

2020 document is here:

[kubeadm Office Hours Meeting Notes](#)

Join this mailing list to gain access to the agenda document:

<https://groups.google.com/forum/#!forum/kubernetes-sig-cluster-lifecycle>

kubeadm Office Hours Meeting Notes

TOP LEVEL DOC: [SIG Cluster Lifecycle Meeting Notes](#)

Meetings:

- 5pm UK time on Wednesdays, weekly (9am PT) - kubeadm office hours

All SIG Cluster Lifecycle meetings share the same Zoom room at

<https://zoom.us/j/179916854>

Contributing guide:

- <https://github.com/kubernetes/kubeadm/blob/master/CONTRIBUTING.md>
-

Wed 25th December 2019 - 9am PT, noon ET

Canceled (xmas)

Wed 18th December 2019 - 9am PT, noon ET

Canceled

Wed 11th December 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Leigh Capili - Weaveworks
- Ed Bartosh - Intel
- Sandeep Rajan - Infoblox,
- Marek Counts, Rafael Fernández López - SUSE
- Luca Lanziani - NearForm

- Lubomir I. Ivanov, Ross Georgiev, Tim St. Clair, Tunde Olu-Isa, Yassine Tijan, Fabrizio Pandinii - VMware

Agenda:

- New meeting participants?
- PSA:
 - [lubomir] 1.17 was released
 - no major hiccups, release notes needed more tweaks compared to 1.16.
 - alejandrox1 (Jorge Alacron) appointed as Release 1.18 lead
 - created housekeeping issue for 1.18:
 - <https://github.com/kubernetes/kubeadm/issues/1963>

Kubeadm planning for 1.18

(here we collect a list of features/changes to work on during the cycle, we discuss them, prioritize them and assign people to work on them).

priority: P0 highest - P4 lowest

- **[priority: P0]** moving kubeadm to the k/kubeadm repository:
 - **assigned: lubomir, yassine, luca, rosti**
 - <https://github.com/kubernetes/kubeadm/issues/1949>
 - informed both sig-docs and sig-release about some of the actions they need to help with:
 - the sig-docs action is about reference docs and is not that complex
 - on the sig-release side might have to help staffing changes - we need to adapt the release and CI builds.
- **[priority: P2]** component config management / refactors
 - **assigned: rosti**
 - KEP: <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/20190925-component-configs.md>
 - WIP PRs:
 - <https://github.com/kubernetes/kubernetes/pull/86075>
 - <https://github.com/kubernetes/kubernetes/pull/86072>
 - <https://github.com/kubernetes/kubernetes/pull/86070>
- **[priority: P2]** addon installer
 - **assigned: stealthybox, dholbach**
 - PR: <https://github.com/kubernetes/kubernetes/pull/85224>
 - KEP (provisional): <https://github.com/kubernetes/enhancements/pull/1308>
- **[priority: P2]** machine readable output for more commands
 - **assigned: bart0sh, Klaven**
 - tracking issue: <https://github.com/kubernetes/kubeadm/issues/494>

- KEP:
 - <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/20190506-kubeadm-machine-output.md>
- There is a pending PR here:
 - <https://github.com/kubernetes/kubernetes/pull/83941>
- **[priority: P1] kubeadm upgrade apply phases**
 - **assigned: Klaven, stealthybox, Tunde**
 - issue: <https://github.com/kubernetes/kubeadm/issues/1318>
 - KEP:
 - <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/0029-20180918-kubeadm-phases-beta.md>
- **[priority: ?] v1beta3**
 - **assigned: rosti, ?**
 - issues:
 - <https://github.com/kubernetes/kubeadm/labels/kind%2Fapi-change>
 - KEP:
 - <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/0023-kubeadm-config.md>
 - there are pending (on hold) PRs in k/k for some of this work
- **[priority: P3] kubeadm operator**
 - **assigned: fabrizio, SataQui, Rafa, Rosti**
 - issues:
 - <https://github.com/kubernetes/kubeadm/issues/1698>
 - KEP (provisional):
 - <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/20190916-kubeadm-operator.md>
- **[priority: P3] Windows support**
 - **assigned: lubomir, ?**
 - issues:
 - <https://github.com/kubernetes/kubeadm/issues/1393>
 - KEP:
 - <https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/kubeadm/20190424-kubeadm-for-windows.md>
- **[priority: P2] CoreDNS changes and deprecate kube-dns**
 - **assigned: Rafa, Yassin, Rajansandeep**
 - tracking issues:
 - <https://github.com/kubernetes/kubeadm/issues/1931>
 - <https://github.com/kubernetes/kubeadm/issues/1954>
 - <https://github.com/kubernetes/kubeadm/issues/1943>
 - problems:
 - our upgrades stomp existing user addon changes instead of preserving them.
 - fixing the replica count to 2 is problematic, we might have to dynamically scale up and down with new kubeadm nodes joining.
 - kube-dns is a maintenance burden, we should mark it as deprecated in 1.18

- **[priority: P1]** kubeadm survey for 1.18 / kubecon eu 2020
 - **assigned: fabrizio, tim, rafa, marek, lubomir**
 - tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1915>
- **[priority: P2]** Remove ClusterStatus from kubeadm-config
 - **assigned: fabrizio, rafa**
 - KEP:
 - <https://github.com/kubernetes/enhancements/pull/1380>
- **[priority: P2]** Kubelet/KubeProxy Node/NodeGroup Specific Config
 - **assigned: stealthybox, sraghunathan, fabrizio, rosti**
 - **Needs design**
 - Looking for general approach for multiple components
 - Issue:
 - related: <https://github.com/kubernetes/kubeadm/issues/1682>
 - TBD (pre-existing related issues probably exist)
 - KEP:
 - TBD

please add your items following the above format.

Wed 4th December 2019 - 9am PT, noon ET

Canceled.

Wed 27th November 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Rafael Fernández López - SUSE
- Sandeep Rajan - Infoblox
- Luca Lanziani
- Lubomir I. Ivanov, Ross Georgiev, Tunde Olu-Isa, Yassine Tijani, Fabrizio Pandini - VMware

Agenda:

- New meeting participants?
- PSA:
 - [lubomir] Release 1.17 timeline:
 - Cherry pick deadline - 2nd of December
 - Release - 9th of December

- <https://github.com/kubernetes/sig-release/tree/master/releases/relase-1.17>
- [lubomir] Planning session for 1.18 - When, perhaps next week?
 - 11th of December
- [lubomir] Moving kubeadm out of tree proposal document:
 - <https://docs.google.com/document/d/1xFAGAjfPQZrcqmapTwpJR5aUA7-2ku1asVeq3dBrjnQ/edit>
 - requesting review by kubeadm contributors
- [lubomir/rafael/rajansandeep]: For 1.18 make CoreDNS into a DaemonSet that targets CP nodes instead of a Deployment with 2 replicas?:
 - <https://github.com/kubernetes/kubeadm/issues/1931>
 - DaemonSet is not preferred due to auto-scanner issues.
 - undecided how to proceed: everyone should comment on the tickets
- [lubomir] kubeadm token discovery optimization:
 - <https://github.com/kubernetes/kubernetes/pull/85468>
 - <https://docs.google.com/document/d/1lf-DsWoJiZAXm9iVeO4MoP0ku0z8S30lzm9uLYigaUw/edit#heading=h.ee22he69cqq>
- [ereslibre] CRI socket annotation on upload-config kubelet
 - <https://github.com/kubernetes/kubeadm/issues/1924>
 - <https://github.com/kubernetes/kubernetes/pull/85419>
 - undecided how to proceed: everyone should comment on the ticket
- [Tunde] kubeadm upgrade apply and node command behavior
 - <https://github.com/kubernetes/kubeadm/issues/1916>
 - <https://github.com/kubernetes/kubernetes/pull/85412>
 - undecided how to proceed: everyone should comment on the tickets
- [fabrizio] Update from KubeCon NA
 - We had questions about the operator
 - Make kubeadm vendorable
 - The roadmap was well received
- [lubomir] we need to get this KEP on addon installer to “implementable”
 - <https://github.com/kubernetes/enhancements/pull/1308>
- [tunde] are there any live-coding VOD / sessions for kubeadm?
 - [fabrizio] i’m going to record something in those lines.
 - existing VOD of code walkthrough
<https://www.youtube.com/watch?v=N0OJW0nlEpQ>
-

Wed 20th November 2019 - 9am PT, noon ET

Skipping due to KubeCon NA.

Wed 13th November 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Leigh Capili - Weaveworks
- Rafael Fernández López - SUSE
- Lubomir I. Ivanov, Ross Georgiev, Fabrizio Pandini - VMware
- Luca Lanziani
- Sandeep Rajan - Infoblox

Agenda:

- New meeting participants?
- PSA:
 - [lubomir] release 1.17 timeline:
 - <https://github.com/kubernetes/sig-release/tree/master/releases/release-1.17#timeline>
 - [lubomir] k/k presubmit CI is failing for a variety of reasons. folks broke it right before code freeze.
 - [lubomir] v1beta3 is punted for the 1.18 release
 - [lubomir] kubeadm no longer depends on k/k
 - closing <https://github.com/kubernetes/kubeadm/issues/1600>
 - TODO[lubomir] we need to create an issue for moving kubeadm and enumerate tasks
 - unclear picture around test, build, release
 - [lubomir] whitelisting of mirror pods is punted to 1.18; sig-auth are still having discussions:
 - <https://github.com/kubernetes/kubeadm/issues/1835>
- [Fabrizio] remove or update
 - <https://kubernetes.io/docs/setup/production-environment/tools/kubeadm/setup-ha-etcd-with-kubeadm/> if <https://github.com/kubernetes/kubeadm/issues/1850#issuecomment-551373827>
 - [lubomir] we added a note that etcdadm will be used in the future:
 - <https://github.com/kubernetes/kubeadm/issues/1907>
- [leigh] kubeadm addon installer PR
 - Can we make alpha (feature-gated) for 1.17?
 - Do we need an exception if we can't get merged before freeze EOD tomorrow?
 - **Group agrees (+3):**
 - too tight for this complexity**
 - punt to 1.18 cycle**
 - <https://github.com/kubernetes/kubernetes/pull/85224>
 - needs rebase
 - needs upgrade (see TODO)
 - no default behavior yet

- Action Items:
 - Add e2e similar to networking: IPv6 alpha
 - Update KEP test plan
 - KEP merged
 - KEP promoted to provisional
 - Implement kube-proxy
 - Implement coredns
 - Demonstrate node-local DNS
- [lubomir/rostiti] release a new kubeadm survey for KubeCon NA 2019?
 - we can copy-paste the old one with minor changes?
 - <https://docs.google.com/forms/d/1A855D054FZfdvNbx6mAa54tcZ204GmPClhxe1UrXKd8/edit>
 - [lubomir] log ticket for KubeCon EU
- [lubomir] walk pending PRs
 - <https://github.com/kubernetes/kubernetes/pulls?q=is%3Apr+is%3Aopen+label%3Aarea%2Fkubeadm>
- [lubomir] walk the 1.17 backlog
 - <https://github.com/kubernetes/kubeadm/milestone/14>
 - fyi document about pki for external CA
<https://kubernetes.io/docs/setup/best-practices/certificates/>
 - fyi tracking issue for housekeeping activities
<https://github.com/kubernetes/kubeadm/issues/1807>

Mon 11th November 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Ross Georgiev - VMware
- Rafael Fernández López - SUSE

Agenda:

- [Ross] [punt] v1beta3 in 1.17?
 - <https://github.com/kubernetes/kubernetes/pull/84352>
 - <https://github.com/kubernetes/kubernetes/pull/84357>
 - <https://github.com/kubernetes/kubernetes/pull/84991>
 - TODO: Remove useHyperKubeImage PR
- [Ross] Component config patches for 1.17
 - <https://github.com/kubernetes/kubernetes/pull/79223>
- [Fabrizio] etcd certificate management
 - <https://github.com/kubernetes/kubeadm/issues/1850>
- [ereslibre] UseHyperKubeImage warning PR

- <https://github.com/kubernetes/kubernetes/pull/85094>

Wed 6th November 2019 - 9am PT, noon ET

Recording TODO

Attending:

- Lubomir I. Ivanov, Yassine Tijani, Ross Georgiev, Fabrizio Pandini - VMware
- Marek Counts - SUSE

Agenda:

- New meeting participants?
 - PSA:
 - [lubomir] we published a kubeadm onboarding video ("how to contribute"):
 - <https://www.youtube.com/watch?v=N0OJW0nIEpQ>
 - [lubomir] e2e test failures for kind/kinder:
 - <https://github.com/kubernetes/kubernetes/issues/84690>
 - <https://github.com/kubernetes/kubernetes/issues/84676>
 - [lubomir] the work on whitelisting mirror pods has begun:
 - <https://github.com/kubernetes/enhancements/issues/1314#issuecomment-548956137>
 - [Ross] Component config internals removal PR is rebased and updated
 - <https://github.com/kubernetes/kubernetes/pull/79223>
 - [Ross] PTAL at v1beta3 PRs
 - [lubomir] enable the TaintNodesByCondition admission controller?
 - <https://github.com/kubernetes/kubeadm/issues/1877>
 - [ereslibre/lubomir] sunsetting useHyperKubeImage from configuration (late addition to v1beta3?)
 - Hyperkube has moved out of tree and our config might need some changes:
 - <https://github.com/kubernetes/kubernetes/issues/81760>
 - <https://github.com/kubernetes/kubernetes/blob/a8220072eca6f7ae9972340422dfb46a57fa777b/cmd/kubeadm/app/apis/kubeadm/v1beta2/types.go#L167>
 - we are considering removing the flag.
 - [lubomir] to log an issue in k/kubeadm
 - [thelastproject] support a new way of creating a full join command on demand?
 - <https://github.com/kubernetes/kubeadm/issues/1872>
 - we discussed that probably adding --certificate-key to the `kubeadm create token` that can be used to also form a join command for control-planes.
 - [fabrizio] thread from kubernetes dev
 - <https://groups.google.com/forum/#!topic/kubernetes-dev/8ad1puP8hxo>
- Should we do something? My understanding:
- for single node clusters —> no
 - for single other clusters —> yes, but before 1.18

```
node.kubernetes.io/exclude-from-external-load-balancers=true \
node.kubernetes.io/exclude-disruption=true
```

- Will attempt to do backlog grooming next week - e.g. Monday.

Wed 30th October 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Fabrizio Pandini, Yassine Tijani - VMware
- Leigh - Weaveworks

Agenda:

- New meeting participants?
- PSA:
- [fabrizio] simplify wait for etcd in kubeadm join --control-plane
 - From wait all the etcd members report status to only the new etcd member report status
 - [tim] should we support parallel join
 - [leigh] it could make sense to wait for single node status
 - [fabrizio] I will do some more investigation and open an issue
- [fabrizio] Test Grid status update
 - Suffered for flakes due to external reasons
 - CoreDNS patch release
 - Delay in availability of artifacts
 - Spamming on the SCL mailing list. Change the distribution list
 - [Tim] if you are in the list, you care. It helps also for new contributors. Let's document how to implement filters for people who don't care
 - Kinder need attentions
 - to be upgraded to the latest kind release
 - [leigh] Will come back on the pending PR
 - to find a solution for kind re-architect
 - [Tim] We still need docker in our CI
- Backlog grooming
 - New issues
 - Pending PRs

Wed 23th October 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Tim St. Clair, Fabrizio Pandini, Ross Georgiev, Yassine Tijani, Andy Goldstein - VMware
- Marek Counts, Rafael Fernández López - SUSE

Agenda:

- New meeting participants?
- PSA...
 - [Fabrizio] Names for new phases
https://github.com/kubernetes/kubernetes/pull/84118#discussion_r336755506
 - [Fabrizio] Renew certs issue
<https://github.com/kubernetes/kubeadm/issues/1850#issuecomment-545099961>
- [Tim] Issues with type imports and trying to make it more consumable for v1beta3
 - 1. clusterapi copied them into the code b/c they wanted them to be omitted.
 - I believe this is planned to be fixed in v1beta3.
 - 2. Multiple api-versions (beta1/beta2)
 - How does this get exported out of a field in CAPI.
 - AI: (Tim + Andy) to open and issue.
 - 3. Object Meta
 - Being fixed.
- [Marek] vote on phases for upgrade apply
 - https://docs.google.com/document/d/1Nicy27rt9jm9tOzZ_MEcQsnmyOFI2ZB6k110aRblv2M/edit?usp=sharing
- [Marek] retries
 - <https://github.com/kubernetes/kubeadm/issues/1844>

Wed 16th October 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Ross Georgiev, Fabrizio Pandini, Yassine Tijani - VMware
- Rafael Fernández López, Marek Counts - SUSE
- Yago Nobre - Nubank
- Pawel Bek
- Irvi Aini

Agenda:

- New meeting participants?
- [lubomir] PSA: all tests are flaky due to missing artifacts:
 - <https://github.com/kubernetes/release/issues/897>
 - <https://k8s-testgrid.appspot.com/sig-cluster-lifecycle-kubeadm>
- [lubomir] PSA: moving ownership for SIG Cluster Lifecycle tests to SIG CP GCP:
 - <https://github.com/kubernetes/kubernetes/pull/83744>
- [lubomir] we have a kubeadm bug with GCP and manual Node object creation:
 - <https://github.com/kubernetes/kubernetes/issues/83923>
 - <https://github.com/kubernetes/kubernetes/pull/83926>
- [lubomir] sig-auth KEP that affects kubeadm and static pods:
 - <https://github.com/kubernetes/kubeadm/issues/1835>
- [leigh] feedback & discussion for [AddonInstaller KEP](#)
- [ross] Send an exception request for component config KEP?
 - <https://github.com/kubernetes/enhancements/pull/1256>
- [fabrizio] trying to wrap my mind on impacts on users of ongoing KEPs see [doc](#)
 - going to ask Leigh to fill the doc
- [marek] Started kubeadm upgrade change doc
https://docs.google.com/document/d/1Nicy27rt9jm9tOzZ_MEcQsnmyOFI2ZB6k110aRblv2M/edit?usp=sharing
-
- audit phases -
https://docs.google.com/document/d/1FZA873LBXf-aG4UULOk96MhT8xm_nPHq5Vu0z61S8pk/edit#heading=h.vldlop62qkot
-
- k/kubeadm PR triage
-

Wed 9th October 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Fabrizio Pandini, Ross Georgiev, Yassine Tijani - VMware
- Rafael Fernández López, Marek Counts - SUSE

Agenda:

- New meeting participants?

- [lubomir]: PSA/reminder: deadline for KEPs the 15th of Oct.
 - 1.17 release schedule is here:
 - <https://github.com/kubernetes/sig-release/tree/master/releases/relase-1.17#timeline>
- [lubomir]: PSA: there is a report that multi-CP does not work with ipv6 related to etcd:
 - <https://github.com/kubernetes/kubernetes/issues/83550>
 - we are backporting a fix for 1.16 in the etcd client.
- [lubomir]: PSA: got contacted by SIG Windows that they will probably not have the time to work on kubeadm / Windows support for going beta in 1.17.
- [lubomir]: VOTE deprecating `--config` for `upgrade`:
 - <https://github.com/kubernetes/kubernetes/pull/83530>
 - <https://github.com/kubernetes/website/pull/16703>
 - will change the PR to print warning.
- [lubomir/rostiti] VOTE vote on overriding apiserver auth modes:
 - <https://github.com/kubernetes/kubernetes/pull/82616>
 - voted +1
- [lubomir/rostiti/ed]: VOTE merge the initial structured output PR and possibly iterate later if needed?
 - <https://github.com/kubernetes/kubernetes/pull/78764>
 - voted +1
- [rostiti/tim]: KEP for component config management:
 - discussion requested by Tim
 - <https://github.com/kubernetes/enhancements/pull/1256>
- [rostiti/lubomir/yassine]: mixture of --control-plane and --certificate-key on join:
 - <https://github.com/kubernetes/kubernetes/pull/83661>
- [tim/lubomir] bump docker in the kubeadm validators:
 - <https://github.com/kubernetes/kubeadm/issues/1821>
- [lubomir] reset nukes the cni directory; do we want that?:
 - <https://github.com/kubernetes/kubeadm/issues/1822>
 - Yassine promised to have a look.
- [fabrizio] VOTE on /hold cancel for
 - <https://github.com/kubernetes/kubeadm/pull/1824>
 - <https://github.com/kubernetes/test-infra/pull/14673>

Wed 2nd October 2019 - 9am PT, noon ET

Recording: **file got corrupted**

Attending:

- Lubomir I. Ivanov, Ross Georgiev, Fabrizio Pandini, Tim St. Clair, Yassine Tijani - VMware

- Rafael Fernández López - SUSE
- Arvinderpal Wander - Ericsson
- Yago Nobre - Nubank
- Daniel Lipovetsky - D2IQ
-

Agenda:

- New meeting participants?
- PSA?
- [Ross] v1beta3 proposed changes (max 5 min discussion for each)
 - **Y P0** +OPTIONAL omitempty fields in the config
 - <https://github.com/kubernetes-sigs/cluster-api-bootstrap-provider-kubeadm/issues/254>
 - KEP by Justin for optional fields: <https://github.com/kubernetes/enhancements/pull/1173>
 - **Y P1** Add ObjectMeta to kubeadm GVKs
 - <https://github.com/kubernetes/kubeadm/issues/1727>
 - **Y P1** Make AddOns opt-in by supplying an array of names in a field of the ClusterConfiguration
 - <https://github.com/kubernetes/kubeadm/issues/1756>
 - Example: strings as new addon list


```
apiVersion: kubeadm.k8s.io/v1beta3
kind: ClusterConfiguration
addons:
- CoreDNS
- KubeProxy
```
 - Example 2: bools in new addon map


```
apiVersion: kubeadm.k8s.io/v1beta3
kind: ClusterConfiguration
addons:
  CoreDNS: false
  KubeProxy: true
```
 - Example 3: bools on existing structs


```
apiVersion: kubeadm.k8s.io/v1beta3
kind: ClusterConfiguration
dns:
  enabled: false
  type: CoreDNS
kubeProxy:
  enabled: true
```
 - AddonInstaller can use these bools/list-items to structure default the component-config for the proper kubernetes version

