



Alerts Reference

This chapter lists the default alerts available in the SMI Cluster Manager.

- [Alerts Reference, on page 1](#)

Alerts Reference

Alerts Summary

Table 1: Alerts Summary

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
server-alert	UCS server hardware alerts	Always	Depending on the hardware component, it can impact the full node functionality and might cause all pods to go down on the node.	k8s-node-not-ready kvm-node-not-ready	Verify the CEE alerts. If there is a k8s-node-not-ready or kvm-node-not-ready alert, then it means the hardware issue might have caused the OS to crash.	Check the CIMC of the node to troubleshoot the problem and decide if an RMA is required.
server-not-reachable-alert	UCS server hardware alert	Always	CIMC server is not reachable.	—	Check CEE and if the alert is present, escalate immediately.	—
k8s-node-not-ready	K8s node is in NotReady state for more than 1 minute	30 minutes	The node is not reporting health-check. There could be multiple reasons for this alert, for example, node down, OS crash, network issues, hardware, or a bug.	-	Verify the node status on the cluster. Investigate potential reason for the node being not ready.	This may be expected during MW. If node stays in the non-ready state outside MW, then contact your Cisco Account representative to troubleshoot.

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
k8s-node-status-change	K8s node Ready status changed during the last 5 minutes	30 minutes	The node changed state from Ready to Not Ready and back due to network issues, reboots, or other issues.	-	Verify the node status on the cluster. Investigate potential reason for the node changing status.	This may be expected during MW. If a node continues to change state outside MW, then contact your Cisco Account representative to troubleshoot.
kvm-node-not-ready	KVM node is not reachable for more than 1 minute	1 hour	The node is not reporting health-check. There could be multiple reasons for this alert, for example, node down, OS crash, network issues, hardware, or a bug.	-	Verify the CEE alerts to ensure that the alerts are active.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
kvm-tunnels-flapping	KVM node reachability changed during the last 5 minutes	1 hour	The connectivity to KVM node is not stable.	-	Verify the CEE alerts to ensure that the alerts are active.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-pod-restarting	Pod restarting one or more times during the last 5 minutes	1 hour	The overall impact is minimal and is expected during MW. Verify the service alerts.	-	Verify the CEE alerts to ensure that the alerts are active.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-pod-crashing-loop	Pod restarting two or more times during the last 5 minutes	1 hour	The overall impact is minimal and is expected during MW. Verify the service alerts. This alert indicates a continuous problem that needs investigation.	-	Verify the CEE alerts to ensure that the alerts are active.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
k8s-pod-pending	Pod is in pending state and cannot be scheduled for more than 1 minute	1 hour	This alert is expected during MW or when a node is down. Otherwise, this alert might be because of deployment misconfiguration.	k8s-node-not-ready	Verify the CEE alerts to ensure that the alerts are active. Check if nodes are in the NotReady state.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-pod-not-ready	Pod is not able to get into ready state for more than 1 minute	1 hour	This alert is expected for a short duration during MW, otherwise it points to application issues that need to be investigated.	-	Verify the CEE alerts to ensure that the alerts are active.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-deployment-generation-mismatch	An upgrade or change to a deployment failed	1 hour	Some upgrade failed to run properly.	-	-	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-deployment-replica-mismatch	Not all the pods in a replica are running for more than two minutes	1 hour	This alert is triggered when the pods are crashing or are unable to be deployed as part of a deployment. If too many application pods in a replica are not running, then this can cause a service impact.	k8s-pod-not-ready k8s-pod-pending k8s-pod-crashing-loop	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-ss-mismatch	Not all the pods in a stateful set are running for more than two minutes	1 hour	This alert is triggered when the pods are crashing or are unable to be deployed as part of a stateful set. If too many application pods in a stateful set are not running, then this can cause a service impact.	k8s-pod-not-ready k8s-pod-pending k8s-pod-crashing-loop	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
k8s-ss-generation-mismatch	An upgrade or change to a stateful set failed	1 hour	Some upgrade failed to run properly.	-	-	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-ss-update-not-rolled-out	The upgrade of a stateful set is stuck	1 hour	Some upgrade failed to run properly.	-	-	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-daemonset-rollout-stuck	Not all the pods in a daemonset are running for more than five minutes	1 hour	Some daemonset pods are having issues to run on nodes. Daemon pods are critical for different functionalities. If they are not running, then there can be a short or long term impact to the service.	k8s-pod-not-ready k8s-pod-pending k8s-pod-crashing-loop	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-daemonset-not-scheduled	Not all the pods in a daemonset are scheduled for more than 5 minutes	1 hour	This is a rare case where daemon sets can not be scheduled, daemonsets have toleration for almost all taints and hence, the pods should get scheduled. This alert points to an issue with K8s or node resources.	k8s-pod-pending	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
k8s-daemonset-mischeduled	Daemonset pods are running where these are not supposed to run	1 hour	This alert is a rare case and can be triggered during some upgrade. This alert indicates a larger K8s issue.	-	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
pod-oom-killed	Pod got killed and restarted due to OOM in the last five minutes	Always	The pod is crossing the memory limit. This is an application issue and indicates a misconfiguration or a memory leak issue. Depending on the functionality, there can be a service impact.	-	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	Contact your Cisco Account representative to troubleshoot.
container-memory-usage-high	Pod crossed the 80% memory limit in the last five minutes	Always	The pod is crossing the memory limit. This is an application issue and indicates a misconfiguration or a memory leak issue. Depending on the functionality, there can be a service impact.	-	Verify the CEE alerts to investigate the reason for the pods not getting deployed or running.	Contact your Cisco Account representative to troubleshoot.
pod-not-ready-but-all-containers-ready	Pod is not ready, but all its containers are ready	1 hour	This is a K8s issue and is caused when the pod is not correctly marked as running. Self-healing restarts the pod. If the alert stays on for too long, then it means either the self-healing is not working or there is a bigger issue.	-	Verify the CEE alerts and investigate the pod with the issue.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
vm-deployed	UPF VM deployed	None	This is a notification that the VM is deployed	-	-	-
vm-alive	UPF VM running	None	This is a notification that the VM is running	-	-	-

