

What's New in Kubernetes 1.18



CLOUD NATIVE
COMPUTING FOUNDATION

Presenters



Jeremy Rickard

1.18 Enhancements Lead

@jerickar



Jorge Alarcon

1.18 Release Lead

@alejandrox1



Karen Chu

*1.18 Communications Lead &
Moderator*

@karenchu



Kubernetes 1.18 - A Bit Quarky



image credit: Maru Lango
@marulango



Agenda

- ★ 1.19 Release Update
- ★ 1.18 Stats
- ★ 1.18 Highlights
- ★ SIG Updates
- ★ Q&A



1.19 Release Dates

Overview

★ 1.19 Release Will Be Extended

- Original Target: June 30th
- New Target: August 4th

★ Read more:

https://groups.google.com/forum/#!msg/kubernetes-dev/IVpilOZ4WcM/I_LGDpNkAgAJ



1.18 Enhancements

Overview

- ★ 38 total enhancements tracked in 1.18
 - 15 Stable Enhancements
 - 11 Graduating to Beta
 - 12 Introduced Alpha features



1.18 Highlights

Fit & Finish

Stability! New Features!



client-go!

- Methods now take context.Context as first param
- Create, Update, and Patch methods now accept CreateOptions, UpdateOptions and PatchOptions
- Delete and DeleteCollection methods now accept DeleteOptions by value
- Check out <https://github.com/kubernetes-sigs/clientgofix>



Stability

- ★ Taint Based Eviction
- ★ Kubectl Diff
- ★ API Server Dry Run
- ★ CSI Block Storage Support
- ★ Windows Improvements



New Things!

- ★ Priority and Fairness for API Server Requests
- ★ Kubectl Debug
- ★ Configurable HPA Scale Velocity
- ★ Immutable Secrets and ConfigMaps



SIG Updates

API MACHINERY

Priority and Fairness for API Server Requests

Priority and Fairness for API Server Requests

- Allow for self-maintenance under heavy traffic
- Keep components/tenants from crowding out others

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



API Server Network Proxy KEP to Beta

Move API Server Network Proxy to Beta

- separates user-initiated network traffic from API server-initiated traffic
- Cloud providers want to control how API server to pod, node and service network traffic is implemented

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



API Server Dry-Run

Add apiserver "dry-run" query-parameter so that requests can be validated and "processed" without actually being persisted.

```
kubectl apply --dry-run=server
```

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



client-go signature refactor to standardize options and context handling

Add context and missing options to client-go interfaces.

```
ctx := req.Context()
pod, err := k8client.
    CoreV1().
    Pods(ns).
    Get(ctx, podName, ...)
```

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



ARCHITECTURE

Enable running conformance tests without beta REST APIs or features

Establish a process for ensuring that the Kubernetes components and the Kubernetes conformance tests have no dependencies on beta REST APIs or features.

Status: Phase 1

Identify beta requirements

[Tracking Issue](#)

[Enhancement Proposal](#)



Auth

Provide OIDC discovery for service account token issuer

Provide OIDC discovery endpoints for the API server's service account token issuer.

- enables federation of Kubernetes service account tokens issued by a cluster (the identity provider) with external systems (relying parties).

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



CertificateSigningRequest API

The Certificates API enables automation of x509 credential provisioning by providing a programmatic interface for clients of the Kubernetes API to request and obtain x509 certificates from a Certificate Authority (CA).

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Autoscaling

Configurable scale velocity for HPA

Add scale velocity configuration parameters to the HPA to control the rate of scaling.

behavior:

scaleUp:

policies:

- type: percent

value: 900%

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



CLI

kubectl debug

Add a new command to debug pods.

```
kubectl debug mypod  
--image=debian
```

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



Kubectl Diff

Users can run a kubectl command to view the difference between a locally declared object configuration and the current state of a live object.

```
kubectl diff -f ./my-manifest.yaml
```

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Moving kubectl package code to staging

Moving kubectl package code to staging.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Cloud Provider

Support Out-of-Tree vSphere Cloud Provider

Implemented a well-tested version of the cloud-controller-manager that has feature parity to the kube-controller-manager.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



CLUSTER LIFECYCLE

Kubeadm for Windows

support for Windows.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



NETWORK

Adding AppProtocol to Services and Endpoints

Added support for allowing Services and Endpoints to specify TCP, UDP, or SCTP.