- [HOLD] Y/N Pn Support supplying API Server Authorization modes in the ClusterConfiguration
 - <https://github.com/kubernetes/kubeadm/issues/1787>
 - Example:


```
apiVersion: kubeadm.k8s.io/v1beta3
kind: ClusterConfiguration
apiServer:
  authorizationModes:
    - Node
    - Webhook
    - RBAC
```
- [HOLD] Y/N Pn Use proper lists for Service and Pod subnets (IPv6 dual stack support)
 - <https://github.com/kubernetes/kubeadm/issues/1612>
 - Example:


```
apiVersion: kubeadm.k8s.io/v1beta3
kind: ClusterConfiguration
networking:
  serviceSubnet:
    - 10.96.0.0/12
    - 3ffa::/24
  podSubnet:
    - 10.97.0.0/12
    - 3ffb::/24
```
- [fabrizio] Kubeadm Operator
 - Demo
 - Next steps: provisional KEP approved
 - Discussion of alpha feature over multiple cycles:
 - [tim] Can we hide this somehow from users?
 - [stealthybox] Docker has DOCKER_CLI_EXPERIMENTAL=enabled
 - [stealthybox] Should have a way to ship CLI help, docs, and other U/X without unreasonable fear of breaking users of clearly marked alpha features
 - Discussion of imperative vs declarative API's
 - breakout of imperative details into multiple sub objects is good for debugging, testing, and U/X
 - most k8s API's are declarative
 - we should be careful about exposing things we don't want people to use
 - often, status objects are used for operator/controller state machines
 - we can refactor to whatever design is needed
 - it's good that the operator POC is being built this way because it forces good patterns (typed state is being stored into the apiserver -- can be picked up)
- [Arvinder]: Dual-Stack List of Service CIDRs

- <https://github.com/kubernetes/kubernetes/pull/82473>
- [stealthybox] Functional POC of kubeadm vendoring addon-installer:
 - Demoed on the addons call yesterday:
 - Recording: <https://youtu.be/a-FIhcXrnp0?t=520>
 - Meeting notes: https://docs.google.com/document/d/10_t1_SXcFGb-2109OpcFVrdxfnVEuO05MBrXtasB0vk/edit?pli=1#heading=h.7unkz13d5r8p
 - Diff: <https://github.com/stealthybox/kubernetes/compare/kubeadm-addon-installer>

Wed 25th September 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Tim St. Clair, Ross Georgiev, Yassine Tijani, Fabrizio Pandini - VMware
- Rafael Fernández López - SUSE
- Jorge Alarcon
- Yago Nobre - Nubank

Agenda:

- New meeting participants?
- [lubomir] PSA: backporting a fix for 1.16 for feature-gates/dualstack:
 - <https://github.com/kubernetes/kubernetes/pull/83025>
 - we don't have e2e/integration tests for the feature gate itself.
 -
- [lubomir] PSA: adding Yassine as reviewer. Welcome!
 - <https://github.com/kubernetes/kubernetes/pull/83116>
 - <https://github.com/kubernetes/kubeadm/pull/1802>
- [Ross] PSA: kubeadm component config management KEP PR is up
 - <https://github.com/kubernetes/enhancements/pull/1256>
- [Yassine] PSA - ipallocator work ongoing
 - ipallocator merged in k8s.io/utils <https://github.com/kubernetes/utils/pull/112>
 - bump utils version in k/k <https://github.com/kubernetes/kubernetes/pull/83113>
- [lubomir] flannel issue reports
 - we need to point our flannel installation guide at this commit:
 - <https://github.com/coreos/flannel/pull/1181>
 - [rafa] to send PR to update.
- [fabrizio] Check of v1.17 priority with Tim and Lubomir
- [fabrizio] - If there is time - KEP / POC Kubeadm Operator
- [your agenda items here]

Wed 18th September 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Fabrizio Pandini, Ross Georgiev - VMware
- Rafael Fernández López, Marek Counts - SUSE
- Yago Nobre - Nubank
- Leigh Capili - Weaveworks
- Sandeep Rajan - Infoblox

Agenda:

- New meeting participants?
- PSAs
 - 1.16 release is landing / kubeadm test grid looks ok
 - Kubeadm repo code organization
<https://github.com/kubernetes/kubeadm/issues/1792>
- **kubeadm planning for 1.17**
- 1.17 milestone
 - <https://github.com/kubernetes/kubeadm/milestone/14>
- **Carry over activities for 1.17**
- [P0] [Rosti, Yassine, Marek] preparing to move kubeadm out of tree
 - <https://github.com/kubernetes/kubernetes/issues/78534>
 - <https://github.com/kubernetes/kubeadm/issues/1600>
 - [P0] use the external APIs of kube-proxy and kubelet
 - [P1] move system validators
 - [P1] remove ipallocator
 - [P1.5] Component Config handling in kubeadm
 - [P2] check remaining import-boss deps to remove?
- [P1] [Arvinder] dual-stack support
 - <https://github.com/kubernetes/kubeadm/issues/1612>
- [P2] [Iubomir, [ksubrmnn](#)] windows support
 - <https://github.com/kubernetes/kubeadm/issues/1393>
- **New feature/activities for 1.17**
- [P2] [Fabrizio, Yago, Rafael] kubeadm operator

- <https://github.com/kubernetes/kubeadm/issues/1698>
- merge KEP
- MVP with
 - Action/TaskGroup/Task Orchestration
 - Upgrade workflow
- [P1][Ross] deprecate v1beta1 - it is just starting the deprecation cycle
- [P1.5][Ross, Rafael] adding v1beta3
 - <https://github.com/kubernetes/kubeadm/issues?q=is%3Aopen+api+label%3Akind%2Fapi-change>
 - collect a list of things we want to change in a tracking issue
 - CABPK friendly (avoid rendering/requiring default or empty values)
 - Kustomize friendly (add object meta)
 - dual stack (CIDR/addresses from csv string to list of strings)
 - addons
 - make the list of addons fully customizable (e.g. use local DNS, skip kube proxy)
 - unify the semantic for addon configuration
 - Fix authorization modes
- [P2.5] [Fabrizio, Leigh, Yago] make addons more flexible
 - [P1.5] [Leigh/Fabrizio] KEP
 - make addons customizable in the kubeadm config?
 - cluster-addon / addon-installer integration?
 - Requirements for alpha/feature gated.
 - CoreDNS and KubeProxy
 - Airgapped support (custom image repository)
 - Component config
 - DryRun?
 - Upgrades? -- Prune comparing with clusterconfig ?
 - Kustomize?
- **Backlog activities for 1.17**
- [P3] [Marek] `kubeadm upgrade apply` phases?
 - <https://github.com/kubernetes/kubeadm/issues/1318>
- [P3] security improvements
 - [P?] finish work on the kubernetes security audit items
 - <https://github.com/kubernetes/kubeadm/issues/1717>
 - [P?] other certificate management improvements?
 - <https://github.com/kubernetes/kubeadm/labels/area%2Fsecurity>
 - [P?] CIS benchmark
 - <https://github.com/kubernetes/kubeadm/issues/683>
 - <https://github.com/kubernetes/kubeadm/issues/1648>
 - add tests
 - not tracked; needs help from the CIS folks.
 - document workarounds
 - <https://github.com/kubernetes/kubeadm/issues/1649>
- [P3] Machine readable output (<https://github.com/kubernetes/kubeadm/issues/494>)
- **your items here**

Wed 11th September 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Fabrizio Pandini, Tim St. Clair - VMware
- Arvinderpal Wander -- Ericsson
- Marek Counts - SUSE
- Yago Nobre - Nubank

Agenda:

- New meeting participants?
- [lubomir] PSA: we are in code-thaw for 1.16. release is planned for Monday. the master branch is open for 1.17 PRs.
- [lubomir] PSA: planning for 1.17 is next week
 - separate doc or inline in this one?
- [lubomir] PSA: kube-proxy currently incompatible with iptables ≥ 1.8
 - <https://github.com/kubernetes/kubernetes/issues/71305>
 - <https://github.com/kubernetes/website/pull/16271>
 - we asked sig-network to see if we have CI tests that catches such problems.
 - follow up, potentially raise to steering
- [lubomir] kubeadm cert / CA discussion:
 - there is a request that CA should expire after more than 10years
 - <https://github.com/kubernetes/kubeadm/issues/1783>
 - support checking CA expiration time:
 - <https://github.com/kubernetes/kubeadm/issues/1784>
 - <https://github.com/kubernetes/kubernetes/pull/82524>
 - support kubelet client cert rotation on `kubeadm init` nodes:
 - <https://github.com/kubernetes/kubeadm/issues/1753>
- [lubomir] applying extraArgs vs static pod flag overrides:
 - <https://github.com/kubernetes/kubeadm/issues/1787>
- [arvinder] Dual-Stack List of Service CIDRs
 - <https://github.com/kubernetes/kubernetes/pull/82473>
- [fabrizio] kubeadm operator discussion
 - should it be in new repo? should be in k/kubeadm as a start
 - should the KEP be tracked in k/enhancements? possible, yes.

Wed 4th September 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Ross Georgiev, Yassine Tijani - VMware
- Rafael Fernández López - SUSE
- Sandeep Rajan - Infoblox

Agenda:

- New meeting participants?
- [lubomir] PSA: testgrid shows our -master jobs have not run since yesterday:
 - <https://k8s-testgrid.appspot.com/sig-cluster-lifecycle-kubeadm>
 - notified #sig-testing
 - seems like a testgrid job only problem and the jobs are actually running
- [lubomir] PSA: coredns / etcd bump updates:
 - kubeadm at master is already using coredns 1.6.2
 - etcd bump is WIP: <https://github.com/kubernetes/kubernetes/pull/82199> waiting on the image to be pushed.
- [lubomir] PSA: we sent a couple of PRs attempting to deflake kind/kinder test jobs:
 - <https://github.com/kubernetes/test-infra/pull/14143>
 - <https://github.com/kubernetes/kubeadm/pull/1767>
 - quick question about [this code](#) in kinder
- [lubomir] PSA: all our docs for 1.16 passed tech review:
 - <https://github.com/kubernetes/kubeadm/issues/1739>
- [lubomir] PSA: concurrent join of control plane nodes is seeing rare flakes, possibly we need to increase timeouts:
 - <https://github.com/kubernetes-sigs/kind/pull/685>
- [ross] Optional addons (kube-proxy in particular)
 - <https://github.com/kubernetes/kubernetes/pull/82248>
 - <https://github.com/kubernetes/kubeadm/issues/1733>
 - <https://github.com/kubernetes/kubeadm/issues/1756>
- [fabrizio] PSA: kubeadm bootstrap provider (CABPK) v0.1.0 is out
 - <https://github.com/kubernetes-sigs/cluster-api-bootstrap-provider-kubeadm>
- [fabrizio] PSA: Discussion about etcdadm integration document: [recording](#)
- [fabrizio] WIP on [kubeadm operator](#)
- [Yassine] discuss proxy/kubelet config
- Your agenda items here

Wed 28th August 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Yassine Tijani, Ross Georgiev - VMware
- Rafael Fernández López - SUSE

Agenda:

- New meeting participants?
- [lubomir] PRs to add 1.16 test signal are up; need LGTM (/cc Fabrizio):
 - <https://github.com/kubernetes/kubeadm/pull/1744>
 - <https://github.com/kubernetes/test-infra/pull/14037>
- [lubomir] PSA: updating etcd and coredns is still on hold, mostly waiting on approval and someone from Google to push images to GCR.
 - <https://github.com/kubernetes/kubeadm/issues/1726>
 - <https://github.com/kubernetes/kubeadm/issues/1745>
- [lubomir] PSA: the PR for adding a kubeadm install script for Windows merged:
 - <https://github.com/kubernetes-sigs/sig-windows-tools/pull/1>
 - still targeting alpha for 1.16. tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1393>
- [lubomir] PSA: docs placeholder PRs for 1.16 are WIP:
 - will work on them today + tomorrow and ping for reviews. deadline is next tuesday.
 - <https://github.com/kubernetes/kubeadm/issues/1739>
- [rostiti] working on component config planning.
 - there has to be a distinction between component config created by users and tools like kubeadm.
 - how to do migration of user configs.
 - potentially drop full blown component configs from kubeadm and allow users to only patch.
- [yassine] move ipallocator imports to k/utils.
- [lubomir] PSA: import boss and _test.go files.
 - once we stop using k/k/pkg we should check for _test.go imports using the new import-boss flag.
- [leigh] update on addon work / addon installer.
 - discussion about how to pipe addon component config from user to the addon installer
 - storing addon patches in the cluster vs github repo.
 - [rafael] What is the expected upgrade U/X and how can we best preserve User intent keeping in mind the history of kubeadm is to store all of the config in the cluster?

- [leigh]
 - It's good to maintain a record of all configs in the cluster to preserve the user intent
 - It's bad for kubeadm to, automatically on upgrade, fetch and use this record as the authority on what the user intends for the upgrade. This makes what the source of truth for the upgrade is really hard to determine
 - It's reasonable to expect the user to provide the inputs for their customization of the cluster install during every operation (install / upgrade)
 - I think we should improve the U/X of kubeadm to fetch all of the configs to make it easier for users who do not have them saved in git or other source control. We should consider taking a hash of everything to make it easy to determine the provenance of the config. `kubeadm upgrade diff` is an important U/X
 - This is related to the U/X of user-managed ComponentConfig upgrades in preparation for a cluster upgrade
- [ross]
 - We need to be careful about what we do when we store configs in the cluster
 - A transparent/automatic behavior of re-using previously used patches is not desirable. There could be edge cases when ComponentConfigs are changing fields between APIVersions.
- [lubomir] walk the backlog / pending PRs.
 - kubeadm join fails on kube-proxy config missing
 - <https://github.com/kubernetes/kubeadm/issues/1733>
 - TODO lubomir to send POC PR
 - use-case is to support Cilium/kube-router kube-proxy replacement
 - [ross] this is a side-effect of our ComponentConfig interface pulling all supported config types despite join only requiring the KubeletConfig
 - fix is to make ComponentConfig pull optionally specific

Wed 21th August 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Ross Georgiev, Yassine Tijani, Fabrizio Pandini - VMware

- Rafael Fernández López - SUSE

Agenda:

- New meeting participants?
- [lubomir] PSA: we have a number of PRs that are not merging - “tide” the tool that merges PR seems to have a problem.
- [lubomir] PSA: coredns in kubeadm just updated to 1.5.0:
 - CI is green for 1.15 -> master upgrades
 - <https://k8s-testgrid.appspot.com/sig-cluster-lifecycle-kubeadm#kubeadm-kinder-upgrade-1-15-master>
- [lubomir] PSA: we are updating etcd to 3.3.15 in 1.16:
 - <https://github.com/kubernetes/kubeadm/issues/1726>
- [lubomir] PSA: import boss update:
 - the only remaining items are IPVS checks and internal kube-proxy/kubelet types:
 - <https://github.com/kubernetes/kubeadm/issues/1600>
 - we created a new repository for validators (thanks Yassine):
 - <https://github.com/kubernetes/kubeadm/issues/1638>
 - <https://github.com/kubernetes/system-validators>
- [lubomir] discuss printing stack trace for all kubeadm errors:
 - <https://github.com/kubernetes/kubernetes/pull/80937>
- [ross & fabrizio] Last call for review on kustomize UX PR:
 - <https://github.com/kubernetes/kubernetes/pull/81458>
- [lubomir] we walked the pending PRs in k/k.
- [lubomir] we walked a couple of new tickets in k/kubeadm.

Wed 14th August 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Fabrizio Pandini, Lubomir I. Ivanov, Ross Georgiev, Yassine Tijani, Naadir Jeewa, Tim St. Clair - VMware
- Arvinderpal Wander - Ericsson
- Marek Counts - SUSE
- Andrew Rynhard - Talos Systems
- Tommy Chong

Agenda:

- New meeting participants?
- [fabrizio] update on the Kustomize work:
 - <https://github.com/kubernetes/kubernetes/pull/80905> Merged!

- PR with UX changes for kustomize in flight
- PR for implementing E2E tests for kinder in flight
- TODO: jsonPathes6902 (not blocking for v1.16)
- AI: contact @seh/power users for early feedback
- [fabrizio, lubomir]
 - Please take a look at documents for
 - [kubeadm-etcdadm integration](#)
 - [tim] don't want to force adoption if the project isn't ready
 - [kubeadm-addons integration](#)

The goal is to collect requirements in order to start a discussion with etcdadm team for defining an etcdadm library interface (s/etcdadm/addon-installer)
- [arvinder] Dual-Stack Update
 - <https://github.com/kubernetes/kubernetes/pull/79033> is ready. Waiting on kube-proxy approval.
 - Starting on phase-2 tasks:
 - <https://github.com/kubernetes/kubeadm/issues/1612>
- [lubomir] Kubernetes Security Audit - kubeadm related items:
 - <https://github.com/kubernetes/kubeadm/issues/1715>
 - AI
 - remove shell expansions
 - document kubeadm reset is best effort
 - check that "best effort" behave consistently
 - <https://github.com/kubernetes/kubeadm/issues/1716>
 - <https://github.com/kubernetes/kubeadm/issues/1717>
 - list of all findings:
 - <https://github.com/kubernetes/kubernetes/issues/81146>
- [andrewrynhard]
 - Moving kubeadm out of tree
- [yastij] moving kubeadm out of tree
 - moving validators to a new repo, sig-arch notice at <https://groups.google.com/forum/#!topic/kubernetes-sig-architecture/QKeNaEnB-P>
 -

Wed 7th August 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Tim St. Clair, Lubomir I. Ivanov, Ross Georgiev, Yassine Tijani - VMware
- Rafael Fernández López - SUSE

- Leigh Capili - Weaveworks
- Arvinderpal Wander - Ericsson

Agenda:

- New meeting participants?
- [lubomir] PSA: did big grooming of the backlog at k/kubeadm - closed 30/40 issues.
- [fabrizio] update on the Kustomize work:
 - <https://github.com/kubernetes/kubernetes/pull/80905>
 - Requires top-level approval because kustomize needs to be re-imported into go-modules since kubectrl has moved out-of-tree, so will have to wait a little while longer.
- [timothysc]
 - KubeConNA Talk submitted (fabrizio, and yago)
 - Would like to get others involved for KubeCon CN/Asia
 - [lubomir] we have reps for cn
- [arvinder] Dual-Stack Update - PRs for review
 - Pod-CIDR support: <https://github.com/kubernetes/kubernetes/pull/79033>
 - E2E test for list of pod-cidrs: <https://github.com/kubernetes/kubernetes/pull/80947>
- [yassine] discuss move of validators to a new repo.
 - <https://github.com/kubernetes/kubeadm/issues/1638>
- [lubomir] check backlog for 1.16 and PR list:
 - <https://github.com/kubernetes/kubeadm/milestone/13>
 - <https://github.com/kubernetes/kubernetes/pulls?q=is%3Apr+is%3Aopen+label%3Aarea%2Fkubeadm>

Wed 31th July 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Fabrizio Pandini, Lubomir I. Ivanov, Tim St. Clair, Ross Georgiev, Yassine Tijani - VMware
- Yago Nobre - Nubank
- Marek Counts - SUSE
- Leigh Capili - Weaveworks

Agenda:

- New meeting participants?
- [lubomir] PSA: need to patch 1.15's minimal control-plane versions
 - <https://github.com/kubernetes/kubernetes/pull/80794>
 - [fabrizio] we can transform this in something automated like client version
 - [leigh] we should make sure it is developer friendly
 - [leigh/lubomir] we should be careful to tight to K8s release stuff
 - AI: update and transform into an issue for better visibility
<https://github.com/kubernetes/kubeadm/blob/master/docs/release-cycle.md>
- [Fabrizio] action required and cherry pick on [#80676](#). YES or NO?
 - [Lubomir] NO (see comments)
 - [Tim, Ross] NO
- [Fabrizio] PSA: [Kustomize POC](#)
- [lubomir] PSA: logged ticket for “best practices for upgrades” related to kubeadm:
 - <https://github.com/kubernetes/kubernetes/issues/80774>
- [timothysc] PSA: Collecting behavioral features for a kubeadm operator.
 - <https://github.com/kubernetes/kubeadm/issues/1698>
- [arvinder] Dual-Stack Update:
 - Merged FG propagation: <https://github.com/kubernetes/kubernetes/pull/80531>
 - Updated PR for '--pod-network-cidr' will be out shortly:
<https://github.com/kubernetes/kubernetes/pull/79033>
 - [Lubomir] should we document this?
 - [Tim] we should not document alpha feature
 - [arvinder] probably when it will reach Beta we will need some docs
- [leigh] Addons installer demo/code-walkthrough
 - <https://github.com/kubernetes-sigs/addon-operators/pull/25/>
- [lubomir] look at pending PRs and backlog?
 - *we looked at some pending k/k PRs*

Wed 24th July 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Fabrizio Pandini, Tim St. Clair, Ross Georgiev, Andrew Kutz - VMware
- Arvinderpal Wander - Ericsson

Agenda:

- New meeting participants?