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
vm-error	UPF VM is in error state	1 hour	The VM is in error state	-	Verify the CEE alerts to investigate the KVM VMs.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
vm-recovering	UPF VM is recovering	None	This alert is a notification that the VM is recovering.	-	-	-
vm-recovery-failed	UPF VM recovering failed	1 hour	The VM failed to recover.	-	Verify the Ops Center to assess the reason for the failure.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
ops-system-sync-running	The Ops Center is running a sync operation	None	This alert is triggered during an upgrade. The Ops Center is trying to either upgrade the application or apply recently changed configuration.	-	-	-
ops-latest-sync-failed	The Ops Center sync operation failed	1 hour	The Ops Center sync operation failed to complete. This alert indicates an issue with either new release or with the latest configuration.	-	Verify the Ops Center to assess the reason for the failure.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
helm-deploy-failure	An application helm deployment failed	1 hour	This alerts is triggered when one of the applications fails to deploy the helm chart, usually due to misconfiguration or release issues.	-	Verify the Ops Center to assess the reason for the failure.	If the alert stays active, then contact your Cisco Account representative to troubleshoot.
node-disk-running-full-24hours	The node disk is estimated to run out space in less than 24 hours	Always	The node is projected to run out of space in one of the partitions in less than 24 hours.	-	Verify the CEE alert and escalate immediately.	-

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
node-disk-running-full-2hours	The node disk is estimated to run out space in less than two hours	Always	The node is projected to run out of space in one of the partitions in less than 2 hours. If the K8s partition is affected, this alert might severely impact the services on the node.	-	Verify the CEE alert and escalate immediately.	-
k8s-persistent-volume-usage	The persistent volume has less than 3% free space	Always	The volume is almost full or full and most likely is impacting the application.	-	Verify the CEE alert and escalate immediately.	-
k8s-persistent-volume-usage-projected-full	The persistent volume will reach 100% usage in less than four days and is currently at 85% or more	Always	The volume is projected to get full in less than four days, potentially impacting the application.	-	Verify the CEE alert and escalate immediately.	-
kube_certificate_expiring	Alert that the k8s certificates are about to expire	Always	The K8s certificates have to be renewed every year. With the automated process, this alert is not triggered.	-	Verify the CEE alert and escalate immediately.	-
kubelet-too-many-pods	Too many pods attempted to be deployed on one node	Always	The applications are designed to not exceed the max pod limit. If the alert is seen, then it indicates a misconfiguration.	-	Verify the CEE alert and escalate immediately.	-
clock-skew-detected	Clock skew detected on a node	Always	The node is having a wrong NTP configuration or the NTP servers have issues. This issue might be seen on new systems.	-	Verify the CEE alert and escalate immediately.	-

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
clock-is-not-in-synch	Clock is not in synch with NTP for last five minutes	Always	The node is not able to get the clock in sync with NTP. Either the NTP configuration is incorrect or there is some network issue.	-	Verify the CEE alert and escalate immediately.	-
network-receive-errors	Specific network is seeing receive errors in the last two minutes	Always	Networking issue with received packets.	-	Verify the CEE alert and escalate immediately.	-
network-transmit-errors	Specific network is seeing transmit errors in the last two minutes	Always	Networking issue with sent packets.	-	Verify the CEE alert and escalate immediately.	-
network-interface-flapping	Specific network up/down status is changing in the last two minutes	Always	Networking issue with specific interface.	-	Verify the CEE alert and escalate immediately.	-
k8s-cpu-overcommit	The CPU is overcommitted compared to quota on namespaces	Always	This alert indicates a deployment misconfiguration.	-	Verify the CEE alert and escalate immediately.	-
k8s-mem-overcommit	The memory is overcommitted compared to quota on namespaces	Always	This alert indicates a deployment misconfiguration.	-	Verify the CEE alert and escalate immediately.	-
k8s-quota-exceeded	A namespace is using more than 90% of its allocated CPU/memory quota	Always	This alert indicates a deployment misconfiguration.	-	Verify the CEE alert and escalate immediately.	-

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
cpu-throttling-high	K8s is throttling the pod CPU for more than 25% of the time in last five minutes	1 hour	Application is running too hot and CPU throttling is too high. If the application pods are affected, then this can result in a service impact.	-	Verify the CEE alert and escalate immediately.	-
cndp-ha-switchover	Cluster Manager switchover from primary to standby	1 hour	This alert indicates a CM failover either caused by MW, RMA or a hardware issue.	-	Verify the CEE alert and escalate immediately.	-
backup-node-down	Cluster Manager Backup node is not reachable from primary	1 hour	The backup CM is not reachable from primary either because of network issues or a hardware issue, RMA, or MW.	-	Verify the CEE alert and escalate immediately.	-
user_password_expiring	User password will expire in less than <configured number of days>	Always	User password will expire and must be updated	-	Verify the CEE alert and escalate immediately.	-
user_password_expired	User password expired	Always	User password expired	-	Verify the CEE alert and escalate immediately.	-
k8s-persistent-volume-errors	Persistent volume has issues	Always	Persistent volume has issues and can impact the application using it. Most likely, there is a hardware issue or it can be caused by a failed install or upgrade.	-	Verify the CEE alert and escalate immediately.	-
k8s-version-mismatch	The system has K8s components with different versions	Always	The components should always run on the same K8s version. This alert can be triggered by a failed upgrade.	-	Verify the CEE alert and escalate immediately.	-

Alarm Name	Description	Active Alert Duration *	System Impact	Associated Alerts	Validation	Resolution
k8s-client-errors	Specific K8s API server is having issues communicating with the API server	Always	This alert is triggered by either a failed upgrade, connectivity or incorrect/expired certificates. It indicates a major impact to applications or service.	-	Verify the CEE alert and escalate immediately.	-
prometheus	Various errors related to Prometheus	1 hour	This alert indicates issues with Prometheus or monitoring and must be investigated.	-	Verify the CEE alert and escalate immediately.	-
node-disk-running-Low-24hours	The node disk partition is > 75%, and will be >80 in less than 24 hours	Always	Node disk partition is projected to be 80% full within the next 24 hours.	-	Verify the CEE alert and escalate immediately.	-
node-disk-running-Low-2hours	The node disk partition is > 75%, and will be >80% in less than 2 hours	Always	Node disk partition is projected to be 80% full within the next 2 hours.	-	Verify the CEE alert and escalate immediately.	-

* Escalate to investigate when the alert is active for longer than the specified time period.

Alert Details

cndp-ha

Rules:

- **Alert:** cndp-ha-switchover
- **Annotations:**
 - **Type:** Switching Over To Primary
 - **Summary:** "CNDP-HA is switched {{ \$labels.hostname }} over to primary."

- **Expression:**

```
ha_is_failed_over == 1
```

- **For:** 1m
- **Labels:**
 - **Severity:** major
- **Alert:** backup-node-down
 - **Annotations:**
 - **Type:** Backup node down
 - **Summary:** "The Backup CM node of {{ \$labels.hostname }} is down."
 - **Expression:**

```
|
      backup_node_status == 0
```
- **For:** 1m
- **Labels:**
 - **Severity:** major

kubernetes-apps

Rules:

- **Alert:** pod-oom-killed
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** 'Pod {{ \$labels.namespace }}/{{ \$labels.pod }} got OOM Killed.'
 - **Expression:**

```
|
      sum_over_time(kube_pod_container_status_terminated_reason{reason="OOMKilled"}[5m])
      > 0
```
 - **For:** 1m
 - **Labels:**
 - **Severity:** critical
- **Alert:** container-memory-usage-high
 - **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** 'Pod {{ \$labels.namespace }}/{{ \$labels.pod }}/{{ \$labels.name }} uses high memory {{ printf "%.2f" \$value }}%'.

