace:** Use the namespaces to filter the events.

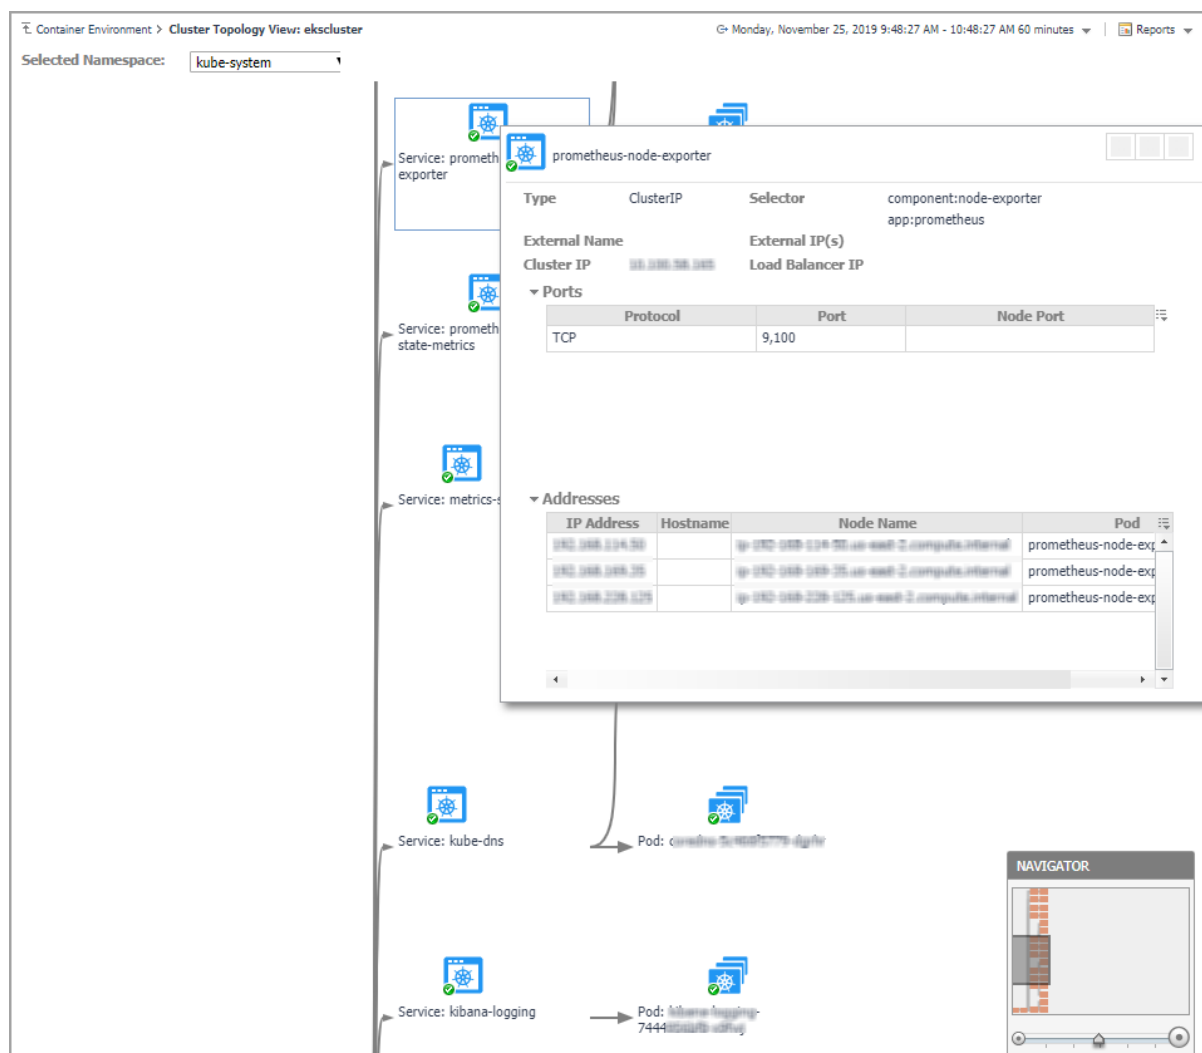
- **Type:** Indicates the severity of the event, including warning and normal.
- **NOTE:** The color in the heatmap indicates the severity of component alarms.
Green: indicates normal. Yellow: indicates warning. Orange: indicates critical. Red: indicates fatal.

Figure 4. Kubernetes Clusters Explorer view Events tab



Cluster Topology view

Figure 5. Kubernetes Clusters Topology view



The *topology* view visualizes the relationships between the objects from the pods accessible aspect in your environment through an interactive dependency map. The map illustrates how different components relate to each other, and the levels of the available resources available to them. Click on Pod, another sub topology view will popup to show the relationship from pods controller to storage for the selected Pod. Click other components or click the Pod inside the sub topology view, an information view will popup to show alarms, basic information, some metrics. From the information popup view of Pod, Node, and Cluster, click the Explore button will navigate to the explorer view of the selected Pod/Node/Cluster.

The **NAVIGATOR** in the bottom-right corner allows you to easily set the zoom level by dragging the slider into the appropriate position.

Monitoring Kubernetes Nodes

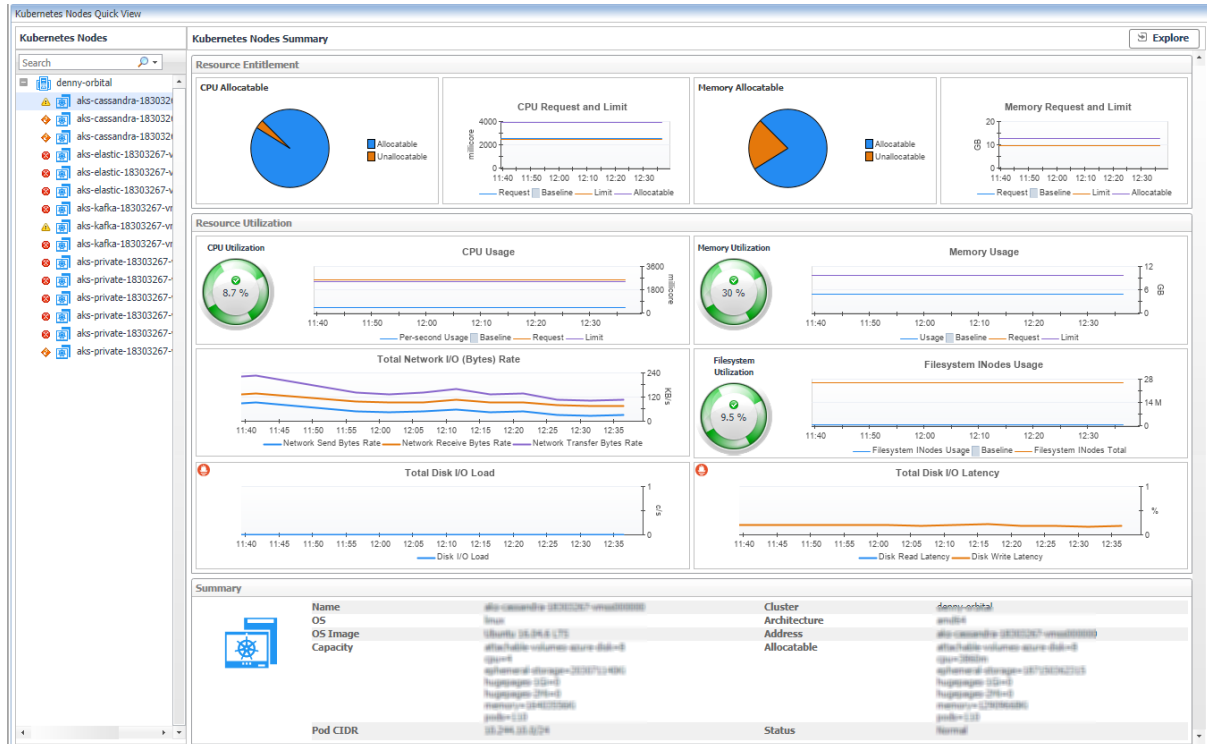
A node, previously known as a minion, is a worker machine in Kubernetes. A node may be a VM or physical machine, depending on the cluster. Each node has the services necessary to run pods and is managed by the master components. The *Kubernetes Nodes Quick View*, which appears after clicking **Monitoring > Nodes**, shows the data collected about the selected clusters and namespaces. This view consists of the following two panes:

- The **Kubernetes Nodes** tree view, which appears on the left of *Kubernetes Nodes Quick View*, lists the nodes existing in the monitored Kubernetes environment.
- The **Kubernetes Nodes Summary view**, which appears on the right after you select an individual node in the **Kubernetes Nodes** tree view.

Kubernetes Nodes Summary view


The **Kubernetes Nodes Summary** view appears on the right when you select a node in the **Kubernetes Nodes** tree view.

Figure 6. Kubernetes Nodes Summary view



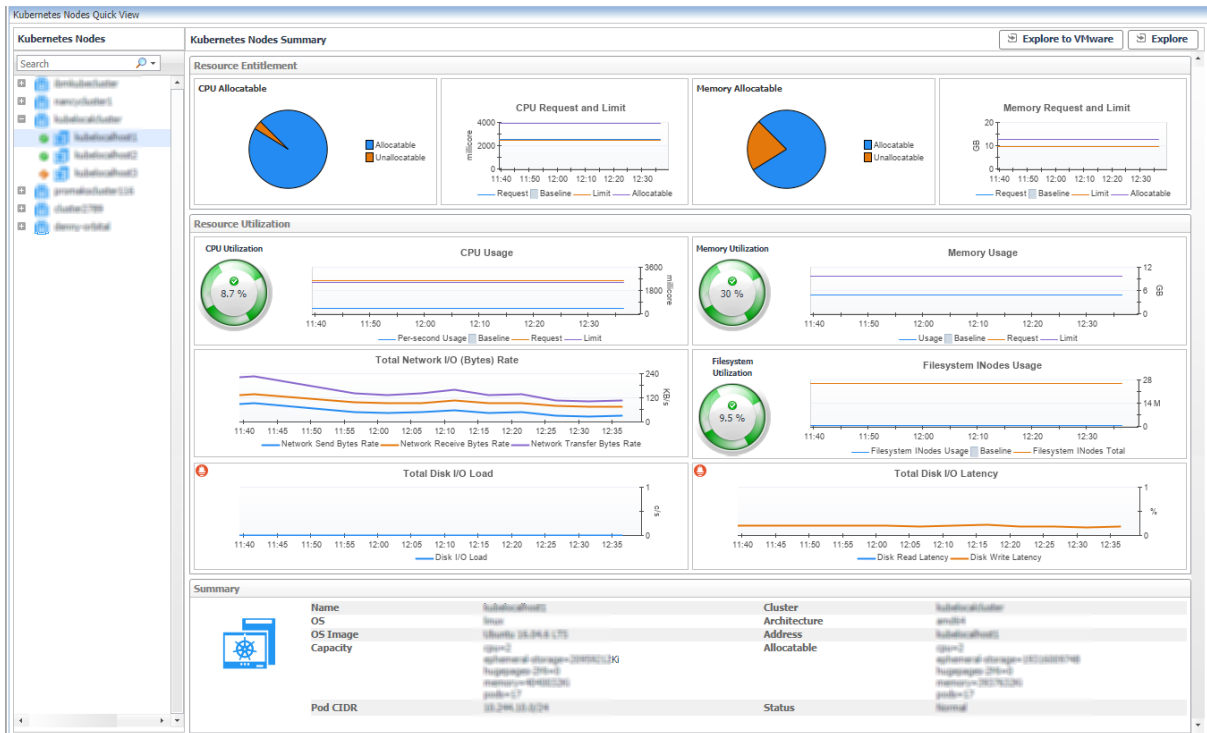
The **Kubernetes Nodes Summary** view displays the following data:

- **Resource Entitlement:** The resource allocation for the selected Kubernetes node over a selected period of time, which includes the following:
 - **CPU Allocatable:** Shows the current allocatable CPU resources of this node.
 - **Memory Allocatable:** Shows the current allocatable Memory resources of this node.
 - **CPU Request and Limit:** Shows the trend of CPU request, limit, and capacity of this node.
 - **Memory Request and Limit:** Shows the trend of Memory request, limit, and capacity of this node.
- **Resource Utilizations:** The resource utilization for the selected Kubernetes node over a selected period of time, which includes the following:
 - **CPU Utilization:** shows the percentage of CPU usage divide CPU capacity.
CPU Usage: shows the usage, request, and limit of CPU.
 - **Memory Utilization:** shows the percentage of Memory usage divide Memory capacity.
Memory Usage: shows the usage, request and limit of Memory.
 - **Total Network I/O:** shows the sending, receiving, and transferring rate in byte/s, aggregated from all interfaces.

- **Filesystem Utilization:** shows the filesystem utilization.
- **Filesystem Inodes Usage:** shows the inodes usage and total inodes.
- **Total Disk I/O Load:** shows the number of IOs in progress per second, aggregated from all disk devices.
- **Total Disk I/O Latency:** shows the read and write latency in percentage.
- : The icon indicates this metrics is collected by Prometheus metrics collector.
- **Summary:** Displays the detailed information about the selected Kubernetes node, including *Name, Pod CIDR, OS, Architecture, OS Image, Address, Capacity, Allocatable, and Status.*

Click **Explore** on the upper right of the **Kubernetes Nodes Summary** view to open the **Nodes Explorer view**, which shows more detailed information about this Kubernetes node.

Figure 7. Kubernetes Nodes Summary view for VMware



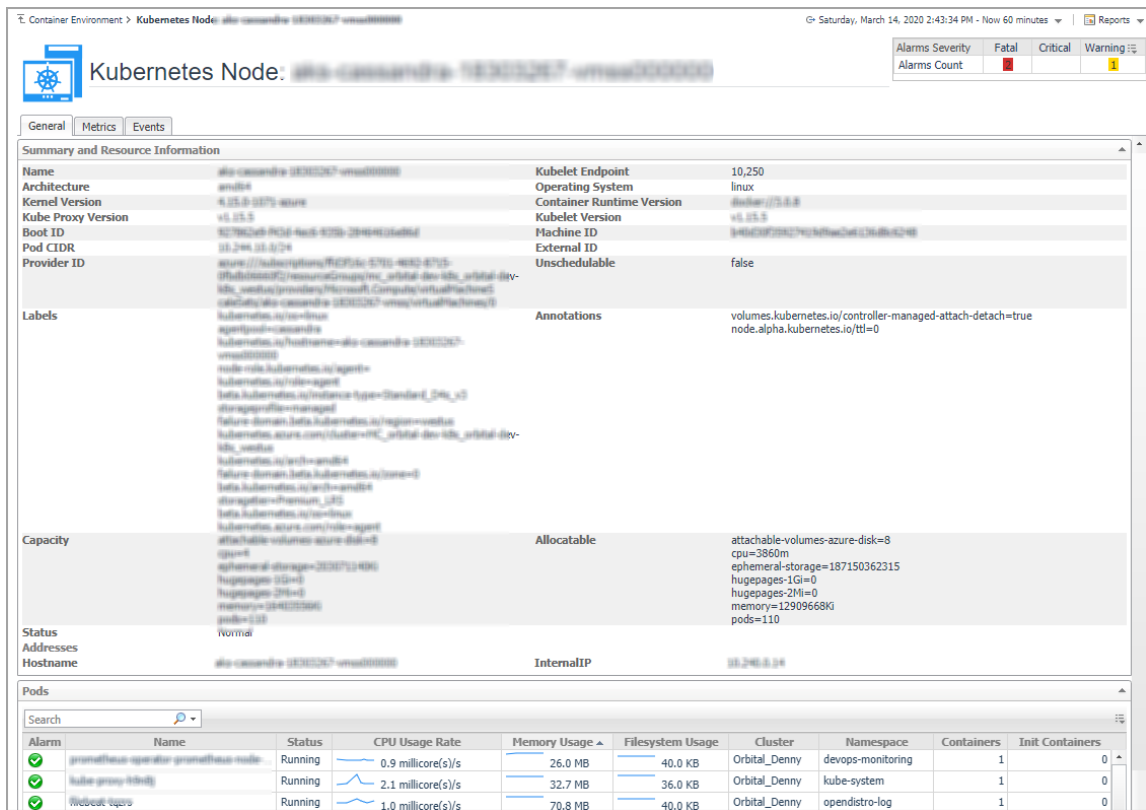
- **Explore to VMware VM:** Click the button to open the **VMware Explorer** view, which is the same view from VMware cartridge. The **Explore to xx** button varies from the cartridge that is monitoring the machines. Currently, the supported cartridges include: VMware, Infrastructure, AWS, and Azure.

Nodes Explorer view

The **Nodes Explorer** view opens when you click **Explore** in the **Kubernetes Nodes Summary** view, which includes the following tabs:

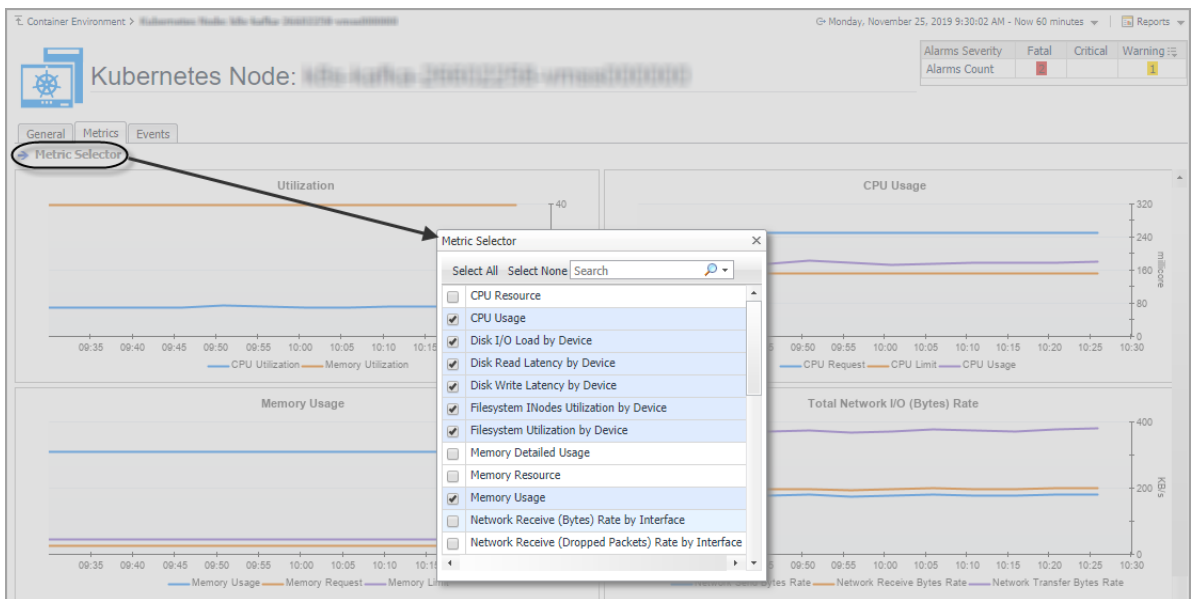
- **General tab:** The **General** tab displays the overall information of the selected Kubernetes node over a selected period of time, including the **Summary and Resource Information** table and the **Pods** table. For more information, see **Node metrics on page 77**.

Figure 8. Kubernetes Nodes Explorer view General Tab




- **Metrics tab:** The *Metrics* tab displays a *Metric Selector* allowing you to choose the metrics to be plotted on this dashboard. Charts of *CPU Usage*, *Utilization*, *Memory Usage*, and *Network I/O* are presented by default.

Figure 9. Kubernetes Nodes Explorer view Metrics Tab



- **Events tab:** The *Events* tab lists all the events occur on the nodes.
 - Name: name of the event.
 - Type: type of the event, Warning or Normal.
 - Namespace: namespace of where this event happens.
 - Kind: type of the Kubernetes component on which this event occurs.
 - Involved Object: name of the Kubernetes component on which this event occurs.
 - Source: where this event has been triggered from.
 - Reason: reason of this event.
 - Message: detailed message of this event.

Figure 10. Kubernetes Nodes Explorer view Events tab



Monitoring Kubernetes Workloads

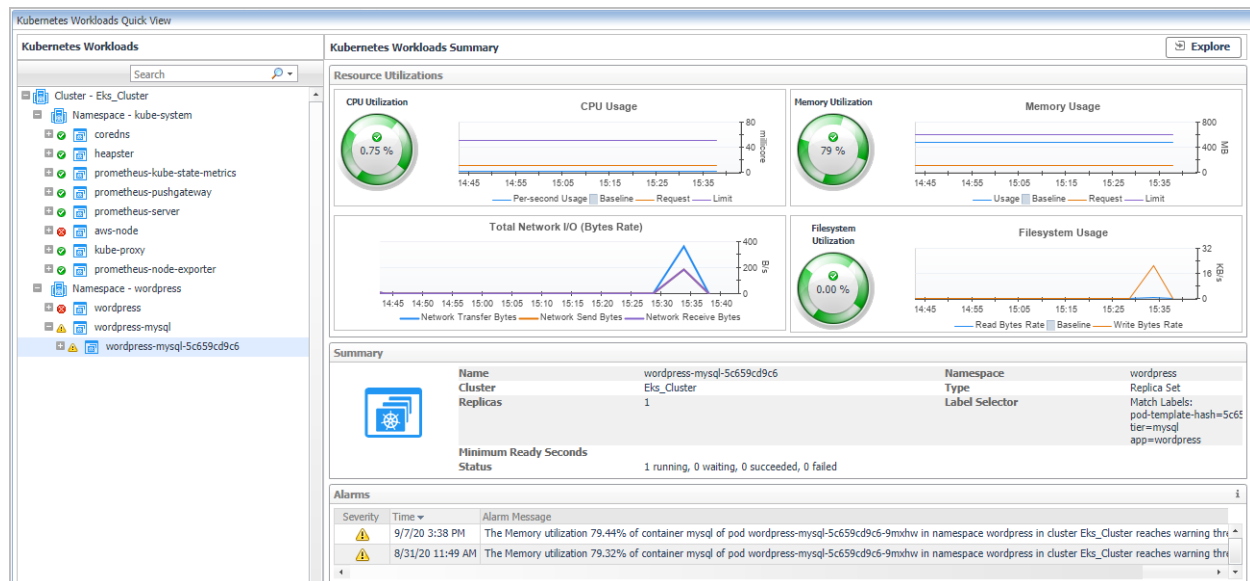
Kubernetes workload is a group of Kubernetes resources. It includes the following Kubernetes resources types: deployment, daemon set, stateful set, replica set, replication controller, cron job, job, and pod. The *Kubernetes Workloads Quick View*, which appears after clicking **Monitoring > Workloads**, shows the data collected about the selected clusters and namespaces. This view consists of the following two panes:

- The **Kubernetes Workloads** tree view, which appears on the left of *Kubernetes Workloads Quick View*, lists the workloads existing in the monitored Kubernetes environment. The menu level differs from different workload types, following are some examples of the menu levels:
 - Cluster -> Namespace -> Deployment -> Replica Set -> Pod
 - Cluster -> Namespace -> Cron Job -> Job -> Pod
 - Cluster -> Namespace -> Daemon Set -> Job -> Pod
 - Cluster -> Namespace -> Stateful Set -> Pod
- The [Kubernetes Pods Summary view](#), which appears on the right after you select an individual pod in the **Kubernetes Workloads** tree view.