- [lubomir] PR to add a note about using control-plane-endpoint instead of advertise address for easy single CP -> HA transition:
 - <https://github.com/kubernetes/website/pull/15524>
- [arvinder] Dual-Stack Update
 - Propagate dual-stack feature-gates
<https://github.com/kubernetes/kubernetes/pull/80531>
 - E2E tests: <https://github.com/kubernetes/kubernetes/pull/80259>
- [Fabrizio] Update on enhancements
 - [Kustomize](#): feedback addressed, open points:
 - [Lubomir comment](#)
 - [Machine readable output](#): agreement on “there are no blocking comments about what the KEP is proposing at high level, while pending discussion on implementation details can be moved/postponed the to implementation PRs”; please make a last pass
 - TODO: we should merge both KEPs before Tuesday next week.
- [Fabrizio] Demo on kustomize: Know limitations
 - #1 two patches for the same object --> error
 - #2 JSON patch through patchesJson6902 --> a wrapper is required

Wed 17th July 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Tim St. Clair, Lubomir I. Ivanov, Yassine Tijani, Ross Georgiev, John Harris, Fabrizio Pandini - VMware
- Marek Counts - SUSE
- Ed Bartosh - Intel

Agenda:

- [lubomir] HPA conformance tests require the metric server now:
 - <https://github.com/kubernetes/kubernetes/pull/79954>
 - How this impacts kubeadm?
 - we should revert the PR or implement a mock (test without any specific implementation of a metric system)
 - [lubomir] log a k/k ticket to ask for revert (EDIT: discussed on slack that revert will be done directly).
- [Fabrizio] Proposal for a new meeting organization
 - Welcome new contributors
 - 1st part of the meeting: Actionable items, status updates, PSA or RFE
 - 5 minutes max each
 - possible each should be closed with a decision/action items
 - if times allows meeting: Alternatively or in separated meetings

- backlog grooming
 - deep dive on specific topics
- [lubomir] new contributors on the call?
 - John Harris: works for VMware, interested in kubeadm/etcdadm integration.
- [Fabrizio] PSA: image name in CI bucket on GCE is changed for master ([#80047](#)); PR in flight to fix kinder tests ([#kubeadm/1673](#))
- [Fabrizio] PSA: [#website/15452](#) upgrade procedure should be changed for HA (thanks Rafal for bringing this out, Lubomir for the fix)
 - CURRENT: upgrade control-plane instances and kubelet on the first CP node, then repeat on other CP nodes, finally upgrade nodes → upgrading 1 kubelet before all the api servers instance violates version skew policy
 - PROPOSAL: first upgrade all the CP instances, then kubelet on CP nodes, finally upgrade nodes
- [Arvinder] Dual-Stack ([#kubeadm/1612](#)) Updates
 - FG merged into master
 - Will send a new PR for propagating dual-stack FG to all k8s components.
 - we should have a Message for the feature gate at the end of the cycle.
- [lubomir] documenting CNAME/etc/hosts usage with control-plane-endpoint
 - will send PR for this to k/website
- [Fabrizio] status upgrade on v1.15 priorities (first deadline EOM)
 - [P0] security improvements -> no KEP/enhancement issue required
 - [P1] cleanup of the kubeadm dependencies - > [KEP PR in flight](#) (update to the API kep)
 - [P1] customization of running clusters
 - advanced configurations with kubeadm --> KEP started in GDOC (I will share soon link,)
 - change a running cluster:
 - no KEP/enhancement issue required (document what we have & use cases)
 - [P2] better support for async node join -> no KEP/enhancement issue required
 - [P2] machine-readable output ([KEP](#) in flight)
 - [P2] window support - It depends by Windows folks (There is a PR in flight by Microsoft folks, needs review)
 - [P2] improve test coverage -> no KEP/enhancement issue
 -

Wed 10th July 2019 - 9am PT, noon ET

Recording TODO

Attending:

- Lubomir I. Ivanov, Fabrizio Pandini, Ross Georgiev, John Harris, Yassine Tijani - VMware
- Leigh Capili - Weaveworks
- Marek Counts, Rafael Fernández López - SUSE
- Jorge Alarcon
- Arvinder Wander - Ericsson
- Yago Nobre - Nubank

Agenda:

- Add your agenda here
- [fabrizio] disambiguate “change the running cluster” (P1) and “customize”
 - “change the running cluster” = a new kubeadm workflow that will allow to modify settings of an existing cluster.
It is meant to replace error prone/manually crafted workarounds based on kubeadm phases, kubeadm upgrade without changing version, manual changes to the static pod manifests.
 - “advanced configurations with kubeadm” = a new kubeadm feature that will allow users to do *advanced* customizations today not possible by kubeadm. It is meant to increase flexibility, without adding more knobs to the kubeadm config API.
 - Do we agree on the definitions above?
 - Do we agree on focusing on the two problems separately?
- [fabrizio] brainstorming on “change the running cluster”
 - what is a change? ideally:
 - changing cluster configurations is the only thing in scope (settings that are in cross components)
 - everything else should “simply” restart with an updated component config
 - In practice component config is not there, and it should be defined what to take in the consideration
 - AI: This is an Epic. Let’s start to list use cases / bin issues
 - AI: For this cycle, let’s try to document how to make changes with the tools already there
- [ereslibre] Self hosting cleanup?
 - leave it there for a while - it’s under alpha and the docs are separate.
- [ereslibre] Upgrade UX
 - using the same command instead of “upgrade apply” and “upgrade node”.
 - [ereslibre] create an issue explaining what is desired and the proposed approach
- [arvinder] Dual-Stack Support in kubeadm
 - kubeadm didn’t have any e2e tests for network configs. I have added some - <https://github.com/kubernetes/kubernetes/pull/79033> - please take a look.

They are part of the dual-stack branch though the tests are for both single and dual stack.

- With dual-stack certain fields that were singular are now slices. These changes will trickle in as PRs are merged into master. Additionally, they are spread across different configs -- e.g. ClusterConfig.Networking, InitConfig.APIEndpoint, KubeproxyConfig, etc. For dual-stack validation, it's straightforward to check individual fields but it becomes tricky to check that if one field is in dual-stack "mode", the rest are all configured accordingly.
- [leigh(stealthybox)] discuss addon-installer integration
 - library to be checked in to k8s.io/addon-operators
 - standalone POC CLI using that lib
 - alpha feature for 1.16 ? TODO: KEP
 - coredns, kube-proxy, node-local dns
 - constraints
 - ordered install
 - shell out to kubectl for customize build and apply
 - demo of POC?
-

Wed 3rd July 2019 - 9am PT, noon ET

Recording TODO

Attending:

- Lubomir I. Ivanov, Fabrizio Pandini, Naadir Jeewa, Ross Georgiev, Yassine Tijani - VMware
- Rafael Fernández López - SUSE
- Jorge Alarcon

Agenda:

- Your agenda here?
- [lubomir] PSA: Kubernetes CI artifacts are broken due to changes in k/release.
- [ross] kubeadm public config API extraction
 - [POC PR](#), kubeadm config KEP to be updated (probably next week)
 - Find location - staging or k/kubeadm w/ manual copy?
 - Staging appears to be a clearer option
 - [yassine] the issue with staging is that we have to erase the contents of the kubeadm repository
 - [lubomir] jason suggested a separate repo for the types but prefer it being in k/kubeadm
 - [tim] +1 - further federation creates release issues
 - [fabrizio] +1

- [ross] AI: We should start with k/kubeadm and then perhaps start with staging, but want to avoid that transition at all costs.
 - [tim] what is the checklist of items for kind so that we don't need kinder
 - [lubomir] the checklist should be on kind's side but it's not being worked on, so kinder is the intermediate solution that works for us.
 - [tim] it's ok to have kinder for a period of time, but we need a checklist so kind can reach parity.
 - [lubomir] things we're doing in kinder are out of scope for the kind project.
 - AI: Take this offline to figure out in detail.
- What to do with defaulting consts?
 - [fabrizio] don't see a use case for consumers to access defaulting
 - [tim] might want to sync with chuck and jason with regards to the kubeadm bootstrapper
 - [fabrizio] checked what they are doing for kubelet & kube-proxy and they are not using defaults
- [lubomir] what is the plan for the files
 - [ross] we will need a couple of a PRs to be done in sync. Have all the things prepared in the kubeadm repo, and then PR to k/k
 - [lubomir] does it need to move to a directory first?
 - [ross] Only if this work doesn't span several releases, because vendoring users will be relying on the package locations.
 - [tim] so how are we planning to migrate vendoring users
 - [ross] the idea is to act as a drop-in replacement. we only expect people to use public types, and that should be fine. if they use constants or defaulting code, they will have problems.
 - [tim] if people are already vendoring today, don't know what the public policy is for breakage. we need to do a series of PSAs
 - [fabrizio] makes sense to do a small KEP with a plan of what to do and then do a public announcement.
 - [tim] AI: check with jordan as to what the policy is.
 - AI: Do a KEP with the migration plan
 - [fabrizio] does it make sense to update the config kep or make a new one for making the types vendorable?
 - [ross] keep everything in the same doc would be better. don't think we should GA kubeadm before this is done.
- v1beta2 only or both v1beta1 & v1beta2
 - preference is for both
 - [tim] I think you have to. bootstrap modification will need to be batched for v1beta3
- [lubomir] concurrent etcd join support needs approval:
 - <https://github.com/kubernetes/kubernetes/pull/79677>
 - We graduated to beta without this fix, so should we backport
 - [tim] I'm ok with not backporting this as we filed the bug as a known issue.
 - [lubomir] we said we were going to fix it in the release notes

- [lubomir] backport new flag:
 - <https://github.com/kubernetes/kubernetes/pull/79734>
 - [ross] don't think we should backport this as it will touch flags on the phases such that they won't be usable
 - +1, same logic.
 - [yassine] as long as it doesn't break something on the existing release, it should be fine.
 - [lubomir] we're currently missing signal for kubeadm above
 - [tim] this is not technically a backportable fix as it's not critical. policy is important so external consumers can reason about the release.
 - [rafael] +1 on not backporting
- [lubomir] walk the k/kubeadm backlog for Milestone 1.15, 1.16 and Next
-

Wed 26th June 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Tim St. Clair, Yassine Tijani - VMware
- Rafael Fernández López, Marek Counts - SUSE
- Ed Bartosh - Intel
- Jorge Alarcon

Agenda:

- vote on OWNERS file update:
 - <https://github.com/kubernetes/kubernetes/pull/79218>
 - got OK from Chuck and Liz to be temp-removed from reviewers.
 - add Rafael as approver?
 - add SataQiu as reviewer?
- Discussion about kubeadm dependency tree and cleanup

[P1] [lubomir, rosti, yassine, ed] cleanup of the kubeadm dependencies

 - Only allowing vendoring on kubeadm types is really P1, the rest is P2
 - tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1600>
 - this work will span across multiple releases and we need to establish sub-priorities:
 - P1 expose the kubeadm public api types. don't depend on the internal kubeadm type in conversions in external type packages.
 - we had a chat with rosti and fabrizio about this and the conclusion is that need a KEP.
 - P2 don't depend on internal kubeproxy kubelet types

- defaulting, validation, conversions?
 - rosti already has a WIP PR:
 - <https://github.com/kubernetes/kubernetes/pull/79223>
 - P3 further decouple from kubernetes/pkg
 - a number of WIP PRs are already up
 - P1 don't use the docker api in validators
 - other?
 - walk over the kubernetes/pkg dependencies at <https://docs.google.com/document/d/1oDp7fXhGHA66cpaKE7Sm4eF3w4G4uZrq8EhtSWBxfus/edit>
-

Wed 19th June 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Tim St. Clair, Yassine Tijani, Ross Georgiev - VMware
- Rafael Fernández López, Marek Counts - SUSE
- Yago Nobre - Nubank
- Fabrizio Pandini
- Ed Bartosh - Intel

Agenda:

- **kubeadm 1.16 planning meeting**
 - 1.16 milestone <https://github.com/kubernetes/kubeadm/milestone/13>
 - Next milestone <https://github.com/kubernetes/kubeadm/milestone/12>
- (if you have topics you want to propose/discuss, add them next to your name and link to a tracking issue)
- [P2] [Fabrizio, Rafael, Yago, Lubomir] better support for async node join (improve stability of cluster creation)
 - HA graduated to beta in 1.15
 - concurrent join of control-plane nodes does not work yet
 - experiments: <https://gist.github.com/neolit123/4ecd657175f17df0a8207e7ffb2c0ec4>
 - possibly needs to be patched in a 1.15 PATCH release
 - if the api-server is not ready we have a workaround to retry fetching configmaps:
 - <https://github.com/kubernetes/kubeadm/issues/1606>
 - LBs without https health check can have problems:

- <https://github.com/kubernetes-sigs/kind/issues/588>
 - [yago] Move Control Planes taint to kubelet config instead of markcontrolplane phase
 - <https://github.com/kubernetes/kubeadm/issues/1621>
- [P1] [Ross, Fabrizio, Marek] customization of running clusters
 - <https://github.com/kubernetes/kubeadm/issues/1379>
 - <https://github.com/kubernetes/kubeadm/issues/1581>
 - this is by far the most requested feature at this point, but we don't have a good set of user stories.
 - “kubeadm customize <target-module>” vs “kubeadm apply <target-module>” new command?
 - should be a replacement for “kubeadm upgrade apply --config”
 - [Fabrizio, Marek] needs a KEP, to establish targets, best practices and caveats.
 - possibly dependent on the next item (kubeadm upgrade apply phases)? the upgrade process refactor can lead to a cleaner shared backend between a customization command and upgrade
- [Punting] upgrade phases
 - <https://github.com/kubernetes/kubeadm/issues/1318>
 - this is for “kubeadm upgrade apply”, as “node” already has phases.
 - the most requested part of this is skipping addons on both init and upgrade.
- [Punting to next cycles when we get critical mass] v1beta3
 - it would be nice to accumulate a list of features to add and bump the config version, yet currently we only have one known/good proposal:
 - <https://github.com/kubernetes/kubeadm/issues/1600>
 - (although, this ticket was originally closed and i don't remember why)
 - anything else?
 - a reminder that we cannot go v1 without component configs such as one for the api-server becoming a reality.
- [P1] [Iubomir, rosti, yassine, ed] cleanup of the kubeadm dependencies
 - Only allowing vendoring on kubeadm types is really P1, the rest is P2
 - we promised a meeting about this - e.g. next week, full hour?
 - this work has already started with priority important/soon:
 - does the priority need adjustment and possibly set sub-priorities?
 - tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1600>
 - this work will span across multiple releases and we need to establish sub-priorities:
 - px? don't depend on internal kubeproxy kubelet types
 - defaulting, validation, conversions?
 - px? don't depend on the internal kubeadm type in conversions in external type packages
 - px? further decouple from pkg

- px? don't use the docker api in validators
 - other?
 - some items require major changes - e.g. use only external componentconfig types
- [P0] [Yago, Marek] security improvements?
 - kubeadm repo label search:
 - Some items are P0, some other are backlog, so we should prioritize first
 - <https://github.com/kubernetes/kubeadm/issues?q=is%3Aopen+is%3Aissue+label%3Aarea%2Fsecurity>
 - CIS compliance?
 - <https://github.com/kubernetes/kubeadm/issues/683>
 - this is a >2 step process:
 - support running kubebench in e2e?
 - integrate with e2e framework
 - needs regular maintenance
 - add e2e test jobs for it and kubeadm
- [P2] [Ed, Andrew] machine-readable output
 - <https://github.com/kubernetes/kubeadm/issues/494>
- [P2] window support
 - It is necessary to get a meeting with window people and try to shape out a plan on this at the beginning of the cycle. Priority will be adjusted accordingly
 - kubeadm side is mostly done for the alpha
 - in the hands of sig-windows at this point
- [P2] [Fabrizio, Lubomir] test coverage
 - upgrade kind dependency in kinder (with containerd)
 - preserve in kinder support for docker
 - consolidate current tests
 - make naming more consistent
 - use always top of branch
 - fix external-etcd tests (make selection of etcd release smarter)
 - add more tests?
 - kube-dns?
 - phases?
 - kubeadm utility commands (token, config upload, alpha certs etc. etc.)
 - [tim, lubomir, Fabrizio] evaluate /cluster directory tests.
- [P3] [tim] evaluate the package install / container story.
- we need more reviewers for kubeadm
 - add @SataQiu as a reviewer
 - <https://github.com/kubernetes/kubernetes/pull/79167>
 - Sponsors: @neolit123 @roster
- we need more contributor for kubeadm
 - Improve issue triage/make help wanted issues more actionable?
 - EU/China friendly meeting?

- organize code review sessions?

Wed 12th June 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Tim St. Clair, Ross Georgiev - VMware
- Ed Bartosh - Intel
- Yassine Tijani - VMware
- Leah Hanson, Danny Berger - Pivotal
- Rafael Fernández López - SUSE
- Arvinderpal Wander -- Ericsson

Agenda:

- [lubomir] share your kubeadm logo ideas if you haven't already:
 - <https://docs.google.com/document/d/1oAzGFaUPdZSDwsFsSzK1WMY8rfUoj6GE2JIN3Yv02G0/edit>
 - <https://github.com/kubernetes/kubeadm/issues/1588>
- [lubomir] State of kubeadm in 1.15:
 - All of our docs PR merged.
 - keeping this issue open until release
 - <https://github.com/kubernetes/kubeadm/issues/1577>
 - We are promoting HA to beta:
 - <https://github.com/kubernetes/kubeadm/issues/1567>
 - The state is good, with the exception of concurrent control-plane node join, which we might resolve in a 1.15 PATCH release.
 - Need retry logic for etcd member joining (error if two try to join at the same time); waiting to see if etcd maintainers have different advice – see next agenda item (“etcd blackout”)
 - Test signal is in a good state:
 - <https://k8s-testgrid.appspot.com/sig-cluster-lifecycle-kubeadm>
 - The new external etcd jobs are failing
 - fix PR:
<https://github.com/kubernetes/kubeadm/pull/1615>
 - upgrade tests are flaky:
 - potential fix:
<https://github.com/kubernetes/kubeadm/pull/1605>
 - Code thaw probably starts today EOD PST (was delayed).
 - Cherry pick deadline is 13th of June (tomorrow)
 - Release is on Monday 17th (delayed)

- [lubomir] etcd blackout and k8s API object fetch timeout handling WRT to concurrent joining of nodes.
 - We are doing experiments when joining multiple control-plane/worker nodes concurrently. It mostly works, but kubeadm needs some patches.
 - Related PRs/issues:
 - <https://github.com/kubernetes-sigs/kind/issues/588>
 - <https://github.com/kubernetes/kubernetes/pull/78915>
 - Add exponential back-off when adding new control-plane member. Seems to solve etcd becoming unhealthy.
 - Seeing occasional random failure when joining worker node, seems related to config map update, also solved by retry logic.
 - [rafael] I see different things here. What happens when we're joining two control planes? They both see the cluster as healthy and then the both start? I have trouble seeing how the work problems have to do with etcd?
 - [ivanov] The worker problems were very rare. Under idempotency – I created a patch that creates clusters, tears them down, creates clusters, and I only caught it a couple again – maybe the load balancer/config map problem. I'm going to write a guide to reproduce this in Kinder. You need a more powerful machine to run this setup with 3 control plane nodes.
 - [rafael] We need to agree on whether we want the retry logic everywhere. Kubeadm is a one-shot application. Failures should be pretty fast. How long should we retry? Should we assume that we shouldn't reach a control plane that isn't healthy yet? Makes our codebase more complex.
 - [tim] This code already exists in the client; the decay-and-try-again on certain conditions is built-in (like too-busy). We may be getting other errors, like weird etcd errors, then we'd have to retry higher up the stack.
 - [rafael] We've been retrying unauthorized errors, so we're discussing retrying all errors. But even retrying unauthorized errors is pretty weird. This comes down to the load balancers being misconfigured.
 - [tim] Jordan and I fundamentally disagree. Fail fast. An error the percolates up that you don't understand means you should fail. Retrying on a join seems ok; retrying on a config map with a broad set of errors seems not cool.
 - [rafael] We want to fix for 1.15. Commit says "Please revert this". We need to find the root cause.
 - [ivanov] Some of the errors from the api-server are very confusing and I think we'll need help from api-machinery to root-cause this.

- [tim] Depending on how fast you create things, there's a different code path in the api-server.
- [ivanov] Depending on the state, it throws an error. I don't know how to implement the logic of concurrent joining and all the errors we might get.
- [tim] You're doing this in a test env where it's very fast and all on the same machine. This is different from most deployments. Doing masters and workers at the same time has not been tested; adding everything at once has not been tested. It's a good user story to add control-plane nodes quickly, but it's not been tested.
- [ivanov] Let's hold this PR and experiment and find the actual problem. We need this workaround, but also find the real problems.
- [tim] I'm find with the patch for the control-plane with the Todos, but the node config map one gives me the willies.
- [ivanov] Maybe we can get help with figuring out what the repro cases are?
- [rafael] I don't know to what extent we can fix this in kubeadm. I don't know if we have the control to fix this. I don't want to be retrying everywhere. Finding a root cause is fine, but if it's the load balancer, I don't want to retry everywhere.
- [ivanov] We'll also need to setup a bunch of load balancers and see how they work. In Kinder, the change from haproxy to nginx broke something and no one knows how or why.
- [lubomir] decouple the public kubeadm API types v1beta1/v1beta2 from importing `k8s.io/kubernetes/cmd/kubeadm`
 - The sig-arch code organization project wants to move kubeadm out of k/k (and into staging) which isn't quite feasible yet.
 - Cluster API providers (and possibly other users too) however have requested us to make the kubeadm types public / importable without k/k.
 - We can get this ready in some sort of state, but it's also interesting what we can do with k/kubeadm as the potential home for that WRT to publishing bot, staging etc. There are some questions there.
 - [tim] We should set import boss on the types to help us not import things from k/k
 - [ivanov] they've asked up to be importable without k/k
 - [tim] I'm ok with the types being importable without k/k; most of the other stuff can't be imported, but the types can be.
 - [ivanov] I have questions.
 - [tim] I think we need to figure out a mishmash of content in how we're going to do this. The dependency graph problems are real. If you import types, you shouldn't drag in other parts of k/k.

- [ivanov] Kube-proxy and kubelet already do that, but yea. I'll ask my questions tomorrow. Rosti - what do think of our api package imports? I think something is wrong?
 - [rosti] I don't think there are that many problems. We can fix them in the next release. We can as them as new constants.
 - [ivanov] Yea, it's just constants. We can start this work to allow the types to be importable separately.
- [Arvinder] <https://github.com/kubernetes/kubeadm/issues/1591>
 - Found bug while deploying a local cluster. Service CIDR not being passed onto controller manager.
 - [tim] I think we've looked at PR, but we're at the end up 1.15, so we can't merge until after code freeze.
 - [ivanov] The change is good, but after tomorrow we'll open the master branch again. Ideally more people will look at this because it's a new use case, but consensus is the change is good.
- [Arvinder] Dual-Stack integration for kubeadm
 - There's been good progress in k8s; I'd like to contribute to kubeadm support. I have a local branch; I'd like to submit a WIP PR for review.
 - [ivanov] IPv6/IPv4?
 - [arvinder] Yes. There was a KEP, and it's now being implemented.
 - [ivanov] The one by Loki from Microsoft? Let's created a ticket in kubeadm to track this. I know there's a developer in the Kind project, which uses Kubeadm, who is integrating the work as well. You can also share you work on that ticket.
 - [yassim] I've been reviewing the dual-stack work, and would love to review your PR. Just ping me.
 - [tim] Just make a WIP PR and assign Yassim. Mark it as WIP to prevent merging.
 - [ivanov] There might be a wall of comments because we're not familiar with the ipv6 changes. I'll also bring in this Kind dev.
- [Ed Bartosh] POC for structured output: <https://github.com/kubernetes/kubernetes/pull/78764>
 - There's a KEP for this. This is for bringing YAML/JSON/GoTemplate output to kubeadm. I have a PR for one of the proposed way to solve this. I want to share my feelings, my findings.
 - The good things is that it's separated. For maintenance, it is most flexible.
 - The not good thing is the duplication. The token list is almost completely duplicated to avoid changing the bootstrap token API, and to make it a runtime object.
 - [tim] We're blocked on (??) PR on to rip out on api-machinery and have been blocked for 2 months now. When that happens, you'd be unblocked because it would be a separate type.
 - [bartosh] Customers of this API – when they get this JSON – they will get the structure of this type, which is separated, so they will end up

converting it into the kubeadm type, to send to the kubeadm api or functions. That seems not good.