- **Expression:**

```
|
    ((container_memory_usage_bytes{pod!="",container!="POD",image!=""} -
    container_memory_cache{pod!="",container!="POD",image!=""}) /
    (container_spec_memory_limit_bytes{pod!="",container!="POD",image!=""} != 0)) * 100
    > 80
```

- **For:** 2m

- **Labels:**

- **Severity:** critical

- **Alert:** pod-not-ready-but-all-containers-ready

- **Expression:**

```
>
    (count by (namespace, pod) (kube_pod_status_ready{condition="true"} == 0))
    and
    (
        (count by (namespace, pod) (kube_pod_container_status_ready==1))
        unless
        (count by (namespace, pod) (kube_pod_container_status_ready==0))
    )
```

- **For:** 5m

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Pod {{ \$labels.namespace }}/{{ \$labels.pod }} ({{ \$labels.container }}) is restarting {{ printf "%.2f" \$value }} times / 5 minutes.

- **Expression:**

```
|
    rate(kube_pod_container_status_restarts_total[5m]) * 60 * 5 > 0
```

- **For:** 1m

- **Labels:**

- **Severity:** minor

- **Alert:** k8s-pod-crashing-loop
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Pod {{ \$labels.namespace }}/{{ \$labels.pod }} ({{ \$labels.container }}) is restarting {{ printf "%.2f" \$value }} times / 5 minutes.
 - **Expression:**

```
|  
rate(kube_pod_container_status_restarts_total[5m]) * 60 * 5 >= 2
```
 - **For:** 1m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-pod-pending
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Pod {{ \$labels.namespace }}/{{ \$labels.pod }} has been in pending state state for longer than 1 minute.
 - **Expression:**

```
|  
sum by (namespace, pod) (kube_pod_status_phase{  
phase=~"Failed|Pending|Unknown"}) > 0
```
 - **For:** 1m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-pod-not-ready
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Pod {{ \$labels.namespace }}/{{ \$labels.pod }} has been in a non-ready state for longer than 1 minute.
 - **Expression:**

```
|
```

```
sum by (namespace, pod) (kube_pod_status_ready(condition="false")) > 0
```

- **For:** 1m

- **Labels:**

- **Severity:** critical

- **Alert:** k8s-deployment-generation-mismatch

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Deployment generation for {{ \$labels.namespace }}/{{ \$labels.deployment }} does not match, this indicates that the Deployment has failed but has not been rolled back.

- **Expression:**

```
|
    kube_deployment_status_observed_generation
    !=
    kube_deployment_metadata_generation
```

- **For:** 5m

- **Labels:**

- **Severity:** critical

- **Alert:** k8s-deployment-replica-mismatch

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Deployment {{ \$labels.namespace }}/{{ \$labels.deployment }} has not matched the expected number of replicas for longer than 2 minutes.

- **Expression:**

```
|
    kube_deployment_spec_replicas
    !=
    kube_deployment_status_replicas_available
```

- **For:** 2m

- **Labels:**

- **Severity:** critical

- **Alert:** k8s-ss-mismatch
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** StatefulSet {{ \$labels.namespace }}/{{ \$labels.statefulset }} has not matched the expected number of replicas for longer than 5 minutes.
 - **Expression:**

```
|
    kube_statefulset_status_replicas_ready
    !=
    kube_statefulset_status_replicas
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-ss-generation-mismatch
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** StatefulSet generation for {{ \$labels.namespace }}/{{ \$labels.statefulset }} does not match, this indicates that the StatefulSet has failed but has not been rolled back.
 - **Expression:**

```
|
    kube_statefulset_status_observed_generation
    !=
    kube_statefulset_metadata_generation
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-ss-update-not-rolled-out
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** StatefulSet {{ \$labels.namespace }}/{{ \$labels.statefulset }} update has not been rolled out.

- **Expression:**

```
|
    max without (revision) (
        kube_statefulset_status_current_revision
        unless
        kube_statefulset_status_update_revision
    )
    *
    (
        kube_statefulset_replicas
        !=
        kube_statefulset_status_replicas_updated
    )
```

- **For:** 5m

- **Labels:**

- **Severity:** critical

- **Alert:** k8s-daemonset-rollout-stuck

- **Annotations:**

- **Type:** Processing Error Alarm
- **Summary:** Only {{ \$value }}% of the desired Pods of DaemonSet {{ \$labels.namespace }}/{{ \$labels.daemonset }} are scheduled and ready.

- **Expression:**

```
|
    kube_daemonset_status_number_ready
    /
    kube_daemonset_status_desired_number_scheduled * 100 < 100
```

- **For:** 5m

- **Labels:**

- **Severity:** critical

- **Alert:** k8s-daemonset-not-scheduled

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** '{{ \$value }} Pods of DaemonSet {{ \$labels.namespace }}/{{ \$labels.daemonset }} are not scheduled.'

- **Expression:**

```
|
    kube_daemonset_status_desired_number_scheduled
    kube_daemonset_status_current_number_scheduled > 0
```

- **For:** 5m

- **Labels:**

- **Severity:** major

- **Alert:** k8s-daemonset-mischeduled

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** '{{ \$value }} Pods of DaemonSet {{ \$labels.namespace }}/{{ \$labels.daemonset }} are running where they are not supposed to run.'

- **Expression:**

```
|
    kube_daemonset_status_number_mischeduled > 0
```

- **For:** 5m

- **Labels:**

- **Severity:** major

- **Alert:** k8s-cronjob-running

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** CronJob {{ \$labels.namespace }}/{{ \$labels.cronjob }} is taking more than 1h to complete.

- **Expression:**

```
|
    time() - kube_cronjob_next_schedule_time > 3600
```

- **For:** 1h

- **Labels:**

- **Severity:** major

- **Alert:** k8s-job-completion
 - **Annotations:**
 - **Type:** Processing Error Alarm

 - **Summary:** Job `{{ $labels.namespace }}`/`{{ $labels.job_name }}` is taking more than one hour to complete.

 - **Expression:**

```
|  
    kube_job_spec_completions - kube_job_status_succeeded > 0
```

 - **For:** 1h

 - **Labels:**
 - **Severity:** major

- **Alert:** k8s-job-failed
 - **Annotations:**
 - **Type:** Processing Error Alarm

 - **Summary:** Job `{{ $labels.namespace }}`/`{{ $labels.job_name }}` failed to complete.