Kubernetes Workloads Summary view

The **Kubernetes Workloads Summary** view appears on the right when you select a workload in the **Kubernetes Workloads** tree view.

Figure 11. Kubernetes Workloads Summary view



The **Kubernetes Workloads Summary** view displays the following data:

- NOTE:** If the selected workload type is a Pod, the metrics values are from the Pod. If the selected workload type is other than a Pod, then the metrics values are the aggregated values from all the related pods of the current workload (e.g. Deployment, Daemon Set, and so on).
- **Resource Utilizations:** The resource utilization for the selected Kubernetes Workloads over a selected period of time, which includes the following:
 - **CPU Utilization:** The CPU utilization is calculated as the percentage of CPU usage rate and limit, and it will be unset or blank if the CPU limit is not configured.
CPU Usage: A comparison of per-second usage, request, and limit.
 - **Memory Utilization:** The Memory utilization is calculated as the percentage of memory usage and limit, and it will be unset or blank if the memory limit is not configured.
Memory Usage: A comparison of usage, request, and limit.
 - **Total Network I/O (Bytes Rate):** transferring, sending, and receiving rate in bytes.
 - **Filesystem Utilization:** temporary filesystem utilization.
Filesystem Usage: temporary filesystem reading or writing rate in bytes.
 - **Summary:** Displays the detailed information about the selected Kubernetes workload. The listed information might vary from different types of workload.
 - **Alarms:** Displays the alarms aggregated from the pod and the current workload related to the pod.

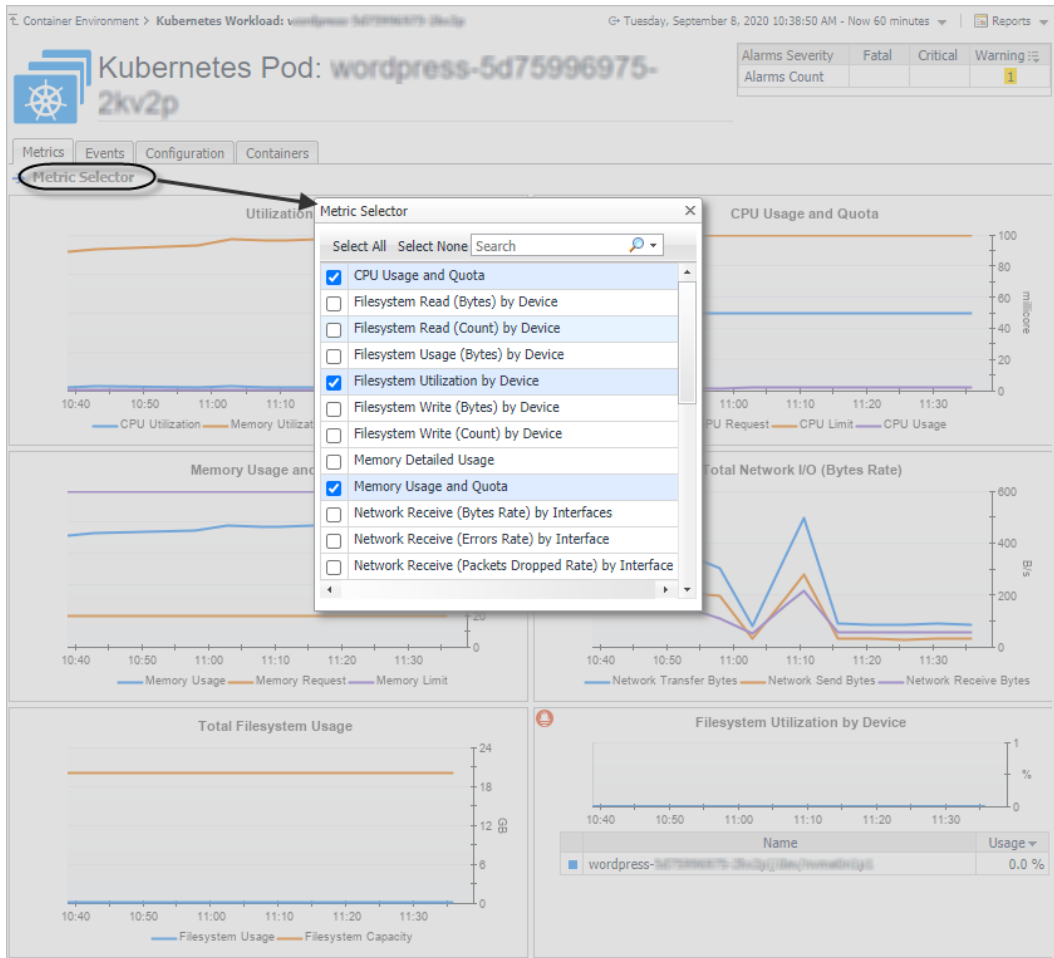
Click **Explore** on the upper right of the **Kubernetes Workloads Summary** view to open the [Workloads Explorer view](#), which shows more detailed information about this Kubernetes workload.

Workloads Explorer view

The *Workloads Explorer* view opens when you click **Explore** in the [Kubernetes Workloads Summary view](#), which includes the following tabs:

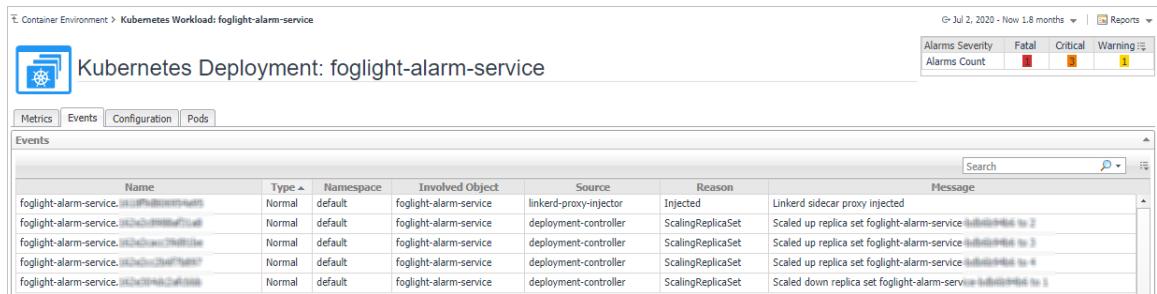
- **Metrics tab:** The *Metrics* tab displays a *Metric Selector* allowing you to choose the metrics to be plotted on this dashboard. Charts of *Utilization, CPU Usage and Quota, Memory Usage and Quota, Total Network I/O, Total Filesystem Usage, and Filesystem Utilization by Device* are presented by default.
 - NOTE:** If the current workload is a Pod, then the metrics values are from the Pod. Otherwise, the metrics values come from aggregated values of the related pods of this workload (e.g. Deployment, Daemon Set).

Figure 12. Kubernetes Workloads Explorer view Metrics Tab



- **Events tab:** The *Event* tab lists all the events occur on the workloads. It will not aggregate the alarms from other related workloads.
 - **Name:** name of the event.
 - **Type:** type of the event, Warning or Normal.
 - **Namespace:** namespace of where this event happens.
 - **Involved Object:** name of the Kubernetes component on which this event occurs.
 - **Source:** where this event has been triggered from.
 - **Reason:** reason of this event.
 - **Message:** detailed message of this event.

Figure 13. Kubernetes Workloads Explorer view Events tab




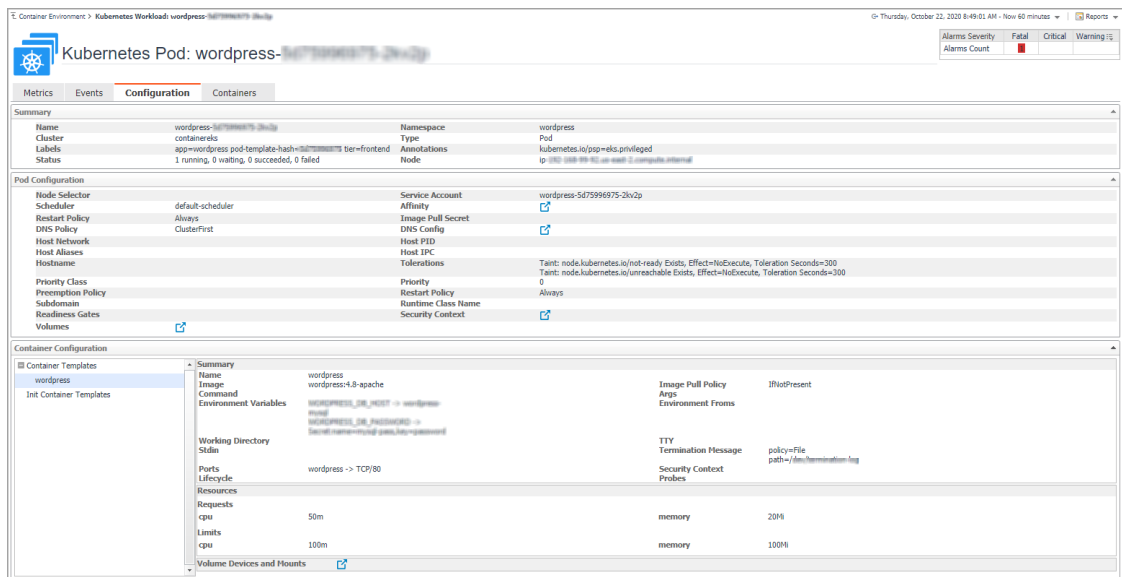
- **Configuration tab:** The *Configuration* tab displays the overall information of the selected Kubernetes Workload over a selected period of time. It includes the following tables:
 - **Summary:** Summary information of this workload.
 - **Pod Configuration:** Pod related configuration. If the current workload is a pod, it displays the configuration of the pod. Otherwise, it displays the pod template information of the current workload to create related pods.
 - **Container Configuration:** Including containers configurations and Init-containers configurations. If the current workload is a pod, it displays the configuration of the pod containers. Otherwise, it displays the container template of the pod containers created by this workload.
 - Click  to see detailed information on this item.

Figure 14. Kubernetes Workloads Explorer view Configuration Tab



- **Pods tab:** The *Pods* tab displays the related Pods of the workload. This tab is not displayed if the workload type is a Pod.



Click the Pod name or the  button will navigate to the related Pod explorer view.

Figure 15. Kubernetes Workloads Explorer view Pods Tab



- **Containers tab:** The *Containers* tab displays the related created by this pod. This tab is displayed only when the workload type is a Pod.



Click the Container name or the  button will navigate to the [Container Dashboard](#).

Figure 16. Kubernetes Workloads Explorer view Containers Tab

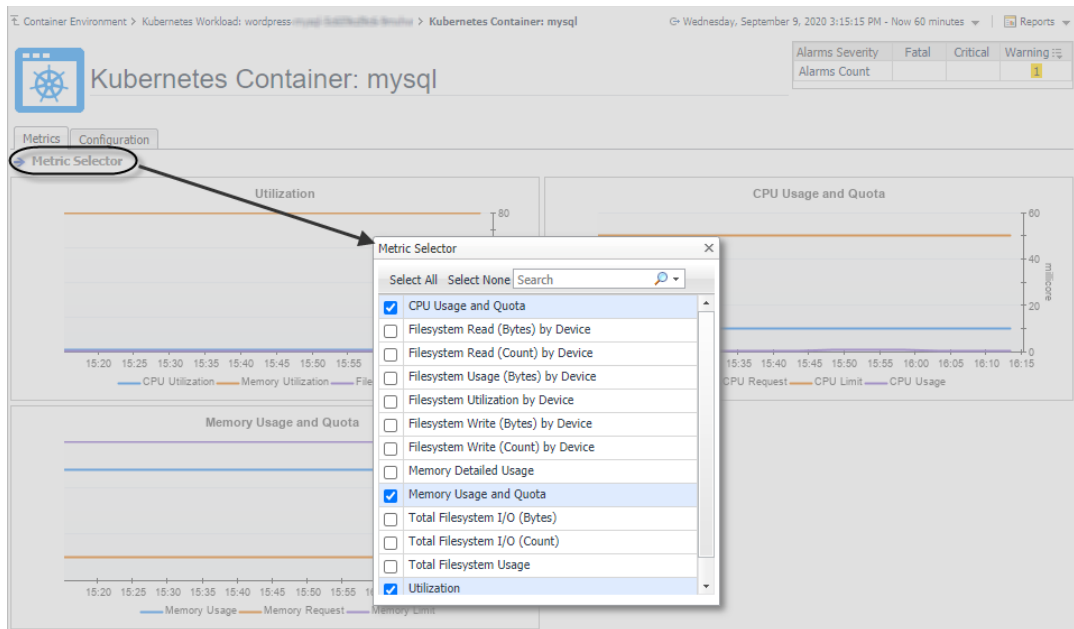


Container Dashboard

- *Metrics* tab on Container Dashboard

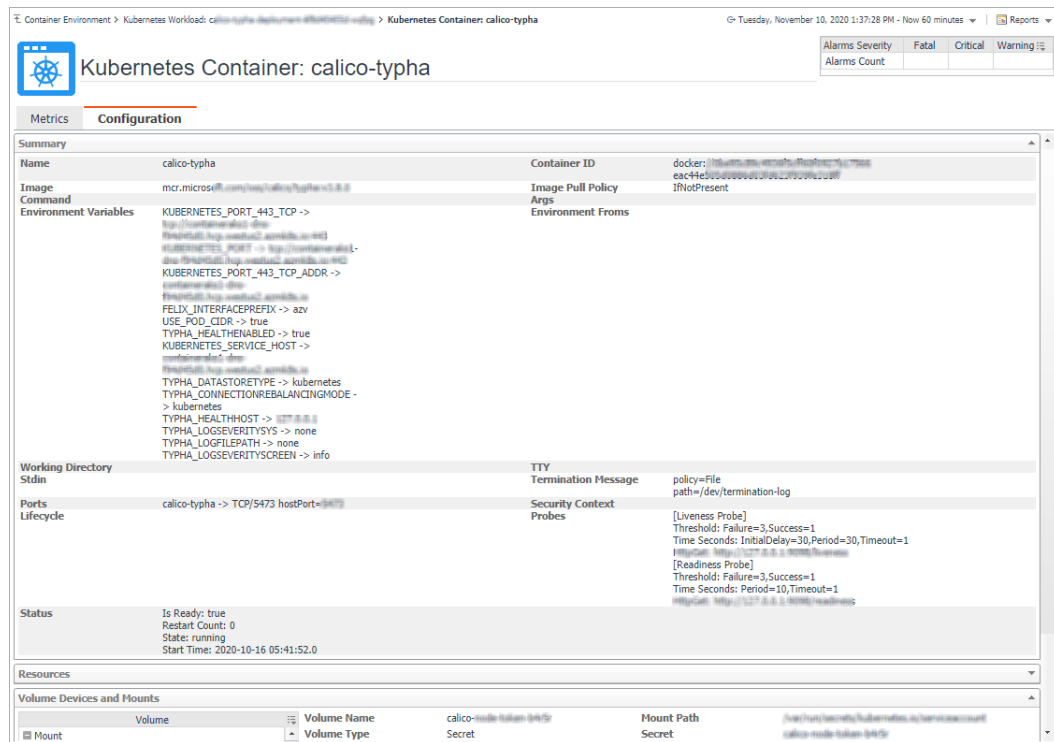
The *Metrics* tab displays a *Metric Selector* allowing you to choose the metrics to be plotted on this dashboard. Charts of *Utilization*, *CPU Usage and Quota*, and *Memory Usage and Quota* are presented by default.

Figure 17. Metrics tab on Container Dashboard



- *Configurations* tab on Container Dashboard
 - **Summary:** summary information of this container.
 - **Resources:** resource request and limit configured for this container.
 - **Volume Devices and Mounts:** volumes mounted by this container.

Figure 18. Configurations tab on Container Dashboard



Monitoring Kubernetes Pods

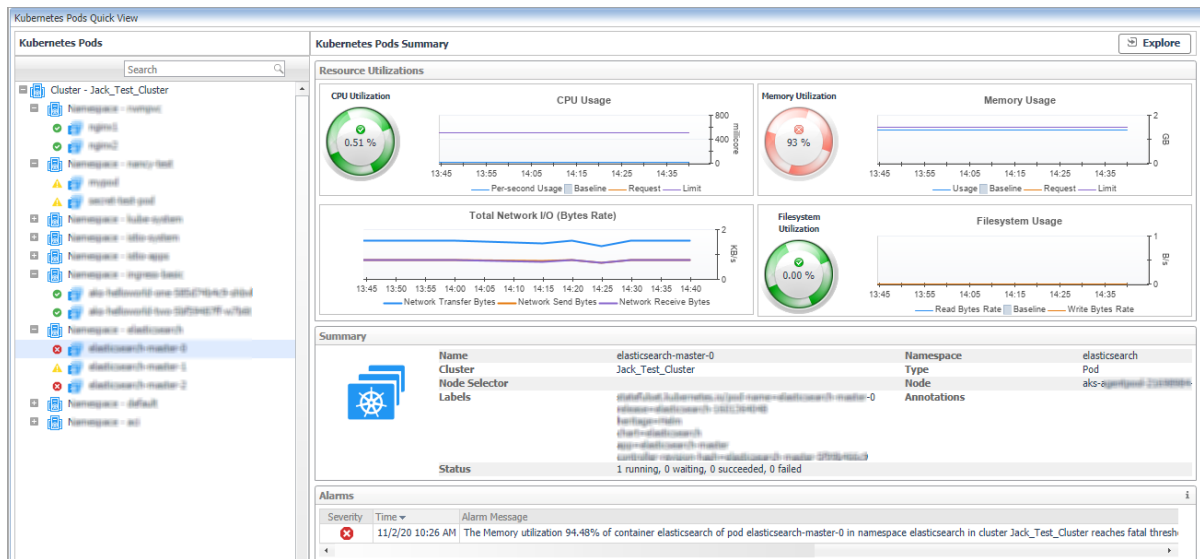
A pod contains one or multiple containers, such as Docker containers, which contains storage/network and the specification about how to run the containers. The *Kubernetes Pods Quick View*, which appears after clicking **Monitoring > Pods**, shows the data collected about the selected clusters and namespaces. This view consists of the following two panes:

- The **Kubernetes Pods** tree view, which appears on the left of *Kubernetes Pods Quick View*, lists the pods existing in the monitored Kubernetes environment.
- The **Kubernetes Pods Summary** view, which appears on the right after you select an individual pod in the **Kubernetes Pods** tree view.

Kubernetes Pods Summary view

The **Kubernetes Pods Summary** view appears on the right when you select a cluster in the **Kubernetes Pods** tree view.

Figure 19. Kubernetes Pods Summary view



The **Kubernetes Pods Summary** view displays the following data:

- **Resource Utilizations:** The resource utilization for the selected Kubernetes Pod over a selected period of time, which includes the following:
 - **CPU Utilization:** shows the percentage of CPU usage divides CPU limit, if the limit is configured for this Pod.
 - **CPU Usage:** shows CPU usage rate, request, and limit, the trend of usage rate, as well as request and limit configuration.
 - **Memory Utilization:** shows the percentage of Memory usage divides Memory limit, if the limit is configured for this Pod.
 - **Memory Usage:** shows Memory usage, request, and limit, the trend of usage rate, as well as request and limit configuration.
 - **Total Network I/O:** shows the sending, receiving, and transferring rate.
 - **Filesystem Utilization:** shows the Filesystem utilization.
 - **Filesystem Usage:** shows the read and write rate in byte/s.
- **Summary:** Displays the detailed information about the selected Kubernetes Pod, including *Name*, *Node*, *Cluster*, *Namespace*, *Owner*, *Pod IP*, *Service Account*, *DNS Policy*, *Restart Policy*, and *Status*.

Click **Explore** on the upper right of the **Kubernetes Pods Summary** view to open the **Pods Explorer** view, which shows more detailed information about this Kubernetes cluster.

Pods Explorer view

The *Pods Explorer* view opens when you click **Explore** in the **Kubernetes Pods Summary** view. For detailed information for each tabs in Pods Explorer view, see **Workloads Explorer** view on page 36.

Monitoring Kubernetes Other Components

Kubernetes other components here including pods controllers, services, ingresses, persistent volumes, secrets and so on. All these components are grouped and displayed in tabs.

- [Configurations](#)
- [Storage](#)
- [Load Balancer](#)

Configurations

Figure 20. Kubernetes Configuration Dashboard

Name	Cluster	Namespace	Labels	Annotations	Configured Data Keys
cluster-info	localckacluster	kube-public			kubeconfig
coredns	localckacluster	kube-system			Corefile
extension-apiserver-authentication	localckacluster	kube-system			client-ca-file, requestheader-extra-headers-prefix, requestheader-die...
kube-flannel-cfg	localckacluster	kube-system	[app=flannel], [tier=node]		net-conf.json, cni-conf.json
kube-proxy	localckacluster	kube-system	[app=kube-proxy]		config.conf, kubeconfig.conf
kubeadm-config	localckacluster	kube-system			ClusterStatus, ClusterConfiguration
kubelet-config-1.13	localckacluster	kube-system			kubelet
metrics-server-config	localckacluster	kube-system	[addonmanager.kubernetes.io/m...		NannyConfiguration
fair-lambkin-elasticsearch-curator-config	nancyakcluster	default	[heritage=Tiller], [app=fair-lamb...		config.yml, action_file.yml
impressive-llama-mariadb-master	nancyakcluster	default	[heritage=Tiller], [app=mariadb...		my.cnf
impressive-llama-mariadb-slave	nancyakcluster	default	[component=slave], [release=im...		my.cnf
impressive-llama-mariadb-tests	nancyakcluster	default			run.sh
metricbeat-config	nancyakcluster	default	[k8s-app=metricbeat], [app=fair...		metricbeat.yml
metricbeat-modules	nancyakcluster	default	[component=fair-lambkin-elastic...		system.yml, kubernetes.yml
sysdig-agent	nancyakcluster	default		[kubectf.kubernetes.io/last-appli...	dragent.yml
understood-zebra-elasticsearch-curator-config	nancyakcluster	default	[release=understood-zebra], [he...		action_file.yml, config.yml
aks-nodepool1-11370379-0-config-5fgt4dhcbf	nancyakcluster	kube-system			kubelet

The *Configurations* dashboard includes Kubernetes Secret and Config Map.

- A Kubernetes Secret is an object that contains a small amount of sensitive data, such as a password, a token, or a key. Such information might otherwise be put in a Pod specification or in an image; putting it in a Secret object allows for more control over how it is used, and reduces the risk of accidental exposure.
- A Kubernetes Config Map binds configuration files, command-line arguments, environment variables, port numbers, and other configuration artifacts to your Pods' containers and system components at runtime. Config maps allow you to separate your configurations from your Pods and components, which helps keep your workloads portable, makes their configurations easier to change and manage, and prevents hardcoding configuration data to Pod specifications.