- [ivanov] Clarify users importing bootstrap token api?
 - [bartosh] Most would require bootstrap token from api group. They would have token from output api group, so would converting from one to another to join cluster, if they want to use api.
 - [ivanov] What happens if we graduate the token to a runtime object and we output to that directly?
 - [bartosh] That would be better, but for somethings, like to output list of images, we don't have the same structure.
 - [ivanov] We should try to keep the api/output group separate. And keep the object there. When we have a kubeadm api someday we can use this.
 - [bartosh] If we keep this api group, and keep everything that's supposed to be here, that's fine, but if we split between this and kubeadm api group, then that would be bad.
 - [ivanov] Maybe only bootstrap token should be exception. Everything else should be under output.
- [rafael] Should we version output, so we can convert to internal type?
 - [bartosh] As we have internal type and external type, and conversion boilerplate, I'm not how to make the boilerplate generate between api groups. I think this would require at least some manual changes.
- [ivanov] In terms of the output group, I think we should deal with not much api-machinery stuff – I think we should keep the conversions minimal. I don't know why we should convert to internal type. Essentially, the functions can use an external type. The functions are going to be public, so they should accept external type.
 - [rafael] Functions should use internal types. Leave it to the boiler plate to convert from external to internal. Functions always use internal.
 - [bartosh] There will be some conversion because there are two API groups.
- [tim] We still need to get the externalized runtime object for bootstrap tokens. We're blocked on that. We need api-machinery to unblock that. I don't know how we get that moving.
- [ivanov] We don't even have a location for the bootstrap token object. Are we even going to have a problem with type conversions between these groups?
 - [bartosh] We have two almost identical structures. If you get the JSON structures, and marshal it, and then you want to use the kubeadm apis, then those want kubeadm types from another group, then you'll end up with some silly code to convert.

- [ivanov] We shouldn't duplicate the types at all. If the bootstrap token type is an exception, fine. Use it as a direct type. But the other types, like image list, they should be only an output.
- [bartosh] I don't know how many fall in this category. Maybe only bootstrap token. Maybe not a problem. Could be lots, could be a problem.
- [ivanov] I don't think we have that many objects. Maybe 7 to 10? This is not a big concern.
- [bartosh] Should I focus on other things, since bootstrap token is blocked?
 - [ivanov] Yes, start with images. I'll show you the bootstrap token comment. You see this repo? They told us to put the type here; the README says don't put API types here.
- [bartosh] I started on upgrade plan, but it doesn't touch the API. I added Andrew to the reviewer list. I'd love more reviewers; thank you to those who already did.
 - [ivanov] I think Andrew had some streaming ideas.
 - [bartosh] Andrew had ideas about avoiding buffered output and only output unbuffered. This code outputs one piece at a time.
 - [ivanov] When I type "kubeadm tokens list" or "images list" with an output of json, do I wait for them to buffer?
 - [bartosh] No, you'll have them as they come.
 - [ivanov] What if I do something that is not a list? Like join or upgrade plan?
 - [bartosh] upgrade is a list.
 - [ivanov] So only for the table is buffered. The rest are unbuffered.
 - [bartosh] Yes, that's true. So far, I've gotten the idea to output only the list. I'm not sure that for init we need structured output at all.
 - [ivanov] For the command at the end.
 - [bartosh] That can be obtained differently. You can't repeat init, but you can run another command any time.
 - [ivanov] But they can share a backend. Kubeadm token does not handle the certificate key.
 - [bartosh] I would propose not to mix certificate output with the kubeadm join output. But we can get it without running init. We can get it anytime. Just a question where to put this code.
 - [ivanov] It's doable for joining worker nodes, but for control plane nodes you're missing the certificate key. Even if you need to create a separate output object specifically for certificate key, that's fine. We don't have one for init output that contains a certificate key.
 - [bartosh] The question is where to put all the unstructured output?

- [ivanov] We can still accumulate the output we don't need in a string, which can be added to a specific field. The rest of the JSON will be what we need, and this field will be a buffered string. Andrew has ideas about this. Maybe there are ways to hack around this to print the output and write to the JSON at the same time. We want to see the output and buffer the lines to print at the end.
 - [bartosh] I'll talk to Andrew.
- 3min early finish! :)
 - Ask questions in #sig-cluster-lifecycle

Wed 5th June 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Ross Georgiev, Naadir Jeewa - VMware
- Marek Counts, Rafael Fernández López - SUSE
- Ed Bartosh - Intel
- Leah Hanson - Pivotal
- Martin Dietze - Dermalog
- Leigh Capili - Weaveworks
- Lucas Källdström - luxas labs (Contracting for Weaveworks)

Agenda:

- [lubomir/naadir] kubeadm CIS report discussion (**full hour?**)
 - <https://discuss.kubernetes.io/t/kubernetes-and-security-benchmarks/6635>
 - link to follow up with tickets: <https://workbench.cisecurity.org/>
 - [naadir] Summarizing discuss thread
 - kubernetes & docker results vs clusters brought up with kubeadm
 - seeing lots of false positives
 - they're checking flags when we've moved to component config
 - there's things we can improve
 - if we want to comply to have a secure cluster by the CIS definition, we need to modify the static pod definitions.
 - Some suggestions in Discuss for how to implement this.
 - Longer discussion on Hard Problems listed in Discuss
 - Control 1.1.1 Disable Anonymous Authentication
 - CIS wants `anonymous-auth=false`; kubeadm does not work with this.

- Fix: add a new port for unauthenticated /healthz endpoint
 - Not a huge risk since we do check auth other than health & metrics
- [tim] But enabling anonymous-auth is fine. We don't recommend that.
- <screenshare of issue 43784>
- [liz] In the latest draft, this test is not scored. If you're using RBAC auth, then it's ok to use anonymous-auth for health/metrics.
 - [tim] Is there a timeline for the update?
 - [liz] I think we're aiming for 1.15, but there's an issue with the CIS spec for unusual envs like OpenSpec. We're adding vendor-spec versions, but only after we add all the common things. It should come out without a few months, but that may be optimistic.
- Control 1.4.12 Ownership of etcd directories
 - /var/lib/etcd should be owned by etcd:etcd
 - kubeadm objection: we run etcd as root
 - [liz] Make an issue on CIS asking for "root or etcd"
 - [tim] Since we don't install the service and run the service as etcd, the whole kubelet runs as root, I think it makes sense to change this. This rule assumes a lot about your deployment.
 - [liz] We may also move this to a different section since not all deployments have etcd on master nodes.
 - Action Item: File issue on CIS for allowing root:root
- Control 1.7.40 Use of host network namespace
 - kubeadm intends controlplane components to live on host network, so this we don't want change anything
 - [tim] how are we treating the CIS test suite?
 - [naadir] sec teams at user companies tend to treat it as gospel. We'll have to make a case for why each test is red to people without a lot of k8s context.
 - [tim] We don't want to change this.
 - [liz] Spec does allow some leeway – "do not generally commit pods to use the host network namespace". Pods that need host network should be defined in a separate PSP. If setting up the PSP is too onerous – or turning on PSPs makes it too difficult to get anything to run at all – ... there is some leeway.
 - [tim] Can you even have a PodSecurityPolicy for static manifests?
 - [ivanov] PSPs for static manifests are not that easy. The admission controllers conflict. You have to start the api-server with the node admission controller, then apply the pod security policies, then add the admission controller, then restart the control plane.

- [tim] That's a lot of complexity for a not of lot of benefit.
 - [naadir] You can do it, it is possible. We do, do that.
- [liz] Another approach is to modify the rule in the spec to carve out specific namespaces like kube-system.
 - [tim] That makes sense. For self-hosting, that makes a lot of sense. Crafting a PSP seems really weird for this user story.
 - [ivanov] This falls in between the MVP cluster and the cluster users want to make one step more secure. We need to evaluate whether the MVP cluster is insecure from the perspective of host network for containers.
 - [tim] It's a control-plane node, it needs to be different.
 - [naadir] You have to apply PSPs if you're using static pods. We don't have to enforce that with kubeadm – all the CNIs that deploy with daemonsets today use hostnetwork as well anyway.
- Control 2.1.10 Specifying TLS cert and private key file paths
 - kubeadm uses TLS bootstrapping; this is mutually exclusive with --tls-cert-file and--tls-prive-key-file
 - CIS is checking for those specific argument flags; bug in CIS
 - [Link to CIS issue](#)
- Control 2.1.13 Kubelet Server Certificate Rotation
 - CIS wants automated Cluster Renewal turned on
 - kubeadm believes this is a security hole and has kubelet self-sign their own certificates
 - some master-nodes actually rubber-stamp all CSRs and that's also insecure
 - [ivanov] self-signing kubelet certs blocks metrics server. We should write a troubleshooting guide.
 - [naadir] If you look at prometheus operator today, they have a thing specifically for kubeadm. We do self-signing once today and never renew. We've definitely seen issues, so troubleshooting.
 - [naadir] The issue is that we can't enable the feature gate for serverlet certificate rotation when the feature is incomplete.
 - [ivanov] I agree.
 - [liz] I don't understand the objection entirely, but I suggest raising a ticket in CIS to discuss in more depth.
- Docker Benchmark 2.8 User Namespace Support
 - Not available in kubernetes
 - [naadir] It's not clear in the output where this is applicable to, so users think k8s is insecure, but this isn't about docker.
 - [tim] So is this a branding/messaging/tool-usage problem?
 - [naadir] Yes, some tests are reasonable, but not all. It would be nice to have a subset of container runtime tests that are appropriate when using k8s.

- [liz] Same people in CIS maintain k8s and docker benchmarks. I'll chat with the authors – maybe we can specify which tests are inappropriate for k8s or split out a docker-in-k8s benchmark.
 - [tim] All the remaining failing “docker” tests are ambiguous due to the Docker Swarm use-case.
 - [liz] Have to run; thanks for looking at CIS benchmark.
- [naadir] Scanning methodology: how do you know the configuration of the control-plane without scrapping the config files? This is not just a kubeadm feature; this is something tools need to be able to do – interrogate what flags are being set in a cluster.
 - [lucas] ComponentConfig. That's one of the benefits. The component can source the config however it likes, and then exposes it at /config on its http endpoint.
 - [naadir] once we make that transition, we'll want to update the CIS benchmark to use that endpoint.
- [tim] When we talk to CIS, should we spell out that the marketing needs to spell out that this is advisory, etc and not gospel? I think we need to help the branding and awareness of the community.
- [naadir] And all of our containers run as root. Most of these don't need to run as root. Is there a plan to fix this?
 - [tim] I think we can do that. Especially if there are user namespaces in the future. I don't think there are restrictions for most components. Only the kubelet needs root.
 - [naadir] We even run pause as root. What mechanism could we provide to add a SecurityContext to the static pods?
 - [lucas] Leah sent a PR; couldn't make it into this release. Looking at it for next release. Doesn't kube-proxy require root?
 - [naadir] Yes – but you might be able to get away with other options.
- [tim] I want to make sure we're doing the right thing. One thing we've done previously is an RBAC constraint analysis – to get a sane list of rules, you basically had to enable *. And we basically want to make sure we're giving users a good feedback loop.
- [ivanov] CIS is complaining that kubeadm isn't configuring various things because we're using ComponentConfig rather than flags. There are so many things that are just extra args.
 - [naadir] There are many things we don't want to enable by default – they're for high-security environments. We just need to make sure people *can* configure them – and we do.
 - [tim] We need to talk to CIS about branding and also write up a doc that summarizes details like the discuss post after a loop of communication with CIS team.
- [lucas] Are we willing to help out with kubebench? What are we going to use to verify and have a set of bench practices? To have an easy tool to use?
 - [tim] It should be built into the end-to-end test if this is the recommendation. It's weird if the tests aren't part of the community.

- [lucas] end-to-end tests. We need to create a tool that supports the use-case we want.
 - [tim] thoughts?
 - [ivanov] I'm sure that there are lots of people that can contribute to this e2e test idea, to help Aqua combine the CIS tests into k8s.
 - [tim] John can help.
- [leigh] User facing documentation on how to set non-defaults to become CIS compliant.
 - [tim] I agree – that's what I meant by the documentation after the loop with CIS.
 - [ivanov] Yep. Some items are high-risk, some are not. We want to clean up before writing the doc.
- [leigh] request for comments on addon management mechanisms
 - Want to allow tools to source add-ons from similar place, in same format.
 - The plan:
 - kustomize for packaging
 - support for templating, namespace selection, standardizing labels/annotations, swapping images, patching
 - Want to build customization for ComponentConfig yaml
 - [tim] You're talking shell out? Or incorporating package? The problem I see is "import universe".
 - [leigh] There's an effort in kustomize to pull everything into pkgd. In terms of simplicity, we don't have serverside ready for everything you'd want it for. I think it's simpler now to shell out to kustomize, but it doesn't have to be that way.
 - [leigh] Maybe the kustomize templating inside of kubeadm and then passing the manifest...
 - [fabrizio] I don't want to import too many things. We're already talking about moving kubeadm out.
 - [ivanov] We have to figure out package situation again. We already recommend kubectl, but it's ok to say you also have to install another tool (kustomize) –
 - [tim, lucas] It's already built-in to kubectl!
 - [ivanov] I mean they have to install kubectl *on the nodes*.
 - [leigh] yes, that's new.
 - [tim, lucas] In the tutorial, we say to install it. As a package.
 - [leigh] The first step is kustomize support, but if we ever do serverside apply, then we might need to vendor in first.
 - [ivanov] Or we might shell-out forever.
 - [ivanov] Before doing this, do you think we should write a prototype tool to play with the templating?

- [leigh] Yes, that makes sense. We need to experiment with the format and install experience. Until we put the CRDs into the cluster, we can't really validate them. There's a lot of stuff to be proof'd, so doing it separately from kubeadm makes sense. Maybe in the add-ons, then somewhere else, then vendored into kubeadm. Doing it externally might make it more reusable.
 - [ivanov] External from the start sounds right.
- [fabrizio] Why is .. not managed?
 - [leigh] That's the second point. Because they're not versioned – they add-ons are commit-by-commit. Kubernetes versions are tested thoroughly with specified add-on/package versions. Docker image versions are bumped. CoreDNS/kube-proxy need to be tested together.
 - [fabrizio] Might help to draw some diagrams between kubeadm/kops and the manager tool to understand the add-on management tool – where kubeadm has to create the CRD and what ever is needed to have the add-on manager running and then install one add-on and then upgrade one add-on.
 - [leigh] The add-on manager is likely to be a library vendored into kubeadm.
 - [ivanov] If we decide to keep kubeadm coupled to kubernetes release, will add-ons bundle also be coupled?
 - [leigh] Yea, you need to pull add-ons in via reference to k8s version.
 - [tim] There's a bunch of logistics that will take time to make clean. Let's start simple – CoreDNS for example. Lots of UX workflow – air-gapped installs, etc.
 - [ivanov] CoreDNS is complicated this release.
 - [leigh] We have existing work for installing things like CoreDNS.
- [leigh] I really wanted to vet these ideas. Thanks for the feedback.
 - [tim] I suggest a KEP. What's the timeframe? 1.16?
 - [leigh] 1.16 for alpha and vendoring the library into kubeadm.
 - [ivanov] [Kubectrl is going out of tree, so shelling out might be more complex.](#)
 - [lucas] Shelling out should be ok – because it's just calling a binary. Please link to the discussion.
 - versioning/lifecycle /w kubeadm dev
- [lucas] kubeadm is a required competency for CKA v1.15 :tada:
 - <https://github.com/cncf/curriculum/pull/30>

Wed 29th May 2019 - 9am PT, noon ET

Attending:

- Lubomir I. Ivanov - VMware
- Andrew Kutz - VMware
- Ed Bartosh - Intel
- Leah Hanson, Danny Berger - Pivotal
- Yago Nobre - Nubank
- Rafael Fernández López - SUSE
- Martin Dietze - Dermalog

Agenda:

- [lubomir] kubecon - quick de-briefing
- [lubomir] feature freeze is on 30th
 - doc PRs are WIP; tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1577>
 - review pending [PRs at k/k](#)
- [lubomir] coredns upgrade broke upgrade test signal:
 - fix PR is here but needs some changes:
 - <https://github.com/kubernetes/kubernetes/pull/78033>
- [lubomir] hardcoding client CA path:
 - <https://github.com/kubernetes/kubernetes/pull/78079>
- [lubomir] expose /etc/kubernetes from config? (too late for this cycle)
- [ereslibre] Are there any e2e tests missing for HA reaching GA?
- [akutz] Touching base on kubeadm machine output kep
 - <https://github.com/kubernetes/enhancements/pull/1054>
-

Tue 21st May 2019 - 6pm GMT F2F @ KubeConf

Attending:

- Fabrizio Pandini
- Lubomir I. Ivanov, Ross Georgiev, Bert, Timothy VMWare
- Ed, Alex Intel
- Raphael, Suse

- Andrew Saber, Anton

Agenda:

- what about the team around kubeadm (brainstorming)
 - Generally good fee but we can always do better
 - People on rotation for faster issue triage/first screening on PR
 - Improve new contributors onboarding
 - Chinese friendly meeting (let's see if it works)
 - Leverage on "good first issue"
 - Give better visibility to contributor tools/documentation
- What's next (brainstorming)
 - Stream in the config file and stream out the output + the artifacts e.g. usage of the phases to create certs for many cluster
 - Kustomize vs dry-run vs machine readable output have a common denominator (all of them requires some sort of control of the output of kubeadm)
 - Cert expiration needs to be documented. Key rotation should be addressed as well
 - Changes the cluster is problem. Let's try to scope it down to some use case and see if using upgrades can be an option (pay)
 - Changing certs
 - Changing static pod flags
 - Kube-proxy configuration
 - Kubelet config
 - Etcd version/client is changing
 - Distroless is going to break how we do etcd health check
 - CIS report is going to be P0
 - Deep dive
 - Get E2E CIS signal
 - Kubeadm as library vs machine readable output (are they in overlap?)

Wed 15th May 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Fabrizio Pandini
- Lubomir I. Ivanov, Ross Georgiev , Tim St. Clair, Yassine Tijani - VMware
- Jorge Alarcon
- Marek Counts - SUSE
- Ed Bartosh, Alexander Kanevskiy - Intel
- Ace Eldeib
- Leigh Capili -- Weaveworks

- Yago Nobre

Agenda:

- [fabrizio] PTAL: What is missing for HA to beta [1567](#)
 - `kubeadm upgrade node config` and `kubeadm upgrade node experimental-control-plane` >> merge into `kubeadm upgrade node`
- [fabrizio] PTAL: Certs renewal [77780](#) / check-expiration [77863](#) improvements
 - What is an acceptable renewal period (time before expiration when automatic certs renewal should be triggered)
 - 3/4 Months
 - kubelet.conf renewal probably won't make 1.15
- [lubomir] PSA: spoke with Naadir about CIS reports. he promised to join the office hours on the 5th of June.
 - Yassin to look at CIS benchmarks and start to investigate
 - Naadir to create a public document

Wed 8th May 2019 - 9am PT, noon ET

Recording: TODO

Attending:

- Rafael Fernández López, Marek Counts - SUSE
- Jorge Alarcon
- Tim St. Clair, Lubomir, Ross - VMware
- Thomas Strömberg - Google
- Ed Bartosh - Intel
- Fabrizio Pandini

Agenda:

- [Fabrizio] kubeadm upgrade tests are up / kubernetes/anywhere deprecation ongoing
 - **YES/NO** run frequency will be harmonized on every **4h** for master, every 12h for an older releases
 - **YES/NO** notification will be harmonized and everything will send notifications after failing for 24h
 - **YES/NO** to create a new "sig-cluster-lifecycle-kubeadm" dashboard
- [Fabrizio] How to debug CI errors for kinder jobs:
 - each job runs a sequence of tasks, where a task is a command + args
 - tasks are defined [here](#) (same name of the jobs)
 - build-log.txt gives you an overview of tasks completed/failed
 - each task generates a separated log under artifacts/
 - under artifacts/ there are also E2E test output and cluster logs
 - you can easily reproduce locally:

- the entire job using `kinder test workflow ci/workflows...`
 - task by task by using `kinder test workflow ci/workflows... --dry-run` and then running one command at time
- AI: Document somewhere
- [akutz, bart0sh] machine/structured output:
 - <https://github.com/kubernetes/enhancements/issues/1053>
- [Fabrizio] How to give users awareness about certs expiration
 - **YES/NO** Add a new kubeadm alpha certs check-expiration command
 - **YES/NO** Upgrade plan will print a table with cert Expiration, highlighting what is "going to expire" (add only a message)
 - **YES/NO** "going to expire" = certs expire in the next 180d (the renewal window)
 - **YES/NO** make upgrade to renew only what is expiring
- [Fabrizio] Other CI tests / kubeadm configuration variants test
 - A or B kubeadm HA test
 - A) change standard test/upgrade test/skew test adding another control-plane node
 - B) create a new dedicated set of tests for HA
 - **YES/NO** add a test for automatic copy certs
 - **YES/NO** add a test for external-etcd
 - **YES/NO** add a test for kube-dns
- [Fabrizio] What is missing for HA going BETA
 - E2E test (see above)
 - **YES/NO** "--experimental-control-plane" flag to be renamed "--control-plane"
 - **YES/NO** "--experimental-upload-certs" flag to be renamed "--upload-certs"
 - **YES/NO** "kubeadm upgrade node experimental-control-plane" to be merged into "kubeadm upgrade node"
 - others?
- [leigh] quick question: Is there a plan to make containerd default? Timeline?
 - no current plan

Wed 1st May 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Ben Moss - Pivotal
- Jorge Alarcon - Vividcortex
- Kalya Subramanian - Microsoft (@ksubrmnn)
- David Schott - Microsoft (@daschott)
- Patrick Lang - Microsoft (@patricklang)
- Michael Michael - VMware
- Lubomir I. Ivanov - VMware
- Marek Counts - SUSE

- Yago Nobre - Nubank

Agenda:

- [Tim] Skip phases on upgrade <https://github.com/kubernetes/kubeadm/issues/1318>
- [lubomir] Disable static pods for worker nodes by default:
 - Proposed by Josh Rosso, +1ed by Duffie
 - <https://github.com/kubernetes/kubeadm/issues/1541>
- [ksubrmnn] Kubeadm for Windows
 - <https://github.com/kubernetes/enhancements/pull/994>
 - tracking issue at k/kubeadm:
 - <https://github.com/kubernetes/kubeadm/issues/1393>
- [patricklang] Q if time: minikube and CNI? Is this implemented for flannel? No clear docs on how to configure it but some binaries are there.

