Kubernetes 1.7

New Features
OpenShift Commons Briefing

Clayton Coleman
Architect, OpenShift and Kubernetes
Red Hat

Runtime

Security

Extensibility

Running Apps

Extensibility

Beta

Feature: Third Party Resources move to beta as Custom Resource Definitions

Description: Fix some of the outstanding storage issues and provide a solid foundation

Changes:

- Ensures storage model is appropriate
- Room for future features like validation
- Semi-automatic migration path
- TPR removed in 1.8

```
apiVersion: extensions/ ~~v1beta1~~ ALPHA
kind: ThirdPartyResource
metadata:
  name: cron-tab.stable.example.com
description: "A specification of a Pod to
run on a cron style schedule"
versions:
- name: v1
```



```
apiVersion: apiextensions.k8s.io/v1beta1
kind: CustomResourceDefinition
metadata:
  name: crontabs.stable.example.com
spec:
  scope: Namespaced
  group: stable.example.com
  version: v1
  names:
    kind: CronTab
    plural: crontabs
    singular: crontab
```

Extensibility

Beta

Feature: Register new API groups

Description: Register API to make it discoverable, automatic proxying

How it Works:

- Admin installs server and “APIService” definition
- /apis/mygroup is now proxied to that server
- kubectl and other clients now detect those resources automatically

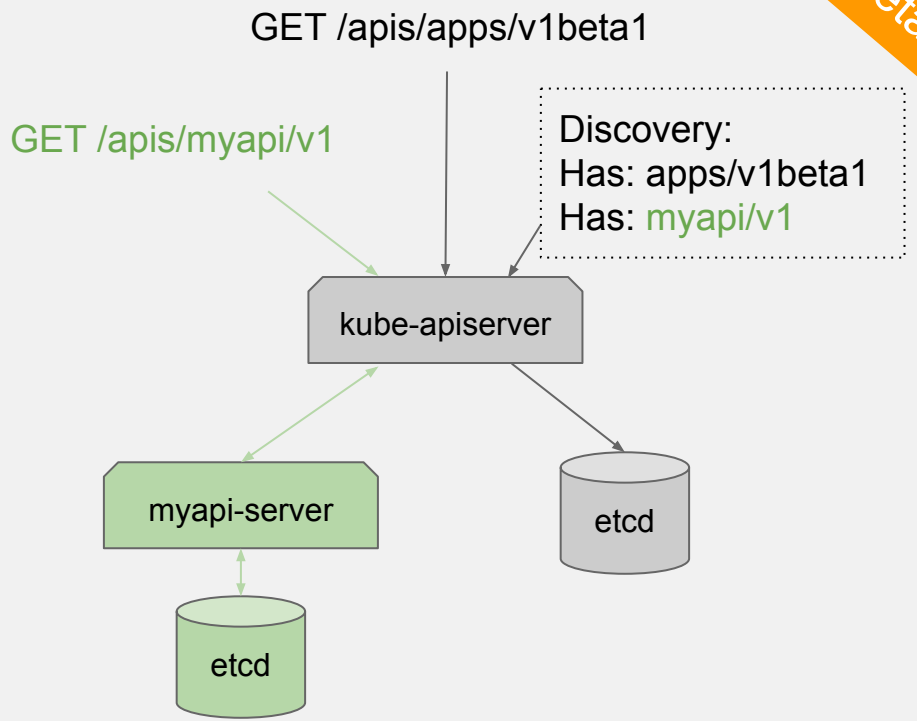

Extensibility

Alpha

Feature: Webhooks for policy

Description: Allow extension of “is this update allowed”

How it Works:

- Admin registers a webhook for making decisions on user changes (create, update, proxy, connect, etc) per resource
- Called in parallel, accept or reject
- Future work will include mutation