Storage

Figure 21. Kubernetes Storage Dashboard

Alarms	Status	Name	Cluster	Reclaim Policy	Claim	Storage Class	Source Type
✓	Bound	pv-sc-no-sc-customize	localckacluster	Retain	pvc-sc-pv-customize-sc	no-sc-customize	HostPath
✓	Available	pv-sc-default	localckacluster	Retain		default	HostPath
✓	Available	pv-sc-invalid-provisioner	localckacluster	Retain		sc-invalid-provisio...	HostPath
✓	Bound	pv-invalid-nfs	localckacluster	Recycle	pvc-invalid-sc-pv	slow	NFS
✓	Available	pv-pvc-oversize	localckacluster	Retain		sc-oversize	HostPath
✓	Available	pv-pvc-acm1	localckacluster	Retain		sc-pvc-acm1	HostPath
✓	Bound	pvc-2b95e22d-dc28-11e8-b2ed-befa22179703	nancyakcluster	Delete	data-mehdb-1	default	AzureDisk
✓	Bound	pvc-45f1fe1e-5f54-11e9-b660-16063de8b09f	nancyakcluster	Delete	data-elasticsearch-2	default	AzureDisk
✓	Bound	pvc-59cb23a5-fd17-11e8-adf4-de8994810bc3	nancyakcluster	Delete	data-elasticsearch-0	default	AzureDisk
✓	Bound	pvc-7049bcb8-fd17-11e8-adf4-de8994810bc3	nancyakcluster	Delete	data-elasticsearch-1	default	AzureDisk
✓	Bound	pvc-90f76a94-2e94-11e9-810c-0a130f143c9f	nancyakcluster	Delete	alertmanager-prom-prometheus-operator-alertmanager-db-al...	default	AzureDisk
✓	Bound	pvc-9727ba6-2e94-11e9-810c-0a130f143c9f	nancyakcluster	Delete	prometheus-prom-prometheus-operator-prometheus-db-prom...	default	AzureDisk

The Kubernetes storage contains volumes, storage class, persistent volume, and persistent volume claim. Volumes are on-disk files used by the containers for persistent their data as well as sharing with other containers. The *Storage* dashboard shows the information about the following storage classes:

- *Storage Class* provides a way for the administrator to describe the "class" of storage they offer.
- *Persistent Volume* subsystem provides an API for users and administrators that abstracts details of how storage is provided from how it is consumed.
- *Persistent Volume Claim* is used for dynamic volume provisioning which allow storage volumes to be created on-demand.

Load Balancer

Figure 22. Kubernetes Load Balancer Dashboard

Alarms	Name	Cluster	Namespace	Type	Cluster IP	External IPs	Load Balancer Addresses	Ports	External Traffic Policy	Endpoint
✓	fve	ekscluster	default	ClusterIP	10.100.211.210	3.142.294.176		[TCP-8080, Target ...]	Cluster	fve
✓	guestbook	ekscluster	default	LoadBalancer	10.100.240.79		elb-93292721-148963762871338-13237621.us-east-2.elb.amazonaws.com	[TCP-3000, Target ...]	Cluster	guestbook
✓	kubernetes	ekscluster	default	ClusterIP	10.100.0.1			[TCP-443, Target P...]		kubernetes
✓	nginx-statefulset	ekscluster	default	ClusterIP	None			[TCP-80, Target Po...]		nginx-statefulset
✓	nginx-stateless	ekscluster	default	ClusterIP	None			[TCP-80, Target Po...]		nginx-stateless
✓	redis-master	ekscluster	default	ClusterIP	10.100.94.188			[TCP-6379, Target ...]		redis-master
✓	redis-slave	ekscluster	default	ClusterIP	10.100.94.227			[TCP-6379, Target ...]		redis-slave
✓	elasticsearch-logging	ekscluster	kube-system	LoadBalancer	10.100.188.40			[TCP-9200, Target ...]	Cluster	elasticsearch-logging
✓	heapster	ekscluster	kube-system	ClusterIP	10.100.212.40			[TCP-80, Target Po...]		heapster
✓	kibana-logging	ekscluster	kube-system	LoadBalancer	10.100.212.173			[TCP-5601, Target ...]	Cluster	kibana-logging
✓	kube-dns	ekscluster	kube-system	ClusterIP	10.100.0.10			[UDP-53, Target P...]		kube-dns
✓	metrics-server	ekscluster	kube-system	ClusterIP	10.100.170.200			[TCP-443, Target P...]		metrics-server
✓	prometheus-kube-state-metrics	ekscluster	kube-system	ClusterIP	10.100.278.81			[TCP-80, Target Po...]		prometheus-kube-stat
✓	prometheus-node-exporter	ekscluster	kube-system	ClusterIP	10.100.58.185			[TCP-9100, Target ...]		prometheus-node-exp
✓	prometheus-pushgateway	ekscluster	kube-system	ClusterIP	10.100.15.38			[TCP-9091, Target ...]		prometheus-pushgate
✓	prometheus-server	ekscluster	kube-system	ClusterIP	10.100.288.135			[TCP-80, Target Po...]		prometheus-server
✓	prometheus-server-lb	ekscluster	kube-system	LoadBalancer	10.100.71.138		ap-0054-0000000112eab2776a779f00000-000770000.us-east-2.elb.amazonaws.com	[TCP-80, Target Po...]	Cluster	prometheus-server-lb
✓	tiler-deploy	ekscluster	kube-system	ClusterIP	10.100.124.53			[TCP-44134, Targe...]		tiler-deploy
✓	dashboard-metrics-scraper	ekscluster	kubernetes-dashboard	ClusterIP	10.100.188.227			[TCP-8000, Target ...]		dashboard-metrics-scr

The *Load Balancer* dashboard includes information about Kubernetes service, endpoint, and ingress. A Kubernetes ingress can provide load balancing, SSL termination, and name-based virtual hosting. A Kubernetes service is an abstraction which defines a logical set of pods and a policy by which to access them - sometime called micro-services. Kubernetes will update the endpoint whenever the set of pods in a service changes.

Alarms

Figure 23. Kubernetes Alarms Dashboard

Time	Severity	Ack'd	Impacting	Source	Rule Name	Alarm Message
5/15/19 1:22 AM	Warning	false	etcd-kubeckamaster	KubePod	Kubernetes Pod Memory	The cpu usage 131.2mb of pod etcd-kubeckamaster in namespace kube-system in cluster localcluster over threshold
5/15/19 1:11 AM	Warning	false	postgresql-6558cdf45c-4s59k	KubePod	Kubernetes Pod Memory	The cpu usage 245.9mb of pod postgresql-6558cdf45c-4s59k in namespace questfive in cluster nancyakscluster over three
5/14/19 11:36 PM	Critical	false	fve-app-76ccb864c-968z6	KubePod	Kubernetes Pod Memory	The memory utilization 99.32% of pod fve-app-76ccb864c-968z6 in namespace default in cluster localcluster reaches
5/14/19 11:36 PM	Critical	false	kube-flannel-ds-amd64-j5n62	KubePod	Kubernetes Pod Memory	The memory utilization 94.75% of pod kube-flannel-ds-amd64-j5n62 in namespace kube-system in cluster localcluster158
5/14/19 11:36 PM	Critical	false	kube-apiserver-kubeckamaster	KubePod	Kubernetes Pod Memory	The memory usage 549.9mb of pod kube-apiserver-kubeckamaster in namespace kube-system in cluster localcluster c
5/14/19 11:36 PM	Critical	false	fglam-f7795b8-65k2c	KubePod	Kubernetes Pod Memory	The memory usage 868.2mb of pod fglam-f7795b8-65k2c in namespace default in cluster localcluster over threshold
5/14/19 11:36 PM	Critical	false	fve-app-76ccb864c-968z6	KubePod	Kubernetes Pod Memory	The memory usage 3.9gb of pod fve-app-76ccb864c-968z6 in namespace default in cluster localcluster over threshol
5/14/19 11:36 PM	Critical	false	tqa-loadtest-jmeter-slaves-768d8c4dc9-sk	KubePod	Kubernetes Pod Memory	The memory usage 1003.4mb of pod tqa-loadtest-jmeter-slaves-768d8c4dc9-sk in namespace tqa in cluster localclust
5/14/19 11:36 PM	Critical	false	jmeter-operator-784bd76967-kfvrx	KubePod	Kubernetes Pod Memory	The memory usage 386.6mb of pod jmeter-operator-784bd76967-kfvrx in namespace kube-system in cluster localcluster.

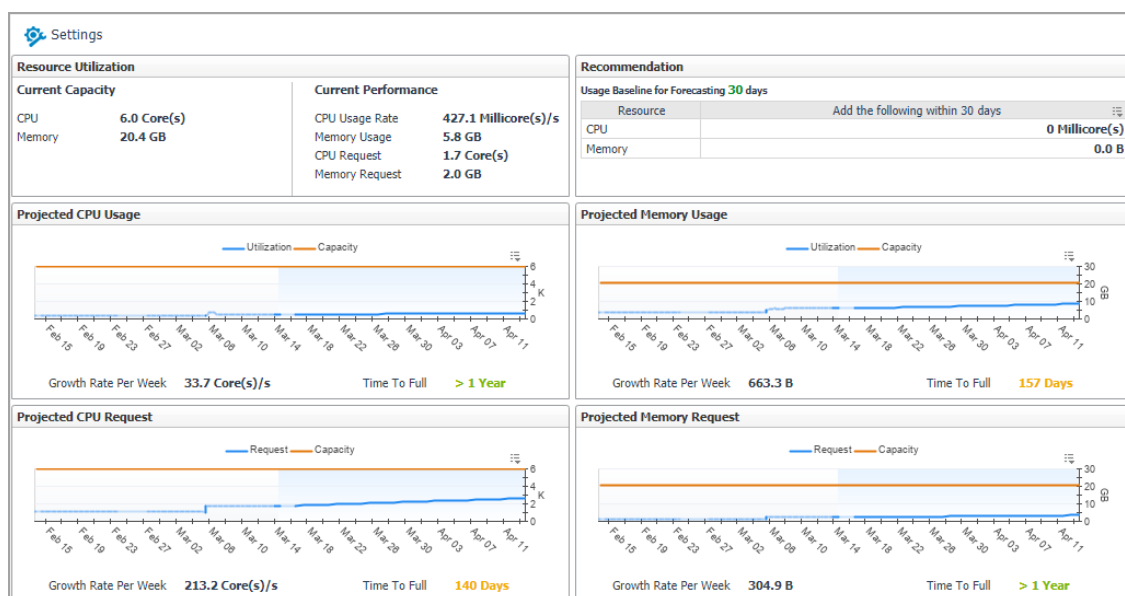
The *Alarms* dashboard displays a list of alarms generated against the monitored Kubernetes environment. Use this view to quickly identify any potential problems related to a specific Kubernetes component.

Capacity Management

Foglight for Container Management provide capacity management feature for Kubernetes. This feature uses historical data to predict the trend and usage within a specific future period.

- NOTE:** If the Capacity Management tab is not displayed, ensure the following:
1. You have purchased a license for Capacity Management. If not, contact Quest Support to purchase a license.
 2. You have the Container Administrator role.

Figure 24. Capacity Management for Kubernetes



The Capacity Management dashboard contains the following fields:

- **Setting:** Click to change the following values:
 - **Baseline for Forecasting:** Defines the historical period used for the calculations of metric views, current capacity, and recommended resources in the Resource Utilization view. The default value is *60 Days Trending*.
 - **Time Frame:** Defines the predicted period for calculating metric views, current capacity, and recommended resources in the Resource Utilization view. The default value is *Next 30 Days*.
- **Resource Utilization:**
 - **Current Capacity:** current resource capacity.
 - **Current Performance:** current resource usage.
- **Recommendation:**

In this section, it shows how many resources are recommended to be added in the current trend, so as to meet the predicated usage.
- **Projected CPU/Memory Usage:** Shows the historical data and the predicted usage trend within the configured future period.
- **Projected CPU/Memory Request:** Shows the historical data and the predicted request trend within the configured future period.
 - **Utilization:** usage.
 - **Capacity:** upper bound which the usage might reach.
 - **Growth Rate per Week:** growth amount of the resource.

- *Time to Full*: how many days the resource usage/request will reach the capacity.
- i** | **NOTE:** If a value *Never* is displayed at *Time to Full*, which means the usage/request trend is declining and the usage/request will never reach the capacity.

Cost

Foglight for Container Management provide cost management feature for Kubernetes. The *Cost* view appears after clicking **Container** > **Kubernetes** > **Cost**. Currently, the Kubernetes costs for AKS and EKS environment are calculated.

- i** | **NOTE:** If the Cost tab is not displayed, ensure the following:
1. You have purchased a license for Container Cost. If not, contact Quest Support to purchase a license.
 2. You have the Container Administrator role.

Prerequisites

Container cost does not collect costs from cloud providers, but uses the cost data collected by cloud agents for parsing the Container environment cost. The cost start date is the first date of the current month.

To use the cost data, ensure the following:

- For AKS cost, an Azure agent with the version 1.9.0 or above is required. The agent should be enabled to collect the inventory data of the AKS cluster resource group and that of the related infrastructure resource group, as well as the subscription cost related to the AKS cluster.
 - For EKS cost, an AWS agent with the version 1.9.0 or above is required. The agent should be enabled to collect the regional inventory data related to the EKS cluster and the account cost related to the EKS cluster.
- i** | **NOTE:** Cost metrics should be configured for cloud agents to collect cost data. Refer to the *Cost Tab* section in *Foglight Hybrid Cloud Manager User and Administration Guide* for detailed information.
- The Kubernetes agent collects the inventory and metrics data. Otherwise, the cost data won't be displayed if there is no Kubernetes data on that day.
 - The clusters have been assigned to BUs. Otherwise, the clusters costs won't be displayed in Cost Dashboard.

Administration dashboard

Administration dashboard enables you to manage the relationships between your cluster and BU, configure budgets at cluster level, and view summarized information about cluster, BU, and cost.

You need to assign cluster to BU here in order to see cost metrics aggregated at different levels in Container Cost dashboards.

Business Units tab

Business Units tab enables you to manage your BU for your Organizations. You can create different BUs to manage your clusters.

Figure 25. *Business Units* tab

Cost data will not be effected by the selected cluster on above.

Kubernetes Administration

Clusters Business Units

+ Add - Delete Assign Organization Search

<input type="checkbox"/>	Edit	Name ^	Location	Organization	Clusters
<input checked="" type="checkbox"/>		Foglight	n/a	Quest Organization	1
<input type="checkbox"/>		Rainbow	n/a	Quest Organization	1
<input type="checkbox"/>		Unallocated	n/a	Default Organization	0

The *Business Units* tab displays the following information:

- *Name*: BU name.
- *Location*: Geographic location of the BU.
- *Organization*: Organization of this BU.
- *Clusters*: numbers of the Kubernetes clusters that have been assigned to this BU.
- *Add*: Add a new Business Unit and assign it to a specific organization.
- *Delete*: Delete a Business Unit.
- : Click the button to edit the information of this Business Unit.
- *Assign Organization*: click the button to assign Business Unit to another Organization.

Clusters tab

Kubernetes clusters are collected by Kubernetes agents. By default, clusters will not be assigned to any BU and do not have a default monthly budget.

- i** **NOTE:** For EKS cluster, Container cost will automatically detect the EKS cluster collected by AWS agent for the cluster. However, if different AWS agent collects the same EKS cluster resources, then you need to assign the AWS account for this cluster.

Figure 26. *Clusters* tab

Cost data will not be effected by the selected cluster on above.

Kubernetes Administration

Clusters Business Units

Set Budget Assign Business Unit Remove from Business Unit Assign AWS Account


<input type="checkbox"/>	Edit	Name	Business Unit	Platform	Spending	Monthly Budget	Last Month(Oct. 2020)	Month-to-Date(Nov. 2020)	Forecast(Dec. 2020)	Last Collected Date
<input type="checkbox"/>		jack_single_vm	-	AKS			\$ 228.94	\$ 0.00	\$ 228.94	2020-10-26
<input checked="" type="checkbox"/>		certmanaks	Foglight	EKS			\$ 76.94	\$ 34.05	\$ 80.77	2020-11-06
<input type="checkbox"/>		Jack_Test_Cluster	Rainbow	AKS		\$ 229.00	\$ 229.05	\$ 82.37	\$ 247.99	2020-11-07

The Clusters tab displays the following information:

- *Name*: The name of the Kubernetes cluster.
- *Business Unit*: The name of the BU.
- *Platform*: AKS, EKS, or On-Prem. If Container cost failed to detect the AKS or ESK cluster, it will display On-Prem by default.
- *Spending*: The trend of cost spending.
- *Monthly Budget*: Monthly budget, which can be configured through *Set Budget*.
- *Last Month (Month. Year)*: The cost of last month.
- *Month-to-Date (Month. Year)*: The cost of current month.

- **Forecast (Month, Year):** The cost prediction for next month.
- **Last Collected Date:** The latest date till when the cost data is collected. It depends on the Last collected time from different platforms.

When mousing over the time, the start date and the end date of container cost collection are displayed.

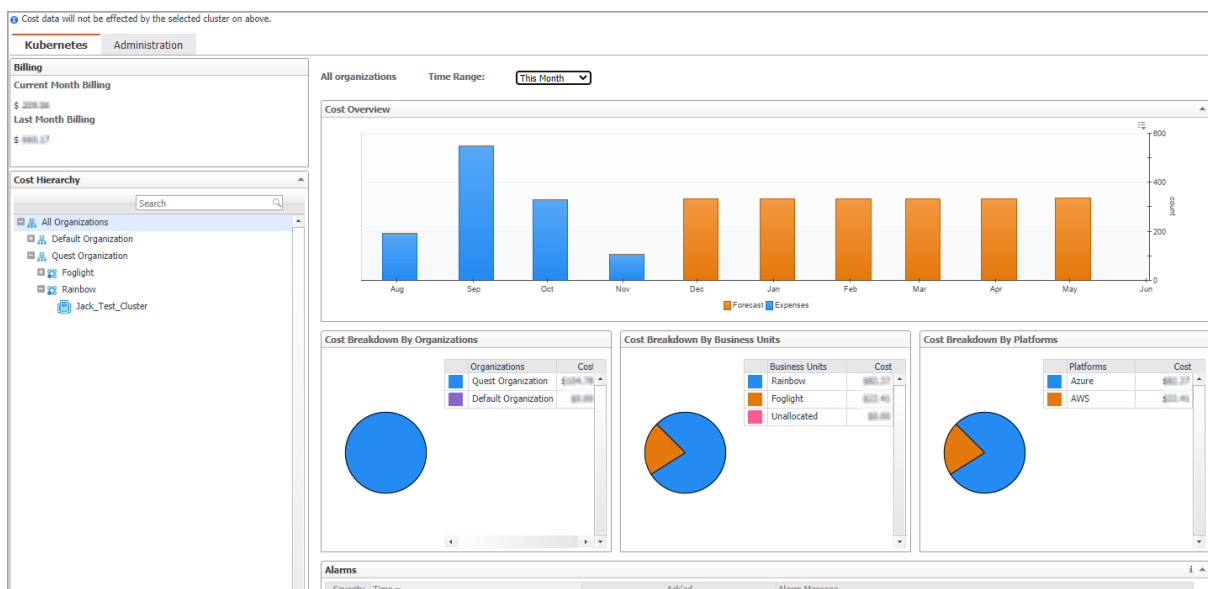
- **Set Budget:** Set monthly budget for the cluster
- **Assign Business Unit:** Assign the cluster to a specific BU.
- **Remove from Business Unit:** Remove the cluster from the selected BU.
- **Assign AWS Account:** This button will be enabled for editing only when an EKS cluster is selected.
-  : Click the button to change the monthly budge for the cluster.

Kubernetes dashboard

The *Kubernetes* dashboard displays a *Cost Overview* and several cost breakdown charts. The cost breakdown charts vary when choosing All Organizations, a single organization, a single BU, or a single cluster.

Kubernetes Overview

Figure 27. Kubernetes Overview



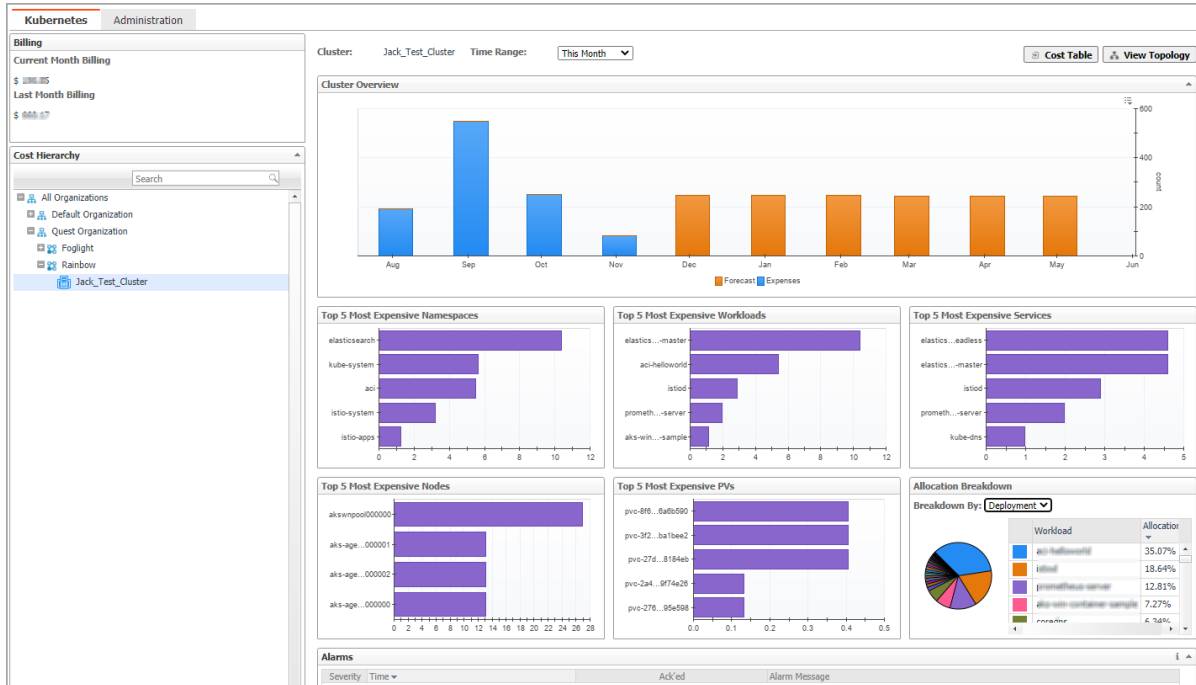
The Kubernetes overview shows the following information:

- **Time Range:** Enables you to change the time range for displaying the cost, including *This Month*, *Last Month*, or *Last N months*.
- **Cost Overview:** Monthly cost aggregated from all the clusters and the forecast aggregated cost for the continuing months.
- **Cost Breakdown By Organizations:** Cost categories by different organizations. Displayed when choosing All Organization.
- **Cost Breakdown By Business Units:** Cost categories by different BUs cross Organizations. Displayed when choosing All Organization or an organization.
- **Cost Breakdown by Platforms:** Cost categories by different platforms, such as AKS, EKS, and so on.
- **Most Expensive Clusters:** Displays the top N clusters that costs the most. Displayed when choosing an organization.

- **Clusters Cost:** Cost of the clusters of this BU. Displayed when choosing a BU.
- **Alarms:** all the cluster cost alarms. Currently we have over-budget alarm.