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



IPv6 support added

Full Kubernetes capabilities using IPv6 networking instead of IPv4 networking.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



New Endpoint API

New Endpoint API to replace the current core/v1 endpoint API to mitigate current performance/scalability issues.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Graduate NodeLocal DNSCache to GA

NodeLocal DNSCache is an add-on that runs a dnsCache pod as a daemonset to improve clusterDNS performance and reliability.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Ingress Changes

- Support wildcard hostnames
 - Better Path Matching
 - Ingress Class Declaration

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Ingress Changes

- Support wildcard hostnames
 - Better Path Matching
 - IngressClass Declaration

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



NODE

Pod Overhead: account resources tied to the pod sandbox, but not specific containers

Sandbox runtimes introduce a non-negligible overhead at the pod level which must be accounted for effective scheduling, resource quota management, and constraining.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Node Topology Manager

Kubelet component to coordinate fine-grained hardware resource assignments for different components in Kubernetes.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Add pod-startup liveness-probe holdoff for slow-starting pods

adds **initializationFailureThreshold** to probes allowing a greater number of failures during the initial start of the container.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Extending Hugepage Feature

enable applications running in a Kubernetes cluster to use huge pages.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



SCHEDULING

Even pod spreading across failure domains

Allow scheduler to spread a group of pods across failure domains.

The existing hard inter-pod anti-affinity does not allow more than one pod to exist in a failure domain.

Status: Beta

[Tracking Issue](#)

[Enhancement Proposal](#)



Taint Based Eviction

Automatically taint nodes with **NoExecute** when they become unready or unreachable.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Add a configurable default Even Pod Spreading rule

Add default spreading rules to pods that don't define one and allow operators to define it.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Run multiple Scheduling Profiles

- Users may choose to run more heterogeneous workloads in a single cluster.
- Running multiple schedulers can cause race conditions

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



STORAGE

Raw block device using persistent volume source

Make raw block devices
(non-networked) available for
consumption via a Persistent
Volume Source

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Skip Volume Ownership Change

NewPermissionChangePolicy
string field in
PersistentVolumeClaimVolumeSource struct and related conditional
volume ownership change
behavior

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



CSI Block storage support

Support for raw block storage in
CSI

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Pass Pod information in CSI calls

Pass Pod information to CSI drivers in **NodePublish** request as CSI volume attributes.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Skip attach for non-attachable CSI volumes

Allow CSI drivers that don't support attach to not require the CSI external-attacher

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Immutable Secrets and ConfigMaps

Allow for specifying Secrets and ConfigMaps as immutable

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



Generic data populators

The PVC spec contains a **DataSource** field to specify a data source for new volumes.

This will enable users to create pre-populated volumes in a manner consistent with current practice

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



Extend allowed PVC DataSources

Enable specifying an existing PVC as a **DataSource** parameter for creating a new PVC (Cloning)

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



WINDOWS

Support CRI-ContainerD on Windows

Currently CRI support for windows is planned for ContainerD 1.3.

users will be able to take advantage of the latest container platform improvements that shipped in Windows Server 2019 / 1809 and beyond.

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



Implement RuntimeClass on Windows

Use RuntimeClass and labels to support multiple Windows versions in the same cluster, and same node with Hyper-V isolation

Status: Alpha

[Tracking Issue](#)

[Enhancement Proposal](#)



Support GMSA for Windows workloads

Enable support for Group Managed Service Account (gMSA) for Windows workloads on Kubernetes.

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



RunAsUserName for Windows

Add **RunAsUserName** field to set the user for a container's initial process on Windows

Status: Stable

[Tracking Issue](#)

[Enhancement Proposal](#)



Release Team Shadow Program

Release Team Shadow Program

★ Release Team Roles

- Release Team Lead
- Enhancements
- CI Signal
- Bug Triage
- Docs
- Release Notes
- Communications

★ 1 lead : 3 - 4 shadows

★ ~3 months // weekly workload varies depending on team

★ [Release Team Shadows Github repo](#)



Questions?

Thank You