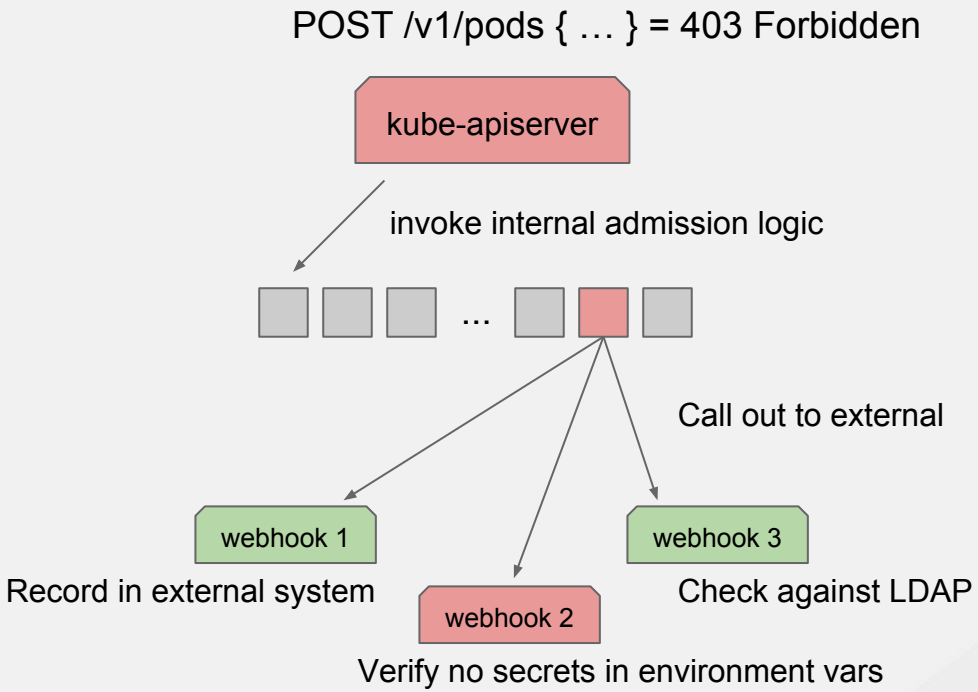

Extensibility

Alpha

Feature: Register initialization logic for resources

Description: Make it easier to hook in new defaults

How it Works:

- Register a new initializer name
- Initializer watches for new resources via the API, updates them
- After all initializers complete, other clients see that resource

```
apiVersion: v1
kind: Pod
metadata:
  name: mypod
  initializers:
    pending:
    - name: k8s.io/VerticalAutosize

Returns 201 Accepted, not visible yet

apiVersion: v1
kind: Pod
metadata:
  name: mypod
  initializers: null
spec:
  containers:
  - name: app
    resources:
      request:
        cpu: 1

Returns 200 Created, visible to all
```

Extensibility

Alpha

Feature: CLI extensions

Description: New commands can be plugged into the CLI

How it Works:

- Add a local yaml definition to specify command
- kubectl myplugin finds plugins and runs them
- Future work includes making it easy to connect to server, get credentials

```
$ cat ~/.kube/plugins/myplugin/plugin.yaml
name: "myplugin"
shortDesc: "My plugin's short description"
command: "echo Hello plugins!"

$ kubectl myplugin
Hello plugins!
```

Security

Alpha

Feature: Encrypt secrets at rest

Description: Secrets are stored in etcd in encrypted form

How it Works:

- Admin configures preferred encryption type and keys via config file
- When writing and reading from etcd, whole object is encrypted
- Support for AES-CBC, Secretbox
- Alternative to whole disk encryption (still recommended)

```
kind: EncryptionConfig
apiVersion: v1
resources:
  - resources:
    - secrets
    providers:
    - aescbc:
        keys:
        - name: key1
          secret: <BASE 64 ENCODED SECRET>
    - identity: {}
```

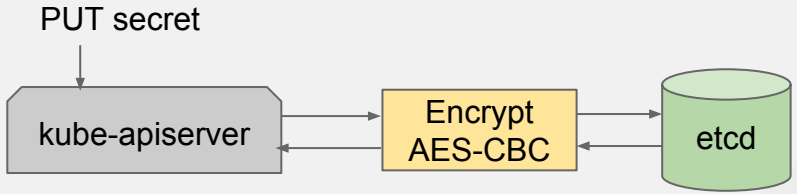

Security

Alpha

Feature: Limit node access to resources

Description: Nodes are limited to only resources scheduled onto them

How it Works:

- Nodes are granted access only to resources referenced by pods scheduled on them
- Can be given additional permissions
- Not represented by RBAC, part of basic support