Cluster Overview dashboard

Figure 28. Cluster Overview



The Cluster overview dashboard shows the following information:

- **Time Range:** Enables you to change the time range for displaying the cost, including *This Month*, *Last Month*, or *Last N months*.
- **Cluster Overview:** The monthly cost aggregated from all the clusters of the organization and the forecast aggregated cost for the continuing months.
- **Top 5 Most Expensive Namespaces:** Displays the top 5 most expensive namespaces whose cost is aggregated from its workloads.
- **Top 5 Most Expensive Workloads:** Displays the top 5 most expensive workloads whose cost is aggregated from its related pods.
- **Top 5 Most Expensive Services:** Displays the top 5 most expensive services whose cost is aggregated from its related pods.
- **Top 5 Most Expensive Nodes:** Displays the top 5 most expensive nodes whose cost is collected from Azure or AWS.
- **Top 5 Most Expensive PVs:** Displays the top 5 most expensive persistent volumes whose cost is collected from Azure or AWS.
- **Allocation Breakdown:** Displays the Allocation Breakdown cost by Deployment or Stateful Set.
- **Alarms:** Displays all the cluster cost over-budget alarms.
- **Cost Table** button: Clicking the button will navigate to a *Cost table grid layout* dashboard.
- **View Topology** button: Clicking the button will navigate to a *Cluster Topology View*.

Cost Table

The Cost Table lists the detailed cost information calculated by different categories, including Infrastructure, Workload, and Service.

Figure 29. Cost Table

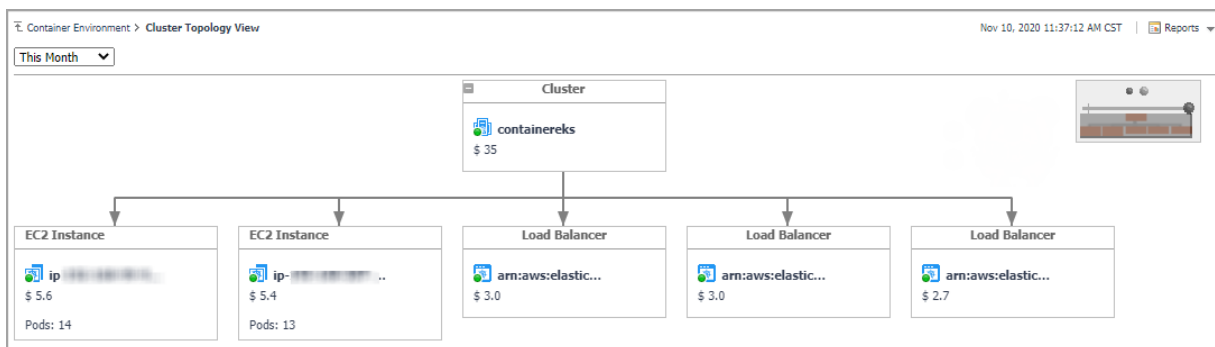
Name	Type	Cost	Resource Group
kubern...	ContainerAzurePublicIpAddress	\$ 0.00	/subscriptions/...
6eab06e7...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzurePublicIpAddressCost	\$ 0.00	/subscriptions/...
e916c8...	ContainerAzureDNSZoneCost	\$ 0.00	/subscriptions/...
akswnp...	ContainerAzureScaleSetCost	\$ 27	/subscriptions/...
aks-agent...	ContainerAzureScaleSetCost	\$ 29	/subscriptions/...
aks-agent...	ContainerAzureScaleSetCost	\$ 0.00	/subscriptions/...
aks-agent...	ContainerAzureScaleSetCost	\$ 0.00	/subscriptions/...
akswnp...	ContainerAzureScaleSetCost	\$ 0.00	/subscriptions/...
aks-agent...	ContainerAzureScaleSetCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.40	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.40	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
kubern...	ContainerAzureDiskCost	\$ 0.00	/subscriptions/...
f245efd...	ContainerAzureStorageAccountCost	\$ 0.00	/subscriptions/...
f245efd...	ContainerAzureStorageAccountCost	\$ 0.00	/subscriptions/...

- Infrastructure: Lists the Azure or AWS costs collected.
- Workload: The calculated workload cost is aggregated from its related pods. The workload types include deployment, statefulset, daemonset, and so on.
- Service: The service cost includes related workload and load balancer costs in AWS.

View Topology

Click View Topology button to open the Cluster Topology View.

Figure 30. Cluster Topology View



Clicking virtual machine, a popup view will be displayed to show the pods cost which run on this node and also the persistent volume that used by some pod.

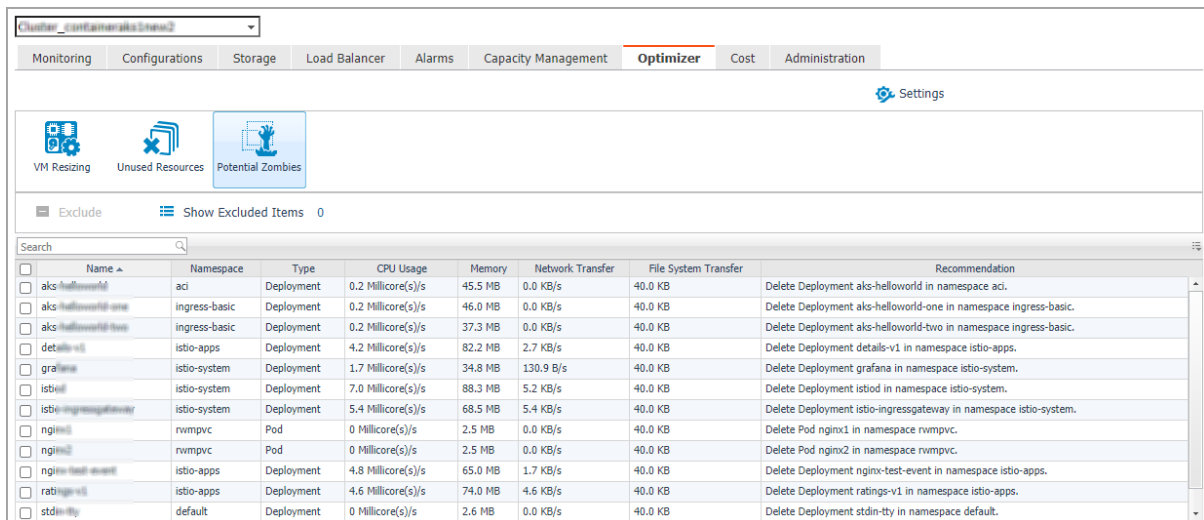
Optimizer

The *Optimizer* view appears after clicking **Container > Kubernetes > Optimizer**.

- NOTE:** If the Optimizer tab is not displayed, ensure the following:
 1. You have purchased a license for Optimizer. If not, contact Quest Support to purchase a license.
 2. You have the Container Administrator role.

NOTE: The displayed views are varied if the cluster hosts are monitored by the agents of VMware, AWS, or Azure.

Figure 31. Kubernetes Optimizer Dashboard



The Optimizer view includes the following elements:

- **Settings.** The Settings dialog box is used to change the time period and properties that are used for calculation. For more information, see [Settings](#) on page 51.
 - **Automate.** Use the Automate menu to set the criteria for automatically sending recommendations for improvements. Currently, this button only functions for **CPU** and **Memory** when a VMware cluster is selected.
 - **Reclaim Now** and **Reclaim Later** buttons. System administrator can select a VM from the list and review the Reclaiming Savings bar for information about how many resources can be reclaimed.
- NOTE:** The **Reclaim Now** and **Reclaim Later** buttons are enabled only after selecting a checkbox from the table. Currently, the two buttons only function for **VM Resources**, **CPU**, and **Memory** when a VMware cluster is selected.
- The **Automate**, **Reclaim Now**, and **Reclaim Later** buttons are displayed only when a VMware cluster is selected.
- **Exclude** button. Select an object you want to exclude from the table to enable the Exclude button, and click Exclude. Then, this object is added to the list of excluded objects under a specific category.
 - **Show Excluded Items** button. Click the Show Excluded Items button to view the excluded objects. The Settings dialog box appears. For more information, see [Excluded tab](#) on page 53.
 - **VM Resources/VM Resizing.** Shows instance or virtual machine name, utilization, recommendations for both CPU and memory resources, and savings.
 - **CPU.** Shows instance or virtual machine name, utilization, recommendations for CPU resource, and estimated savings.
 - **Memory.** Shows instance or virtual machine name, utilization, recommendations for memory resource, and estimated savings.

- **Storage.** Shows virtual machine name, utilization, storage and modify recommendations, and savings.

i | **NOTE: VM Resizing** will be displayed when a cloud cluster is selected. **VM Resources, CPU, Memory,** and **Storage** will be displayed when a VMware cluster is selected.

- **Unused Resources** table. Detects and shows those unused resources in container environment.

For example, persistent volume stays unused for more than 3 months. persistent volume stays in unbound status. This is due to the Unused Resources configuration in Settings.

- **Potential Zombies** table. Detects and shows the potential pod controllers in container environment, including Deployment, Daemon Set, Stateful Set, Replication Controller, as well as Pod that is not managed by any Pod Controller.

For example, if all pods managed by a pod controller are zombies, then we might suggest you to delete the whole pod controller.

Settings

Use the Settings menu to define the default optimization settings for your environment. The Settings Dialog box provides information about the following components:

- [Configuration tab](#)
- [Waste tab](#)
- [Excluded tab](#)
- [Credentials tab](#)
- [Constraints tab](#)

Configuration tab

Figure 32. Configuration tab

Settings Dialog

Configuration Waste Excluded Credentials Constraints

These settings are for CPU, Memory and Storage Optimization.

Thresholds

CPU	Memory	Storage
Warning: 75% Critical: 83%	Warning: 85% Critical: 90%	Warning: 90% Critical: 95%

Recommendation Calculation

Resource	CPU	Memory	Storage
Reserve Margin	5 %	5 %	5 %
Acceptable Variation	3 % 50 MHz	3 % 50 MB	3 % 1024 MB
Recommended Basis	Maximum Peak Utilization	Maximum Peak Utilization	Maximum Peak Utilization

Peak analysis period: 15 minute(s) Threshold for merging peaks: 5%

Evaluate calculation over this period of time 30 Day(s) History Period 30 Day(s)

Save Cancel

The **Configuration** tab provides the recommended settings for CPU, memory, and storage optimization.

- **Thresholds.** Provides the values of a resource metric that define the Warning and Critical levels (for CPU, memory, and storage).
- **Recommendation Calculation** area. Allows you to define the following parameters for optimizing the CPU, memory resources in your environment, Storage resources not supported at current version:

To save any changes made to the **Configuration** settings, click **Save** at the bottom of the tab.

Waste tab

Figure 33. Waste tab

Settings

Configuration Waste Excluded

These settings are for Unused Resources and Potential Zombies

Determine as waste if :

Resource has been created 1 Days

Persistent Volume Status

- Available
- Released
- Failed

Determine as a potential zombie if :

Time period used for average calculation is 7 Days

Average resource utilization

100.0 Millicore/s for CPU 100.0 KB/s for Network

100.0 MB for Memory 100.0 MB for Filesystem

Excluded Namespace

Save Cancel

The **Waste** tab allows you to configure the settings for determining resources wasted in your environment. These include unused resources and potential zombie Pod controllers.

- *Determine as waste if:* used to filter Unused Resources.
 - *Resource has been created [time] days:* Resources that has been created more than the set days will be considered here.
 - *Persistent Volume Status:* By default, select *Available* and *Failed*. For detailed information, go to <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#phase>.
- *Determine as a potential zombie if:* used to filter Potential Zombies.

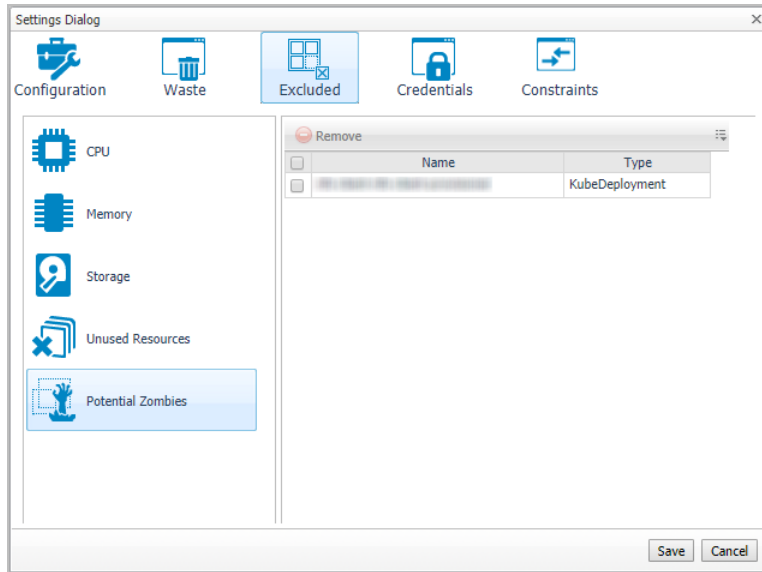
In container environment, Potential Zombie Pod Controller is considered here, including Deployment, Daemon Set, Stateful Set, Replication Controller, and Pod that is not managed by any Pod Controller. Settings work for single pod managed by Pod Controller. If all pods or partial pods of a Pod Controller are considered as zombies, different recommendations will be generated.

- *Time period used for average calculation is [time] days:* The average metrics for the pods are calculated, so a time range should be set to calculate the average value.
- *Average resource utilization:* only if a pod's metrics satisfy all the conditions, it will be considered to be a potential zombie pod.
- *Excluded Namespace:* pods in the namespace can be excluded in the Potential Zombies check.

To save any changes made to the **Waste** settings, click **Save** at the bottom of the tab.

Excluded tab

Figure 34. Excluded tab



The **Excluded** tab allows you to remove a resource from the list of excluded objects. The **Excluded** tab includes the following information:


- On the left side, a navigation tree, that allows you to select the resource category.
- On the right side, the list of resources excluded from the selected category.

To remove resources from the list of **Excluded** objects, select the check boxes for these resources and click **Remove**. To save any changes made to the **Excluded** settings, click **Save** at the bottom of the tab.

The **Excluded** tab can also be accessed by clicking **Show Excluded Items** on the **Optimizer** tab.

Credentials tab

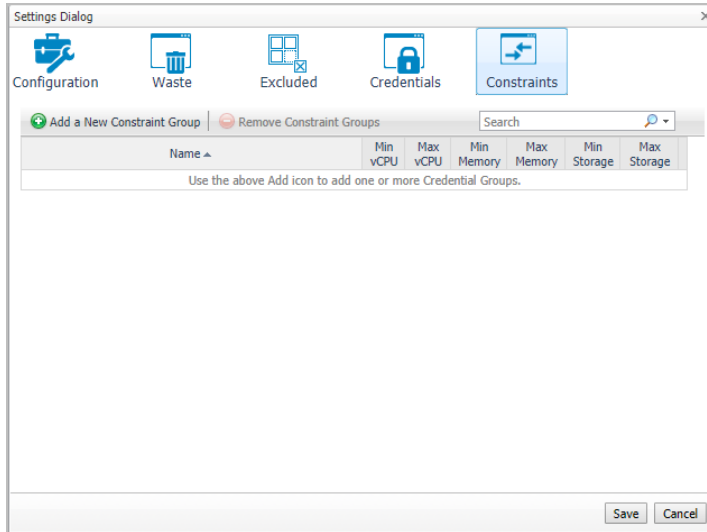
Figure 35. Credentials tab



This tab is available in VMware environment. The **Credentials** tab allows you to add, edit, and remove credential groups. This tab is only for the Storage rightsizer.

Constraints tab

Figure 36. Constraints tab

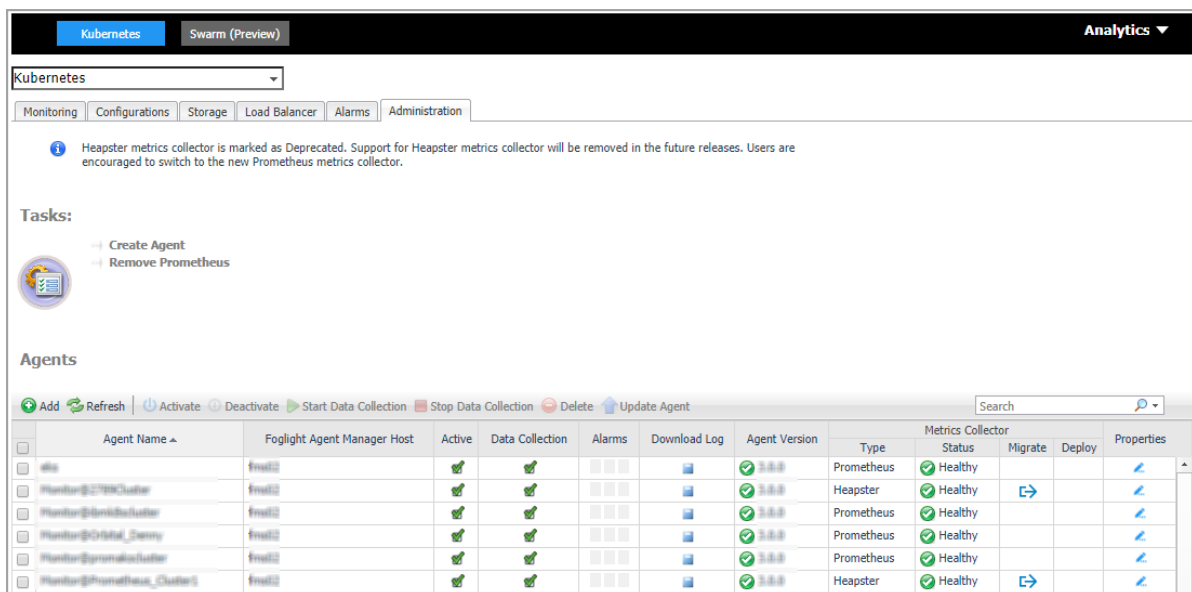


This tab is available in VMware environment. The **Constraints** tab allows you to set custom thresholds for select objects in the environment. These recommendations are displayed in the **Optimizer** tab > **VM Configuration/ CPU/ Memory/ Storage** views > **Modify Recommendation** column. Use this tab to add, edit, and remove constraints groups.

- IMPORTANT:** A virtual machine may have several partitions. VM environment makes recommendations for each partition separately, but the custom constraints can be set only for the entire VM (not for individual partitions). Therefore, the custom constraint for storage are applied to all partitions on the selected VM.

Administration

Figure 37. Kubernetes Administration Dashboard



i | **NOTE:** The *Administration* dashboard can be accessed only when the users have both the Administrator role and the Container Administrator role. To grant the users with the Container Administrator role, go to **Administration > Users & Security** management under *Administer Server > Manage Users, Groups, Roles > Roles* tab.

i | **IMPORTANT:** Heapster metrics collector is marked as Deprecated. Support for Heapster metrics collector will be removed in the future release. Users are encouraged to switch to the new Prometheus metrics collector.

The *Administration* dashboard supports new agent creation and existing agents management. It contains the following features:

- [Create Agent task](#)
- [Remove Prometheus task](#)
- [Agents table](#)

Create Agent task

Use *Create Agent* or the *Add* button to start a wizard to create a new agent. See [Creating and Activating a Kubernetes Agent](#) on page 19 for more information.

Remove Prometheus task

Remove Prometheus will start a wizard to remove the Prometheus successfully deployed through our platform.

i | **NOTE:** Prometheus that are not deployed through our platform are not listed here.

	Agent Name ▲	Metrics Collector	
		Type	Status
<input type="checkbox"/>	Prometheus@k8s1001001001	Prometheus	Healthy

Select the Prometheus you want to remove and click **Remove**.


i | **NOTE:** All the related Kubernetes components will also be removed according to the template that you used to deploy the Prometheus through our platform.

If the Prometheus service is deleted successfully, a result message will be display. Otherwise, the *Remove Prometheus* wizard will appear again with error messages to guide you with further operations. Follow the messages and delete them again later.

Agents table


Use *Activate*, *Deactivate*, *Start Data Collection*, *Stop Data Collection*, *Delete*, and *Update Agent* to manage the agent.


Metrics Collector

- *Type*: the type of Metrics Collector, supports Heapster, Prometheus, or OpenShift Prometheus.
- *Status*:
 - **Healthy**: discovered the Kubernetes service of the Metrics Collector and the result for the health check of the service is successful.
 - **Unhealthy**: discovered the Kubernetes service of the Metrics Collector, however, the result for the health check of the service is failed.
 - **Discover failed**: failed to discover the Kubernetes service in your cluster. The *Deploy*  icon is enabled.
 - **Waiting for data update**: waiting for status update after performing some operations for Metrics Collector.
 - **Failed to update status**: connection failed or other known issues. Contact the support.



- *Migrate* : start a wizard to migrate Heapster to Prometheus Metrics Collector.

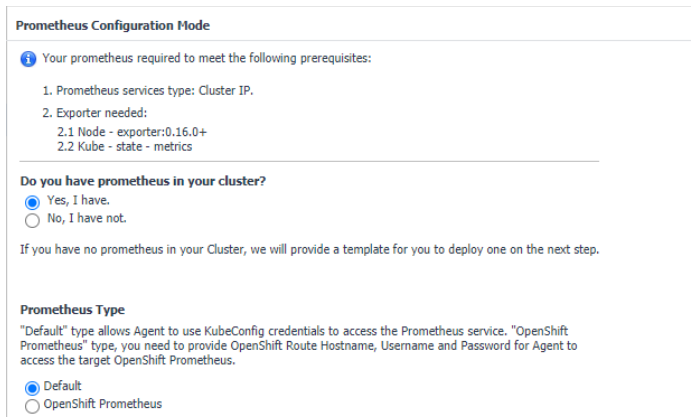
If the agent version is 3.0.0 and is using Heapster as the metric collector, you can Migrate from Heapster to Prometheus metrics collector.

- *Deploy* : start a wizard to configure or guide to deploy Prometheus to your cluster.


If our platform failed to discover Prometheus in your cluster and the *Deploy*  icon is enabled.

Migrate or Deploy process

Clicking *Migrate*  or *Deploy*  will launch the *Configure Prometheus Metrics Collector Wizard*.



Prometheus Configuration Mode

 Your prometheus required to meet the following prerequisites:

1. Prometheus services type: Cluster IP.
2. Exporter needed:
 - 2.1 Node - exporter:0.16.0+
 - 2.2 Kube - state - metrics

Do you have prometheus in your cluster?

Yes, I have.

No, I have not.

If you have no prometheus in your Cluster, we will provide a template for you to deploy one on the next step.

Prometheus Type

"Default" type allows Agent to use KubeConfig credentials to access the Prometheus service. "OpenShift Prometheus" type, you need to provide OpenShift Route Hostname, Username and Password for Agent to access the target OpenShift Prometheus.

Default

OpenShift Prometheus

- Have existing Prometheus in your cluster.

i | **NOTE:** Ensure that you have Prometheus in your cluster before this step. We will check the service existence and health status after you click Next.

- Select the Prometheus Type.

- Default type

Allows Agent to use KubeConfig credentials to access the Prometheus service. Enter the existing Prometheus service namespace and name to configure the metrics collector.

Prometheus Configuration

Please enter prometheus service information.

Namespace: ⓘ

Name: ⓘ

- OpenShift Prometheus type: you need to provide OpenShift Route Hostname, Username, and Password for Agent to access the target OpenShift Prometheus.

Prometheus Configuration

ⓘ OpenShift Prometheus uses Route Hostname to access the OpenShift OAuth secured Prometheus. You need to provide OpenShift Route Hostname, Username and Password for Agent to access the target OpenShift Prometheus.

Route Hostname: ⓘ

OpenShift Username: ⓘ

OpenShift Password: ⓘ

ⓘ **NOTE:** To find the Route Hostname, go to the **OPENSHIFT ORIGIN** dashboard, **Application > Routes**. Your OpenShift Login User should have the privilege to access to the Prometheus.

- Do not have existing Prometheus in your cluster.