 - **Expression:**

```
|  
    kube_job_status_failed > 0
```

 - **For:** 1h

 - **Labels:**
 - **Severity:** major

- **Alert:** k8s-pod-cpu-usage-high
 - **Annotations:**
 - **Type:** Processing Error Alarm

 - **Summary:** '`{{ $labels.namespace }}`.`{{ $labels.pod }}` pod cpu usage is above 80%.'

 - **Expression:**

```
|
```

```

sum(rate(container_cpu_usage_seconds_total{container!="POD", pod!="",
image!=""}[5m])) by (namespace, pod) * 100 /
sum(kube_pod_container_resource_limits_cpu_cores) by (namespace, pod) > 80

```

- **For:** 1m
- **Labels:**
 - **Severity:** major

kubernetes-resources

Rules:

- **Alert:** k8s-cpu-overcommit
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Cluster has overcommitted CPU resource requests for Namespaces.
 - **Expression:**

```

|
sum(kube_resourcequota{ type="hard", resource="cpu"})
/
sum(kube_node_status_allocatable_cpu_cores)
> 1.5

```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major
- **Alert:** k8s-mem-overcommit
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Cluster has overcommitted memory resource requests for Namespaces.
 - **Expression:**

```

|
sum(kube_resourcequota{ type="hard", resource="memory"})
/
sum(kube_node_status_allocatable_memory_bytes)
> 1.5

```

- **For:** 2m

- **Labels:**

- **Severity:** major

- **Alert:** k8s-quota-exceeded

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Namespace {{ \$labels.namespace }} is using {{ printf "%0.0f" \$value }}% of its {{ \$labels.resource }} quota.

- **Expression:**

```
|
  100 * kube_resourcequota{ type="used" }
    / ignoring(instance, job, type)
    (kube_resourcequota{ type="hard" } > 0)
    > 90
```

- **For:** 2m

- **Labels:**

- **Severity:** major

- **Alert:** cpu-throttling-high

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** '{{ printf "%0.0f" \$value }}% throttling of CPU in namespace {{ \$labels.namespace }} for container {{ \$labels.container }} in pod {{ \$labels.pod }}.'

- **Expression:**

```
"100 * sum(increase(container_cpu_cfs_throttled_periods_total{container!="\"},
  [5m])) by (container, pod, namespace)\n
/\nsum(increase(container_cpu_cfs_periods_total{} [5m]))
  by (container, pod, namespace)\n > 25 \n"
```

- **For:** 2m

- **Labels:**

- **Severity:** major

kubernetes-storage**Rules:**

- **Alert:** k8s-persistent-volume-usage
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** The PersistentVolume claimed by `{{ $labels.persistentvolumeclaim }}` in Namespace `{{ $labels.namespace }}` is only `{{ printf "%0.2f" $value }}`% free.
 - **Expression:**

```
|
    100 * kubelet_volume_stats_available_bytes
      /
    kubelet_volume_stats_capacity_bytes
    < 3
```
 - **Labels:**
 - **Severity:** critical
- **Alert:** k8s-persistent-volume-usage-projected-full
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Based on recent sampling, the PersistentVolume claimed by `{{ $labels.persistentvolumeclaim }}` in Namespace `{{ $labels.namespace }}` is expected to fill up within four days. Currently `{{ printf "%0.2f" $value }}`% is available.
 - **Expression:**

```
|
    100 * (
      kubelet_volume_stats_available_bytes
      /
      kubelet_volume_stats_capacity_bytes
    ) < 15
    and
    predict_linear(kubelet_volume_stats_available_bytes[6h], 4 * 24 * 3600) < 0
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-persistent-volume-errors
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** The persistent volume {{ \$labels.persistentvolume }} has status {{ \$labels.phase }}.
 - **Expression:**

```
|
kube_persistentvolume_status_phase{phase=~"Failed|Pending",namespace=~"(kube-.*|default|logging)"}
> 0
```
 - **Labels:**
 - **Severity:** critical

kubernetes-system

Rules:

- **Alert:** k8s-node-not-ready
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** '{{ \$labels.node }}' has been unready for more than 1 minutes.'
 - **Expression:**

```
|
kube_node_status_condition{condition="Ready",status="true"} == 0
```
 - **For:** 1m
 - **Labels:**
 - **Severity:** critical
- **Alert:** k8s-node-status-change
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** '{{ \$labels.node }}' status was changed in the past 5 minutes.'
 - **Expression:**

```
|
```

```
changes(kube_node_status_condition{condition="Ready",status="true"}[5m]) > 0
```

- **For:** 0m
- **Labels:**
 - **Severity:** major

- **Alert:** k8s-version-mismatch

- **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** There are {{ \$value }} different semantic versions of Kubernetes components running.

- **Expression:**

```
|
    count(count by (gitVersion)
(label_replace(kubernetes_build_info,"gitVersion","$1","gitVersion","(v[0-9]*.[0-9]*.[0-9]*).*"))
> 1
```

- **For:** 5m
- **Labels:**
 - **Severity:** major

- **Alert:** k8s-client-errors

- **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Kubernetes API server client '{{ \$labels.job }}/{{ \$labels.instance }}' is experiencing {{ printf "%0.0f" \$value }}% errors.'

- **Expression:**

```
|
    (sum(rate(rest_client_requests_total{code=~"5.."}[5m])) by (instance, job)
/
sum(rate(rest_client_requests_total[5m])) by (instance, job))
* 100 > 1
```

- **For:** 2m
- **Labels:**
 - **Severity:** major

- **Alert:** kubelet-too-many-pods
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Kubelet {{ \$labels.instance }} is running {{ \$value }} Pods, close to the limit of 110.
 - **Expression:**

```
|
    kubelet_running_pod_count > 110 * 0.9
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** critical

- **Alert:** k8s-client-cert-expiration
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** A client certificate used to authenticate to the apiserver is expiring in less than 30 days
 - **Expression:**

```
|
    apiserver_client_certificate_expiration_seconds_count > 0 and
    histogram_quantile(0.01, sum by (job, le)
    (rate(apiserver_client_certificate_expiration_seconds_bucket[5m]))) < 2592000
```
 - **Labels:**
 - **Severity:** warning

- **Alert:** k8s-client-cert-expiration
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** A client certificate used to authenticate to the apiserver is expiring in less than 24.0 hours.
 - **Expression:**

```
|
```



```
apiserver_client_certificate_expiration_seconds_count > 0 and
  histogram_quantile(0.01, sum by (job, le)
    (rate(apiserver_client_certificate_expiration_seconds_bucket[5m]))) < 86400
```