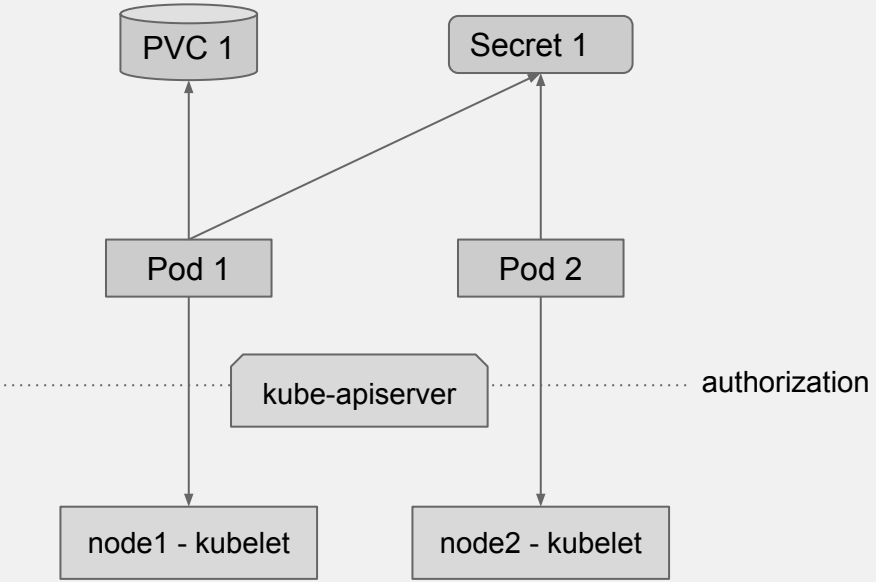

Security

Alpha

Feature: Kubelet certificate rotation

Description: Bootstrapped nodes automatically rotate client/server certificates

How it Works:

- The Kubelet, when bootstrapping, will try to refresh certificates if they are going to expire
- Adds support on the masters to help automate approving expiration refresh

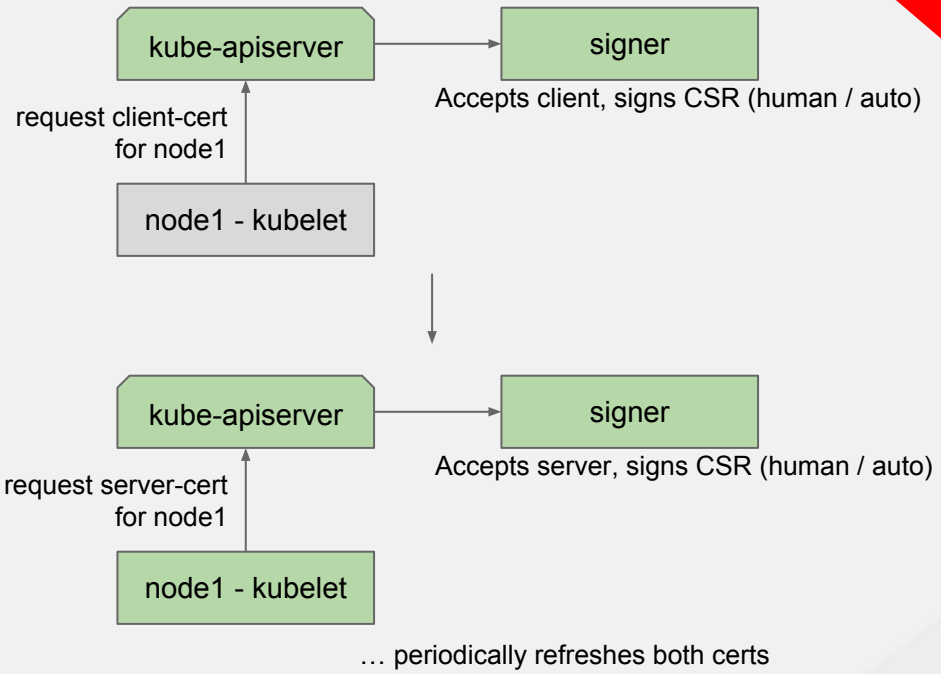

Security

Alpha

Feature: Audit improvements

Description: Provide more control and flexibility to audit logging

How it Works:

- Support selectively filtering audit events
- Send events for different phases of request lifecycle
- Send audit to a remote sink