A *Prometheus Configuration* wizard will appear. To configure the metrics collector, either use the default Prometheus template or upload your own Prometheus deployment .yaml files through *Load from files*.

ⓘ **NOTE:** Whenever you change the content in Configure file (template), ensure that the *Namespace* and *Name* fields are consistent with the Prometheus service configurations in your Configure file (template). If you are using the default template, there is need to change anything.

Prometheus Configuration

Please edit the prometheus configuration file. If you have your own prometheus configuration files, load the files from local filesystem. Ensure that the namespace and the name of the Prometheus service are consistent with those in the configuration files you edited or loaded. And then click Deploy to deploy prometheus.

Configure file (template):

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  labels:
    app: prometheus
    component: server
    name: prometheus-server
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: prometheus-server
subjects:
- kind: ServiceAccount
  name: prometheus-server
  namespace: kube-system
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  labels:
    app: prometheus
    component: kube-state-metrics
    name: prometheus-kube-state-metrics
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: prometheus-kube-state-metrics
subjects:
- kind: ServiceAccount
  name: prometheus-kube-state-metrics

```

Load from files

Namespace: ⓘ

Name: ⓘ

Deploy Reset from template

Reset from template: Helps you to reset the Configure file (template) content to the default template.

Deploy: Creates Prometheus components to your cluster with the Configure file.

After clicking *Deploy*, a progress message will be displayed.

- If deployed successfully, a succeeded message will be displayed. Close the *Validation Result* page and then click **Next** to finish your agent creation process.

- If failed to deploy the Prometheus configuration, a *Validation Result* page will be displayed with possible solutions.

Agent Edit properties

Click  in Properties column of the Agents table to edit the property of the agent.

- For agent using Heapster as metric collector, a wizard similar as below will be displayed. Update the agent properties and save the changes. A new data collection process will be initiated.

The screenshot shows a dialog box titled "Kubernetes Agent [Monitor 2789Cluster]: Edit Properties". It contains the following fields and options:

- Cluster Name *: 2789Cluster
- Kubernetes API Service End Point *: https://172.17.0.1:443
- Collected Event Level *: ALL ABNORMAL
- Metrics Collector *: **Heapster**
- Service Namespace: kube-system
- Service Name: heapster
- Proxy: Enable Proxy

Buttons for "Save" and "Cancel" are at the bottom right.

- For agent using Prometheus as metric collector and the Prometheus is successfully deployed through our platform, a wizard similar as below will be displayed.

The screenshot shows a dialog box titled "Kubernetes Agent [Monitor ibmk8scluster]: Edit Properties". It contains the following fields and options:

- Cluster Name *: ibmk8scluster
- Kubernetes API Service End Point *: https://172.17.0.1:443
- Collected Event Level *: ALL ABNORMAL
- Metrics Collector *: **Prometheus** (with a "Change to Heapster" button)
- Service Namespace: kube-system
- Service Name: prometheus-server
- Config File: (with an "Export" button)
- Proxy: Enable Proxy

Buttons for "Save" and "Cancel" are at the bottom right.

i **NOTE:** If the Prometheus is deployed through our platform, the Service Namespace and Service Name cannot be updated through *Kubernetes Agent Edit Properties* wizard. You can use *Remove Prometheus* to remove Prometheus and deploy a new one.

Export: export the deployed .yml file of the Prometheus.

Change to Heapster: change from Prometheus Metrics Collector to Heapster. However, it is not recommended.

- For agent using existing Prometheus as Metrics Collector, a wizard similar as below will be displayed.

Kubernetes Agent [eks]: Edit Properties

Cluster Name *: Ekscluster ⓘ

Kubernetes API Service End Point *: <https://XXXXXXXXXXXXXXXXXXXXXXXXXXXX.us-east-1.amazonaws.com> ⓘ

Collected Event Level *: ALL ABNORMAL ⓘ

Metrics Collector *: **Prometheus** Change to Heapster ⓘ

Service Namespace *: kube-system ⓘ

Service Name *: prometheus-server ⓘ

Proxy Enable Proxy ⓘ

i | **NOTE:** If you are using your existing Prometheus as Metrics Collector, we will not help you to manage your Prometheus. Both the Service Namespace and Service Name can be updated in *Kubernetes Agent Edit Properties* wizard.

- For agent using OpenShift Prometheus as Metrics Collector, a wizard similar as below will be displayed.

Kubernetes Agent [Monitor@akscluster]: Edit Properties

Cluster Name *: akscluster ⓘ

Kubernetes API Service End Point *: <https://XXXXXXXXXXXXXXXXXXXXXXXXXXXX.us-east-1.amazonaws.com> ⓘ

Collected Event Level *: ALL ABNORMAL ⓘ

Metrics Collector *: **Prometheus** Change to Heapster ⓘ

Route Hostname <https://prometheus.kube-system> ⓘ

OpenShift Username openshift ⓘ

OpenShift Password ⓘ

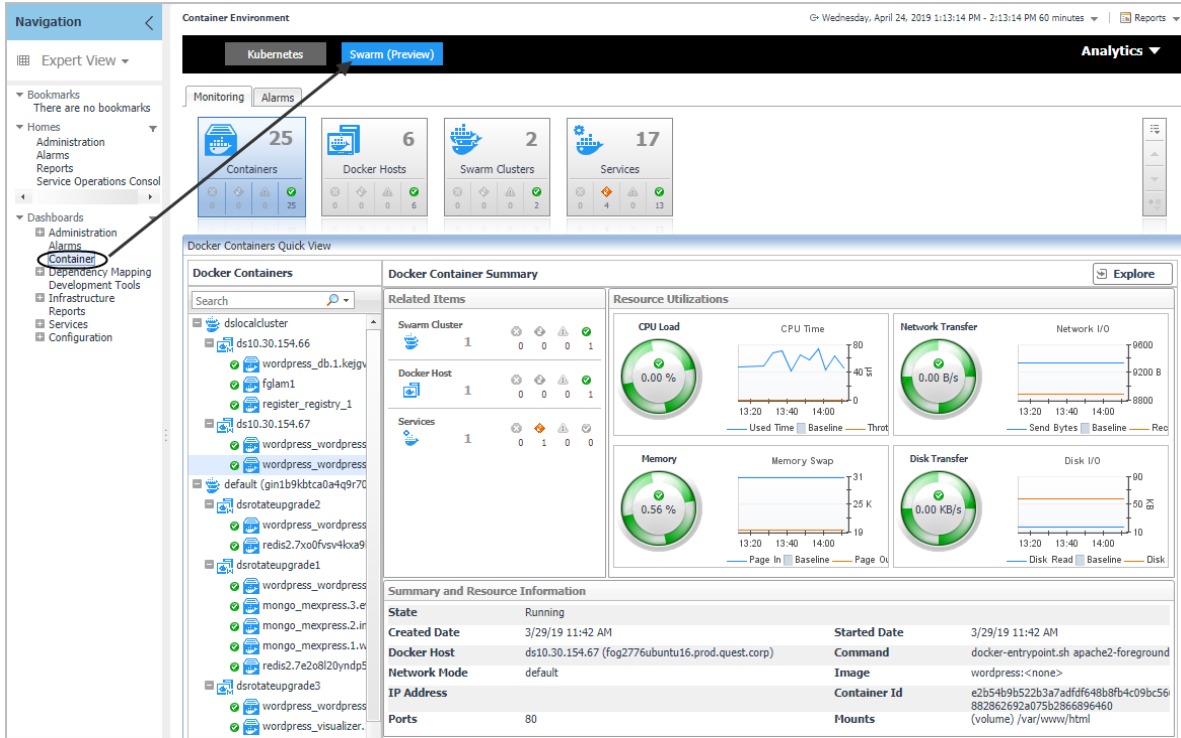
Proxy Enable Proxy ⓘ

Save Cancel

Docker Swarm

The *Docker Container Quick View* appears after clicking **Monitoring > Containers**. Click **Swarm (Preview)** from the header on top to switch to Docker Swarm dashboard.

Figure 38. Docker Swarm Dashboard



Monitoring Docker Containers

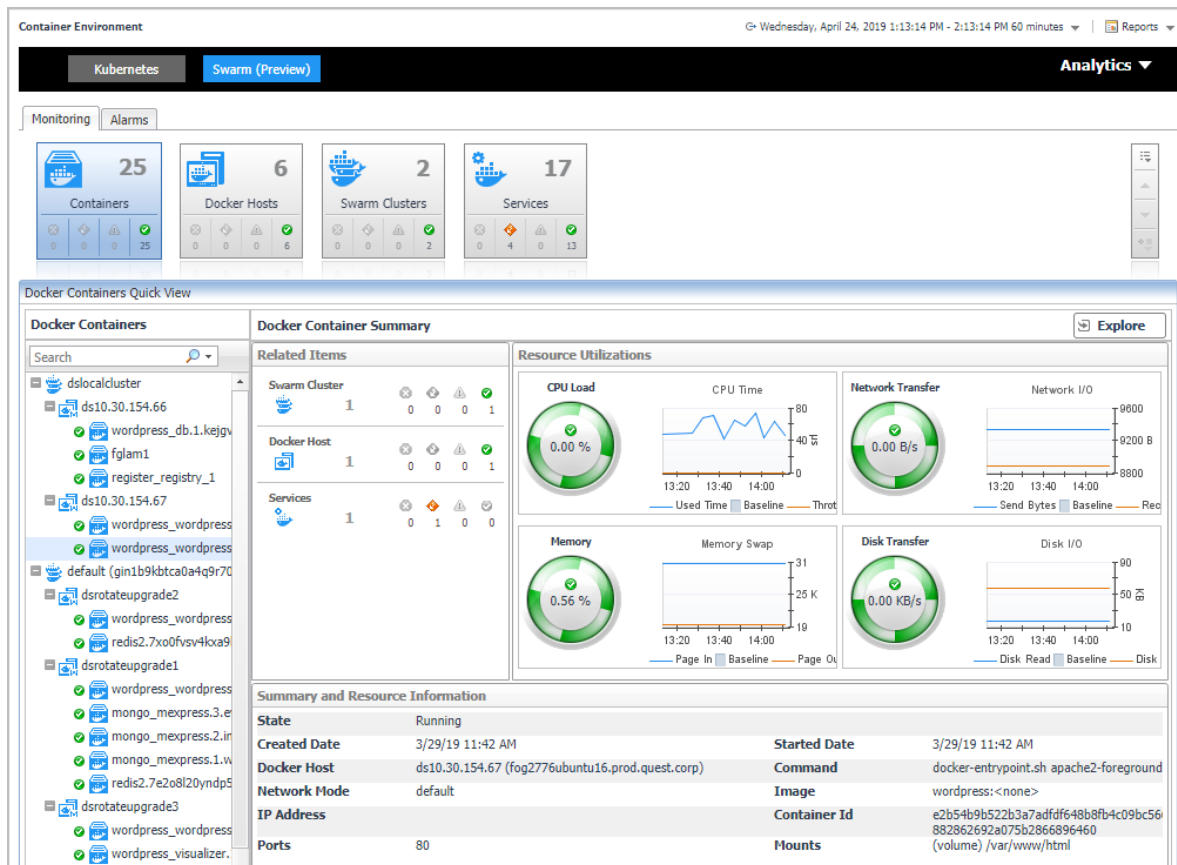
This view consists of the following two panes:

- The **Docker Containers** tree view, which appears on the left of *Docker Containers Quick View*, lists the containers existing in the monitored *Docker* environment. The containers in the tree view are grouped by **cluster > docker host > container**.
- The **Docker Container Summary** view, which appears on the right after you select an individual container in the **Docker Containers** tree view.

Docker Container Summary view

The **Docker Container Summary** view appears on the right when you select a container in the **Docker Containers** tree view.

Figure 39. Docker Container Summary view



The **Docker Container Summary** view displays the following data:

- **Related Items:** Shows the related Docker components grouped by type as well as the associated alarms.
- **Resource Utilizations:** The resource utilization for the selected Docker Container over a selected period of time, which includes the following:
 - **CPU Load:** Shows the CPU utilization of the selected container.
 - **CPU Time:** Shows the used time and throttled time of the selected container.
 - **Network Transfer:** Shows the transfer bytes rate of the selected container over a selected period of time.
 - **Network I/O:** Shows the total send/receive bytes of the selected container.
 - **Memory:** Shows the memory utilization of the selected container.
 - **Memory Swap:** Shows the mounts of memory pages that are swapped to disk.
 - **Disk Transfer:** Shows the disk transfer bytes rate of the selected container over a selected period of time.
 - **Disk I/O:** Shows the disk read/write bytes of the selected container.
- **Summary and Resource Information:** Displays the detailed information about the selected Container, including *State*, *Command*, *Created Time*, *Started Time*, *Image*, and so on.

Click **Explore** on the upper right of the **Docker Container Summary** view to open the **Container Explorer** view, which shows more detailed information about this container.

Container Explorer view

The *Container Explorer* view opens when you click **Explore** in the [Docker Container Summary view](#), which includes the following tabs:

Monitoring tab: The *Monitoring* tab displays the overall information of the selected container over a selected period of time, including the *Summary and Resource Information* table, Resource Management table as well as the Metrics list. To set the Metrics list displayed, go to **Action > General > Metric Selector**. For more information, see [Container metrics](#) on page 79.

Figure 40. Docker Container Explorer view Monitoring Tab

The screenshot displays the Docker Container Explorer view for a container named 'wordpress_wordpress...'. The interface is divided into several sections:

- Monitoring Tab:** Shows the container's state as 'Running' and provides a detailed 'Summary and Resource Information' table. This table includes fields for State, Created Date, Docker Host, Image, Network Mode, Ports, Environment Variables (such as WORDPRESS_DB_HOST, WORDPRESS_DB_PASSWORD, etc.), Labels, Mounts, and IP Address.
- Swarm Information:** Displays the container's location within a Swarm cluster, including the Swarm Cluster (localdscluster), Swarm Service (wordpress_v), and Swarm Node (fog2776ubuntu16).
- Resource Management:** Contains two line graphs: 'CPU Utilization' and 'Memory Utilization', both showing 0% usage over a 2-hour period.
- Actions Panel:** Located on the right, it features a 'Metric Selector' dropdown menu. The selected metrics include CPU Utilization, Memory Utilization, Network Bytes, Network Packets, Network Packets Dropped, Network Transfer Bytes, Disk Bytes, and Disk Transfer Bytes.

Monitoring Docker Hosts

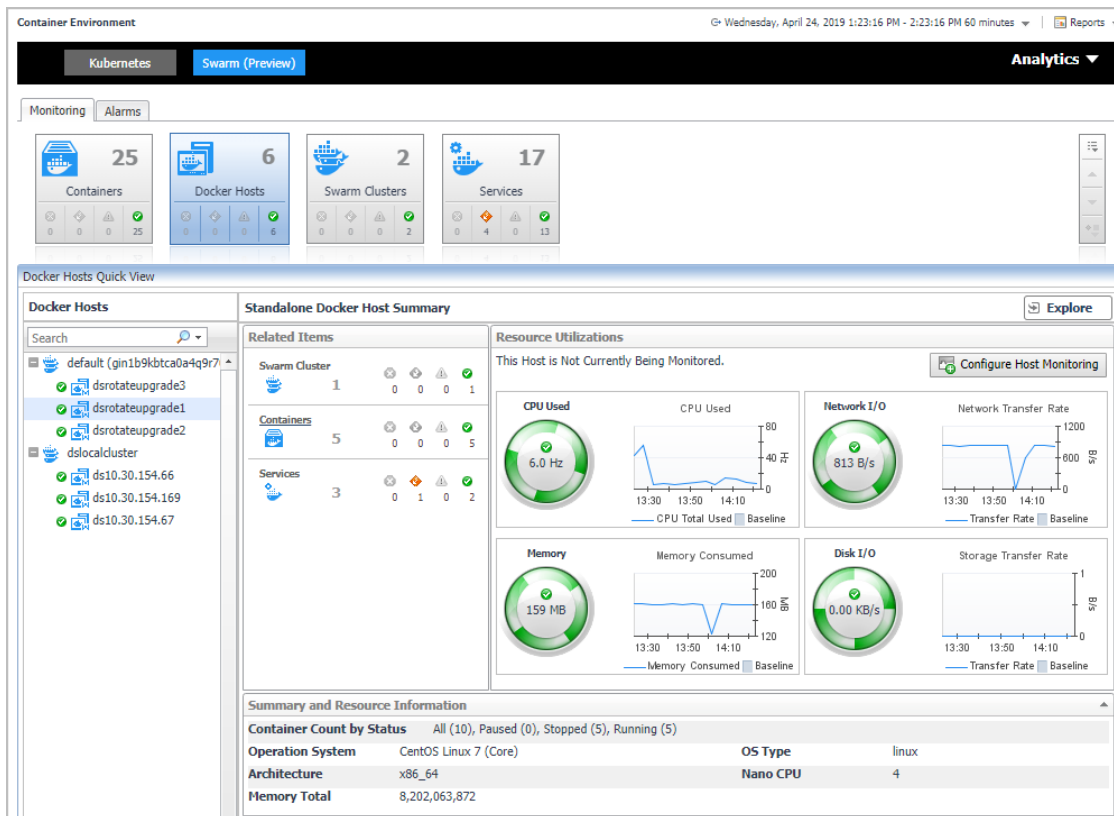
This view consists of the following two panes:

- The **Docker Hosts** tree view, which appears on the left of *Docker Hosts Quick View*, lists the docker hosts existing in the monitored *Docker* environment. The docker hosts in the tree view are grouped by **cluster > docker host**.
- The [Docker Host Summary view](#) appears on the right after you select an individual docker host in the **Docker Hosts** tree view.

Docker Host Summary view

The **Docker Host Summary** view appears on the right when you select a docker host in the **Docker Hosts** tree view.

Figure 41. Docker Host Summary view



The **Docker Host Summary** view displays the following data:

- **Related Items:** Shows the related Docker components grouped by type as well as the associated alarms.
- **Resource Utilizations:** The resource utilization for the selected docker host over a selected period of time, which includes the following:
 - **CPU Load:** Shows the CPU utilization of the selected docker host.
 - **CPU Used:** Shows the used CPU resources aggregated from the containers running on the docker host.
 - **Network I/O and Network Transfer Rate:** Shows the transfer bytes rate of the selected docker host aggregated from the containers running on the docker host over a selected period of time.
 - **Memory and Memory Consumed:** Shows the memory consumed bytes aggregated from the containers running on the docker host.
 - **Disk I/O and Disk Transfer:** Shows the disk transfer bytes rate of the selected docker host aggregated from the containers running on the docker host over a selected period of time.
- **Summary and Resource Information:** Displays the detailed information about the selected docker host, including *Container Count by Status*, *Operating System*, *Memory Total*, and so on.

Click **Explore** on the upper right of the **Docker Host Summary** view to open the [Docker Host Explorer view](#), which shows more detailed information about this container.

Docker Host Explorer view

The *Docker Host Explorer* view opens when you click **Explore** in the [Docker Host Summary view](#), which includes the following tabs:

- **Monitoring tab:** The *Monitoring* tab displays the overall information of the selected docker host over a selected period of time, including the *Summary and Resource Information* table, *Containers* table, *Images* table, and *Volumes* table.

NOTE: All the docker host metrics are calculated from the aggregated metrics of the containing containers on the docker host.

Figure 42. Docker Host Explorer view Monitoring Tab

The screenshot shows the Docker Host Explorer interface for a host named 'dsrotateupgrade3'. It is divided into three main sections: Summary and Resource Information, Containers, and Images.

Summary and Resource Information:

- Container Count by Status:** All (22), Paused (0), Stopped (9), Running (13)
- Operation System:** CentOS Linux 7 (Core)
- Memory:** 1.8 GB
- Docker Version:** 18.09.2
- Host:** dsrotateupgrade3 (10.4.117.155)
- Nano CPU:** 1
- Group Driver:** cgroupfs
- Docker Host Swarm Information:**
 - Swarm Cluster:** default (gin1b9kbtca0a4q9r70aym1je)
 - Swarm Node:** dsrotateupgrade3 (xypq2vtrwx4cztzyo16cjg9v3)
 - Swarm Role:** Manager
 - Swarm Node Status:** [ready]

Containers Table:

Alarms	Name	Id	Image	Netw Mod
✓	voting_result.1.m1ab9gkqai8l1krq24oudykd	aa25171d29a8076e9a0a076340a967927a237cccd066b6e460c4e44d1ad2979	dockersamples/examplevotingapp_result:<none>	defau
✓	voting_vote.2.ua654033tdvqp7fn3pvyqo0rb	133ae0269871f9e90b894be92719216e9beef2d46328c5d849a5acec025dd4b4	dockersamples/examplevotingapp_vote:<none>	defau
✓	wordpress_wordpress.3.3d78lku8t9evlza3mpt8ehr	7b6b79b24d5c7efa67e4ac707f1c1efcc696daf36c6205dae4fd1645123d3dec2	wordpress:<none>	defau
✓	voting_redis.1.awis611hu72gydx5h5epjy6v	288e8a1dde3a0e893913559c7e2202e96f52a88fadfe5b862d4b6988a735d996	redis:<none>	defau
✓	voting_vote.1.deivdnggt0nlrcb8qo0b8r1hd	f5eca4b5aec9294745fe742a2eb9e12ee26b13699e9db4c473cc6b334d9e78bc5d	dockersamples/examplevotingapp_vote:<none>	defau
✓	redis2.xypq2vtrwx4cztzyo16cjg9v3.vrubg5q34jlaq8fh21mg6vex	149fa7578f49f332fed8ac292d61cd19d4c8eafecdf45a1b3482201312fadf0e6	redis:<none>	defau
✓	mongo_mongo.1.95ak9qe8gp158etdue4pwwu2w	78a7e75c36a9056c9abc9eba4d636f2a1558f236873d67bd114f21df12aad2d	mongo:<none>	defau
✓	mysql_phpmyadmin.1.0c2dfj4k6ghuo4n463mf8q7m	f89221041f8ddcd2c594bf9a304af3924d8b1f26700ecc794f7173b52e47bd7	phpmyadmin/phpmyadmin:<none>	defau
✓	wordpress_db.1.7c7ftk36ojtcsqb6zr1tjoh	1711d54eef15255a20e37685ccf85ad6a3656beae84bf9ce44cb763d29d734a2	mysql:<none>	defau
✓	wordpress_visualizer.1.drldgdbtulabszbeSeriegdiq7	d1af858af95546c4db25fc7218dc89bfe4c99b73980d746a2eea9dd74a489719	dockersamples/visualizer:<none>	defau
✓	voting_visualizer.1.nwq4Hjmbwpyvmh9wn30qddy	16b274edf0a972a313b61f785e3473571cfa6f380ef1a57991cc64fc53dddbb9	dockersamples/visualizer:<none>	defau
✓	mysql_mysql.1.mroqoc2f57qm5pjk9y4bnjke	2c27187424e082060de9cb9da6d2e9da440cdbc50ffa626ed45157698cc722	mysql:<none>	defau
✓	voting_db.1.8e03l95f71050jrvpq4j9hu	7a60a959b3afa27006ed3383d69e966ae25270c15029f95127c6696ae226db	postgres:<none>	defau

Images Table:

State	Name	Id	Size	Virtual Size	Comment	Containers	Not Updated Duration
✓	dockersamples/examplevotingapp_result:<none>	sha256:e10df791f13c3ac17efa123dfce57e3297fcea05a34b3bbf305749f22a9b3c83	216.0 MB	216.0 MB		1	2 minute(s)
✓	phpmyadmin/phpmyadmin:<none>	sha256:c6ba363e7c9bba3bc96aa490e31da3e266e67e5d8c525fb8a36df2544c2aa54	158.2 MB	158.2 MB		1	2 minute(s)
✓	mongo:<none>	sha256:0d183f48c313d863d26aed97c27a0fb73833674c87da2576b6282de2439a144c	389.5 MB	389.5 MB		1	2 minute(s)
🗑️	nate/dockviz:latest	sha256:93b5259c1e18862e143ae39678640cbdd555d1b8e2742bc6f4da9c2b78acd8ab	6.3 MB	6.3 MB		0	
✓	mysql:<none>	sha256:7bb2586065cd50457e315a5dab0732a87c45c5fad619c017732f5a13e58b51dd	454.8 MB	454.8 MB		2	2 minute(s)
✓	redis:<none>	sha256:d4deec2c521cdae0450218bd53c69611bacd2eb10838057a5de7dcb341c66cf5	144.2 MB	144.2 MB		1	2 minute(s)
✓	postgres:<none>	sha256:d7cf98b297166b40efca50ff11e19c7e801d45a0f6c1ba316854984229667578	214.9 MB	214.9 MB		1	2 minute(s)
🗑️	alpine:<none>	sha256:5cb3aa0f89934411fba5c063a9bc98ace875d8f92e77d0029543d9f2ef4ad0	5.2 MB	5.2 MB		0	

- **Containers table:** Includes the containers on this docker host.
- **Images table:** Includes the images pulled onto this docker host.
 - : Indicates this image is using by a container.
 - : Indicates no container is using this image and the image can be recycled.
- **Volumes table:** Includes the volumes created on this docker host.
 - : Indicates this volume is using by a container.
 - : Indicates no container is using this volume and the volume can be recycled.

Figure 43. Docker Host Explorer view Images table and Volumes table under Monitoring tab

The screenshot shows the Docker Host Explorer interface. At the top, there are two tables: 'Images' and 'Volumes'. The 'Images' table has columns for State, Name, Id, Size, Virtual Size, Comment, Containers, and Not Updated Duration. The 'Volumes' table has columns for State and Name. A modal window titled 'Docker Host Explore Containers' is open, showing a table with columns for Alarms, Name, and Id. A red circle highlights the 'Containers' column in the 'Images' table, and a red arrow points to the modal window.

By clicking the number in the *Containers* column, a *Docker Host Explore Containers* view will open, which lists the containers using this image or this volume. Click the Name or ID of the container and an explore page of the container will appear.

- **Metrics** tab: The *Metrics* tab displays the Metrics list. To set the Metrics list displayed, go to **Action > General > Metric Selector**. For more information about the description of the metrics, see [Container metrics](#) on page 79.

Figure 44. Docker Host Explorer view Metrics Tab

The screenshot shows the Docker Host Explorer Metrics tab. It features four line graphs: 'Total CPU Used', 'Total Memory Consumed', 'Total Network Transfer Rate', and 'Total Disk Transfer Rate'. A 'Metric Selector' panel on the right allows selecting which metrics to display. The 'Metric Selector' panel has a search bar and a list of metrics with checkboxes. The selected metrics are: Total CPU Used, Total Memory Consumed, Total Network Packets Rate, Total Network Transfer Rate, and Total Disk Transfer Rate.

Monitoring Docker Swarm Clusters


This view consists of the following two panes:

- The **Swarm Clusters** tree view, which appears on the left of *Swarm Clusters Quick View*, lists the docker swarm clusters existing in the monitored *Docker* environment.
- The **Docker Swarm Cluster Summary** view, which appears on the right after you select an individual docker swarm cluster in the **Swarm Clusters** tree view.

Docker Swarm Cluster Summary view

The **Docker Swarm Cluster Summary** view appears on the right when you select a docker swarm cluster in the **Swarm Clusters** tree view.

Figure 45. Docker Swarm Cluster Summary view



The **Docker Swarm Cluster Summary** view displays the following data:

- **Related Items:** Shows the related Docker components grouped by type as well as the associated alarms.
- **Resource Utilizations:** Shows CPU Utilization, Memory Utilization, Network Transfer Rate, Disk Transfer Rate metrics of the containers running in this docker swarm cluster in descending order.

Monitoring Docker Swarm Services

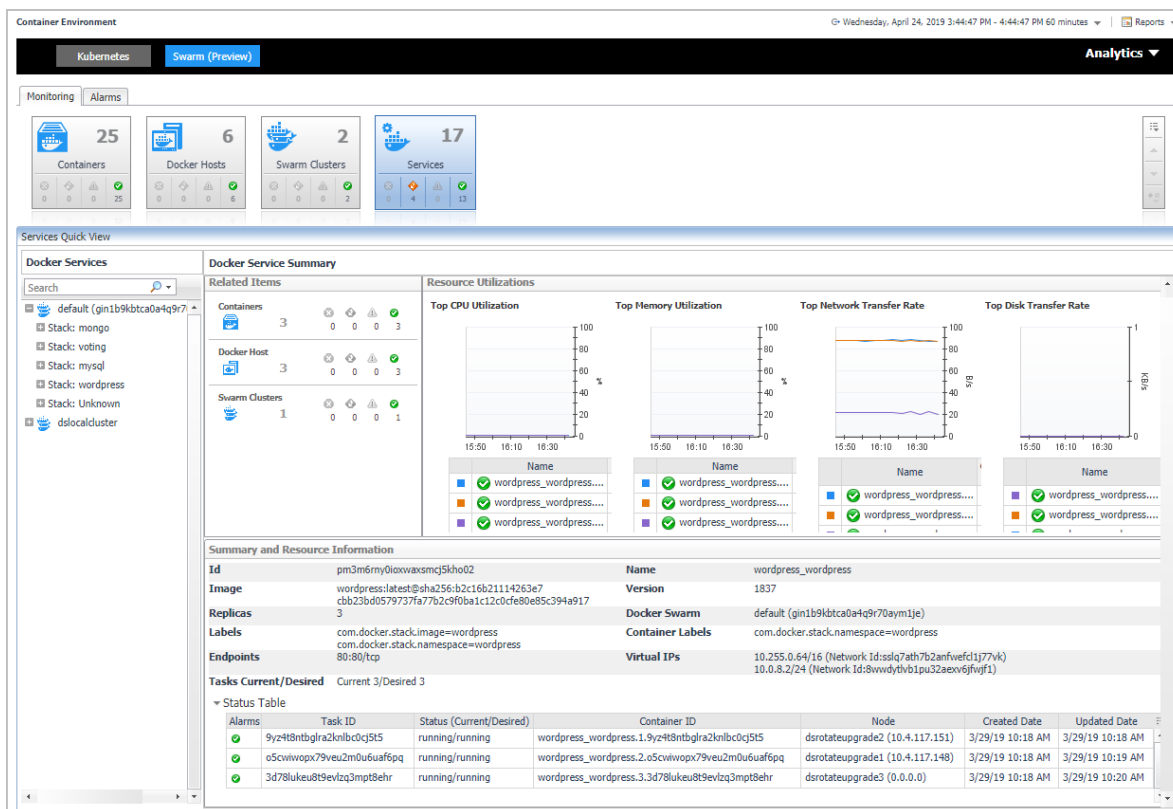
This view consists of the following two panes:

- The **Swarm Services** tree view, which appears on the left of *Swarm Services Quick View*, lists the docker swarm services existing in the monitored *Docker* environment.
- The **Docker Swarm Service Summary** view, which appears on the right after you select an individual docker swarm service in the **Swarm Services** tree view.

Docker Swarm Service Summary view

The **Docker Service Summary** view appears on the right when you select a docker swarm service in the **Docker Services** tree view.

Figure 46. Docker Service Summary view

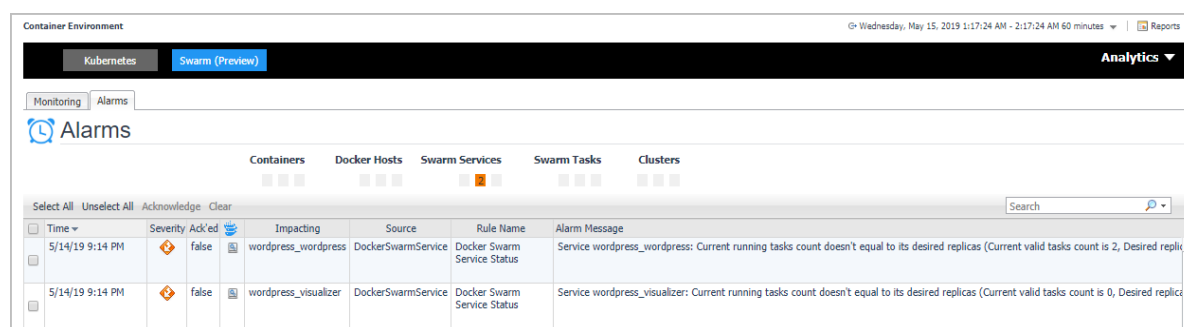


The **Docker Service Summary** view displays the following data:

- **Related Items:** Shows the related Docker components grouped by type as well as the associated alarms.
- **Resource Utilizations:** Shows CPU Utilization, Memory Utilization, Network Transfer Rate, Disk Transfer Rate metrics of the containers running in this docker swarm service in descending order.
- **Summary and Resource Information:** Shows the summary information of the docker swarm service, including Labels, Image, Mount Volumes, Ports, Container Status and so on.

Alarms

Figure 47. Docker Swarm Alarms Dashboard



The *Alarms* dashboard displays a list of alarms generated against the monitored Docker environment. Use this view to quickly identify any potential problems related to a specific Docker component.

Analytics

Foglight for Container Management provide analytics feature for Kubernetes and Docker Swarm.

Heat Map is a two-dimensional representation of data in which values are represented by colors. Showing collected metrics with elaborate heat maps allows you to understand complex data sets and the monitored cluster environment well.

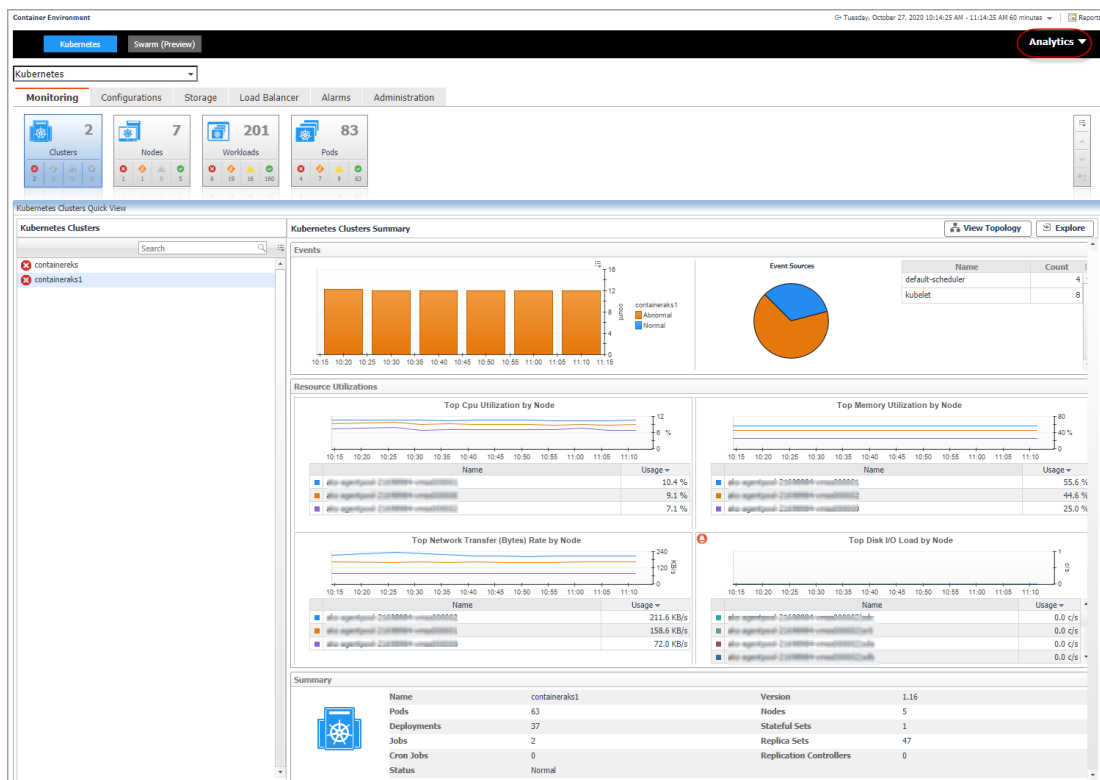
Scatter Plot is used to display values in points using two variables for a set of data. The points is color-coded also, Color Metric can be used to display one additional variable.

- [Kubernetes analytics](#)
 - [Heatmap analytics](#)
 - [Scatter Plot analytics](#)
- [Docker Swarm analytics](#)
 - [Heatmap analytics](#)
 - [Scatter Plot analytics](#)

Kubernetes analytics

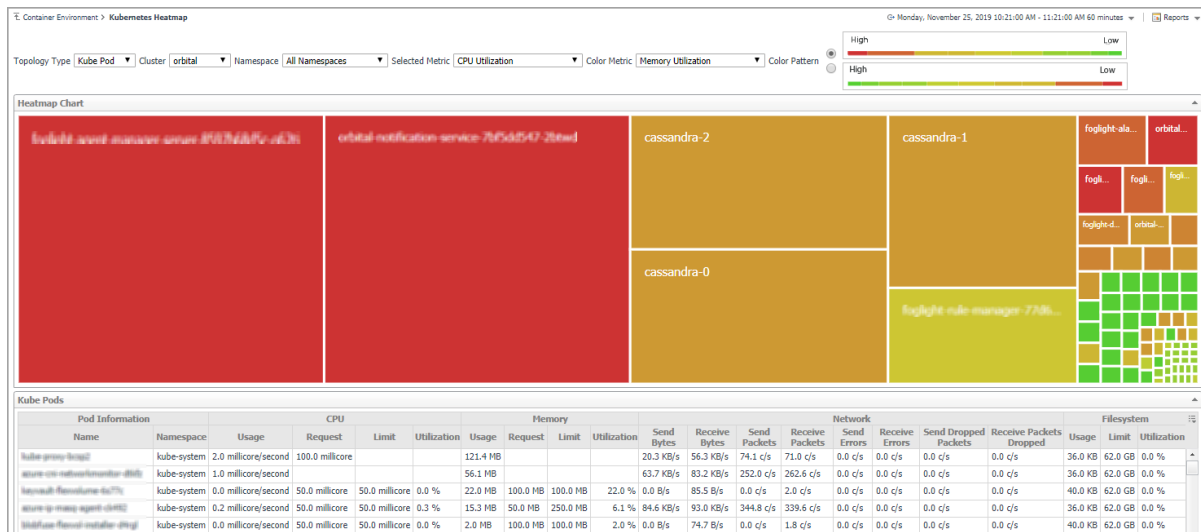
In the Container dashboard, choose **Kubernetes** from the header. Then click **Analytics** from the header, a drop down view will display with **Heatmap** and **Scatter** on it. Click **Heatmap** will navigate to the Kubernetes **Heatmap Analytics** dashboard, while click **Scatter** will navigate to the Kubernetes **Scatter Plot Analytics** dashboard.