- **Labels:**
 - **Severity:** critical

general.rules

Rules:

- **Alert:** watchdog
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** | This is an alert meant to ensure that the entire alerting pipeline is functional. This alert is always firing, therefore it should always be firing in Alertmanager and always fire against a receiver. There are integrations with various notification mechanisms that send a notification when this alert is not firing.
 - **Expression:**

```
vector(1)
```
 - **Labels:**
 - **Severity:** minor

sync.rules

Rules:

- **Alert:** ops-system-sync-running
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** | ops center system upgrade for {{ \$labels.namespace }} is in progress
 - **Expression:**

```
system_ops_upgrade_running > 0
```
 - **Labels:**
 - **Severity:** minor
- **Alert:** ops-latest-sync-failed
 - **Annotations:**

- **Type:** Communications Alarm
- **Summary:** | ops center latest system sync for {{ \$labels.namespace }} failed

- **Expression:**

```
system_synch_error > 0
```

- **Labels:**

- **Severity:** major

kube-prometheus-node-alerting.rules

Rules:

- **Alert:** node-disk-running-full-24hours
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Device {{ \$labels.device }} of node-exporter {{ \$labels.namespace }}/{{ \$labels.pod }} will be full within the next 24 hours.
 - **Expression:**

```
|
(node:node_filesystem_usage: > 0.85) and
(predict_linear(node:node_filesystem_avail:[6h], 3600 * 24) < 0)
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** major
- **Alert:** node-disk-running-full-2hours
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Device {{ \$labels.device }} of node-exporter {{ \$labels.namespace }}/{{ \$labels.pod }} will be full within the next 2 hours.
 - **Expression:**

```
|
(node:node_filesystem_usage: > 0.85) and
(predict_linear(node:node_filesystem_avail:[30m], 3600 * 2) < 0)
```
 - **Labels:**

- **Severity:** critical

node-time

Rules:

- **Alert:** clock-skew-detected
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Clock skew detected on hostname {{ \$labels.hostname }} . Ensure NTP is configured correctly on this host.
 - **Expression:**

```
|
    abs(node_timex_offset_seconds) > 0.03
```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major
- **Alert:** clock-is-not-in-synch
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Clock not in synch detected on hostname {{ \$labels.hostname }} . Ensure NTP is configured correctly on this host.
 - **Expression:**

```
|
    min_over_time(node_timex_sync_status[5m]) == 0
    and
    node_timex_maxerror_seconds >= 16
```
 - **For:** 10m
 - **Labels:**
 - **Severity:** major

node-network**Rules:**

- **Alert:** network-receive-errors
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** Network interface "{{ \$labels.device }}" showing receive errors on hostname {{ \$labels.hostname }}"
 - **Expression:**

```
|
    rate(node_network_receive_errs_total{device!~"veth.+"}[2m]) > 0
```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major

- **Alert:** network-transmit-errors
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** Network interface "{{ \$labels.device }}" showing transmit errors on hostname {{ \$labels.hostname }}"
 - **Expression:**

```
|
    rate(node_network_transmit_errs_total{device!~"veth.+"}[2m]) > 0
```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major

- **Alert:** network-interface-flapping
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** Network interface "{{ \$labels.device }}" changing it's up status often on hostname {{ \$labels.hostname }}"
 - **Expression:**

```
|
    changes(node_network_up{device!~"veth.+"}[2m]) > 2
```

- **For:** 2m
- **Labels:**
 - **Severity:** major

- **Alert:** kvm-tunnels flapping

- **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** Pod {{ \$labels.namespace }}/{{ \$labels.pod }} tunnel to ({{ \$labels.ip }}:{{ \$labels.port }}) is flapping {{ printf "%.2f" \$value }} times / 5 minutes.

- **Expression:**

```
|
    changes(kvm_metrics_tunnels_up[5m]) > 2
```

- **For:** 5m
- **Labels:**
 - **Severity:** major

- **Alert:** kvm-node-not-ready

- **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** KVM node {{ \$labels.hostname }} ({{ \$labels.ip }}) is not reachable.

- **Expression:**

```
|
    changes(kvm_metrics_tunnels_up[2m]) > 0
```

- **For:** 0m
- **Labels:**
 - **Severity:** major

fluentbit.rules

Rules:

- **Alert:** fluent-proxy-output-retries-failed
 - **Annotations:**
 - **Type:** Communications Alarm
 - **Summary:** 'Fluent-proxy {{ \$labels.namespace }}/{{ \$labels.pod }} output retries failed for target: {{ \$labels.name }}'
 - **Expression:**

```
|
    rate(fluentbit_output_retries_failed_total{pod=~"fluent-proxy.*"}[5m]) > 0
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** major

prometheus.rules

Rules:

- **Alert:** prometheus-config-reload-failed
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Reloading Prometheus' configuration failed
Reloading Prometheus' configuration has failed for {{ \$labels.namespace }}/{{ \$labels.pod }}
 - **Expression:**

```
|
    prometheus_config_last_reload_successful == 0
```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major
- **Alert:** prometheus-notification-q-running-full
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Prometheus' alert notification queue is running full

Prometheus' alert notification queue is running full for `{{ $labels.namespace }}`/`{{ $labels.pod }}`

- **Expression:**

```
|
    predict_linear(prometheus_notifications_queue_length[5m], 60 * 30) >
prometheus_notifications_queue_capacity
```

- **For:** 10m

- **Labels:**

- **Severity:** major

- **Alert:** prometheus-error-sending-alerts

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Errors while sending alert from Prometheus

Errors while sending alerts from Prometheus `{{ $labels.namespace }}`/`{{ $labels.pod }}` to Alertmanager `{{ $labels.Alertmanager }}`

- **Expression:**

```
|
    rate(prometheus_notifications_errors_total[5m]) /
rate(prometheus_notifications_sent_total[5m]) > 0.01
```

- **For:** 2m

- **Labels:**

- **Severity:** major

- **Alert:** prometheus-error-sending-alerts

- **Annotations:**

- **Type:** Processing Error Alarm

- **Summary:** Errors while sending alerts from Prometheus

Errors while sending alerts from Prometheus `{{ $labels.namespace }}`/`{{ $labels.pod }}` to Alertmanager `{{ $labels.Alertmanager }}`

- **Expression:**

```
|
    rate(prometheus_notifications_errors_total[5m]) /
rate(prometheus_notifications_sent_total[5m]) > 0.03
```

- **Labels:**
 - **Severity:** critical

- **Alert:** prometheus-not-connected-to-alertmanagers
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Prometheus is not connected to any Alertmanagers
Prometheus {{ \$labels.namespace }}/{{ \$labels.pod }} is not connected to any Alertmanagers
 - **Expression:**

```
|
    prometheus_notifications_alertmanagers_discovered < 1
```
 - **For:** 2m
 - **Labels:**
 - **Severity:** major

- **Alert:** prometheus-tsdbs-reloads-failing
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Prometheus has issues reloading data blocks from disk
'{{ \$labels.job }} at {{ \$labels.instance }} had {{ \$value | humanize }} reload failures over the last four hours.'
 - **Expression:**

```
|
    increase(prometheus_tsdbs_reloads_failures_total[2h]) > 0
```
 - **For:** 5m
 - **Labels:**
 - **Severity:** major

- **Alert:** prometheus-tsdbs-compactions-failing
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** Prometheus has issues compacting sample blocks

'{{ \$labels.job }}' at '{{ \$labels.instance }}' had '{{ \$value | humanize }}' compaction failures over the last four hours.'