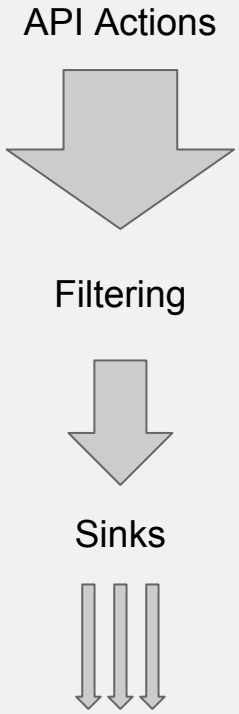

Running Apps

Beta

Feature: Stateful Set Updates

Description: Automatically roll-out updates to stateful sets

How it Works:

- Like Deployments, when updated a roll-out can be performed on each pod
- The new “partition” field can be used to control how far the rollout goes for canary testing or staged rollout
- Only updates pods

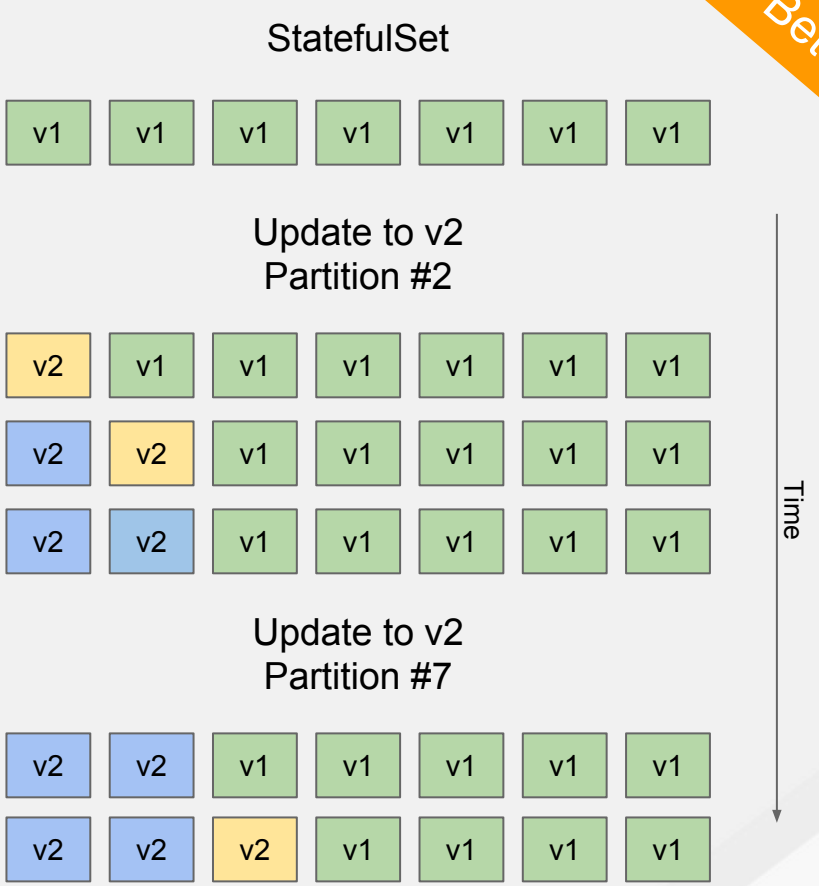

Running Apps

Alpha

Feature: Stateful Set Parallel Scaling

Description: Allow stateful sets to “burst” and create all pods at once

How it Works:

- Adds new pod management policy field
- OrderedReady = default, same as 1.6
- Parallel = create / update / scale fully
- Doesn’t affect rolling updates