Figure 48. Kubernetes analytics Navigation



Heatmap analytics

Figure 49. Kubernetes Heatmap Analytics Dashboard



Heat maps will be refreshed automatically when you change either of the following fields:

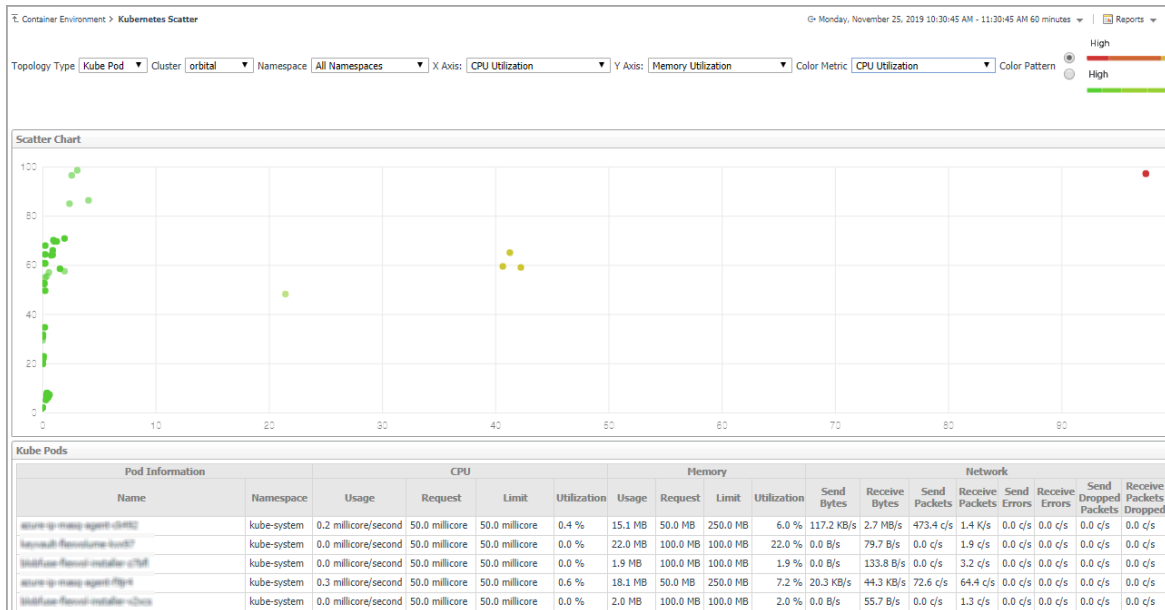
- **Topology Type:** Indicates the monitored topology object, including Kubernetes Pod and Kubernetes Node.
- **Cluster:** Lists all clusters available in the monitored Kubernetes environment.
- **Namespace:** Lists all namespaces available in the monitored Kubernetes environment.
- **Selected Metric:** Populates a rectangle based upon the selected metrics. For example, if you select **Memory Usage** from the **Selected Metric** drop-down list, the rectangle area will be populated based on the used memory for the selected topology object. For more information about metrics, refer to [Kubernetes metrics on page 75](#).
- **Rendering related metrics:** For example, if you select **CPU Usage Rate** and Red to Green, the rectangle of the topology object that has larger value of CPU Usage Rate will be rendered in red.
 - **Color Metric:** Renders the color of rectangle based upon the selected color metric.
 - **Color Pattern:** Offers two patterns, Red to Green (larger value shows in red) or Green to Red (larger value shows in green).

Figure 49 shows an example of heat map. Clicking the object name on the heat map directs you to the relevant object Explorer dashboard. For more information, see:

- [Workloads Explorer view on page 36](#)
- [Workloads metrics on page 75](#)
- [Nodes Explorer view on page 33](#)
- [Node metrics on page 77](#)
- [Cluster Explorer view on page 29](#)
- [Docker Swarm metrics on page 79](#)

Scatter Plot analytics

Figure 50. Kubernetes Scatter Plot Analytics Dashboard



The points on the chart will be refreshed automatically when you change either of the following fields:

- *Topology Type*: Indicates the monitored topology object, including Kubernetes Pod and Kubernetes Node.
- *Cluster*: Lists all clusters available in the monitored Kubernetes environment.
- *Namespace*: Lists all namespaces available in the monitored Kubernetes environment.
- *X Axis*: Indicates which metrics will be plotted on X axis.
- *Y Axis*: Indicates which metrics will be plotted on Y axis.
- Rendering related metrics:
 - *Color Metric*: Renders the color of circle based upon the selected metrics.
 - *Color Pattern*: Offers two patterns, Red to Green (larger value shows in red) or Green to Red (larger value shows in green).

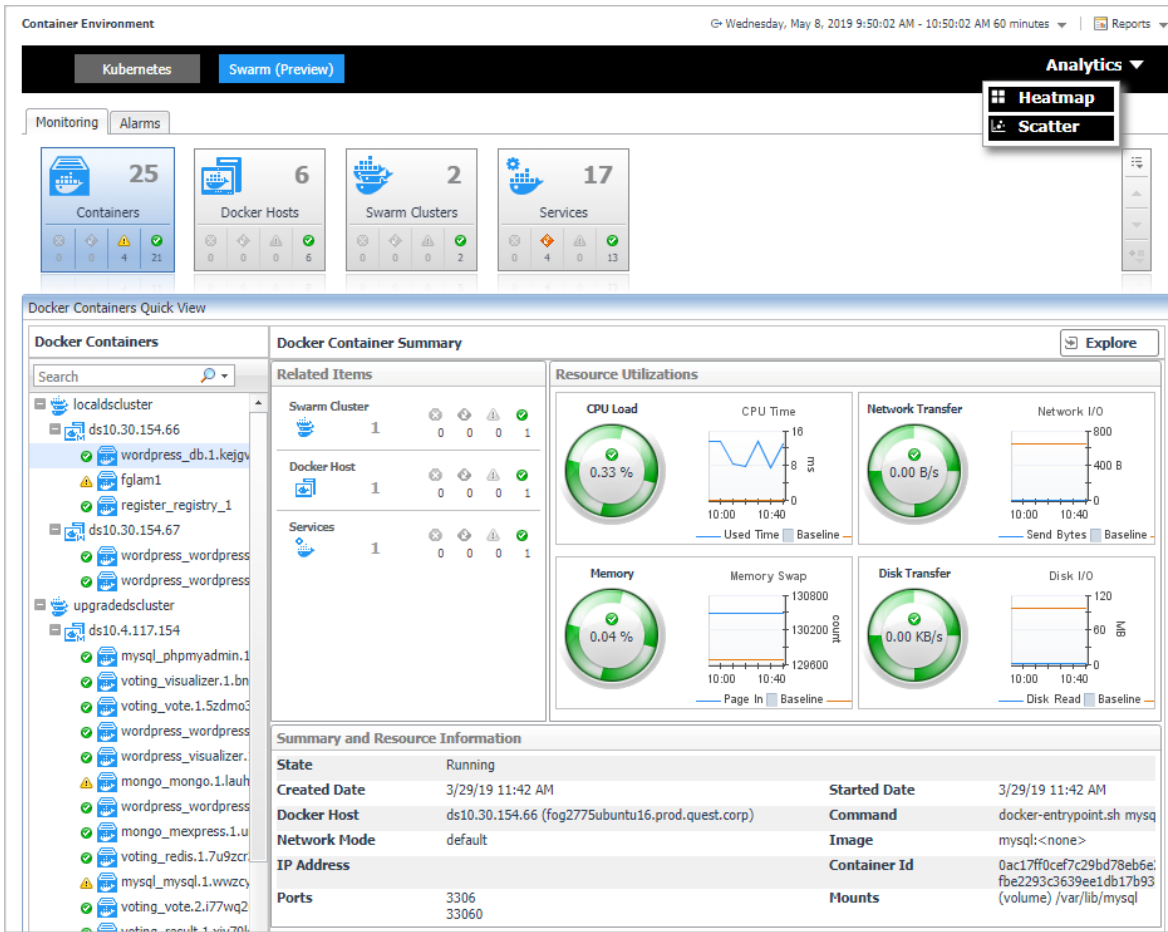
Figure 50 shows an example of Scatter Plot analytics. For more information, see:

- [Workloads Explorer view](#) on page 36
- [Workloads metrics](#) on page 75
- [Nodes Explorer view](#) on page 33
- [Node metrics](#) on page 77
- [Cluster Explorer view](#) on page 29
- [Docker Swarm metrics](#) on page 79

Docker Swarm analytics

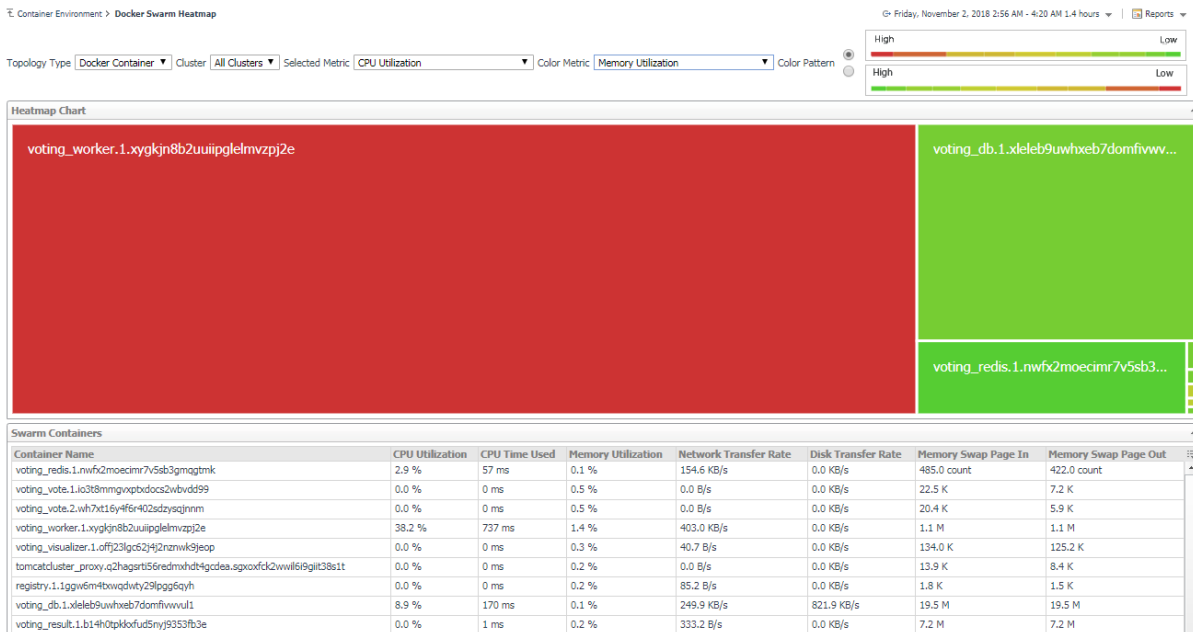
In the Container dashboard, choose **Docker Swarm** from the header. Then click **Analytics** from the header, a drop down view will display with **Heatmap** and **Scatter** on it. Click **Heatmap** will navigate to the Docker Swarm **Heatmap Analytics** dashboard, while click **Scatter** will navigate to the Docker Swarm **Scatter Plot Analytics** dashboard.

Figure 51. Docker Swarm Analytics Navigation



Heatmap analytics

Figure 52. Docker Swarm Heatmap Analytics Dashboard



Heat maps will be refreshed automatically when you change either of the following fields:

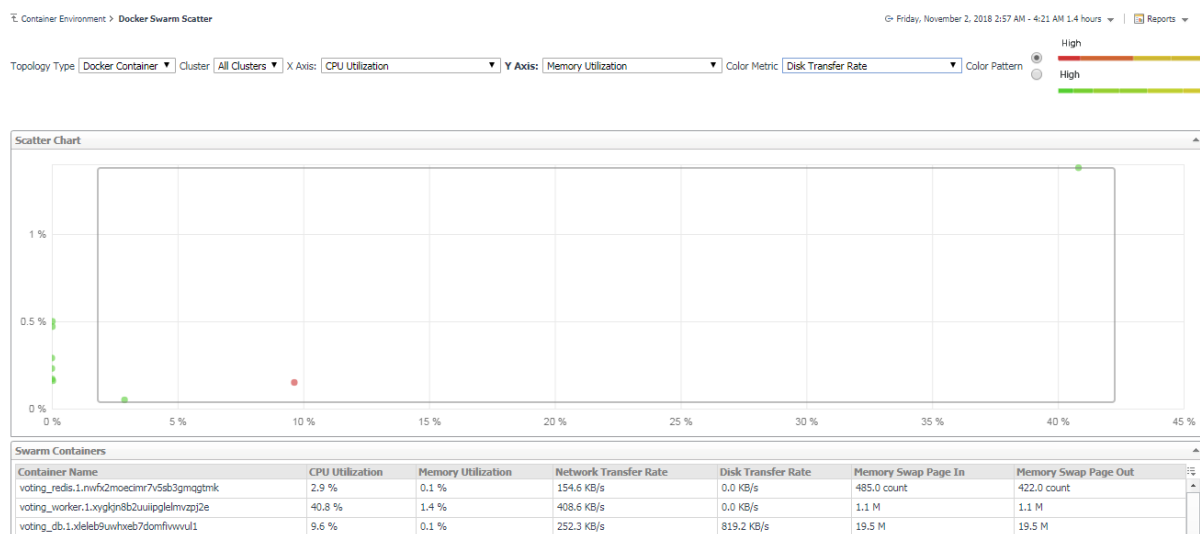
- **Topology Type:** Indicates the monitored topology object, including Docker Container and Docker Host.
- **Cluster:** Lists all clusters available in the monitored Docker Swarm environment.
- **Selected Metric:** Populates a rectangle based upon the selected metrics. For example, if you select *Memory Time Used* from the *Selected Metric* drop-down list, the rectangle area will be populated based on the used CPU time for the selected topology object. For more information about metrics, refer to [Docker Swarm metrics on page 79](#).
- **Rendering related metrics:** For example, if you select *CPU Utilization* and Red to Green, the rectangle of the topology object that has larger value of CPU Utilization will be rendered in red.
 - **Color Metric:** Renders the color of rectangle based upon the selected color metric.
 - **Color Pattern:** Offers two patterns, Red to Green (larger value shows in red) or Green to Red (larger value shows in green).

Figure 52 shows an example of heat map. This sample diagram represents the “voting_redis.1.nwfx2moecimr7v5sb3gmqgtmk” has the maximum amounts of CPU Utilization which is the largest in size, and also it has the higher Memory Utilization since it is in Red. If you switch the Color Pattern, then “voting_redis.1.nwfx2moecimr7v5sb3gmqgtmk” will turn to green. Clicking the object name on the heat map directs you to the relevant object *Explorer* dashboard. For more information, see:

- [Container Explorer view on page 62](#)
- [Docker Host Explorer view on page 63](#)
- [Container metrics on page 79](#)

Scatter Plot analytics

Figure 53. Docker Swarm Scatter Plot Analytics Dashboard



The points on the chart will be refreshed automatically when you change either of the following fields:

- **Topology Type:** Indicates the monitored topology object, including Docker Container and Docker Host.
- **Cluster:** Lists all clusters available in the monitored Docker Swarm environment.
- **X Axis:** Indicates which metrics will be plotted on X axis.
- **Y Axis:** Indicates which metrics will be plotted on Y axis.
- **Rendering related metrics:**
 - **Color Metric:** Renders the color of circle based upon the selected metrics.

- *Color Pattern*: Offers two patterns, Red to Green (larger value shows in red) or Green to Red (larger value shows in green).

Figure 53 shows an example of Scatter Plot analytics. The purple circle in the middle represents the following: “voting_redis.1.nwfx2moecimr7v5sb3gmqgtmk” CPU Utilization is 2.9%, its Memory Usage is 0.1%, and its value of Network Transfer Bytes is not high. For more information, see:

- [Container Explorer view](#) on page 62
- [Docker Host Explorer view](#) on page 63
- [Container metrics](#) on page 79

Domains and Object Groups

Domains

A domain is a group of monitored components organized by monitoring technology. This dashboard shows a summarized view of your monitored enterprise organized by domain. Click on a sub-domain for detailed information about the contents and health of the domain.

To access the Domains dashboard, on the Navigation panel, click **Dashboards > Services > Domains**.

Click the + icon to display the components under Container.

Figure 54. Container Components in Domains dashboard

Name	State	History	Alarms	Agents
Container	🔴			
Docker Swarm	🟢	████████████████████		
Kubernetes	🔴	████████████████████	27	15
Custom Applications	🟢			🟢
Databases	🟢			
End User	🟢			
Infrastructure	🟢			
Packaged Applications	🟢			

Click the State, History, Alarms, and Agents column, for detailed alarms and health information.

Object Groups

An object group is a mapping to a certain set of data types of the objects you are interested in.

To access the Object Groups dashboard, on the Navigation panel, click **Dashboards > Services > Object Groups**.

Figure 55. Object Groups for Container

Object Groups Monday, June 24, 2019 10:19:50 AM - 11:19:50 AM 60 minutes Reports

Search

Actions		Name ^	Description	Is Disabled	Created by Foglight
Edit	Test				
<input type="checkbox"/>	<input type="checkbox"/>	Agents	All Agent objects.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	All Models	All model objects.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	App	The Application Tier includes all objects from .NET, Siebel, SAP, PeopleSoft and...	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Azure	All Azure Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	DB	The Database Tier includes all objects from Oracle, SQLServer, Sybase and Db2 ca...	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Docker Swarm	All Docker Swarm Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Geo	All Geo Aware Applications.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Global Services	All Global Services.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Hosts	All Host objects.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Hyper-V	All Hyper-V Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Java EE	The Java EE Tier includes all objects from Weblogic, WebSphere, JBoss, OracleAS,...	false	true
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Kubernetes	All Kubernetes Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	OpenStack	All OpenStack Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	User	The User Tier includes all the objects for end user performance.	false	true
<input type="checkbox"/>	<input type="checkbox"/>	VMware	All VMware Objects	false	true
<input type="checkbox"/>	<input type="checkbox"/>	VMware Clusters	All VMware Clusters	false	true

Kubernetes has 19 subgroup(s)

Search

Actions		Name ^	Description	Data Type	Query Conditions	Is Disabled	Created by Foglight
Edit	Test						
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Clusters	All Kubernetes Clusters	KubeCluster	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Config Maps	All Kubernetes Config Maps	KubeConfigMap	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Cron Jobs	All Kubernetes Cron Jobs	KubeCronJob	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Daemon Sets	All Kubernetes Daemon Sets	KubeDaemonSet	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Deployments	All Kubernetes Deployments	KubeDeployment	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Endpoints	All Kubernetes Endpoints	KubeEndpoint	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Ingresses	All Kubernetes Ingresses	KubeIngress	n/a	false	true
<input type="checkbox"/>	<input type="checkbox"/>	Kubernetes Jobs	All Kubernetes Jobs	KubeJob	n/a	false	true

Select *Docker Swarm* or *Kubernetes* to display the subgroups.

Reference

- [Metrics](#)
 - [Kubernetes metrics](#)
 - [Docker Swarm metrics](#)
- [Rules](#)
 - [Kubernetes](#)
 - [Docker Swarm](#)
 - [Customization](#)

Metrics

Kubernetes metrics

Workloads metrics

Table 4. Workloads metrics

Metric name	Description
CPU Usage Rate	CPU usage rate on all cores in millicores/second.
CPU Request	CPU request (the guaranteed amount of resources) in millicores.
CPU Limit	CPU hard limit in millicores.
CPU Utilization	Percentage of CPU usage / CPU limit if user configured CPU limit for this pod.
Memory Usage	Total memory usage in bytes.
Memory Working Set	Total working set usage. Working set is the memory being used and not easily dropped by the kernel.
Memory RSS	RSS memory usage.
Memory Cache	Number of bytes of page cache memory.
Memory Swap	Container swap usage in bytes.
Memory Request	Memory request (the guaranteed amount of resources) in bytes.
Memory Limit	Memory hard limit in bytes.
Memory Utilization	Percentage of Memory usage / Memory limit if user configured Memory limit for this pod.
Network Send Bytes Rate	Network send bytes per second.
Network Receive Bytes Rate	Network receive bytes per second.

Table 4. Workloads metrics

Metric name	Description
Network Transfer Bytes Rate	Network send and receive bytes per second.
Network Send Errors Rate	Network send errors count per second.
Network Receive Errors Rate	Network receive errors count per second.
Network Transfer Errors Rate	Network send and receive errors count per second.
Network Send Packets Rate	Network send packets count per second.
Network Receive Packets Rate	Network receive packets count per second.
Network Transfer Packets Rate	Network send and receive packets count per second.
Network Send Dropped Packets Rate	Network send dropped packets count per second.
Network Receive Dropped Packets Rate	Network receive dropped packets count per second.
Network Transfer Dropped Packets Rate	Network send and receive dropped packets count per second.
Filesystem Usage	Number of bytes that are consumed by the container on this filesystem.
Filesystem Capacity	Number of bytes that can be consumed by the container on this filesystem.
Filesystem Utilization	Percentage of Filesystem Usage / Filesystem Capacity.
Filesystem Read Bytes Rate	Filesystem read bytes per second.
Filesystem Write Bytes Rate	Filesystem write bytes per second.
Filesystem Read Rate	Filesystem read counts per second.
Filesystem Write Rate	Filesystem write counts per second.

i | **NOTE:**

- Workloads metrics are for the following workloads types, Deployment, Daemon Set, Stateful Set, Replica Set, Replication Controller, Cron Job, Job, and Pod.
- For workloads types except Pods, the metrics values come from the aggregated metrics values of the related Pods.

Container metrics

Table 5. Container metrics

Metric name	Description
CPU Usage Rate	CPU usage rate on all cores in millicores/second.
CPU Request	CPU request (the guaranteed amount of resources) in millicores.
CPU Limit	CPU hard limit in millicores.
CPU Utilization	Percentage of CPU usage / CPU limit if user configured CPU limit for this pod.
Memory Usage	Total memory usage in bytes.
Memory Working Set	Total working set usage. Working set is the memory being used and not easily dropped by the kernel.
Memory RSS	RSS memory usage.
Memory Cache	Number of bytes of page cache memory.
Memory Swap	Container swap usage in bytes.
Memory Request	Memory request (the guaranteed amount of resources) in bytes.
Memory Limit	Memory hard limit in bytes.
Memory Utilization	Percentage of Memory usage / Memory limit if user configured Memory limit for this pod.
Filesystem Usage	Number of bytes that are consumed by the container on this filesystem.
Filesystem Capacity	Number of bytes that can be consumed by the container on this filesystem.
Filesystem Utilization	Percentage of Filesystem Usage / Filesystem Capacity.
Filesystem Read Bytes Rate	Filesystem read bytes per second.
Filesystem Write Bytes Rate	Filesystem write bytes per second.
Filesystem Read Rate	Filesystem read counts per second.
Filesystem Write Rate	Filesystem write counts per second.

Node metrics

Table 6. Node metrics

Metric name	Description
CPU Usage Rate	CPU usage rate on all cores in millicores/second.
CPU Request	CPU request (the guaranteed amount of resources) in millicores.
CPU Limit	CPU hard limit in millicores.
CPU Utilization	CPU utilization as a share of node allocatable.
CPU Allocatable	Available CPU to allocate to workloads.
CPU Capacity	Hard CPU capacity of node.
Memory Usage	Total memory usage in bytes.
Memory Working Set	Total working set usage. Working set is the memory being used and not easily dropped by the kernel.
Memory RSS	RSS memory usage.
Memory Cache	Number of bytes of page cache memory.
Memory Swap	Container swap usage in bytes.
Memory Request	Memory request (the guaranteed amount of resources) in bytes.
Memory Capacity	Hard memory capacity of node.
Memory Limit	Memory hard limit in bytes.

Table 6. Node metrics

Metric name	Description
Memory Allocatable	Available Memory to allocate to workloads.
Memory Utilization	Memory utilization as a share of memory allocatable.
Network Send Bytes Rate	Network send bytes per second.
Network Receive Bytes Rate	Network receive bytes per second.
Network Transfer Bytes Rate	Network send and receive bytes per second.
Network Send Errors Rate	Network send errors count per second.
Network Receive Errors Rate	Network receive errors count per second.
Network Transfer Errors Rate	Network send and receive errors count per second.
Network Send Packets Rate	Network send packets count per second.
Network Receive Packets Rate	Network receive packets count per second.
Network Transfer Packets Rate	Network send and receive packets count per second.
Network Send Dropped Packets Rate	Network send dropped packets count per second.
Network Receive Dropped Packets Rate	Network receive dropped packets count per second.
Network Transfer Dropped Packets Rate	Network send and receive dropped packets count per second.
Filesystem Usage	Number of bytes that are consumed by the container on this filesystem.
Filesystem Capacity	Number of bytes that can be consumed by the container on this filesystem.
Filesystem Utilization	Percentage of Filesystem Usage / Filesystem Capacity.
Filesystem Inodes Usage	Number of Inodes that are consumed by the container on this filesystem.
Filesystem Inodes Total	Number of Inodes that can be consumed by the container on this filesystem.
Filesystem Inodes Utilization	Percentage of Filesystem Inodes Usage / Filesystem Inodes Capacity.
Filesystem Read Bytes Rate	Filesystem read bytes per second.
Filesystem Write Bytes Rate	Filesystem write bytes per second.
Filesystem Read Rate	Filesystem read counts per second.
Filesystem Write Rate	Filesystem write counts per second.

Docker Swarm metrics

Container metrics

Table 7. Container metrics

Metric name	Description
CPU Utilization	CPU utilization.
CPU Time Used	Total CPU time that a container used.
CPU Throttled Time	Total time that a container's CPU usage was throttled.
Memory Page Fault	Total page fault count of a container's Memory.
Memory Consumed	Total memory consumed of a container in bytes.
Memory Utilization	Memory utilization.
Memory PageIn Rate	Total page in count of a container's Memory.
Memory PageOut Rate	Total page out count of a container's Memory.
Disk Read Bytes	Total disk read bytes.
Disk Write Bytes	Total disk write bytes.
Disk Transfer Rate	Sum of total disk read and write bytes.
Network Send Packets	Total network send packets count.
Network Receive Packets	Total network receive packets count.
Network Send Bytes	Total network send bytes.
Network Receive Bytes	Total network receive bytes.
Network Inbound Dropped Packets	Total dropped packet count of all the packets coming into the container.
Network Outbound Dropped Packets	Total dropped packet count of all the packets going out from the container.
Network Transfer Rate	Sum of network send bytes and receive bytes per seconds during a specific period.

Rules

i | **NOTE:** Foglight uses Alarm Templates to gather alarm rules into a domain-specific template that is easily modified and applied to targets. Alarm rules can be modified using the Alarm Templates dashboard.

Foglight for Container Management allows you to create flexible rules that can be applied to complex interrelated data from multiple sources within your clusters. You can associate several different actions with a rule, configure a rule so that it does not fire repeatedly, and associate a rule with schedules to define when it should be evaluated or not.

Different types of data can be used in rules, including registry variables, raw metrics, derived metrics, and topology object properties.

There are two types of rules: simple rules and multiple-severity rules. A simple rule has a single condition, and can be in one of three states: *Fire*, *Undefined*, or *Normal*. A multiple-severity rule can have up to five severity levels: *Undefined*, *Fatal*, *Critical*, *Warning*, and *Normal*.

Rule conditions are regularly evaluated against monitoring data (metrics and topology object properties collected from your monitored environment and transformed into a standard format). Therefore, the state of the rule can change if the data changes. For example, if a set of monitoring data matches a simple rule's condition, the rule enters the *Fire* state. If the next set does not match the condition, the rule exits the *Fire* state and enters the *Normal* state.

Rules can be configured to send emails, pager messages, or perform other actions you define. Performance data can be viewed and analyzed using Foglight for Container Management.

Foglight for Container Management includes a number of predefined rules used to monitor the health of your container clusters. You are allowed to modify these rules to satisfy your different requirements. Many of these rules listed and described in this section have thresholds defined within them. Those thresholds include standard deviations, utilization percentages, and so on, are default values predefined in the registry.

For information on customizing alarms templates and rule, refer to [Modifying alarms and threshold values and Customization](#).

Kubernetes

All rules are controlled by registry variable Kubernetes:AlertSensitivity. If the value is 0, then no alarm can be fired. If the value is 1, warning level alarm can be fired. If the value is above 1, then all level alarm can be fired.

Kubernetes Administrator email address can be configured in Registry Variable KubernetesAdmin.

Health Check

Kubernetes Pod Health Check

Purpose

This rule detects abnormal Pod health status and fires alarm for different severity abnormal health status.

Scope

KubePod

Conditions and Severities

Conditions	Severity	Action
Pods that is in Failed or Unknown status. Or the node which is running the pod gets disconnected.	Critical	Send email to Kubernetes Administrator.