- **Expression:**

```
|  
    increase(prometheus_tsdb_compactions_failed_total[2h]) > 0
```

- **For:** 5m

- **Labels:**

- **Severity:** major

- **Alert:** prometheus-tsdb-wal-corruptions

- **Annotations:**

- **Type:** Processing Error Alarm
- **Summary:** Prometheus write-ahead log is corrupted

'{{ \$labels.job }}' at '{{ \$labels.instance }}' has a corrupted write-ahead log (WAL).'

- **Expression:**

```
|  
    prometheus_tsdb_wal_corruptions_total > 0
```

- **For:** 5m

- **Labels:**

- **Severity:** major

- **Alert:** prometheus-not-ingesting-samples

- **Annotations:**

- **Type:** Processing Error Alarm
- **Summary:** Prometheus isn't ingesting samples

Prometheus {{ \$labels.namespace }}/{{ \$labels.pod }} isn't ingesting samples.

- **Expression:**

```
|  
    rate(prometheus_tsdb_head_samples_appended_total[5m]) <= 0
```

- **For:** 5m

- **Labels:**

- **Severity:** major

- **Alert:** prometheus-target-scrapes-duplicate

- **Annotations:**

- **Type:** Processing Error Alarm
 - **Summary:** Prometheus has many samples rejected

'{{ \$labels.namespace }}/{{ \$labels.pod }} has many samples rejected due to duplicate timestamps but different values'

- **Expression:**

```
|
    increase(prometheus_target_scrapes_sample_duplicate_timestamp_total[5m]) > 0
```

- **For:** 10m

- **Labels:**

- **Severity:** warning

- **Alert:** prometheus-remote-write-behind

- **Annotations:**

- **Type:** Processing Error Alarm
 - **Summary:** Prometheus remote write is behind

'Prometheus {{ \$labels.namespace }}/{{ \$labels.pod }} remote write is {{ \$value | humanize }} seconds behind for target: {{ \$labels.url }}.'

- **Expression:**

```
|
    (
        max_over_time(prometheus_remote_storage_highest_timestamp_in_seconds[5m])
    ignoring(remote_name, url) group_right
    max_over_time(prometheus_remote_storage_queue_highest_sent_timestamp_seconds[5m])
    )
    > 120
```

- **For:** 15m

- **Labels:**

- **Severity:** major

- **Alert:** ssl-earliest-cert-expiry
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** SSL certificate expires in 30 days
'{{ \$labels.namespace }}/{{ \$labels.pod }} ssl certificate expires in 30 days'
 - **Expression:**

```
|  
    probe_ssl_earliest_cert_expiry - time() < 86400 * 30
```
 - **Labels:**
 - **Severity:** major

- **Alert:** ssl-earliest-cert-expiry
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** SSL certificate expires in 7 days
'{{ \$labels.namespace }}/{{ \$labels.pod }} ssl certificate expires in 7 days'
 - **Expression:**

```
|  
    probe_ssl_earliest_cert_expiry - time() < 86400 * 7
```
 - **Labels:**
 - **Severity:** critical

- **Alert:** helm-deploy-failure
 - **Annotations:**
 - **Type:** Processing Error Alarm
 - **Summary:** 'Helm chart failed to deploy for 5 minutes'
'Helm chart {{\$labels.chart}}/{{\$labels.namespace }} deployment failed'
 - **Expression:**

```
|  
    helm_chart_deploy_success < 1
```
 - **Labels:**

- **Severity:** critical

- **For:** 5m

server

Rules:

- **Alert:** server-alert
- **Annotations:**
- **Type:** Equipment Alarm
- **dn:** "{{ \$labels.cluster }}/{{ \$labels.server }}/{{ \$labels.fault_id }}/{{ \$labels.id }}"
- **Summary:** "{{ \$labels.description }}"
- **Expression:**

```
|
sum(server_alert) by ( id, description, fault_id, server, cluster, severity, affectedDn)
== 1
```

- **For:** 1m
- **Alert:** server-not-reachable-alert
 - **Annotations:**
 - **Type:** Equipment Alarm
 - **Summary:** "{{ \$labels.description }}"

- **Expression:**

```
|
sum(cimc_server_not_reachable_alert) by (server, cluster, description) == 1
```

- **For:** 1m
- **Labels:**
- **Severity:** critical

k8s.rules

Rules:

- **Expression:**

```
sum(rate(container_cpu_usage_seconds_total{ image!="", container!="POD"}[5m]))
by (namespace)
```

- **Record:** namespace:container_cpu_usage_seconds_total:sum_rate

- **Expression:**

```
sum by (namespace, pod, container) (
  rate(container_cpu_usage_seconds_total{ image!="", container!="POD"}[5m])
)
```

- **Record:** namespace_pod_container:container_cpu_usage_seconds_total:sum_rate

- **Expression:**

```
sum(container_memory_usage_bytes{image!="", container!="POD"} -
container_memory_cache{image!="", container!="POD"}) by (namespace)
```

- **Record:** namespace:container_memory_usage_bytes:sum

- **Expression:**

```
sum(
  label_replace(
    label_replace(
      kube_pod_owner{ owner_kind="ReplicaSet"},
      "replicaset", "$1", "owner_name", "(.*)"
    ) * on(replicaset, namespace) group_left(owner_name) kube_replicaset_owner,
    "workload", "$1", "owner_name", "(.*)"
  )
) by (namespace, workload, pod)
```

- **Labels:**

- **workload_type:** deployment

- **Record:** mixin_pod_workload

- **Expression:**

```
sum(
  label_replace(
    kube_pod_owner{ owner_kind="DaemonSet"},
    "workload", "$1", "owner_name", "(.*)"
  )
) by (namespace, workload, pod)
```

- **Labels:**

- **workload_type:** daemonset

- **Record:** mixin_pod_workload

- **Expression:**

```
sum(
  label_replace(
    kube_pod_owner{ owner_kind="StatefulSet"},
    "workload", "$1", "owner_name", "(.*)")
  )
) by (namespace, workload, pod)
```