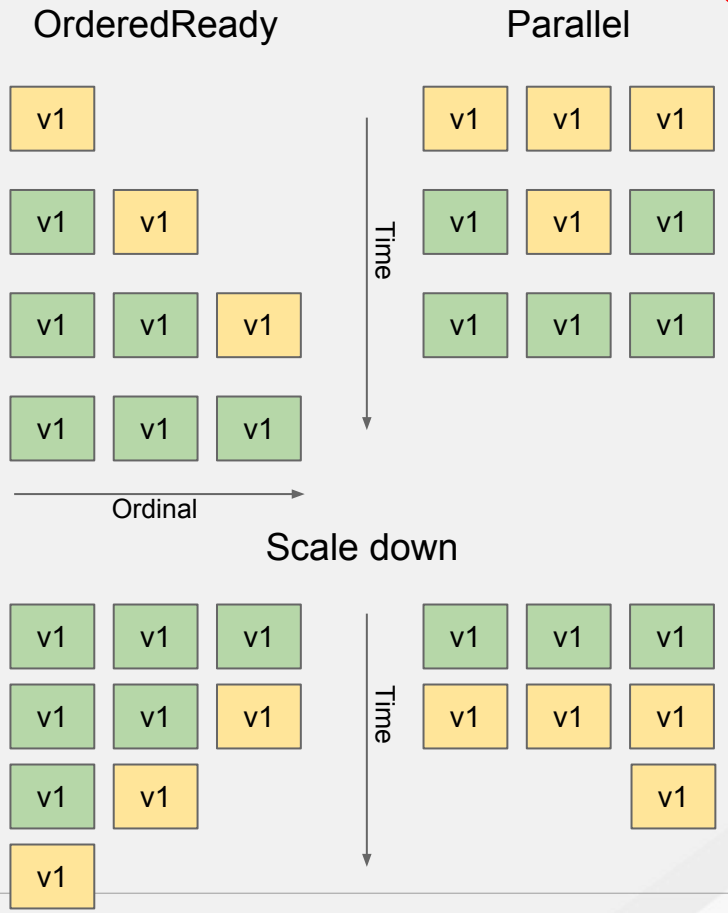

Running Apps

Beta

Feature: Daemon Set Updates

Description: Automatically roll-out updates to daemon sets

How it Works:

- Update by changing the daemon set
- MaxUnavailable rate limits change
- Dead nodes / pods considered unavailable

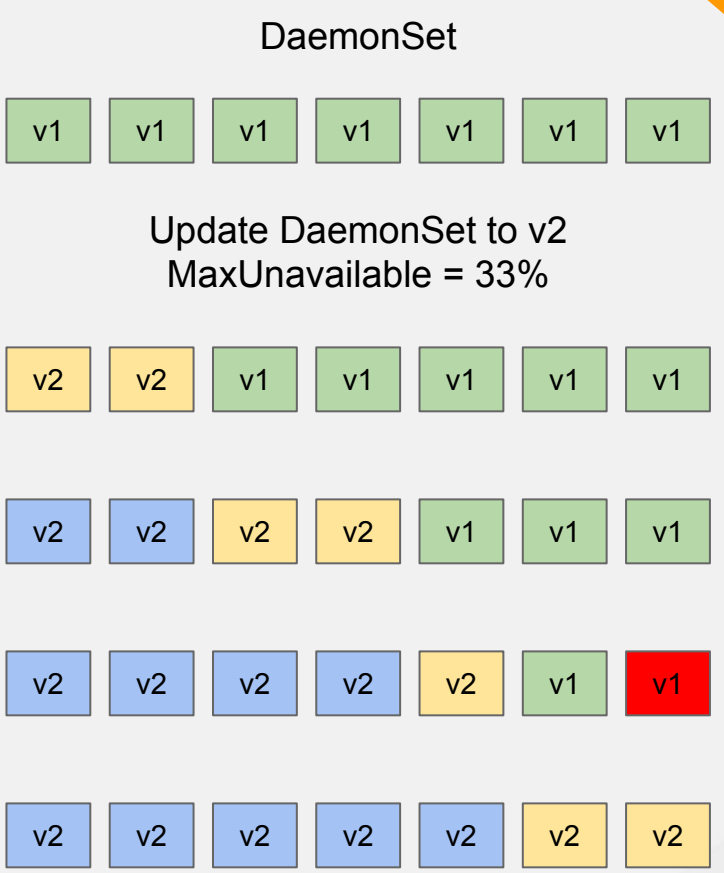

Runtime

Alpha

cri-o

Feature: CRI-o

Description: More flexible and efficient container runtime

How it Works:

- Built around systemd + runc + OCI standard libraries
- Supports exactly what Kube needs, no more
- Minimal overhead
- Future: support other OCI runtimes like runv, clear containers, etc