Pods that is in CrashLoopBackOff status.	Warning	None

Kubernetes Pod Health Check (Pending Phase)

Purpose

This rule detects Pods that stays in pending phase for an abnormal long time.

Scope

KubePod

Conditions and Severities

Conditions	Severity	Action
Pods that is pending for two continuous data submission periods because of Failed to schedule to Node.	Critical	Send email to Kubernetes Administrator.
Pods that is pending for two continuous data submission periods because container is not ready.	Warning	None

Kubernetes Container Health Check

Purpose

This rule detects abnormal Container health status and fires alarm for different severity abnormal health status.

Scope

KubeContainer

Conditions and Severities

Conditions	Severity	Action
Container that is terminated for abnormal reasons.	Critical	Send email to Kubernetes Administrator.

Kubernetes Node Health Check

Purpose

This rule detects abnormal Node health status and fires alarm for different severity abnormal health status.

Scope

KubeNode

Conditions and Severities

Conditions	Severity	Action
Nodes that is not Ready or out of disk or network unavailable.	Critical	Send email to Kubernetes Administrator.
Nodes whose memory or disk is under pressure.	Warning	None

Kubernetes Deployment Health Check

Purpose

This rule detects abnormal Deployment health status and fires alarm for different severity abnormal health status.

Scope

KubeDeployment

Conditions and Severities

Conditions	Severity	Action
Deployment is not available.	Critical	Send email to Kubernetes Administrator.
Deployment has failed to create some of the replicated pods.	Warning	None

Kubernetes Daemon Set Health Check

Purpose

This rule detects abnormal Daemon Set health status and fires alarm for different severity abnormal health status.

Scope

KubeDaemonSet

Conditions and Severities

Conditions	Severity	Action
Some of the pods created by the Daemon Set is not available or mis-scheduled.	Critical	Send email to Kubernetes Administrator.
The daemon set doesn't have enough replicated pods running that meets its desired replicated pods count.	Warning	None

Kubernetes Job Health Check

Purpose

This rule detects abnormal Job health status and fires alarm for different severity abnormal health status.

Scope

KubeJob

Conditions and Severities

Conditions	Severity	Action
Job that is failed.	Warning	None

Kubernetes Persistent Volume Health Check

Purpose

This rule detects abnormal Persistent Volume health status and fires alarm for different severity abnormal health status.

Scope

KubePersistentVolume

Conditions and Severities

Conditions	Severity	Action
Persistent Volume that is in failed status.	Warning	None

Kubernetes Persistent Volume Claim Health Check

Purpose

This rule detects abnormal Persistent Volume Claim health status and fires alarm for different severity abnormal health status.

Scope

KubePersistentVolumeClaim

Conditions and Severities

Conditions	Severity	Action
Persistent Volume Claim that is in failed status.	Warning	None

Kubernetes Persistent Volume Claim Health Check (Long Pending)

Purpose

This rule detects abnormal long pending Persistent Volume Claim and fires alarm for different severities.

Scope

KubePersistentVolumeClaim

Conditions and Severities

Conditions	Severity	Action
Persistent Volume Claim that is pending for two continuous data submission periods.	Critical	None

Kubernetes Container Restart for Termination with Invalid State

Purpose

This rule checks for container that restarted and previously terminated in an invalid state within the last 12 hours.

Scope

KubeContainer

Conditions and Severities

Conditions	Severity	Action
Container that restarted and previously terminated in an invalid state within the last 12 hours.	Critical	Send email to Kubernetes Administrator

Usage

i | **NOTE:** All the Pod usage related rules are disabled by default, these rules are used as default values or examples for the customers to customize their different kinds of workloads. For how to customize the rules, refer to [Customization](#) on page 95.

Kubernetes Pod CPU Utilization

Purpose

This rule detects abnormal CPU Utilization for Pods, and fires alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities. This rule only works for those Pods that configures CPU limit.

Scope

KubePodCpu

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodCpu UtilizationFatal	Send email to Kubernetes Administrator
Pods whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes:PodCpu UtilizationCritical	None
Pods whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes:PodCpu UtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Pod Memory Utilization

Purpose

This rule detects abnormal Memory Utilization for Pods, and fires alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities. This rule only works for those Pods that configures Memory limit.

Scope

KubePodMemory

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodMemory UtilizationFatal	Send email to Kubernetes Administrator
Pods whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes:PodMemory UtilizationCritical	None
Pods whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes:PodMemory UtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Pod CPU Usage

Purpose

This rule detects abnormal CPU Usage for Pods, and fires alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubePodCpu

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodCpu UsageFatal	Send email to Kubernetes Administrator
Pods whose usage is above the value configured in critical Threshold.	Critical	Kubernetes:PodCpu UsageCritical	None
Pods whose usage is above the value configured in warning Threshold.	Warning	Kubernetes:PodCpu UsageWarning	None

*Note: the unit is millicores/second.

Kubernetes Pod Memory Usage

Purpose

This rule detects abnormal Memory Usage for Pods, and fires alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubePodMemory

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodMemory UsageFatal	Send email to Kubernetes Administrator
Pods whose usage is above the value configured in critical Threshold.	Critical	Kubernetes:PodMemory UsageCritical	None
Pods whose usage is above the value configured in warning Threshold.	Warning	Kubernetes:PodMemory UsageWarning	None

*Note: the unit is bytes.

Kubernetes Pod Network Receive

Purpose

This rule detects abnormal Network Receive Rate in bytes/second for Pods, and fire alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubePodNetwork

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodNetwork ReceiveFatal	Send email to Kubernetes Administrator
Pods whose usage is above the value configured in critical Threshold.	Critical	Kubernetes:PodNetwork ReceiveCritical	None
Pods whose usage is above the value configured in warning Threshold.	Warning	Kubernetes:PodNetwork ReceiveWarning	None

*Note: the unit is bytes/second.

Kubernetes Pod Network Send

Purpose

This rule detects abnormal Network Send Rate in bytes/second for Pods, and fire alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubePodNetwork

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodNetworkSendFatal	Send email to Kubernetes Administrator
Pods whose usage is above the value configured in critical Threshold.	Critical	Kubernetes:PodNetworkSendCritical	None
Pods whose usage is above the value configured in warning Threshold.	Warning	Kubernetes:PodNetworkSendWarning	None

*Note: the unit is bytes/second.

Kubernetes Pod Filesystem Utilization

Purpose

This rule checks kubernetes pod filesystem utilization to see if it reaches the defined threshold. The filesystem utilization rule only works for pod that configures filesystem limit.

Scope

KubePodStorage

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Pods whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes:PodFilesystemUtilizationFatal	Send email to Kubernetes Administrator
Pods whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes:PodFilesystemUtilizationCritical	None
Pods whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes:PodFilesystemUtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Node CPU Utilization

Purpose

This rule detects abnormal CPU Utilization for Nodes, and fire alarm on different severities. It is enabled by default. You can change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubeNodeCpu

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose utilization is above the value configured in fatal Threshold.	Fatal	Kubernetes:NodeCpu UtilizationFatal	Send email to Kubernetes Administrator
Nodes whose utilization is above the value configured in critical Threshold.	Critical	Kubernetes:NodeCpu UtilizationCritical	None
Nodes whose utilization is above the value configured in warning Threshold.	Warning	Kubernetes:NodeCpu UtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Node Memory Utilization

Purpose

This rule detects abnormal Memory Utilization for Nodes, and fire alarm on different severities. It is enabled by default. You can change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubeNodeMemory

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose utilization is above the value configured in fatal Threshold.	Fatal	Kubernetes:NodeMemory UtilizationFatal	Send email to Kubernetes Administrator
Nodes whose utilization is above the value configured in critical Threshold.	Critical	Kubernetes:NodeMemory UtilizationCritical	None
Nodes whose utilization is above the value configured in warning Threshold.	Warning	Kubernetes:NodeMemory UtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Node Network Receive

Purpose

This rule detects abnormal Network Receive Rate in bytes/second for Nodes, and fire alarm on different severities. It is enabled by default. You can change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubeNodeNetwork

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose utilization is above the value configured in fatal Threshold.	Fatal	Kubernetes:NodeNetworkReceiveFatal	Send email to Kubernetes Administrator
Nodes whose utilization is above the value configured in critical Threshold.	Critical	Kubernetes:NodeNetworkReceiveCritical	None
Nodes whose utilization is above the value configured in warning Threshold.	Warning	Kubernetes:NodeNetworkReceiveWarning	None

*Note: the unit is bytes/second.

Kubernetes Node Network Send

Purpose

This rule detects abnormal Network Send Rate in bytes/second for Nodes, and fire alarm on different severities. It is enabled by default. You can change value of the registry variables or use your own value to change the threshold of each severities.

Scope

KubeNodeNetwork

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose utilization is above the value configured in fatal Threshold.	Fatal	Kubernetes:NodeNetworkSendFatal	Send email to Kubernetes Administrator
Nodes whose utilization is above the value configured in critical Threshold.	Critical	Kubernetes:NodeNetworkSendCritical	None
Nodes whose utilization is above the value configured in warning Threshold.	Warning	Kubernetes:NodeNetworkSendWarning	None

*Note: the unit is bytes/second.

Kubernetes Node Network Transfer

Purpose

Periodically check Kubernetes node Network Transfer Rate in bytes/second, if the value is too high and changes too much, then an alarm will be triggered.

Scope

KubeNodeNetwork

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes: NodeNetworkTransferFatal	Send email to Kubernetes Administrator
Nodes whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes: NodeNetworkTransferCritical	None
Nodes whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes: NodeNetworkTransferWarning	None

*Note: the unit is bytes/second.

Kubernetes Node Filesystem Utilization

Purpose

Periodically check Kubernetes node Filesystem Utilization, if the value is too high and changes too much, then an alarm will be triggered.

Scope

KubeNodeFilesystem

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Nodes whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	KubeNodeFilesystemUtilizationFatal	Send email to Kubernetes Administrator
Nodes whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	KubeNodeFilesystemUtilizationCritical	None
Nodes whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	KubeNodeFilesystemUtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Container CPU Utilization

Purpose

This rule checks Kubernetes container CPU utilization to see if it reaches the defined threshold. The CPU utilization rule only works for container that configures CPU limit.

Scope

KubeContainerCpu

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Containers whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes:ContainerCpu UtilizationFatal	Send email to Kubernetes Administrator
Containers whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes:ContainerCpu UtilizationCritical	None
Containers whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes:ContainerCpu UtilizationWarning	None

*Note: the unit is percentage.

Kubernetes Container Memory Utilization

Purpose

This rule checks Kubernetes container Memory utilization to see if it reaches the defined threshold. The Memory utilization rule only works for container that configures Memory limit.

Scope

KubeContainerMemory

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Containers whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Kubernetes: ContainerMemoryUtilizationFatal	Send email to Kubernetes Administrator
Containers whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Kubernetes: ContainerMemoryUtilizationCritical	None
Containers whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Kubernetes: ContainerMemoryUtilizationWarning	None

*Note: the unit is percentage.

Container Cost Cluster Budget Over Spending

Purpose

Detecting container cluster whose budget is over spending.

Scope

AbstractKubernetesClusterCost

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Cluster current month cost which is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	ContainerCost:OverBudgetFatal ContainerCostAdmin	Send email to Kubernetes Administrator
Cluster current month cost which is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	ContainerCost:OverBudgetCritical	None
Cluster current month cost which is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	ContainerCost:OverBudgetWarning	None

*Note: the unit is percentage.

Docker Swarm

All rules are controlled by registry variable Docker:AlertSensitivity. If the value is 0, then no alarm can be fired. If the value is 1, warning level alarm can be fired. If the value is above 1, then all level alarm can be fired.

Docker Swarm Administrator email address can be configured in Registry Variable Docker:DockAdmin.

Health Check

Docker Container Status

Purpose

This rule detects abnormal Container health status and fires alarm for different severity abnormal health status.

Scope

DockerContainer

Conditions and Severities

Conditions	Severity	Action
Container that is already stopped for abnormal reason.	Critical	Send email to Docker Swarm Administrator

Docker Container Status - Paused

Purpose

This rule detects abnormal long-time paused Container and fires alarm for different severity abnormal health status.

Scope

DockerContainer

Conditions and Severities

Conditions	Severity	Action
Container paused for two continuous data submission periods.	Warning	None

Docker Service Status

Purpose

This rule detects abnormal Docker Swarm Service health status and fires alarm for different severity abnormal health status.

Scope

DockerService

Conditions and Severities

Conditions	Severity	Action
Missing some of the replicated task running for this service.	Critical	Send email to Docker Swarm Administrator

Docker Task Status

Purpose

This rule detects abnormal Docker Swarm Task health status and fires alarm for different severity abnormal health status.

Scope

DockerTask

Conditions and Severities

Conditions	Severity	Action
Task that is in failed, orphaned or remove status.	Critical	Send email to Docker Swarm Administrator

Docker Task Status -- pending

Purpose

This rule detects abnormal long-time pending Docker Swarm Task and fires alarm for different severity abnormal health status.

Scope

DockerTask

Conditions and Severities

Conditions	Severity	Action
Task that is in pending status for two continuous data submission periods.	Warning	None

Usage

Docker Swarm Container CPU Utilization

Purpose

This rule detects abnormal CPU Utilization for Docker Swarm Containers, and fire alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities. This rule only works for those Containers that configures CPU limit.

Scope

DockerContainerCPU

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Container whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Docker:ContainerCpu UtilizationFatal	Send email to Kubernetes Administrator
Pods whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Docker:ContainerCpu UtilizationCritical	None
Pods whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Docker:ContainerCpu UtilizationWarning	None

*Note: the unit is percentage.

Docker Swarm Container Memory Utilization

Purpose

This rule detects abnormal Memory Utilization for Docker Swarm Containers, and fire alarm on different severities. It is disabled by default. You can customize it and enable it based on your different requirements. For more details about customization, refer to [Customization](#) on page 95. You can also change value of the registry variables or use your own value to change the threshold of each severities. This rule only works for those Containers that configures Memory limit.

Scope

DockerContainerMemory

Conditions and Severities

Conditions	Severity	Threshold (Registry Variable)*	Action
Container whose usage is about to reach the limit, the ration is above the value configured in fatal Threshold.	Fatal	Docker:ContainerMemory UtilizationFatal	Send email to Kubernetes Administrator
Pods whose usage is about to reach the limit, the ration is above the value configured in critical Threshold.	Critical	Docker:ContainerMemory UtilizationCritical	None
Pods whose usage is about to reach the limit, the ration is above the value configured in warning Threshold.	Warning	Docker:ContainerMemory UtilizationWarning	None

*Note: the unit is percentage.

Modifying alarms and threshold values

i | **IMPORTANT:** Avoid editing rules in the **Administration > Rules & Notifications > Rule Management** dashboard. Default rules may be modified during regular software updates and your edits will be lost. Always use the **Alarm Templates** dashboard.

You can and should modify the thresholds associated with alarms to better suit your environment. If you find that alarms are firing for conditions that you consider to be acceptable, you can change the threshold values that trigger the alarm. You can also enable or disable severity levels to better suit your environment.

When a rule has severity levels, a Threshold section appears in the Alarm Settings tab showing the severity levels and bounds by agent. Many rules, such as Baseline rules, do not have severity levels and thresholds.

When editing thresholds, ensure that the new values make sense in context with the other threshold values. For most metrics, threshold values are set so that Warning < Critical < Fatal. However, in metrics where normal performance has a higher value, the threshold values are reversed: Warning > Critical > Fatal.

To change alarm status and threshold values:

i | **IMPORTANT:** The procedure below is a summary. Refer to **Viewing, Creating, and Managing Alarm Templates** in the *Foglight User Guide* for more information on editing alarm templates for more information on working with alarm templates.

- 1 In the **Navigation** panel, click **Alarm Templates**.
- 2 If you have previously configured an alarm template, select that template.
Otherwise, click the Factory Template to view the default rules. Duplicate the factory template to make an editable copy, selecting the appropriate domains.
- 3 Click the appropriate domain tab.
- 4 Scroll or search to find the alarm rule you want to edit. Click the rule to select it.

Edit the rules using the procedure described in **Viewing, Creating, and Managing Alarm Templates** in the *Foglight User Guide*

Customization

To customize a rule, *Rule Scope* and *Condition* will be used frequently.

To access *Rule Scope* and *Condition*, do the following:

- 1 Under **Dashboards**, click **Administration > Rules & Notifications > Rules**, then click on the rule and select *View and Edit*.
- 2 Click **Rule Editor** on the *Rule Detail* popup dialog box. Then click **Continue** on the *Confirm Edit Rule* popup dialog box.
- 3 On the **Rule Editor** dashboard, *Rule Scope* can be located on the **Rule Definition** tab and *Condition* can be located on the **Condition & Actions** tab.

Figure 56. Rule Scope

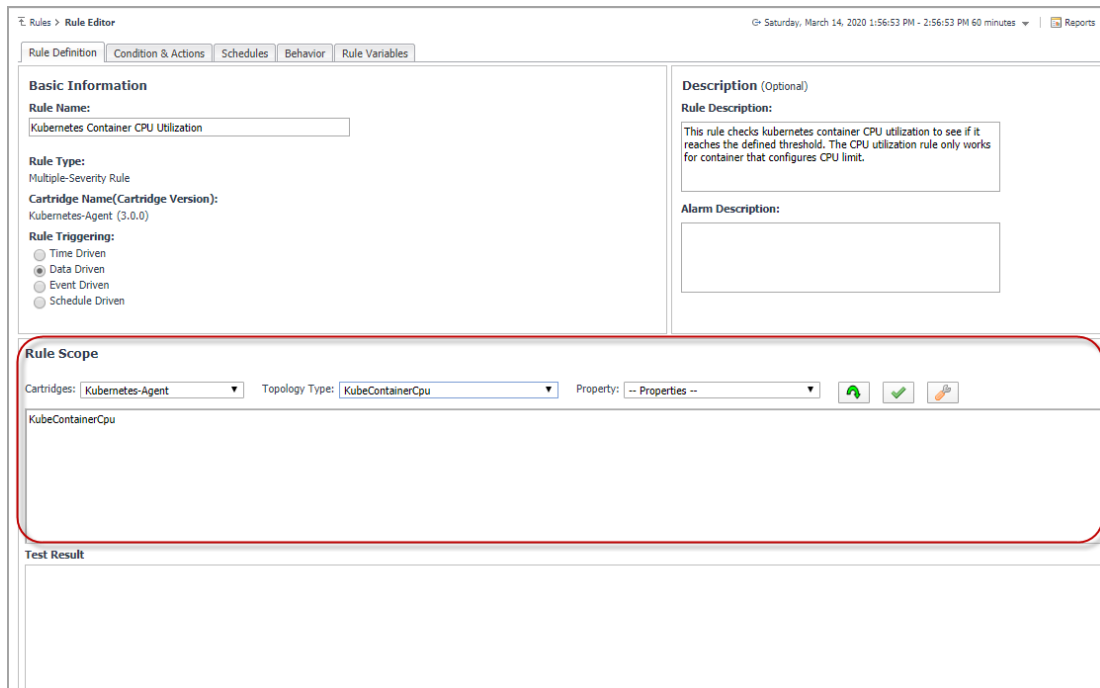
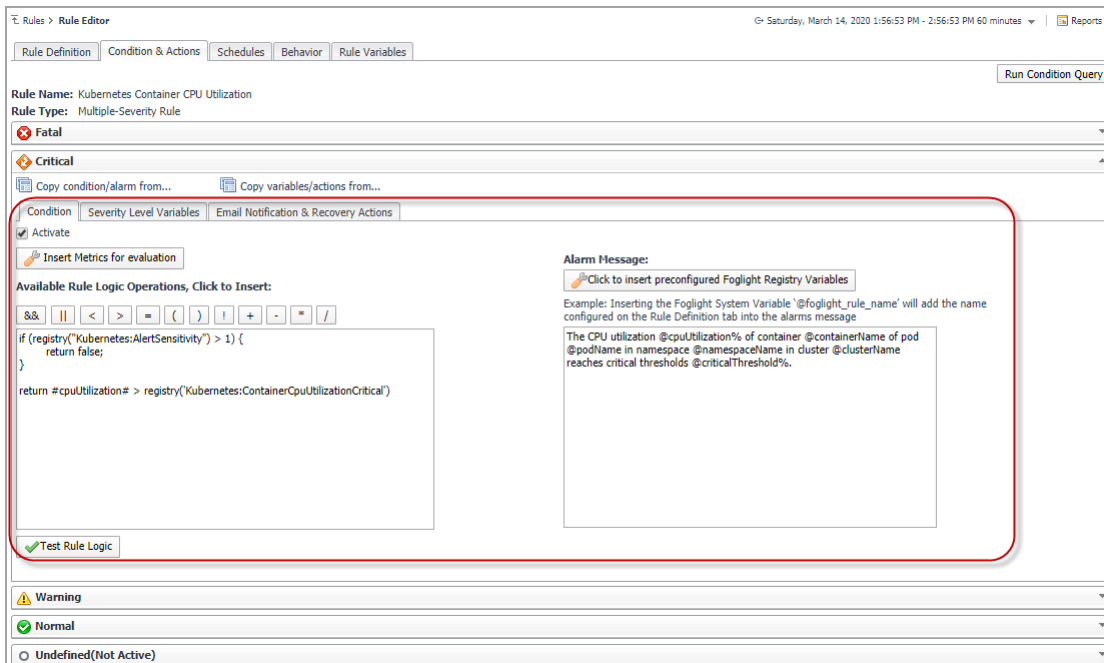


Figure 57. Condition



Kubernetes

Filter Pods by Cluster

Finding Pods inside cluster “kubecuster”, enter following statement in the Scope of a rule, and choose KubePod as the Topology Type in the Rule Scope.