- **Labels:**

- **workload_type:** statefulset

- **Record:** mixin_pod_workload

node.rules

Rules:

- **Expression:**

```
max(label_replace(kube_pod_info, "pod", "$1", "pod", "(.*)")) by (node, namespace, pod)
```

- **Record:** 'node_namespace_pod:kube_pod_info:'

- **Expression:**

```
1 - avg(rate(node_cpu_seconds_total{mode="idle"}[1m]))
```

- **Record:** :node_cpu_utilisation:avg1m

- **Expression:**

```
1 -
  sum(node_memory_MemFree_bytes + node_memory_Cached_bytes +
    node_memory_Buffers_bytes)
  /
  sum(node_memory_MemTotal_bytes)
```

- **Record:** ':node_memory_utilisation:'

- **Expression:**

```
sum(node_memory_MemFree_bytes + node_memory_Cached_bytes +
  node_memory_Buffers_bytes)
```

- **Record:** :node_memory_MemFreeCachedBuffers_bytes:sum

- **Expression:**

```
sum(node_memory_MemTotal_bytes)
```

- **Record:** :node_memory_MemTotal_bytes:sum

- **Expression:**

```
avg(irate(node_disk_io_time_seconds_total{device=~"nvme.+|rbd.+|sd.+|vd.+|xvd.+|dm-.+"} [1m]))
```

- **Record:** :node_disk_utilisation:avg_irate

- **Expression:**

```
avg(irate(node_disk_io_time_weighted_seconds_total{device=~"nvme.+|rbd.+|sd.+|vd.+|xvd.+|dm-.+"} [1m]))
```

- **Record:** :node_disk_saturation:avg_irate

- **Expression:**

```
max by (instance, namespace, pod, device)
((node_filesystem_size_bytes{fstype=~"ext[234]|btrfs|xfs|zfs"}
- node_filesystem_avail_bytes{fstype=~"ext[234]|btrfs|xfs|zfs"})
/ node_filesystem_size_bytes{fstype=~"ext[234]|btrfs|xfs|zfs"})
```

- **Record:** 'node:node_filesystem_usage:'

- **Expression:**

```
max by (instance, namespace, pod, device)
(node_filesystem_avail_bytes{fstype=~"ext[234]|btrfs|xfs|zfs"} /
node_filesystem_size_bytes{fstype=~"ext[234]|btrfs|xfs|zfs"})
```

- **Record:** 'node:node_filesystem_avail:'

- **Expression:**

```
sum(irate(node_network_receive_bytes_total{device!~"veth.+"} [1m])) +
sum(irate(node_network_transmit_bytes_total{device!~"veth.+"} [1m]))
```

- **Record:** :node_net_utilisation:sum_irate

- **Expression:**

```
sum(irate(node_network_receive_drop_total{device!~"veth.+"} [1m])) +
sum(irate(node_network_transmit_drop_total{device!~"veth.+"} [1m]))
```

- **Record:** :node_net_saturation:sum_irate

- **Expression:**

```
max (
  max (
    kube_pod_info{host_ip!=""}
  ) by (node, host_ip)
  * on (host_ip) group_right (node)
  label_replace(
```

```

        (max(node_filesystem_files{ mountpoint="/"}) by (instance)), "host_ip", "$1",
"instance", "(.*):.*"
    )
) by (node)

```

- **Record:** ':node:node_inodes_total!'

- **Expression:**

```

max(
    max(
        kube_pod_info{ host_ip!=""}
    ) by (node, host_ip)
    * on (host_ip) group_right (node)
    label_replace(
        (max(node_filesystem_files_free{ mountpoint="/"}) by (instance)), "host_ip",
"$1", "instance", "(.*):.*"
    )
) by (node)

```

- **Record:** ':node:node_inodes_free!'

- **Expression:**

```

-
sum by (node) (
    (node_memory_MemFree_bytes + node_memory_Cached_bytes +
node_memory_Buffers_bytes)
    * on (namespace, pod) group_left(node)
    node_namespace_pod:kube_pod_info:
)

```

- **Record:** node:node_memory_bytes_available:sum

- **Expression:**

```

-
sum by (node) (
    node_memory_MemTotal_bytes
    * on (namespace, pod) group_left(node)
    node_namespace_pod:kube_pod_info:
)

```

- **Record:** node:node_memory_bytes_total:sum

- **Expression:**

```

max without(endpoint, instance, job, pod, service) (kube_node_labels and on(node)
kube_node_role{role="control-plane"})

```


- **Labels:**

label_node_role_kubernetes_io: control-plane

Record: cluster:master_nodes

kube-prometheus-node-recording.rules

Rules:

- **Expression:**

```
sum(rate(node_cpu_seconds_total{mode!="idle",mode!="iowait"}[3m])) BY
(instance)
```

- **Record:** instance:node_cpu:rate:sum

- **Expression:**

```
sum((node_filesystem_size_bytes{mountpoint="/" } -
node_filesystem_free_bytes{mountpoint="/" }))
BY (instance)
```

- **Record:** instance:node_filesystem_usage:sum

- **Expression:**

```
sum(rate(node_network_receive_bytes_total[3m])) BY (instance)
```

- **Record:** instance:node_network_receive_bytes:rate:sum

- **Expression:**

```
sum(rate(node_network_transmit_bytes_total[3m])) BY (instance)
```

- **Record:** instance:node_network_transmit_bytes:rate:sum

- **Expression:**

```
sum(rate(node_cpu_seconds_total{mode!="idle",mode!="iowait"}[5m])) WITHOUT
(cpu, mode) / ON(instance) GROUP_LEFT() count(sum(node_cpu_seconds_total)
BY (instance, cpu)) BY (instance)
```

- **Record:** instance:node_cpu:ratio

- **Expression:**

```
sum(rate(node_cpu_seconds_total{mode!="idle",mode!="iowait"}[5m]))
```

- **Record:** cluster:node_cpu:sum_rate5m

- **Expression:**

```
cluster:node_cpu:sum_rate5m / ON(cluster) GROUP_LEFT()
count(sum(node_cpu_seconds_total)
BY (cluster, cpu)) BY (cluster)
```

- **Record:** cluster:node_cpu:ratio

kubernetes.rules

Rules:

- **Expression:**

```
sum(container_memory_usage_bytes{container!="POD",container!="",pod!=""} -
container_memory_cache{container!="POD",container!="",pod!=""})
  BY (pod, namespace)
```

- **Record:** pod:container_memory_usage_bytes:sum

- **Expression:**

```
sum(container_spec_cpu_shares{container!="POD",container!="",pod!=""})
  BY (pod, namespace)
```