Wed 24th April 2019 - 9am PT, noon ET

Attending:

- Lubomir, Jason DeTiberus, Tim St. Clair, Ross Georgiev - VMware
- Jorge Alarcon
- Ed Bartosh - Intel

Agenda:

- [fabrizio] the first [kubeadm upgrade test](#) is in test grid! there are still some refinement to do; next week we will demo how to debug/reproduce locally
- [fabrizio] certs renewal. how to proceed?
 - **YES/NO** During upgrade kubeadm should upgrade all certs in pki folder <https://github.com/kubernetes/kubernetes/pull/76862>
 - renewal driven “by component”
 - @maullion comments
 - **YES/NO** During upgrade kubeadm should upgrade all certs embedded in kubeconfig files
 - What about admin.conf (renew, notify the user)
 - What about kubelet.conf (skip, see kubelet cert rotation)
 - **YES/NO** Upgrade certs during upgrade should be optional
 - **YES/NO** kubeadm certs renew should update everything, but user is responsible to restart components
- [ross] PR for HA cert key in v1beta2 config:
 - <https://github.com/kubernetes/kubernetes/pull/77012>
 - this PR should be merged first:
 - <https://github.com/kubernetes/kubernetes/pull/76710>

- UX: Single option in NodeRegistrationOptions vs couple of options in InitConfiguration & JoinControlPlane
- [lubomir] there is a PR that enables wildcards for apiserver cert SANs
 - <https://github.com/kubernetes/kubernetes/pull/76920>
 - do we want that?
- [lubomir] we merged a JoinConfiguration validation bug fix:
 - ticket: <https://github.com/kubernetes/kubeadm/issues/1525>
 - PR: <https://github.com/kubernetes/kubernetes/pull/76923>
 - should we cherry pick - vote?
- [lubomir] (if there is time) a brief overview of the Windows / kubeadm KEP:
 - [WIP doc](#)

Wed 17th April 2019 - 9am PT, noon ET

Attending:

- Tim St. Clair, Lubomir, Rosti - VMware
- Fabrizio
- Yago
- Merek
- Alexander Kanevskiy, Ed Bartosh - Intel

Agenda:

- [lubomir] Reviewer promotion proposal from @SataQiu
 - <https://github.com/kubernetes/kubernetes/pull/76591>
- [tim] Blocked by release team changes.
 - ...
- [ross] v1beta2 KEP and initial PR are up
 - <https://github.com/kubernetes/enhancements/pull/969>
 - <https://github.com/kubernetes/kubernetes/pull/76710>
- [fabrizio]
 - Completed all the changes needed in kind and are working on the 1st job.
 - Can work on the upgrade issues.

Wed 10th April 2019 - 9am PT, noon ET

Attending:

- Fabrizio Pandini
- Tim St. Clair
- Marek Counts - SUSE

- Jorge Alarcon
- Leigh Capili - Weaveworks
- Dylan Cooper - SUSE

Agenda:

- [fabrizio] SCL meeting for new contributors
- [fabrizio] update on the work for creating E2E test for kubeadm-upgrade
 - [overview](#)
- [Dylan] Security Engineer for Suse
 - KubeBench issues would be worthwhile to evaluate.
 - Security guide would also be useful.
 - Reach out to sig-auth group as well
 - We should discuss with the CNCF as well.
- [tim]

Wed 4th April 2019 - 9am PT, noon ET

Attending:

- Tim St. Clair, Ross Georgiev - VMware
- Marek Counts, Rafael Fernández López - SUSE
- Alexander Kanevskiy - Intel
- Jorge Alarcon

Agenda:

- [tim+Merek] PSA walking through packaging tomorrow.
 - 9:30 CDT - kubeadm zoom room
 - POC: https://github.com/Klaven/kubernetes/tree/artifact_generation
- [ross] Walk through v1beta2 proposal changes.
 - <https://docs.google.com/document/d/1XnP67oO1i9VcDIpw42IzptnJsc5OQM-HTf8cVcjCR2w/edit#>
- [lubomir] proposal coming from cluster-api folk to taint worker nodes:
 - <https://github.com/kubernetes/kubeadm/issues/1487>
 - [tim] no consistency here between labels, probably Cluster API should do this; should check for security issues
 - AI: [tim] sync with Lucas and Taufen on this one
- [lubomir] PSA, kubeadm/windows proposal is out in gdocs form:
 - https://docs.google.com/document/d/1yaT7K85qMvZD7Q-ejWHBko1fgGaeGtjGLEZZ_Bz63VA/edit?usp=sharing
- [lubomir] packages are still broken:
 - PR to fix CNI: <https://github.com/kubernetes/release/pull/688>
 - user reports
 - <https://github.com/kubernetes/kubernetes/issues/75701>

- <https://github.com/kubernetes/kubeadm/issues/1481>
 - ^ 1.13.5 kubeadm installs 1.14.0 kubelet
 - /shrug
 - AI: Sync after this meeting
- [fabrizio] fix-kubeadm-upgrade 1.12->1.13->1.14 [75956](#)

Wed 27th Mar 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Lubomir I. Ivanov, Jason DeTiberus, Ross Georgiev, Tim St. Clair- VMware
- Marek Counts - SUSE
- Ed Bartosh, Alexander Kanevskiy - Intel
- Danny Berger - Pivotal
- Jorge Alarcon

Agenda:

- [Fabrizio] PR [75600](#) : How to manage join when ClusterStatus contains not responsive etcd endpoints?
 - a) silently ignore and move to another endpoint
 - b) print a warning, then ignore and move to another endpoint
 - c) blocking preflight check (user can skip this preflight check and force ignore/move to another endpoint)
 - [lubomir] this was reported as a related problem today: <https://github.com/kubernetes/kubeadm/issues/1469>
- [lubomir] upgrade plan doesn't work correctly (**potential cherry pick**):
 - <https://github.com/kubernetes/kubeadm/issues/1470>
 - workaround is to skip "upgrade plan" and just "upgrade apply".
- [lubomir] packages are quite broken for PATCH versions = latest-1
 - TLDR is that the existence of kubernetes-cni 0.7.5 in the repos made the old packages try to pull this version instead and the pinning to `=0.6.0` doesn't work because of semantics related to the fact that the MAJOR is `0` (or something like that).
 - tickets:
 - <https://github.com/kubernetes/kubernetes/issues/75701>
 - <https://github.com/kubernetes/kubernetes/issues/75683>
 - slack discussion:
 - <https://kubernetes.slack.com/archives/C2C40FMNF/p1553612852825500>
 - known workarounds are pinning CNI version for apt/yum.

- sig-release proposed to override the existing packages, but probably just revisions would be pushed instead e.g. `kubelet-1.13.5-00-1`
- [marek] kind of related, can we set up a planning meeting today or tomorrow for those interested?
<https://github.com/kubernetes/enhancements/blob/master/keps/sig-cluster-life-cycle/20190227-artifact-generation.md>
- [ross] Reminder to vote for config changes in the proposal doc
 - <https://docs.google.com/document/d/1XnP67oO1i9VcDIpw42IzptnJsc5OQM-HTf8cVcjCR2w/edit#>
- [lubomir] what can we merge in master and/or cherry pick for 1.14?
 - <https://github.com/kubernetes/kubernetes/pulls?q=is%3Apr+is%3Aopen+label%3Aarea%2Fkubeadm>

Wed 20th Mar 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Tim St. Clair, Ross Georgiev, Lubomir I. Ivanov - VMware
- Ed Bartosh, Alexander Kanevskiy - Intel
- Fabrizio Pandini
- Jorge Alarcon
- Pawel Bek -Codahead
- Marek Counts - SUSE
- Yago Nobre

Agenda:

- Planning (what's your wish list)
 - [Fabrizio] I moved [Marek Counts] and [Ross] lines into my list below
 - (P0) k/k as unique source of truth for Kubernetes release artifacts
 - [Marek Counts, Pavel, Tim]
<https://github.com/kubernetes/enhancements/pull/858>
 - (P0.5) KIND/ Kubeadm test enablement Improve CI [Lubomir, Fabrizio, Jorge, Rafael]
 - Test grid for HA
 - Test grid for cluster variants
 - tickets with e2e tests in k/kubeadm
 - certs-copy/HA:
<https://github.com/kubernetes/kubeadm/issues/1373>

- test-infra items for kind/kubeadm:
 - <https://github.com/kubernetes/kubeadm/issues/1290>
 - tempted to close this one; not enough details/already covered in other tickets:
 - <https://github.com/kubernetes/kubeadm/issues/813>
 - crictl testing?
 - <https://github.com/kubernetes/kubeadm/issues/983>
- kubeadm enhancements
 - (P1) HA to beta [Fabrizio, Tim, *]
 - Certs copy tracking issue:
 - <https://github.com/kubernetes/kubeadm/issues/1373>
 - HA E2E tests (tracked in the above ticket):
 - auto-copy certs
 - manual-copy certs
 - HA upgrade
 - Support ConfigMap locks:
 - <https://github.com/kubernetes/kubeadm/issues/1097>
 - Review final UX, flags for init and join?
 - Docs update (depends on UX changes)
 - (P1.5/2) Config
 - [Ross] <https://github.com/kubernetes/kubeadm/issues/1439>
 - [Google Doc](#)
 - kubeadm as a “toolbox”
 - (P2) Phases UX: fix or re-design ? [proposal/notes](#)
 - (P3) Reset Phases [proposal](#)
 - (Punt) Upgrade phases [proposal](#)
 - (Evaluate) - Kustomize/OtherTools [proposal/notes](#)
 - (P1) Bring back Windows (alpha-grade) support: [Lubomir, SIG-Windows]
 - <https://github.com/kubernetes/kubeadm/issues/1393>
 - CI-Signal
- (P3) Component standard adoption in kubeadm
 - flag vs config - adoption of k8s.io/component-base
 - the “grand unified field theory on component config”
 - component config for kube-proxy (make v1beta1) [#967](#)
 - component config for kube-scheduler (currently v1alpha1)
 - cleanup bootstrap token code in kubeadm [#964](#)
 - cleanup feature-gates code in kubeadm
- wish list
 - (P1.5) Consolidate story about certs (external CA, renewal, CSR-command) [Yago, Rafael]
 - Renew certificates in the kubeconfig... Our current document is not sufficient. (User journeys)
 - Talk with Steve Harris, and Kubespray folks.
 - (P3) “Change the cluster” [proposal](#)
 - (P3) Cleanup code / Review package structure

- [Fabrizio] new appointments
 - milestone maintainer: Lubomir
 - approver: Ross
 - reviewer: Rafael

Wed 13th Mar 2019 - 9am PT, noon ET

[Recording](#)

Attending:

- Tim St. Clair, Lubomir I. Ivanov, Ross Georgiev, V A Amarnath - VMware
- Marek Counts - SUSE
- Paweł Bek - Codahead
- Ed Bartosh - Intel
- Jorge Alarcon - Vividcortex

Agenda:

- [lubomir] status of docs for 1.14
 - <https://github.com/kubernetes/kubeadm/issues/1422>
 - Lubomir and Yago to send remaining 2 PRs later.
 - We are slightly behind on schedule and reviews should start ASAP.
- [fabrizio/lubomir] kinder to be merged in k/kubeadm?
 - kinder is to be used for testing development cases between kubeadm and kind as a library. kinder is to be pushed in the kubeadm repo for now.
 - [tim] Over time I'd like to minimize this surface area and continue to push kind.
 - [lubomir] We can't use kinder for e2e testing currently with the current kubetest/kind deployer.
- [lubomir] CRD discussion and brief future plan for kubeadm?
 - ... we went all meta talking
- [lubomir] Config map locking should be a priority for v1.15
- [lubomir] tracking issues in k/enhancements. we should edit the OPs of tracking issues and better maintain them because they are out of date.

Wed 6th Mar 2019 - 9am PT, noon ET

[Recording:](#)

Attending:

- Tim St. Clair, V A Amarnath, Andy Goldstein, Ross Georgiev - VMware
- Marek Counts, Rafael Fernández López - SUSE

- Daniel Lipovetsky - Platform9
- Yago Nobre - Nubank
- Dmitry Rozhkov - Intel

Agenda:

- [dlipovetsky] Taints marshalling: <https://github.com/kubernetes/kubeadm/issues/1358>
- [fabrizio pandini] [Code freeze](#) is tomorrow!
 - HA (tracking issue [1373](#))
 - [74671](#) Yago / Merged
 - [72886](#) Marek / no updates 😞😞
 - [lubomir] existing PR closed. will send another PR later.
 - Join phases (tracking issue [1204](#))
 - [74873](#) Lubomir / Merged
 - [74870](#) Lubomir / Merged
 - [75013](#) Lubomir / in merge queue
 - args validation moved to v1.15
 - Other
 - [74530](#) Ross / Closed / v1.15
 - [72685](#) Yago / Merged
- 1.14 Docs <https://github.com/kubernetes/kubeadm/issues/1422>
- [ereslibre] Integration tests for <https://github.com/kubernetes/kubeadm/issues/1373>

Wed 27th Feb 2019 - 9am PT, noon ET

Recording:

Attending:

- Fabrizio Pandini
- Yago Nobre
- Marek Counts, Rafael Fernández López - SUSE
- Ed Bartosh - Intel
- Jason DeTiberus, Tim St. Clair - VMware
- Paweł Bek - blind3dd (Codahead)

Agenda/Notes

- [fabrizio] Auditing join phases; sorry I will be traveling and I don't know how connection will be. If I cannot connect, PTAL @ [this doc](#)
- [fabrizio] Update on HA [#1373](#); we are in good shape. we need to get more people starting to executing manual test because there will be no CI signal:
 - Certs copy

- PSA Upload certs ([#73907](#)) and Download certs ([#74168](#)) merged.
Test it, break it, send feedback!
 - TODO Allow user to provide the certificate key on upload-certs (Yago)
 - TODO Change init final message (Yago)
 - TODO Add More tests (Yago/Rafael)
 - [yago] AI - should we add a command to generate certificate keys?
like `kubeadm token`
- Upgrade improvements punted to v1.15
- Reset improvements
 - PSA Remove etcd member merged [#74112](#) and [#74410](#)
Test it, break it, send feedback!
 - WIP Cleanup ClusterStatus is pending [#72886](#) (Marek)
- [marek] <https://github.com/kubernetes/enhancements/pull/858> packaging KEP is now a PR.
 - Pawel would help join on execution in 1.15.

Wed 20th Feb 2019 - 9am PT, noon ET

[Recording:](#)

Attending:

- Tim, Ross, Marek ...?
- Fabrizio Pandini
- Ed Bartosh - Intel

Agenda/Notes

- [lubomir] i'm AFK today, sorry. PTAL at the items below)
- [lubomir] fix config overwrite for RPMs:
 - <https://github.com/kubernetes/release/pull/658>
 - what is blocking this PR? there was a related bug report recently (can't find it).
- [lubomir] kubeadm kubeconfig precedence:
 - <https://github.com/kubernetes/kubernetes/pull/71874>
 - fixes a bug where if you pass `--kubeconfig=/etc/kubernetes/admin.conf` it still searches in default locations like KUBECONFIG.
 - decide if it's OK to not apply precedence on init, join, reset, upgrade (the PR does that).
- [fabrizio] Update on join phases [#1204](#)
 - All the phases are now merged !
 - [lubomir] A last second change introduced a regression on master, but fix already merged [#74277](#)
 - We should plan an audit session final

- [fabrizio] Update on HA [#1373](#)
 - Certs copy
 - Upload certs is in merge queue [#73907](#) with two follow up prs
 - Allow user to provide the certificate key on upload-certs
 - Add More tests
 - Download certs is in good shape [#74168](#) please take a look!
 - Upgrade improvements is pending ()
 - Reset improvements
 - Cleanup ClusterStatus is pending [#72886](#)
 - Remove etcd member is in merge queue [#74112](#) (Merging soon, but requires more tests.)
- [fabrizio] Remove kubeadm-playground from the kubeadm-repo. Welcome kinder! kinder is an example of kind used as a library; All the kind commands will be available in kinder, side by side with additional commands for supporting use cases that are difficult to support in the main CLI.

Roadmap

- [] Enable usage of kinder(kind) as local integration test environment for kubeadm/cluster api development
 - [x] Allow creation of nodes "ready for installing Kubernetes"
 - [x] Provide "topology aware" wrappers for `docker exec` and `docker cp`
 - [x] Provide pre built "developer" workflows for init & join. e.g. init with phases
 - [] Provide pre built "developer" workflow for kubeadm upgrades
 - [] Allow test of kubeadm cluster variations (e.g. External etcd)

Non-Goals include:

- [] Supporting every possible use cases that can be build on top of kind as a library

Next step:

- get merged 1 change in kind, so we can avoid fork
- find a home for kinder
- find volunteers using kinder for v1.14 manual tests

Wed 13th Feb 2019 - 9am PT, noon ET

[Recording:](#)

Attending:

- Tim St. Clair, Lubomir I. Ivanov, Jason DeTiberus, Chuck Ha, Ross Georgiev - VMware
- Marek Counts, Rafael Fernández López, Dario Maiocchi, Pablo Chacin - SUSE
- Dmitry Rozhkov, Ed Bartosh - Intel

- Yago Nobre - Nubank
- Fabrizio Pandini

Agenda/Notes

- [Marek] Discuss the state of packaging
 - notes:
 - https://docs.google.com/document/d/1f5l_plsuysDiQPTnPctjUWINxafaPcbZML1oByuBlu8/edit
 - Going to be working on a KEP for this, this week.
 - Building out an example yaml.
 - Help welcome
 - Research into how to support these functions in yaml
 - [tsc] need to sync with hannes. may be hard pressed to flip the switch on the old packages, but could get to generating new packages side-by-side, and orchestrate the switch in 1.15.
- State of HA
 - [Work for the copy of certs in HA has started](#)
 - Finishing up the PRs
 - State of tests (requirement for beta)
 - Need to update kind to support this
 - Need to keep the manual way for the older versions.
 - There is already a switch linked to the kubernetes version which we can use to determine code path in KIND
 - Will stay on alpha for this release
 - It was decided to add a flag for a user provided key this cycle.
- State of join phases
 - [Slightly blocked on the last PR for join phases](#)
 - This is the final phase before join is complete.
- Other updates
 - Fixed bug related to our diff lib and also a bug in the defaults for the kubelet config (these are now different on Linux and non-Linux):
 - <https://github.com/kubernetes/kubernetes/pull/73945>
 - <https://github.com/kubernetes/kubernetes/pull/73941>
 - [lubomir] do we backport?
 - [tsc] backports are usually for regressions, but not vendor dependencies
 - [our kubeadm pre-submit test was removed](#)
 - Certs cache on init certs phase
 - <https://github.com/kubernetes/kubernetes/pull/73775>
- [Tim] If time backlog grooming for 1.14

Wed 6th Feb 2019 - 9am PT, noon ET

[Recording:](#)

Attending:

- Tim St. Clair, Jason DeTiberus, Andy Goldstein, Ross Georgiev , Chuck Ha, Ruben Orduz, Lubomir I. Ivanov - VMware
- Marek Counts, Dario Maiocchi, Rafael Fernández López, Pablo Chacin - SUSE
- RA489 - NEC
- Hannes Hörl - Pivotal
- Paweł Bek - Codahead
- Ed Bartosh - Intel
- Fabrizio Pandini
- Nicolas Trangez - Scalify
- Davi Garcia (davivcgarcia)

Agenda/Notes

- [Marek] Discuss the state of packaging
 - Spent time looking at nfpm, and fpm, to build both rpms/debs.
 - Run into limitations (hasn't been actively developed) - last commit was in Nov
 - FPM is more active and maintained, FPM and is in ruby.
 - hasn't currently found any bugs in nfpm.
 - Are there any open issues there?
 - Doesn't currently support Changelogs.
 - wrt to signing we can do this in a separate step.
 - Notes
 - https://docs.google.com/document/d/1f5l_plsuysDiQPTnPctjUWINxafaPcbZML1oByuBlu8/edit
 - There are also some dockerfile available to the create build-container.
 - Could use some help here. (@andy)
 - Revisit next.
- [Hannes] Discuss the state of release
 - Is there some change to 1.14?
 - Can work on it and depends on progress.
 - [tim] Have requirements around having separate stable and testing repos
 - Currently no one can test an alpha release externally as alpha and beta releases are not cut.
 - [lubomir] They are produced, but not published to repos.
 - [tim] Also would like to refactor anago to use k/k instead of separate packaging in k/release
 - k/release contains the debian and rpm spec

- We can get a PR together for anago to publish artifacts out of k/k
 - [hannes] anago doesn't actually handle the rpm/deb artifacts, it's done by other scripts used by googlers
 - [andy] the branch managers have some private repository they work with for updating the spec files, which are not generally checked back in to k/release.
 - [andy] the current tools are hard to consume for publishing to anything that isn't the official kubernetes public repositories
 - [tim] There are people who have requirements to rebuild binaries and sign with their own signing keys.
- [nicolas trangez in chat] So, given (n)fpm, is the intent to build infra to take actual binaries produced by... 'some other infra', and wrap those up in an RPM/Deb? As in, in the RPM case, don't build a spec which builds 'from source'
 - [andy] yes, we currently take binary artifacts from gcs storage and wrap them into the rpm/deb
 - [nicolas] this complicates being able to produce custom signed packages
 - [andy] current use case is to bring our own source code for kubernetes, build it and produce rpm/debs. if sig-release wants to continue having the package building process consume precompiled binaries, we could most likely live with that, as long as we can bring our own binaries, but it's more typical to have the rpm/deb packaging process do its own compilation
- [ed bartosh] have you considered using native build tools, like [git-buildpackage](#)
 - [tim] up to those evaluating fpm to see what works
- AI: Take a look at FPM and revisit next week
- [Tim] PSA - Weird critical bug
 - <https://github.com/kubernetes/kubernetes/issues/72102>
- [Q:Tim, A:Lubomir] What's the state of windows?
 - Should we only target Server 2019?
 - This is the minimum version that would want to be tested.
 - Discussed with SIG Windows on the testing mechanic.
 - Changes are being made to support Windows around linux-only tests.
 - kube-up is being used by Google, but SIG Windows is using Azure, as used in test-infra.
 - Spins up an Azure cluster and obtains resultant logs.
 - [tim] Is this federated testing or actually happening in prow?
 - the signal is green except for dns
 - [tim] As soon as Windows goes GA, SIG CL will be asked how to stand up Windows
 - [lubomir] We made linux only assumptions in kubeadm. For the next cycle, we are going to investigate best effort support for windows in kubeadm.
 - Testing showed it failed to bring up kubelet