```
KubePod where namespace.cluster.name='kubecuster'
```

Filter Pods by Namespace

Finding Pods inside namespace “default” of Cluster “kubecuster”, enter following statement in the Scope of a rule, and choose KubePod as the Topology Type in the Rule Scope.

```
KubePod where namespace.cluster.name='kubecuster' and namespace.name='test'
```

Filter Nodes by Cluster

Finding Nodes inside cluster “kubecuster”, enter following statement in the Scope of a rule, and choose KubeNode as the Topology Type in the Rule Scope.

```
KubeNode where cluster.name='nancyakscluster'
```

Filter Pod by Labels

Find Pods with labels “run=nginx” and “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose KubePod as the Topology Type in the Rule Scope.

```
KubePod where labels.key='run' and labels.value='nginx-rollingupdate' and labels.key='env' and labels.value='prod'
```

If you want to find Pods by labels in namespace “test” of cluster “kubecuster”, you can append `and namespace.cluster.name='kubecuster' and namespace.name='test'` to the end of above statement.

Filter Node by Labels

Find Nodes with labels “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose KubeNode as the Topology Type in the Rule Scope.

```
KubeNode where labels.key='env' and labels.value='prod'
```

If you want to find Nodes by labels in cluster “kubecuster”, you can append *and cluster.name='kubecuster'* to the end of above statement.

Filter Pod Metrics by Pod Labels

Find Pods Metrics with labels “run=nginx” and “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose KubeHeapsterMetrics as the Topology Type in the Rule Scope.

```
KubePod.metrics where object.labels.key='run' and object.labels.value='nginx' and object.labels.key='env' and object.labels.value='prod'
```

If you want to find Pods by labels in namespace “test” of cluster “kubecuster”, you can append *and namespace.cluster.name='kubecuster' and namespace.name='test'* to the end of above statement.

Filter Nodes Metrics by Node Labels

Find Node Metrics with labels “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose KubeHeapsterMetrics as the Topology Type in the Rule Scope.

```
KubeNode.metrics where object.labels.key='env' and object.labels.value='prod'
```

If you want to find Nodes by labels in cluster “kubecuster”, you can append *and cluster.name='kubecuster'* to the end of above statement.

Docker Swarm

Filter Container by Swarm Cluster

Find Containers in cluster “dockercluster”, enter following statement in the Scope of a rule, and choose DockerContainer as the Topology Type in the Rule Scope.

```
DockerContainer where dockerSwarm.service.cluster.name='kicakdscluster'
```

Filter Container by Labels

Find Containers with labels “com.docker.stack.namespace=nginx” and “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose DockerContainer as the Topology Type in the Rule Scope.

```
DockerContainer where labels.key='com.docker.stack.namespace' and labels.value='nginx' and labels.key='env' and labels.value='prod'
```

If you want to find Containers by labels in cluster “swarmcluster”, you can append *and dockerSwarm.service.cluster.name='kicakdscluster'* to the end of above statement.

Filter Docker Host by Swarm Cluster

Find Docker Hosts in cluster “dockercluster”, enter following statement in the Scope of a rule, and choose DockerHost as the Topology Type in the Rule Scope.

```
DockerHost where dockerSwarmNodeInfo.node.cluster.name='kicakdscluster'
```

Filter Container CPU Usage by Container Labels

Find Container CPU Usage by container labels “com.docker.stack.namespace=nginx” and “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose DockerContainerCPU as the Topology Type in the Rule Scope.

DockerContainerCPU where *container.labels.key='com.docker.stack.namespace'* and *container.labels.value='nginx'* and *container.labels.key='env'* and *container.labels.value='prod'*

If you want to find Containers by labels in cluster “swarmcluster”, you can append *and container.dockerSwarm.service.cluster.name='kicakdscluster'* to the end of above statement.

Filter Container Memory Usage by Container Labels

Find Container CPU Usage by container labels “com.docker.stack.namespace=nginx” and “env=prod” among all clusters, enter following statement in the Scope of a rule, and choose *DockerContainerMemory* as the Topology Type in the Rule Scope.

DockerContainerMemory where *container.labels.key='com.docker.stack.namespace'* and *container.labels.value='nginx'* and *container.labels.key='env'* and *container.labels.value='prod'*

If you want to find Containers by labels in cluster “swarmcluster”, you can append *and container.dockerSwarm.service.cluster.name='kicakdscluster'* to the end of above statement.

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit <https://www.quest.com/>.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.