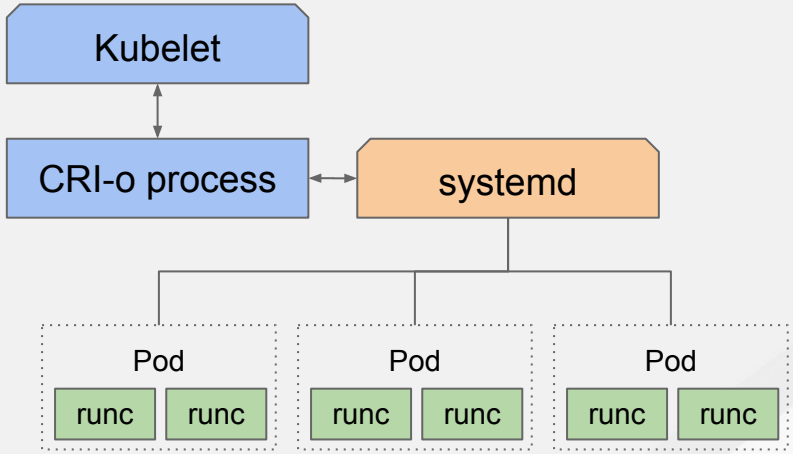

Runtime

Feature: Network Policy is GA

Description: No major changes

How it Works:

- Limit who can access services and pods within a namespace
- Implemented in many network plugins for kubernetes, expect more in the coming months
- Foundational to future policy and tenancy

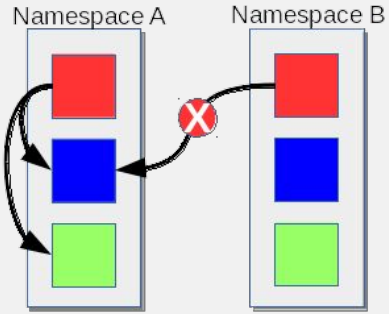


```
kind: NetworkPolicy
apiVersion: extensions/v1beta1
metadata:
  name: allow-from-red
spec:
  podSelector:
  ingress:
  - from:
    - podSelector:
        matchLabels:
          type: red
```

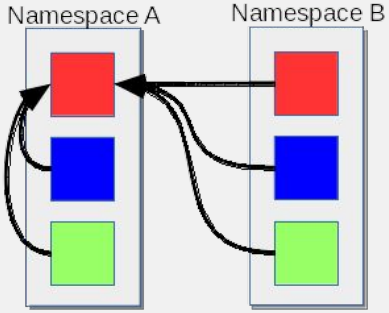


```
kind: NetworkPolicy
apiVersion: extensions/v1beta1
metadata:
  name: allow-to-red
spec:
  podSelector:
    matchLabels:
      type: red
  ingress:
  - {}
```

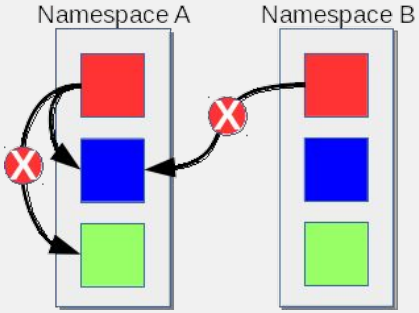


```
kind: NetworkPolicy
apiVersion: extensions/v1beta1
metadata:
  name: allow-from-red-to-blue
spec:
  podSelector:
    matchLabels:
      type: blue
  ingress:
  - from:
    - podSelector:
        matchLabels:
          type: red
```

Other Features

- Local storage volumes (early alpha)
- Lots of CLI improvements
- Performance fixes to kube-proxy
- Service catalog

Additional Resources

- K 1.7 Features:
  - https://github.com/kubernetes/features/tree/master/release-1.7
- K 1.7 Release Notes:
  - https://github.com/kubernetes/features/blob/master/release-1.7/release-notes-draft.md
- Upcoming OpenShift Commons Briefing related to K1.7
  - Kubernetes 1.7 Release Update (Aparna Sinha, Google)
  - Kubernetes Service Catalog Deep Dive (Paul Morie, Red Hat)
  - Distributed Tracing with Jaeger & Prometheus on Kubernetes (Yuri Shkuro,Uber & Greg Brown, Red Hat)
  - Full calendar here: http://commons.openshift.org/events.html
- Register now for Upcoming OpenShift Commons Gathering Dec 5 Austin, Texas

Next OpenShift Commons Gathering Dec 5th Austin https://commons.openshift.org


Questions
redhat.