- Prow doesn't support grabbing windows logs
 - Right now, we will say Windows Containers is GA, but kubeadm is not ready for Windows
 - [tim] we should create tracking issues for windows support, and then target these for planning in 1.15
 - Also would need to sort out packaging
 - [tim] we will not claim multiarch support
- [lubomir] our kubernetes-anywhere tests are red again:
 - need to refactor the k-a log dumping logic to not use the `cluster` folder
 - implementing kind jobs is a WIP
 - we need to plan for jobs that test deb/rpm artifacts (k-a is currently doing that, but kind is not)
- [Chuck] Nearly [duplicate](#) kubeadm commands cause a very confusing experience.
 - non-config flag deprecation?
 - [fabrizio] flags are the first thing users see for customisation, so we cannot get rid of them entirely
 - [tim] tools like kustomize my offer a better escape hatch in the future
 - [lubomir] a kflag library is being proposed to handle overrides in a unified way.
 - [fabrizio] as this is a cross binary problem, then we shouldn't reinvent this wheel
 -
- [Tim] Backlog state
 - <https://github.com/kubernetes/kubeadm/milestone/10>

Wed 30 Jan 2019 - 9am PT, noon ET

Recording:

Attending

- Fabrizio Pandini
- Rafael Fernández López, Marek Counts, Dario Maiocchi, Pablo Chacin - SUSE Linux
- Ed Bartosh - Intel
- Lubomir I. Ivanov, Ross Georgiev, Chuck Ha, Tim St. Clair, Ruben Orduz, Jason DeTiberus - VMware
- RA489 - NEC

Agenda/Notes

- [bart0sh] discuss how kubeadm should deal with parameters specified in the user-defined config. Currently at least some of the parameters are simply ignored, which confuses users. Example issues: [1254](#), [1252](#), [1028](#)
Here are 3 possible solutions proposed by [neolit123] in [1028](#):
 - allow changes by the user but show warnings for deviations from the defaults.
 - always override with the security defaults, but document it.
 - always override but show a warning "[kubelet] Applying kubeadm security defaults."
 - exit with an error if something is dangerous.I'd add another option - to exit with an error when user specifies dangerous parameters, but give user a possibility to force them (--force?).
- [lubomir] the extra unnamed arg in `kubeadm join <some-endpoint>` has opened a can of worms for us and the `join phases`:
 - [lubomir] [here is how it's handled](#)
 - [lubomir] my vote here is to add a flag `--endpoint` and roll the GA deprecation on the unnamed arg ([GA: 12 months or 2 releases \(whichever is longer\)](#))
 - tracking issue: <https://github.com/kubernetes/kubeadm/issues/1375>
- [lubomir] we are now using 18.09 as the recommended Docker:
 - PR1: <https://github.com/kubernetes/kubernetes/pull/72823>
 - PR2: <https://github.com/kubernetes/kubernetes/pull/73457>
 - [help wanted for updating the docs](#)
 - test-infra, release and kind are still at 18.06:
<https://github.com/kubernetes/test-infra/issues/9305>
- [fabrizio] PSA: Green light for Certificates copy for join --control-plane: Help wanted!
 - Tracking issue <https://github.com/kubernetes/kubeadm/issues/1373>
 - Things to agree asap:
 - Config changes for encryption key (init and join)
 - New command for generating encryption key

Wed 23 Jan 2019 - 9am PT, noon ET

Recording: TODO

Attending

- Dario Maiocchi, Pablo Chacin - SUSE Linux
- Ross Georgiev, Chuck Ha, Lubomir I. Ivanov, Vince Prignano, Jason Scarano - VMware
- Fabrizio Pandini
- Pawel Bek
- Jorge Alarcon - VividCortex (
- Ed Bartosh - Intel

Agenda/Notes

- [fabrizio] we should define what are kubeadm official supported cluster variants vs custom cluster
 - <https://github.com/kubernetes/kubeadm/issues/1349>
 - how do we make this more clear to users?
 - [chuck] making what we support documented/explicit is a good thing
 - do we want to “without kube-proxy” an “officially supported variant? rif. kube router support
 - <https://github.com/cloudnativelabs/kube-router/blob/master/docs/architecture>
 - <https://github.com/cloudnativelabs/kube-router/blob/master/docs/kubeadm.md> (`--pod-network-cidr` when you run `kubeadm init // verification?.`)
 - Nat vs NotNat
- [fabrizio] PSA: questions about certs rotation are becoming more frequent. If someone want to help there is room of improvement here!
 - lack of user facing documentation
 - rif [1350](#) (CA rotation: controller-manager needs a separate ca.crt file)
 - rif [1223](#) (kubeadm generate kubelet certs with master(s) CA)
- [fabrizio] backlog cleanup:
 - PSA [69366](#) detect CRIs automatically is merging; manual test will be great
 - search logic for kubeconfig files: ? ([71874](#))
- [lubomir] deprecate placeholder docs in k/k:
 - <https://github.com/kubernetes/kubernetes/pull/70052>
 - most folks are +1 on this
 - it seems to be used in a “linkchecker” that is no longer active:
 - https://github.com/kubernetes/kubernetes/pull/70052#discussion_r239686687
 - also sig-docs are not using these.
- [lubomir] local DNS cache?
 - <https://github.com/kubernetes/kubernetes/pull/71642>
 - do we want this in kubeadm and if yes in what form - `alpha` cmd or f-gate?
 - sig-network folks (+thockin) told me that the `k8s-dns-node-cache` image is going to be maintained and is not going to become abandonware anytime soon.
 - <https://github.com/kubernetes/kubernetes/issues/70707#issuecomment-444282873>
- [tim] Discussion after the dedicated meeting about DEB/RPM.

Wed 16 Jan 2019 - 9am PT, noon ET

Recording: TODO (not recorded :-()

Attending

- Lubomir I. Ivanov, Ross Georgiev - VMware
- Rafael Fernández López, Dario Maiocchi, Marek Counts - SUSE Linux
- Martin Dietze - Dermalog
- Thomas Strömberg - Google
- Bryson Shepherd - WalmartLabs
- Fabrizio

Agenda/Notes

- [fabrizio] Discuss KEP for automating the copying of certificates on ``join --control-plane``
 - [KEP for ``join --control-plane``](#)
 - [Copy edit document for the new KEP](#)
 - I will quickly iterate on google doc for the first round of feedbacks until friday and the move on a PR
- [lubomir] Discuss phases for ``kubeadm join``:
 - [tracking issue](#)
 - [Google doc with WIP proposal](#)
 - We want to declare the list of phases, so that contributors can work on them in parallel @yago and @ereslibre volunteered for this.
- [fabrizio] Discuss optional kube-proxy that users demand.
 - we should define what a custom cluster is
 - removing kube-proxy is not a supported cluster and we don't have CI singal for it.
 - do we want to make this supported?

Wed 9 Jan 2019 - 9am PT, noon ET

Recording: <https://youtu.be/8Mf6vgXWVvl>

this link seems wrong.

also i found this [4 minute recording](#) uploaded to youtube on the 11. which isn't this meeting either.

Attending

- Tim St. Clair, Jason DeTiberus, Ruben Orduz, Lubomir I. Ivanov, Naadir Jeewa - VMWare
- Rafael Fernández López, Dario Maiocchi, Marek Counts - SUSE
- Paweł Bek (blind3dd) - Codahead
- Jorge Alarcon - vividcortex
- Yago Nobre
- Fabrizio Pandini

Agenda/Notes

Issues

- [fabrizio] let's discuss [result from the survey](#) + feedbacks from kubeconf
 - my takeaway (from the survey):
 - users keep up with release and updates pretty well
 - only a small group of users runs compliance tests
 - number of cluster/nodes is interesting
 - more HA clusters than single node clusters; windows and mixed architecture are not a concern
 - Better HA & Documentation are both paramount; upgrade should be improved as well
 - LTS and "change the cluster" will be more than appreciated by users
 - Well done!
- [xing] Sig-storage is working on moving in-tree CRD's CSIDriver and CSINodeInfo to Beta in 1.14 release: <https://github.com/kubernetes/kubernetes/pull/70367>. We need to figure out how to install these two 2 CRD's through deployment tools such as kubeadm.
- [fabrizio] low priority - let's tackle some old PRs: e.g. :
 - search logic for kubeconfig files (rif [71874](#))

1.14 Prioritization

improve Kubernetes release/PR + Tests management

- [P0] Marek, Tim, Paweł, Naadir: [Artifact build + release](#) + Tests (CI build & release build unification)
 - Improve PR Blocking job reliability/speed [Kind]
 - Move kops to periodic job
 - Select replacement for K/Anywhere: Kubespray vs Kind vs Cluster API* vs
 - improve signal for kubeadm main workflows (upgrades, etc.)
 - [P1] improve signal for cluster variations (e.g, CRI-O, external CA ..)

- [P2] improve signal for kubeadm utilities (e.g. Config migrate, certs renew etc.)
- [P2] e2e tests for kubeadm owned settings (e.g. kubeadm-config ConfigMap, RBAC rules etc.)

kubeadm enhancements

- [P1] HA support to beta
 - [P0] UX improvements/address feedbacks (reset, pre pull, upgrade node etc.)
 - [P0] Test signal
 - [P?] Automate certificate copy (alpha)
- [P2] Lucas, Marek: The “grand unified field theory on component config”
 - [P?] flag vs config - adoption of k8s.io/component-base
 - [P?] component config for kube-proxy (make v1beta1)
 - [P?] component config for kube-scheduler (currently v1alpha1)
 - [P?] cleanup bootstrap token code in kubeadm
 - [wg-component-standard Meeting Notes](#)
- [P2] Lucas (will review): Security improvements
 - [P?] secure serving for controller-manager & scheduler
 - [P?] check/improve CIS compliance / kube-bench
 - [P?] certificate rotation enhancements ??? (kubelet, kubeconfig files etc.)
 - ...
- [P3] kubeadm as a “toolbox”
 - [P?] Join Phases
 - [P?] Machine readable output
 - [P?] Allow easier Vending
 - Config vending
 - “Phases” vending
 - [P?] Kube Router use case (skip kube-proxy in join/upgrade vs config)
- [On Going] Documentation
 - [P?] high level doc for docker/external components recommended versions
 - [P?] document how to run conformance in a cluster (created with kubeadm)
 - [P?] document how certificate rotation works (automatic or DIY)
 - [P?] Advanced tutorials
 - apply customize settings to control plane components / to nodes
 - “Using external CA”
 - ...
- [P3] Cut the technical debt
 - [P?] organize code review sessions
 - [P?] ...

- Ideas/any other
 - [P?] Enforce priorityClass/QoS
 - [P?] “Change the cluster” workflow [proposal](#) (design only)
 - [P?] Improve support for Kubernetes on ARM (timeouts)

Wed 2 Jan 2019 - 9am PT, noon ET

(meeting canceled - b/c a lot of folks are still out)

Attending

- Tim St. Clair, Vince Prignano - VMWare
- Ed Bartosh - Intel

Wed 19 Dec 2018 - 9am PT, noon ET

(meeting canceled - SIG leads are not available today. happy holidays!)

Wed 12 Dec 2018 - 9am PT, noon ET

(meeting canceled due to overlap with KubeCon Seattle)

Wed 5 Dec 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Liz Frost - Heptio
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Marek Counts, Rafael Fernández López - SUSE
- Fabrizio
- Jorge Alarcon
- Yago Nobre
- Alexander Kanevskiy - Intel

Agenda/Notes

- Handling legacy godoc linking.
 - We have an issue with godoc links to v1alpha2 versions
- discussion about building deps and rpms.
 - <https://github.com/kubernetes/kubernetes/issues/71677>
- discussion about kubecon seattle.

Wed 28 Nov 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Chuck Ha - Heptio
- Marek Counts, Rafael Fernández López - SUSE
- Lubomir I. Ivanov - VMware
- Jorge Alarcon
- Lucas Käldestrom - luxas labs (Contracting for Weaveworks)

Agenda/Notes

- [Tim] Outstanding 1.13 issues?
 - None.
- [Tim] Blog post
 - https://docs.google.com/document/d/1YMXdrETV1BpREYCwXwa_N7FA50A301k8AOm5R-W6M7M/edit#heading=h.czy22xu9x0b0
- [chuck] Revisit k/k's linux package builds and the reasons why we are maintaining it
 - For testing.
- [Tim]
 - Promoting Lubomir to approver
 - Lucas: +1
 - Tim: +1
 - Fabrizio: +1
 - Chuck: +1!
 - Jason DeTiberus: +1
- [Lucas] k8s.io/component KEP
 - <https://github.com/kubernetes/community/pull/2977>
- [Tim] KubeCon Seattle
 - Contributor Summit
 - Presentation
 - BoF during the day
 - KubeCon
 - SIG Intro
 - SIG Deep Dive (kubeadm)
 - SIG Deep Dive (Cluster API)
- [Lucas] SIG Cluster Lifecycle 2019
 - https://docs.google.com/document/d/13wMdzkvAvSoT5Mmy5V1OP5hk1ot04DWkNLC_OsjTU4A/edit
 - [Lubomir] - Start a survey and we will chat about it next week.
- [Lucas] <https://github.com/kubernetes/kubeadm/issues/1200>

Wed 21 Nov 2018 - 9am PT, noon ET

Attending

- Lubomir I. Ivanov - VMware
- Dhawal Yogesh Bhanushali - VMware
- Tim St. Clair - Heptio
- Marek Counts, Rafael Fernández López - SUSE
- Lucas Käldestrom - luxas labs (Contracting for Weaveworks)

Agenda/Notes

- [lubomir] announcement for WIP blog post about “kubeadm GA”:
 - https://docs.google.com/document/d/1YMXdrETV1BpREYCwXwa_N7FA50A301k8AOm5R-W6M7M/edit#
 - released on the 3rd of Dec.
 - we need more folks looking at this and possibly outline the strengths that the project has - flexible configuration, upgrades, etc.
- [lubomir] discuss status of docs PRs for 1.13:
 - tracking issue: <https://github.com/kubernetes/website/issues/10937>
 - pending PRs: <https://github.com/kubernetes/website/pulls?q=is%3Apr+is%3Aopen+label%3A%2Fcluster-lifecycle>
 - please tag your PRs with ``/priority critical-urgent`` and ``/sig cluster-lifecycle``.
 - **are we missing docs updates for new features?**
- [lubomir] walk the PRs and backlog for 1.13 and look at critical items to escalate to higher priority:
 - PRs: <https://github.com/kubernetes/kubernetes/pulls?q=is%3Apr+is%3Aopen+label%3Aarea%2Fkubeadm>
 - backlog: <https://github.com/kubernetes/kubeadm/milestone/9>
- [lucas] We really need to do flag parsing consistently
 - <https://github.com/kubernetes/kubernetes/pull/71270>
 - Ties into my new proposal: <https://docs.google.com/document/d/1nZnzJD9dC0xrtEla2Xa-J6zobbC9oltdHlJ3KKSSlkh/edit#>

Wed 14 Nov 2018 - 9am PT, noon ET

Attending

- Lubomir I. Ivanov, Ross Georgiev - VMware
- Tim St. Clair, Chuck Ha - Heptio
- Rafael Fernández López - SUSE
- Fabrizio Pandini

Agenda/Notes

- [lubomir] strict YAML unmarshaling, we can enable it as warnings for 1.13?:
 - the change has to happen on the kubeadm (client) side since api-machinery don't want it in the codecs.
 - warnings are going to complicate our unit tests, errors are easier.
 - WIP PR is in flight: <https://github.com/kubernetes/kubernetes/pull/70901>
- [liz/fabrizio/tim] having all flags under init.go instead of sub-phases. should we keep it like that for 1.13 and possibly refactor in 1.14, or refactor now?
 - CSR PR with discussion: <https://github.com/kubernetes/kubernetes/pull/70809>
 - PR with local flags: <https://github.com/kubernetes/kubernetes/pull/71040>
- [fabrizio] deprecate feature flags <https://github.com/kubernetes/kubeadm/issues/1226>
 - Probably ok removing all deprecated feature flags.
- [fabrizio] there are a lot of issues about setting HA clusters in k/kubeadm.
 - [tim] the steps are hard to follow.
 - [tim] probably because more people are using it.
 - [fabrizio] try to figure out if there are pattern/recurring problems
 - [chuck] kubeadm generally supports external etcd very well and stacked control planes requires more steps. The stacked control plane instructions are where most of the issues pop up

Wed 7 Nov 2018 - 9am PT, noon ET

Attending

- Fabrizio Pandini
- Rafael Fernández López, Marek Counts - SUSE
- Tim St. Clair, Jason DeTiberus, Chuck Ha - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Lucas Källdström - luxas labs (Contracting for Weaveworks)
- Yago Nobre - Nubank

Agenda/Notes

- [fabrizio] Api changes
 - **YES** Remove ClusterName from join configuration. use ClusterName from ClusterConfiguration. [Rafael]
 - **YES** Remove FeatureGates from join configuration. use FeatureGates from ClusterConfiguration. [Issue](#) [Rafael]
 - **YES** Changes UnifiedControlPlaneImage to UseHyperKube (bool) [Ross]
 - **YES** Deprecate featureGate:CoreDNS, considering we are going to implement DNS.Type [Fabrizio]
 - **YES** Remove DNS.ExtraArgs [Fabrizio]
 - **YES** Remove Proxy.ExtraArgs [Fabrizio]
 - **YES** Remove AuditPolicyConfiguration from config and from code [Marek]
- [fabrizio] Final API review: When?
 - [tim] tomorrow?
 - [lucas] rename APIEndpoint ??
- What items remain before slush?
 - [tim] Graduate Dynamic Kubelet Config? (why?)
<https://github.com/kubernetes/kubernetes/pull/70410>
 - Action TBD in slack
 - Addons? <https://github.com/kubernetes/kubernetes/pull/70354>
- [Fabrizio] Upgrade regression, probably introduced by stacked etcd
 - kubeadm 1.13 init → kubeadm 1.13 upgrade works
 - kubeadm 1.12 init → kubeadm 1.13 upgrade does not works
 - What is the current status of test coverage for upgrades

Wed 31 Oct 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Ruben Orduz, Chuck Ha - Heptio
- Marek Counts, Rafael Fernández López - SUSE
- Kenneth Redler
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Santos, Yago Nobre - Nubank

Agenda/Notes

- [kenneth] - Join divergence for worker vs. control plane - [Issue 1189](#)
- [tim] - discuss join phases vs. alpha options
(<https://github.com/kubernetes/kubernetes/pull/70333>)
- [tim] 1.13 backlog
- Talk about having a community session

- AI: Marek to discuss on the #kubeadm to see if they would like a session
- Perhaps we could have a survey to help refine the top 10 list of common questions.

Wed 24 Oct 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Lubomir I. Ivanov - VMware
- Chuck Ha, Tim St. Clair - Heptio
- Fabrizio Pandini
- Ben - Google

Agenda/Notes

- [lubomir - 5min] quickly discuss state of e2e vs failing tests
 - <https://github.com/kubernetes/test-infra/pull/9897>
 - <https://github.com/kubernetes/kubernetes/pull/70141>
- [tim/fabrizio] <https://github.com/kubernetes/kubernetes/pull/69486> - discuss question on stacked masters. For V1:
 - use advertise-address
 - scenario with advertise-address = LB require external etcd
 - always scale etcd when scaling control-plane
- [tim/fabrizio] timeouts
 - close [70025](#)
 - instead add a timeout for waiting on control plane to the api server struct (fabrizio will sync with @roster)
- [tim] new Approver
 - next cycle Lubomir will step up as approver!
- [tim] Kubecon
 - who will be in Kubecon seattle?
 - looking for someone to co-present kubeadm update
 - [fabrizio] we should give an overview about sig-projects and how they are related
 - [tim] we should keep going on having updates from different projects periodically at the sig-meetings

Wed 17 Oct 2018 - 9am PT, noon ET

Attending

- Lubomir I. Ivanov, Ross Georgiev - VMware

- Jason DeTiberus, Liz Frost, Ruben Orduz, Nicholas Lane - Heptio
- Naadir Jeewa - Scale Factory

Agenda/Notes

- [lubomir/fabrizio] we should merge [69666](#) so that we have an example phase of how the new phase runner works.
- [tim] we need to figure out the status of packages for 1.13.
 - [naadir] focus for release-infra is to move repos to cncf before making other changes
 - [tim] the way we do packaging is not how you typically do packaging.
 - [tim] we can talk about this (with richard).
 - [naadir] Beginnings of [Improving packaging UX doc](#)
 - [ross] our deps are deviating.
 - [tim] Have spoken to Canonical re: canonicalising (pun intended) deps
- [lubomir/fabrizio] we need to define the list of init phases for 1.13:
 - [issue](#)
 - [gdoc](#)
- [fabrizio] kubeadm config v1beta1 (new requests)
 - Y Timeouts (generic API design, only 1 timeout configuration implemented)
 - HOLD Ignorepreflight
 - TOBE DEFINED Allow to specify IP address for local/stacked etcd
 - DOUBLE CHECK Kube proxy change config on a node base

Wed 10 Oct 2018 - 9am PT, noon ET

Attending

- Marek Counts, Rafael Fernández López - SUSE
- Tim St. Clair, Jason DeTiberus - Heptio
- Fabrizio Pandini

Agenda/Notes

- [tim] Walk through config/api changes
https://docs.google.com/document/d/1Sk2VW7IKaLjrjf_4a1NEU0I7X0ysZoXq9IVBvH-ruYY/edit#

After v1.12/v1alpha3 pending changes are:

- Nice to have but not critical.
 - Refactor the Discovery joining API types by using substructs
 - PR [67763](#)
 - Improve controlplane configuration by using substructs
 - group apiserver/controllermanager/scheduler args.