- **Record:** pod:container_spec_cpu_shares:sum

- **Expression:**

```
sum(rate(container_cpu_usage_seconds_total{container!="POD",container!="",pod!=""}[5m]))
  BY (pod, namespace)
```

- **Record:** pod:container_cpu_usage:sum

- **Expression:**

```
sum(container_fs_usage_bytes{container!="POD",container!="",pod!=""})
  BY (pod, namespace)
```

- **Record:** pod:container_fs_usage_bytes:sum

- **Expression:**

```
sum(container_memory_usage_bytes{container!=""} -
container_memory_cache{container!=""}) BY (namespace)
```

- **Record:** namespace:container_memory_usage_bytes:sum

- **Expression:**

```
sum(container_spec_cpu_shares{container!=""}) BY (namespace)
```

- **Record:** namespace:container_spec_cpu_shares:sum

- **Expression:**

```
sum(rate(container_cpu_usage_seconds_total{container!="POD",container!=""}[5m]))
  BY (namespace)
```

- **Record:** namespace:container_cpu_usage:sum

- **Expression:**

```
sum(container_memory_usage_bytes{container!="POD",container!="",pod!=""} -
container_memory_cache{container!="POD",container!="",pod!=""})
BY (cluster) / sum(machine_memory_bytes) BY (cluster)
```

- **Record:** cluster:memory_usage:ratio

- **Expression:**

```
sum(container_spec_cpu_shares{container!="POD",container!="",pod!=""})
/ 1000 / sum(machine_cpu_cores)
```

- **Record:** cluster:container_spec_cpu_shares:ratio

- **Expression:**

```
sum(rate(container_cpu_usage_seconds_total{container!="POD",container!="",pod!=""} [5m]))
/ sum(machine_cpu_cores)
```

- **Record:** cluster:container_cpu_usage:ratio

- **Expression:**

```
kube_node_labels and on(node)
kube_node_spec_taint{key="node-role.kubernetes.io/master"}
```

- **Labels:**

- **label_node_role_kubernetes_io:** master

- **Record:** cluster:master_nodes

- **Expression:**

```
sum((cluster:master_nodes * on(node) group_left kube_node_status_capacity_cpu_cores)
or on(node) (kube_node_labels * on(node) group_left
kube_node_status_capacity_cpu_cores))
BY (label_beta_kubernetes_io_instance_type, label_node_role_kubernetes_io)
```

- **Record:** cluster:capacity_cpu_cores:sum

- **Expression:**

```
sum((cluster:master_nodes * on(node) group_left
kube_node_status_capacity_memory_bytes)
or on(node) (kube_node_labels * on(node) group_left
kube_node_status_capacity_memory_bytes))
BY (label_beta_kubernetes_io_instance_type, label_node_role_kubernetes_io)
```

- **Record:** cluster:capacity_memory_bytes:sum

- **Expression:**

```
sum(node:node_memory_bytes_total:sum - node:node_memory_bytes_available:sum)
```

- **Record:** cluster:memory_usage_bytes:sum

- **Expression:**

```
sum(cluster:master_nodes or on(node) kube_node_labels ) BY
(label_beta_kubernetes_io_instance_type,
label_node_role_kubernetes_io)
```

- **Record:** cluster:node_instance_type_count:sum

- **Expression:**

```
sum(etcd_object_counts) BY (instance)
```

- **Record:** instance:etcd_object_counts:sum

user-password-expiry

Rules:

- **Alert:** user_password_expiring

- **Annotations:**

- **Type:** Cluster Node User Password Expiring Alarm

- **Summary:** "{{ \$labels.user_name }}" password on host: "{{ \$labels.node_name }}" is expiring in "{{ \$labels.days_to_expire }}" days."

- **Expression:**

```
|
    User_password_expiration == 1
```

- **Labels:**

- **Severity:** critical

- **Alert:** user_password_expired

- **Annotations:**

- **Type:** Cluster Node User Password Expired Alarm

- **Summary:** "{{ \$labels.user_name }}" password on host: "{{ \$labels.node_name }}" is expired "{{ \$labels.days_to_expire }}" days ago."

- **Expression:**

```
|
    User_password_expiration == 2
```

- **Labels:**
 - **Severity:** critical

VM State Alert

Rules:

- **Alert:** vm-deployed
 - **Annotations:**
 - **Type:** Equipment Alarm
 - **Summary:** "{{ \$labels.vm_name }}" is deployed."
 - **Expression:**

```
|
    upf_state == 2
```
 - **For:** 5s
 - **Labels:**
 - **Severity:** minor
 - **State:** DEPLOYED
- **Alert:** vm-alive
 - **Annotations:**
 - **Type:** Equipment Alarm
 - **Summary:** "{{ \$labels.vm_name }}" is alive."
 - **Expression:**

```
|
    upf_state == 1 and changes(upf_state[2m]) > 0
```
 - **For:** 5s
 - **Labels:**
 - **Severity:** minor
 - **State:** ALIVE
- **Alert:** vm-error

- **Annotations:**

- **Type:** Equipment Alarm
- **Summary:** "{{ \$labels.vm_name }}" is down."

- **Expression:**

```
|  
    upf_state == 0
```

- **For:** 5s

- **Labels:**

- **Severity:** major
- **State:** ERROR

- **Alert:** vm-recovering

- **Annotations:**

- **Type:** Equipment Alarm
- **Summary:** "{{ \$labels.vm_name }}" is recovering."

- **Expression:**

```
|  
    upf_state == 3
```

- **For:** 5s

- **Labels:**

- **Severity:** warning
- **State:** RECOVERING

- **Alert:** vm-recovery-failed

- **Annotations:**

- **Type:** Equipment Alarm
- **Summary:** "{{ \$labels.vm_name }}" failed to recover."

- **Expression:**

```
|  
    upf_state == 4
```

- **For:** 5s
- **Labels:**
 - **Severity:** critical
 - **State:** RECOVERY_FAILED

confd-user-status

Rules:

- **Alert:** confd_user_password_expiring
 - **Annotations:**
 - **Type:** Confd User Status Alarm
 - **Summary:** "Confd user {{ \$labels.namespace }}/{{ \$labels.confidUser }} password is expiring in less than 60 days."
 - **Expression:**

```
|
    confd_user_password_days_to_expiry changelesser 60 and
confd_user_password_days_to_expiry >= 0
```
 - **Labels:**
 - **Severity:** major
- **Alert:** confd_user_password_expired
 - **Annotations:**
 - **Type:** Confd User Status Alarm
 - **Summary:** "Confd user {{ \$labels.namespace }}/{{ \$labels.confidUser }} password is expired."
 - **Expression:**

```
|
    confd_user_password_days_to_expiry < 0
```
 - **Labels:**
 - **Severity:** critical