- Addons
 - Manifest changes + Image repository [1091](#)
 - CoreDNS Config map overrides.
 - [Tim] Need to evaluate overlap with bundles.
 - “Promotion” Stacked etcd - issue [1123](#)
 - PR [69486](#)
 - Support ignoring preflight errors [808](#) (added post meeting/consensus?)
- Important Long Term
 - Move out BootstrapTokens
 - Changes landed, and figure out how to excise our dependencies there.
 - Addons configuration [1017](#)
- Critical
- Make flags override config values - issue [1040](#)
- [fabrizio] PSA: [69622](#) implements the phase runner as defined in the KEP + gets kubeadm init ready for graduating phases; another PR for graduating preflight phase in flight

Wed 3 Oct 2018 - 9am PT, noon ET

Attending

- Flavio Castelli, Marek Counts, Rafael Fernández López - SUSE
- Fabrizio Pandini
- Tim St. Clair - Heptio
- Ross Georgiev, Lubomir I. Ivanov, Dhawal Yogesh Bhanushali- VMware
- Jason Murray - iNNOVO Cloud

Agenda/Notes

- Discuss - <https://github.com/kubernetes/kubeadm/issues/1152>
 - kubeadm config print init-defaults
 - “” join-defaults
 - This will work for 1.13 but we need to fix the issues with 1.12 to not fail.
 - We will need to make it tolerant of init/join for different flows.
- Version skew with stable-1
 - AI:Tim review <https://github.com/kubernetes/kubernetes/pull/69301>
- KEP kubeadm phases to beta
 - <https://github.com/kubernetes/community/pull/1939>
 - show the list of phases
 - kubeadm init --help will print the list of available phases (ordered)

- kubeadm init logs [tag] will be consistent with the name of phases
- kubeadm init phase subcommand will allow user to invoke a single phase
NB. we are not considering to let the user to invoke a list of phases because this open
-
- [No] kubeadm init --phase certs --api-server-advertise-address 127.0.1
- **[Yes] kubeadm init --foo-init phase --foo-phase certs --api-server-advertise-address 127.0.1**
- **[Skip] kubeadm init --skip-phase=proxy, certs**
- [Fabrizio] cleanup self-hosting [1072](#) / with the help of Marek -> dropping everything in the main workflows but moving the pivoting logic to another command + warning
- [Fabrizio] PSA I'm prototyping on stacked etcd, please see command [1123](#)

Wed 26 Sept 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Marek Counts, Rafael Fernández López - SUSE
- Tim St. Clair, Jason DeTiberus, Liz Frost, Chuck Ha - Heptio
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Naadir Jeewa - Scale Factory
- Martin Dietze - Dermalog
- Fabrizio Pandini

Agenda/Notes

- [Tim] 1.13 planning / hit-list
 - p0 - [Fabrizio, Ross, Tim] push config to v1beta1 (🎉)
 - https://docs.google.com/document/d/1Sk2VW7IKaLjrjf_4a1NEU0I7X0ysZoXq9IVBvH-ruYY/edit#
 - Improve controlplane configuration by using substructs ??
 - Refactor the Discovery joining API types ??
 - Addons configuration ??
 - CoreDNS: how to give a better UX considering that it doesn't support component config ??
 - p0 - [Fabrizio, Tim, Lubomir] phases to beta
 - possibly add a subcommand that lists the phases.
 - KEP <https://github.com/kubernetes/community/pull/1939>
 - to use phases capabilities: kubeadm init --only-phases or --skip-phases
 - to get the list of phases: kubeadm init --dry-run
 - p1 - Test automation

- [Jason/Chuck/Liz/Tim] Auto-jig [Cluster-api - aws HA deployment]
 - [lubomir] wil patch kubernetes-anywhere until the above is ready ^
- p1/2 - Packaging + Release
 - [cha + tim + naadir] k8s-infra team
 - [naadir] - Producing a straw man doc on ideal end state to distribute to k8s-infra team & sig-cluster-lifecycle
- p2 - addons / <https://github.com/GoogleCloudPlatform/k8s-cluster-bundle>
- p1 - Cleanup deprecated, and promoted features
 - cleanup self hosting
 - remove feature gate
 - add separate destructive sub-command.
 - kubeadm pivot ...
 - P2 - dynamic kubelet
 - Fix an error on our UX. Give instruction to users
 - Figure out how to test it.
 - AI: [Tim] Setup a time with mtaufen.
- p1 - [Fabrizio, chuck, jason] Join --control-plane to beta
 - UX refinements
 - Automatic stacked etcd
 - Test automation
- p3 - New alpha command for changing existing clusters using kubeadm
 - KEP <https://github.com/kubernetes/community/pull/2688>
 - first priority discuss about possible approaches, drawbacks etc.etc. around a first use case:
 - change control-plane extra args

Wed 19 Sept 2018 - 9am PT, noon ET

Attending

- Martin Dietze - Dermalog
- Jason DeTiberus, Tim St. Clair, Chuck Ha, Liz Frost - Heptio
- Marek Counts, Rafael Fernández López, Richard Brown - SUSE
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Fabrizio Pandini

Agenda/Notes

- [chuck] Assigning tickets to folks outside the org. Is this a github thing or a k8s process thing?
 - You can't, in order to assign they need to apply for the org.
 - [lubomir] the lifecycle/active label helps here. we've been using it for a while with new contributors.

- [tim] We need folks to verify full state + Docs walkthrough
- [lubomir] do we need this PR for 1.12.0:
 - "[kubeadm] warn if the CIDR is missing"

<https://github.com/kubernetes/kubernetes/pull/68682>
- [tim] Walk the 1.12 final list..

Wed 12 Sept 2018 - 9am PT, noon ET

Attending

- Lubomir I. Ivanov, Ross Georgiev - VMware
- Rafael Fernández López, Marek Counts - SUSE
- Fabrizio Pandini

Agenda/Notes

- [lubomir] failing periodic e2e test [kubeadm-gce-stable-on-master](#)
 - started failing on Sep 10, 2018
 - tracking issue: [1109](#)
 - possible commit suspect: [68473](#)
 - related fix PR by a contributor: [68534](#)
 - Action:
 - Fabrizio & Lubomir deep dive on failing test to determine if the issue is related to controller manager. **Update:** [confirmed](#) that we have a "stable on master" issue.
 - Ross: follow up with the contributor on the PR and test it
- [fabrizio] PSA on v1.12 issues: help in testing is more than welcome
 - annotate CRI phase - [68449](#) merged
 - preflight master - [68446](#) merged
 - Dynamic Kubelet - waiting for help from sig-node
 - tracking issue [1102](#)
 - [background info](#)
 - Action: Fabrizio: ping sig-node
 - Etcd certs - [68485](#)
 - Controller manager (see above)
- [fabrizio] Are there constant to change in kubeadm before release cut e.g.


```
`MinimumControlPlaneVersion = 1.10`?
```

 - action: [lubomir] investigate if we need to update this. **update:** sent PR:

<https://github.com/kubernetes/kubernetes/pull/68649>
- [lubomir] enabling logs from pods for kubernetes-anywhere based e2e.
 - we need logs from pods and state of nodes because these e2e are very difficult to debug.

- trying to figure out if the change can be made on the k-a side and not in test-infra.
 - action: [lubomir] investigate the required effort for making this change. discuss with sig-testing and test-infra folks if needed. **update:** investigated seems quite doable on the k-a side.
- Welcome to new a contributor: @Marek from SUSE

Wed 5 Sept 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Liz Frost, Vince Prignano, Chuck Ha - Heptio
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Flavio Castelli, Rafael Fernández López - SUSE

Agenda/Notes

- [Tim] CRI + docs
 - <https://github.com/kubernetes/website/pull/10186/files>
 - Test automation is old
 - Need to double check current version being used in test automation, should be 18.06
 - <https://kubernetes.io/docs/setup/independent/install-kubeadm/#installing-docker>
- [Tim] Walking the backlog

Wed 29 August 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair, Jason DeTiberus, Chuck Ha, Vince Prignano - Heptio
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Marek Counts, Rafael Fernández López, Ty Daines - SUSE
- Alexander Kanevskiy - Intel

Agenda/Notes

- [lubomir] Removing 1.8 and 1.9 e2e tests
 - tracking issue: <https://github.com/kubernetes/kubeadm/issues/863>
 - testgrid: <https://k8s-testgrid.appspot.com/sig-cluster-lifecycle-all>

- [tim] 1.12 triage and backlog burndown
 - we need tests for new features
 - we need to solidify the config before that
- [lubomir] Status for `alpha phases` in 1.12:
 - <https://github.com/kubernetes/kubeadm/issues/454>
 - [tim] we wanted to wait for the beta config to be ready
 - [tim] we can still move these before the config is in place
 - [tim] we can have a PR for that and merge it when 1.13 opens.
- [lubomir] Config discussion
 - UUIDs as keys for "status" in the config-map
https://github.com/kubernetes/kubernetes/pull/67944#discussion_r213485283
 - [tim] we should probably go with the UUID idea and vendor in kubeadm.
 - [tim] we should talk to others about this and get more opinions.
 - [jason] the host name should be unique we can use that
 - [alexander] there are related bugs to host names and join
 - We have FeatureGates in the JoinConfiguration, is this desired?
<https://github.com/kubernetes/kubeadm/issues/1088>
 - [tim] we should file an issue an investigate this.
 - Worker nodes don't have an image repository and they can only obtain it from the ClusterConfiguration (config map).
 - [tim] this is on demand and we might add it eventually.
- [lubomir] see if we can optimize this docs page for 1.12
<https://kubernetes.io/docs/reference/setup-tools/kubeadm/kubeadm-init/>
- [lubomir] we need someone to help with docs about setup of different CRIs. Ed Bartosh offered help.

Wed 22 August 2018 - 9am PT, noon ET\

recording

Attending

- Tim St. Clair, Chuck Ha, Jason Detiberus - Heptio
- Leigh Capili - Beatport
- Mike Bryant - Ocado Technology
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Richard Brown, Rafael Fernández López, Marek Counts - SUSE
- Fabrizio
- Archy - CloudOps

Agenda/Notes

- [Fabrizio] - kubeadm join --control-plane
 - Status: [PR66973](#) comments are going to be addressed;

- How to detect this nodes hosts a control plane?
check if manifest exists?
 - Integration with kubeadm config v1alpha3 probably shifts to v1.13
(workaround already implemented using cli flags)
 - PSA: Help for manual testing needed
- [Leigh] User feedback on SelfHosting deprecation
 - Mike from Ocado can speak to user stories / support
 - Their clusters are created using kubeadm phases with some custom logic. This will be possible also in v1.12 (alpha phase self hosting will be there until code cleanup)
 - Full cluster failures are solved with bootkube
 - Evaluate in the future whether to move alpha phases for self-hosting features into a separate sub-command (where the the user explicitly opts-in to the feature)
 - Should consider implementing un-pivoting -- right now, there is no supported downgrade path for SelfHosted enabled clusters
- [Fabrizio] - Update on kubeadm config v1Beta1 [issues/963](https://github.com/kubernetes/kubernetes/issues/963)

--> For v1.12, it is reasonable to complete as much as possible in v1alpha3, but not cutting v1Beta1

--> Define target for "as much as possible" considering the current status

1) Restructure API

- UNBLOCK <https://github.com/kubernetes/kubernetes/pull/67441> Split InitConfiguration and ClusterConfiguration
- (IN PROGRESS @fabrizio) Improve controlplane configuration by using substructs + n control plane instances
 - Co-design/early review API changes
- (IN PROGRESS @rosti) Refactor the Discovery joining API types
- Improve add-on configuration by using substructs
- Minor improvements e.g. improve how to configure images for control plane components

2) Change how kubeadm store/consumes config @Fabrizio

- **Currently updates are broken! Fix it**
- Do not store component config in kubeadm configMap (instead use existing kubelet/kubeproxy configMaps)
- Do not store node specific info in kubeadm config (instead use node attributes/annotations)
 - How to get the node name/get node ref from kubelet?
- Store info about n>1 control plane instances

3) Code cleanup (TODO in restructure API PRs)

- some of them IN PROGRESS @rosti
- /Hold until all the API changes are complete

Wed 15 August 2018 - 9am PT, noon ET\

[recording](#)

Attending

- Leigh Capili - Beatport
- Jason DeTiberus, Liz Frost, Chuck Ha - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Alexander Kanevskiy - Intel

Agenda/Notes

- [lubomir] - remove / deprecate some feature gates for 1.12:
 - <https://github.com/kubernetes/kubeadm/issues/1058> (created by fabrizio)
 - we can ask someone from api-machinery about the deprecation policy
 - <https://kubernetes.io/docs/reference/command-line-tools-reference/feature-gates/>
 - we got a response that it's safe to remove the flags without deprecation as these are alpha. a release note should suffice.
 - HA is already supported so we can simply remove this flag.
 - Self hosting was discussed before and we should remove it.
 - [leigh] CertsInSecrets stores PKI in a secret and then is used in static pods.
 - [liz] we should remove all features which don't have owners.
- look at the backlog for 1.12:
 - <https://github.com/kubernetes/kubeadm/milestones/v1.12>
 - [lubomir] we went through the backlog for 1.12 and moved some items to the GA milestone. also reduced priority of some items. we reached "Tracking issue for "Config to v1beta1" on the first page.
 - about debs and rpms moving to CNCF hosting, i got a response from the release team that this isn't happening anytime soon.
<https://github.com/kubernetes/kubeadm/issues/969>
- Leigh - Support Core Metrics Server Addon?
Minimum, we should reference docs from the addons page.
 - Required by HPA's / kubectl top / dashboard
 - Are any kube-scheduler features dependent on this yet?

- Slack Thread:
<https://kubernetes.slack.com/archives/C13J86Z63/p1534298957000100>
- [leigh] is going to send a docs/website PR for this.
- not making the cut past feature freeze for 1.12
- consider for 1.13 as metrics server becomes closer to GA
- jason and leigh identify work needs to be done to improve user-facing docs and productionalize deployments to be a safe default component

Wed 8 August 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair, Chuck Ha, Liz Frost, Jason DeTiberus - Heptio
- Ross Georgiev, Lubomir I. Ivanov, Tim Pepper - VMware
- Fabrizio Pandini
- Matt Melgard - Charter Communications
- Alexander Kanevskiy - Intel
- Martin Dietze - Dermalog
- Leigh Capili - Beatport

Agenda/Notes

- [chuck] Let's [close milestones](#) unless there is some reason not to?
 - Close milestones generally after the .1 releases [tim to close existing milestones]
- [chuck] The release code is not tied to the actual code so we have a ton of version check logic and now a bunch of [unhappy users](#).
 - Unknown status of bazel
 - [tim] bazel was blocked on was multi-arch builds
 - [liz] bazel has some issues with cgo
- [Tim] CLI -> flags -> config issues.
- [jorge/lubomir] (continued, low-prio) Jorge Castro notified about the launch of a board, where SIGs can propose what areas they need help with community work (jobs). Maybe we can create a list of such areas?
 - the board would be somewhere under
<https://discuss.kubernetes.io/c/contributors>
 - tracking issue for kubeadm
<https://github.com/kubernetes/kubeadm/issues/1042>

- [\[the issue\]](#) Chuck will write up an issue regarding release & packaging (but mostly release)
 - This could graduate to a KEP that's useful for other people building Go binaries
 - Discuss distro-patched Go's and build provenance
 - os-packages? container? snap?
 - todo; link tpepper
- We should s/master/control-plane

Wed 1 August 2018 - 9am PT, noon ET

(skipped since all the SIG leads are busy).

Wed 25 July 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Liz Frost - Heptio
- Lucas Käldestrom - luxas labs (contracting for Weaveworks)
- Tim Pepper - vmware
- Dims
- Alexander

Agenda/Notes

- [lucas] The ComponentConfig KEP is merged: <https://github.com/kubernetes/community/pull/2354>
 - Sync with Stefan and Di about trying to make the shift.
 - Q: Should we co-maintain? A: Partially at least
 - <https://github.com/kubernetes/kubernetes/pull/66058>
 - <https://github.com/kubernetes/kubernetes/pull/66059>
 - <https://github.com/kubernetes/utils/pull/41>
- [lucas] v1beta1 configuration stuff will be on the table for quite a bit of time.
 - https://docs.google.com/document/d/1Sk2VW7IKaLjrjf_4a1NEU0I7X0ysZoXq9IVBvH-ruYY/edit#
 - Liz, Fabrizio and Tim reviewed the whole document
 - [tim] The plan makes a ton of sense from a software design perspective, not ideal from the UX migration perspective
 - [liz] The biggest problem is docs related to this. This is how the new thing works, this setting has moved here, etc.
 - Split MasterConfiguration to InitConfiguration and ClusterConfiguration <https://github.com/kubernetes/kubernetes/pull/66219>
- [tim] PSA - Bootstrap-token needs the fix we discussed in client-go. I'll bump to P0.
 - Will sync with Stefan to make the change request for moving out.

- [tim] I'd like to rip out self-hosting.
 - We should mark-deprecated 1st, in v1.12, and remove completely in v1.13
- [luxas] <https://github.com/kubernetes/community/pull/1939>
-

Wed 18 July 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Chuck Ha🌟, Liz Frost - Heptio
- Fabrizio Pandini
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Matt Melgard - Charter Communications
- Alexander Kanevskiy - Intel

Agenda/Notes

- [Fabrizio] [Copy edit of KEP for v1Beta1 API](#) is going to be translated into a real KEP soonish
- [Tim] want to remove self-hosting before HA master join scenario.
- [Dims] Splitting the repo
- [Alexander] - has requested to go for session in Shanghai, but it's still TBD.

Wed 11 July 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Chuck Ha🌟 - Heptio
- Lubomir I. Ivanov, Ross Georgiev - VMware
- Hippie Hacker - ii.coop contracting for CNCF

Agenda/Notes

- [leigh] Discuss updates to Test plan issue: <https://github.com/kubernetes/kubeadm/issues/772>
- [lubomir] kubeadm init / join detection on the same node: <https://github.com/kubernetes/kubeadm/issues/974>
 - do a pre-flight check for now - simple solution
 - write an issue / proposal for a "command log" - long term solution (both solutions should use a sentinel file in /var/lib...)
- [HH] Using kubeadm with kubetest / e2e for Packet and beyond
 - kubernetes-anywhere is deprecated (or in the process)
 - cluster-api

- Existing in-tree providers (gcp/vsphere):
<https://github.com/kubernetes-sigs/cluster-api/tree/master/cloud>
- There has been discussion around adding a generic terraform provider as well.
- Working group meeting notes:
<https://docs.google.com/document/d/16ils69KImmE94RlmzjWDrkmFZysgB2J4IGnYMRN89WM/edit#>
- [HH] What's the state of getting all e2es on ARM
 - <https://github.com/kubernetes/kubernetes/issues/38067#issuecomment-403182388>

Wed 4 July 2018 - 9am PT, noon ET

(US holiday. mostly EU attendance)

[Recording](#)

Attending

- Lucas Källdström - luxas labs (contracting for Weaveworks)
- Fabrizio Pandini
- Ross Georgiev, Lubomir I. Ivanov - VMware

Agenda/Notes

- [lucas & fabrizio] v1beta1 Discussion
<https://github.com/kubernetes/kubeadm/issues/963>
-

Wed 27 June 2018 - 9am PT, noon ET

Attending

- [Fabrizio Pandini]
- Tim St. Clair, ☀️ Chuck Ha, Liz Frost, Jason DeTiberus - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Alexander Kanevskiy - Intel
- Leigh Schrandt - 🎧 Beatport

Agenda/Notes

- 1.12 planning continued:
<https://docs.google.com/document/d/1YP8sCtgo58yqQ9sR4nsn9iuVFmgICCBMivCTwjd1ERE/edit#>
 - Define priorities

- Deep dives
 - [Fabrizio] Added design notes to the document about [v1beta1](#) API (@luxas please merge the suggestion)
 - [Fabrizio] Copy edits of the [updated KEP for kubeadm join --master](#) (I will wait the end of the week before sending corresponding PR)
- [leigh] Fill out test plan issue: <https://github.com/kubernetes/kubeadm/issues/772>
-

Wed 20 June 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Chuck Ha 🌟, Liz Frost, Ruben Orduz, Jennifer Rondeau - Heptio
- Leigh Capili - 🎧 Beatport
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Lucas Källdström - luxas labs (contracting for Weaveworks)

Agenda/Notes

- [Jennifer] Proposal for restructuring the kubeadm documentation: https://docs.google.com/document/d/1-pAxBmV7c0DfbJVU33oK4Cozd99Z_16HX1-MLZ-A6OY/edit?usp=sharing
- [Tim] Pushing alpha content to kubeadm repo

Wed 13 June 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair, Jason DeTiberus, Liz Frost, Chuck Ha - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Lucas Källdström - luxas labs (contracting for Weaveworks)
- Alexander Kanevskiy - Intel

Agenda/Notes

- [Tim] Current state of docs
 - Walked through the backlog and pruned the list.
- [Tim] Testing is green for we need additional eyes on certain areas?
 - We should build debs/rpms that people can test e2e
 - [lucas] Install kubeadm on 1.10

- Fetch from head and upgrade to latest 1.11 using upgrade apply
- Fetch latest debs/rpms and install them and verify it's working
 - Check kubelet config / unit file changes
- Added doc issue for upgrade
- [lucas] We should edit the release notes
 - Need to talk with nick-chase about the release notes.
 - We should use google docs, but write in markdown instead of using the gdocs formatting.

Wed 6 June 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Jason DeTiberus, Liz Frost, Jennifer Rondeau, Chuck Ha - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Martin Dietze - Dermalog
- Fabrizio Pandini
- Leigh Capili - Beatport

Agenda/Notes

- Leigh will submit patch to make --cluster-dns kubelet flag match serviceSubnet
- [Jennifer/Lubomir] kubeadm pending docs for 1.11 (Umbrella issue for kubeadm docs): <https://github.com/kubernetes/website/issues/8846>
 - status for CoreDNS docs PRs: <https://github.com/kubernetes/kubeadm/issues/843> (merged). one concern here in terms of docs/UX is that the CoreDNS service is called kube-dns. context from the sig-networking folks: https://github.com/kubernetes/website/pull/8829#discussion_r193170655
 - CoreDNS is not getting renamed, it is still kubedns.
 - This is good because users shouldn't have to care about the underlying technology.
 - All command docs are submitted as PRs, just one that Lucas is working one (<https://github.com/kubernetes/kubernetepull/64624>).
 - kubeadm user stories: <https://github.com/kubernetes/kubeadm/issues/849>
 - [Lubomir] should we do this now or in the next few weeks?
 - [Jennifer] docs folks don't have time for this in 1.11 right now. Best to do this anytime after 1.11 is out the door.
 - kubeadm config: including godocs links vs inlining static configs vs generating configs in the future: <https://github.com/kubernetes/kubeadm/issues/842>
 - [Lubomir] What should we do with configs? link out? generate them? Current idea is to encode current version (v1alpha2) and link to the godoc.

- [Tim] We should have a separate page that links out to all the detail of the config object. Mirroring content is bad. If docs link to godoc, then it forces devs to make sure the docs are up to date.
 - [Jennifer] We should branch for now. But relying on godocs is completely different than how the rest of the docs website works. Could bring it up to sig-docs. Make a proposal and present to sig-docs if we want to change how we reference code.
 - **[Action Item]** Generate data into docs and figure out a long term solution later.
 - [Tim] There has to be a deprecation policy on the old config
 - move milestone to 1.12 for some of those, if needed?
 - anything else?
- [Jason or Lubomir] [List of known bugs in our issue tracker](#) (this possibly overlaps with the above)
 - It was suggested to add it here:
<https://github.com/kubernetes/kubernetes/issues/64806>
- Walk-Through + Enumerate the docs list
 - Under Bootstrapping Clusters we want a Creating single master cluster and a Creating a multi-master cluster (both with kubeadm).
 - [Lubomir] Let's copy pages we want to discuss in detail into google docs so we can comment.
 - Content discussion of [Creating HA clusters with kubeadm](#)
 - Remove [Building High-Availability Clusters?](#)
 - Content discussion of [Using kubeadm to create a cluster](#)
 - Content discussion of [Troubleshooting kubeadm](#)
 - [Tim] Additions to trouble shooting guide:
<https://github.com/kubernetes/kubeadm/issues/887>
 - We need to de-dupe the issues [Tim, Lubomir]
 - We don't have cloud provider instructions for kubeadm - the docs are here, but we should probably only link to them + some brief details, as we don't want to expand on cloud provider specifics:
<https://kubernetes.io/docs/getting-started-guides/cloudstack/> tracker issues is here: <https://github.com/kubernetes/kubeadm/issues/526>
 - List of kubeadm docs topics:
https://docs.google.com/spreadsheets/d/1SQ-_FYq9vDCvizlQUZlhfK5lydPf9loCBYIk5m7HwWw/edit?usp=sharing
 -

Wed 30 May 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair, Jason DeTiberus - Heptio
- Lucas Käldestrom - luxas labs (contracting for Weaveworks)
- Fabrizio Pandini
- Alexander Kanevskiy - Intel
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Leigh Capili - Beatport

Agenda/Notes

- [tim+] cgroup driver
 - The default cgroup driver in both docker and kubelet is cgroupfs
 - We want to catch the case when someone is using the systemd cgroup driver for docker. e.g. on CentOS docker is using systemd by default
- [lucas] Patch annotation for the crisocket on the node
 - <https://github.com/kubernetes/kubernetes/issues/64460>
- [lucas] Bootstrap Token API changes in v1.11
 - <https://github.com/kubernetes/kubernetes/pull/64408>
- [lucas] kubelet dropin changes
 - [leigh] assigned issue for resolved auto-config
<https://github.com/kubernetes/kubeadm/issues/845>
 - dropin example:

...

[Service]

Environment="KUBELET_KUBECONFIG_ARGS=--bootstrap-kubeconfig=/etc/kubernetes/boostrap-kubelet.conf --kubeconfig=/etc/kubernetes/kubelet.conf"

Environment="KUBELET_CONFIG_ARGS=--config=/var/lib/kubelet/config.yaml"

EnvironmentFile=/var/lib/kubelet/kubeadm-flags.env

EnvironmentFile=/etc/default/kubelet

ExecStart=/usr/bin/kubelet_v111 \$KUBELET_KUBECONFIG_ARGS

\$KUBELET_CONFIG_ARGS \$KUBELET_KUBEADM_ARGS \$KUBELET_EXTRA_ARGS

...

- [lucas+tim] Starting to test v1.11
 - Docs + Test sprint next week.

Wed 23 May 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair - Heptio
- Lucas Käldestrom - luxas labs (contracting for Weaveworks)
- Martin Dietze - Dermalog
- Ross Georgiev - VMware
- Matt Kelly - Containership
- Alexander Kanevskiy - Intel

Agenda/Notes

- [luxas] v1beta1 API changes
 - https://docs.google.com/document/d/1Sk2VW7IKaLjrjf_4a1NEU0I7X0ysZoXq9IVBvH-ruYY/edit#
-

Wed 16 May 2018 - 9am PT, noon ET

Recording

Attending

- Tim St. Clair, Jason DeTiberus, Chuck Ha - Heptio
- Ross Georgiev, Lubomir I. Ivanov - VMware
- Lucas Käldestrom - luxas labs (contracting for Weaveworks)
- Alexander Kanevskiy - Intel
- Leigh Capili - Beatport
- Ryan Phillips, Michalis Kargakis - Red Hat

Agenda/Notes

- [luxas] Discuss <https://github.com/kubernetes/kubeadm/issues/822>
 - Lots of conversation(s)...
- [tim] glog all in or what?
- [jason] Backport of external etcd upgrade to 1.10
 - <https://github.com/kubernetes/kubernetes/pull/63925>
- Other topics

Wed 9 May 2018 - 9am PT, noon ET

Attending

- Jason DeTiberus, Chuck, Tim St. Clair - Heptio
- Martin Dietze, Dermalog
- Lucas Kåldström - luxas labs (contracting for Weaveworks)
- Alexander Kanevskiy - Intel
- Michalis Kargakis - Red Hat
- Ross Georgiev, Lubomir I. Ivanov - VMware

Agenda/Notes

- [tim/lucas/dims] CI issue(s)
 - [tim] Try to get a meeting setup at a reasonable time tomorrow together with SIG Testing folks.
- [tim] Creating a new GA tracking issue.
 - Update the kubeadm to GA doc [TODO link] and create issues for all subsections of that doc in Github with the GA milestone
- [mbert] Host *etcd* rpms/debs on Kubernetes repo
 - <https://github.com/kubernetes/website/pull/8301>
- [lucas] kube-dns/CoreDNS on systemd-resolved systems
 - Problem: `/etc/resolv.conf` contains `nameserver 127.0.0.53` which breaks both kube-dns and CoreDNS
 - <https://github.com/kubernetes/kubeadm/issues/787>
 - AI: Create a preflight check for this scenario
 - AI: Take back control over the kubelet config
 - AI: Switch to CoreDNS by default
 - AI: Try to mitigate this issue by specifying {some extra parameter} in the CoreDNS configmap

Wed 25 Apr 2018 - 9am PT, noon ET

Attending

- Jason DeTiberus, Chuck - Heptio
- Liz Frost - Heptio
- Hippie Hacker - ii.coop
- Lubomir I. Ivanov - VMware
- Matt Kelly - Containership

Agenda/Notes

- [leigh] <https://kubernetes.slack.com/archives/C13J86Z63/p1524617195000021>
 - “Kubeadm version v1.10.2-beta.0.77+d842c660d63b5e can only be used to upgrade to Kubernetes versions 1.10”
 - Disabling 1.9.x -> 1.9.y upgrade also disables the 1.10.x -> 1.9.y downgrade.
 - Should we restructure the version check to allow this so users have a downgrade path?
 - kubeadm 1.9 can do 1.9.x -> 1.9.y, but it cannot downgrade from 1.10
- [leigh] should we keep the old manifests around in the backup dir since they're cheap to store?
 - related issue: <https://github.com/kubernetes/kubeadm/issues/489>
 - store etcd backup
 - store old config
- [liz] e2e tests, initial pass to get into place.
 - <https://github.com/kubernetes/test-infra/pull/7833> the kubetest PR / <https://github.com/kubernetes/kubernetes/pull/62665> the kubetest e2e
 - [Tim] - Open an issue about creating a test - plan.
- contrib experience -- we have a kubeadm specific session on Friday next week at KubeCon EU
 - talked with luxas and peterzhao about improving ContribX
 - how to build / run / test
 - project layout / coding standards
 - <https://github.com/stealthybox/vagrant-kubeadm-testing>

Wed 18 Apr 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Liz, Chuck, Jason DeTiberus - Heptio
- Fabrizio Pandini
-

Agenda/Notes

- Point of attention about phases
KEP: (<https://github.com/kubernetes/community/pull/1939>)
 - the init workflow - today written in a single func - will be replaced by a set of functions and by a struct that defines the workflow
-> to be mitigated implementing phases step by step
 - all the phases will have the same flags/no dedicated helps or examples
-> acceptable, because flags are there only for the simplified story; everything else need the config file

- over time other function were added to phases. those functions should be now "relocated"
- > to be mitigated implementing phases step by step
- are we planning to divide other commands into phases?
 - TODO: write down possible workflows / command breakdowns
- Config
 - What are relevant versions for managing upgradability?
 - relevant versions:
 - kubeadm api version
 - kube-proxy configuration api version
 - kubelet configuration api version
 - kubernetesVersion or kubeadmVersion are less important
 - typed fields vs string
 - single node pov (either master or slave) vs cluster pov (HA readiness)
 - internal etcd / external etcd
 - https://docs.google.com/document/d/1LNaGKzjyLtPnLbRbW4QUfwVEJS3XcyQy_JmWyV-bVbU/edit#heading=h.zd8ht9pthabe
 - TODO: collect requirements for the KEP

Wed 11 Apr 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Liz, Craig Tracey, Jason DeTiberus, Eryn Muetzel, Chuck Ha- Heptio
- Cole Mickens, Ryan Phillips, Michalis Kargakis - Red Hat
- Matt Kelly - Containership
- Fabrizio Pandini
- Leigh Capili - Beatport

Agenda/Notes

- Update upgrade issue
 - PR to fix race condition is ready but is pending on the upgrade ordering problem.
 - For 1.10, we will remove tls-enablement as part of the upgrade
 - Will re-enable tls-enablement for the 1.11 upgrade
 - 1.10 -> 1.11 upgrades will need to be able to upgrade tls and non-tls etcd static pods.
 - Started a PR using the kubelet api & etcd client to grab status
 - What about using liveness endpoint, and it doesn't show up.
 - <https://github.com/kubernetes/kubernetes/pull/62358>
 - Will engage with sig-node on enhancements to kubelet api
 - potential way to manage static pods rather than filesystem manipulation

- way to retrieve static pod status through the kubelet api, currently static pods do not receive status updates through the kubelet api.
 - Stability of the parts of the kubelet api we are using
 - Idempotency of tls phases
 - Config file / Config Map.
 - Relied on a bunch of other structs, that changed in an incompatible way to deserialize.
 - Stop gap
 - Link to PR ()
 - We need a KEP for long term support for the configuration file
- Self hosting
 - [Cole] quick intro
 - Discuss a different approach
 - kubeadm places a “sentinel” pod on master nodes that runs a temporary apiserver when the selfhosted controlplane is not running for the kubelet to refetch its Bound Pods for the controlplane DaemonSets
 - after the selfhosted apiserver starts, the sentinel pod stops the temporary apiserver so that the kubelet can pivot to the selfhosted apiserver
 - the sentinel pod continues running and writes the apiserver config to disk
- Talk about HA plan
 - Currently breaking apart the WIP into a smaller PR
 - And sending

Action Item

- [Tim] Follow up with the release team to get the upgrade tests back in place.

Wed 4 Apr 2018 - 9am PT, noon ET

Attending

- Tim St. Clair, Liz, Craig Tracey - Heptio
- Matt Kelly - Containership
- Ryan Phillips - Red Hat

Agenda/Notes

- [Tim] Debugging high priority items atm.
 - We need a better jig and set of tests for upgrades.
- HA testing for kubeadm features will be difficult.

- Happy path might be to setup an external jig that stands up the HA environment and runs the tests that get reported back to testgrid.
- [Tim] State of HA-Config work (Matt)
 - 1st easy ½ done - working on the upgrade portion.
- [Tim] Write down self-hosted issue with just the Controller Manager and Scheduler.

Wed 28 March 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Leigh Capili - Beatport
- Jason DeTiberus, Kris Nova, Patrick Barker - Heptio
- Matt Kelly - Containership
- Stefan, Ilya - Weaveworks
- Alexander Kanevskiy - Intel

Agenda/Notes

- [Stefan] Upgrade issue from 1.9 to 1.10
 - Replicator
 - Download and install 1.9.6 version of kubeadm
 - kubeadm init
 - install console
 - Download and install 1.10 version of kubeadm
 - run kubeadm upgrade
 - Starting from 1.9.6 upgrade initially crashes - featuregate changed from string to map. Documented change, but manual process to update makes it not overly clear.
 - [kad] maybe we can do kubeadm tweak in upgrade phase to do that conversion automatically ?
 - Timeout during etcd upgrade
- [Matt] working on POC for multi-master ConfigMap
- [Patrick] begin discussing proposal on dynamic audit configuration

Wed 21 March 2018 - 9am PT, noon ET

No Recording today, as it was a very lightly attended meeting

Attending

- Robert Bailey, Kenny Leung - Google
- Leigh Capili - Beatport
- Matt Kelly - Containership

- Fabrizio

Agenda/Notes

- [leigh] Propose `ExternalEtcd` FeatureFlag?
 - Currently behavior is gated behind existence of `cfg.Etcd.Endpoints`. This is starting to get sprinkled all throughout the phases code.
 - Difficult to use this property for static pod implementation.
 - Previously was possible to supply some TLS flags but not others
 - Currently running into this with the SelfHosting `StoreCertsInSecrets` implementation
- [leigh] [SelfHosting /w StoreCertsInSecrets](#)
 - Static Pod Etcd TLS be ready for 1.10.1 cherry pick
 - External Etcd TLS support does not look to be implemented
- [leigh] Talk from yesterday on support for containerized installations: <https://kubernetes.slack.com/archives/C13J86Z63/p1521564747000544>
 - Also see <https://github.com/v1k0d3n/gantry>
- [matt] KEP for kubeadm config for HA is up, please review: <https://github.com/kubernetes/community/pull/1915>
 - KEP is fully updated with all outstanding comments addressed
 - Outstanding question is how we minimize impact to users that are overriding configs as we are introducing multiple configs
 - Should we keep the config maps as disjoint (as proposed in the KEP)?
 - User would have to provide two config files which is a bummer
 - Or keep the existing config the same?
 - User could specify one file but as an implementation detail we figure it out and split them up
 - Half working on proof of concept; hoping to have something out soon
- [Fabrizio] PSA <https://github.com/kubernetes/community/pull/1939>
 - Merging phases into init
 - Expect to finish a PR in a day or two to showcase what the proposal is explaining

Wed 14 March 2018 - 9am PT, noon ET

Meeting Cancelled via Slack — Please work on and review in-flight items

Wed 7 March 2018 - 9am PT, noon ET

Recording

Attending

- Tim St. Clair, Jason DeTiberus - Heptio
- Martin Dietze - Dermalog

- John Harris - Docker
- Leigh Capili - Beatport
- Matt Kelly, Norman Joyner - Containership
- Ryan Phillips - CoreOS/Red Hat

Agenda/Notes

- [leigh] Status on Release Tests
(<https://k8s-testgrid.appspot.com/sig-release-1.10-blocking>)
- [leigh] Will be drafting Contributor Doc for kubeadm /w Reviewer Cohort -- will be public -- will update next week
 - Potentially talk to Contributor Experience Working Group
- [tim] status of HA upgrade doc
 - [md] HA upgrade guide PR needs another review:
<https://github.com/kubernetes/website/pull/7557>
- [tim] master --join scenario
 - Matt will do one more review prior

Wed 28 Feb 2018 - 9am PT, noon ET

Recording

Attending

- Leigh Capili - Beatport
- Matt Kelly, Ashley Schuett, Norman Joyner - Containership
- Jason DeTiberus, Tim St. Clair - Heptio
- Ryan Phillips - CoreOS/Red Hat

Agenda/Notes

- [Matt] Discuss [Fabrizio's proposal to make kubeadm upgrade HA ready](#)
 - Is the approach of node-specific ConfigMaps to store master metadata acceptable to move forward with?
 - Tim to comment on issue (linked above), Matt will follow up after
- [leigh] Separate [etcd CA patch](#) is up -- needs 1 more /approve
- [Tim] Discuss more about using sidecar pods for load balancing
 - we should document an approach for cluster load balancing
 - a Nodeport style mesh for the "kubernetes svc" would be a good capability to have -- since routes to all of the apiservers will be distributed to all of the nodes in the cluster, kubelets will be able to use this routing to load balance to the controlplane in an HA manner
 - could make it work with kube-proxy and saving IPTables to disk, but this is ugly -- using a formal loadbalancer has benefits

- see fabrizio's doc for some considerations (needs link)

Wed 21 Feb 2018 - 9am PT, noon ET

Recording

Attending

- Leigh Capili - Beatport
- Martin Dietze - Dermalog
- Matt Kelly - Containership
- Tim St. Clair, Jason DeTiberus - Heptio

Agenda/Notes

- [leigh] [Etcad TLS](#) discuss feedback about providing a separate CA arg.
 - Upstream etcd features block meaningful progress with a shared CA approach
 - Tim suggests moving forward with separate CA for this PR
- Martin - Discuss feedback from PR documenting on prem - deployment.
 - We need to get the documentation in place for folks who have referenced
- [Matt] deprecate old docs related to kubeadm HA proposals

Wed 14 Feb 2018 - 9am PT, noon ET

Recording

Attending

- Jason DeTiberus, Tim St. Clair - Heptio
- Leigh Capili - Beatport
- Martin Dietze - Dermalog
- Matt Kelly - Containership

Agenda/Notes

- [leigh] [Etcad TLS](#) update -- rebased and passing tests
 - not ready for 2nd review yet
- [ilya] Discuss Removal of Kubernetes Anywhere
- [leigh] We've mentioned using dind-cluster for replacing hack/local-cluster-up
 - does it have any place in test-infra?
 - [tim] not to replace core kubernetes-anywhere but really good for integration testing effort /w 3rd party projects
-

Wed 7 Feb 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Martin Dietze - Dermalog
- Rodrigo De Castro - Google
- Leigh Capili - Beatport
- Matt Kelly, Ashley Schuett - Containership
- Tim St. Clair - Heptio

Agenda/Notes

- [Leigh] etcd UX -- looking into this
 - We want it to be easy for operators to work with etcd clusters
 - Is this in scope? You can already hack around with the phases commands today -- we have etcdctl but that doesn't bootstrap etcd nodes
 - Tim suggests `kubeadm etcd` / `kubeadm phases controlplane etcd...`
 - propose a KEP in the community repo? Tim suggests still doing this
 - Talked to Lucas -- we should start small with a docs update on how to setup multi-node etcd with kubeadm and the new cert generation features
 - [tim] request for any etcd management code to be separated into own package -- leverage this in phases
 - [tim] discussion on multi-node/cluster cert distribution
 - CRD's?
 - existing static-pod HA controlplane doc just suggests user driven SCP
 - [tim] suggests looking at JustinSB's work on etcd-manager
 - [fabrizio] etcd could be started on separate nodes from the controlpane
- [Martin] (pre-HA) kubeadm HA upgrade instructions
 - Section added to [guide in google docs](#), ansible code in [github](#) updated, needs some feedback
 - Do we feel any need to add this information to the documentation on kubernetes.io eventually?
- [Tim] Talked about what if we completely punt on self-hosting
 - KEPS for:
 - static-pod, multi-node etcd /w kubeadm
 - kubeadm HA upgrades
 - [Fabrizio] Will write-up an issue multi-master config problem with HA.

- [Fabrizio] Will start drafting a prototype/KEP for conflating phases into init - with the goal of improving UX maintainability in the long term

Wed 31 Jan 2018 - 9am PT, noon ET

Recording

Attending

- Martin Dietze - Dermalog

Agenda/Notes

- <add items here>

Wed 24 Jan 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Leigh Capili - Beatport
- Robert Bailey - Google
- Tim St. Clair - Heptio

Agenda/Notes

- [tim] issue triage (on the call, no meeting notes taken)
- [fabrizio] <https://github.com/kubernetes/kubernetes/pull/58698>

Wed 17 Jan 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair - Heptio
- Robert Bailey, Kenny Leung - Google
- Eric Chiang - CoreOS
- Ryan Phillips - CoreOS
- Leigh Capili - Beatport

Agenda/Notes

- Self hosting of etcd
 - Eric: This was brought up at KubeCon, but those conversations weren't recorded or widely distributed
 - Had issues with apiservers not reconnecting to etcd. Sometimes grpc issues; issues with the nebulousness of static pods
 - Customers (of tectonic) have had bizarre crashes of etcd
 - Won't be exploring using self-hosting etcd with the etcd operator any time soon
 - There appear to be some fundamental loops that happen with self hosting etcd that expose bugs in the apiserver / static pods / etc that make it fragile
 - Will explore making some fixes upstream to make this configuration more reliable
 - Not sure how the applies to kubeadm
 - Jamie shouldn't continue to rebase his PR given this information
- How does CoreOS recommend running etcd?
 - CoreOS recommends dedicated etcd hosts not running in static pods
 - Not sure if there is a specific motivation for not doing static pods
 - Masters are often in autoscaling groups, but don't want etcd to autoscale
 - Makes sense architecturally to separate etcd from the master components
 - CoreOS still runs it in a container to make upgrades easier
- Self-hosting kubeadm
 - Last blocking bug: Need to determine what we want to do with secrets for the apiserver
 - Are we going to pursue host mounting or do secrets and look into checkpointing secrets?
 - Ilya: Could we use a volume to store the secrets?
 - Simplest solution is to use host volume mounts
 - If we had a way to sync secrets into volumes we could store the secrets on encrypted volumes
 - The kubelet no longer mounts volumes into the host -- the volume lifecycle is coordinated by code in the controller manager
 - The problem is bootstrapping the apiserver itself, we need to strip down the pod spec to the bare minimum
 - Ilya: If we think of secret checkpointing as a general use case, would we consider configmap checkpointing as well?
 - This opens up the rabbit hole, as you could end up wanting to checkpoint every resource
 - Tim: Kind-of prefer the host mount solution. You could use a daemonset to do cert rotation.

- Ilya: We could apply similar pattern to doing static pod upgrades. Run a daemonset that rotates the static pods on the master(s) rather than logging in and executing kubeadm upgrade
- Ilya has said that he will investigate this in more detail.

Wed 10 Jan 2018 - 9am PT, noon ET

[Recording](#)

Attending

- Tim St. Clair, Craig Tracey - Heptio
- Martin Dietze - Dermalog
- Leigh Capili - Beatport
- Rodrigo De Castro - Google

Agenda/Notes

- HA
 - [Tim] Follow on up on possible takers for documentation on HA regarding kubeadm GA.
 - [Tim] PR to docs.
 - Fabrizio's HA Prototype
 - Feature gating Fabrizio's work.
 - Going to create a PR for review on updating
- [Leigh] Cert Rotation
 - Peter PR'd a feature for apiserver cert rotation on upgrade <https://github.com/kubernetes/kubeadm/issues/548>
 - What is our strategy for rotating certs? Do we want a phase for this that is separate from upgrades?
 - What about rotating the CA?
 - Need to add this feature for the 3 etcd certs
 - [Tim] Suggest talking to Eric Chang, Lucas, Matt Moyer
- [Tim] Self hosting
 - Looking for volunteers to finish off secrets management.

Wed 3 Jan 2018 - 9am PT, noon ET

Attending

- Robert Bailey, Rodrigo De Castro - Google

- Tim St. Clair - Heptio
- Craig Tracey - Heptio
- David Watson - Samsung
- Leigh Capili - Beatport

Agenda/Notes

[Tim] Two topic areas: secret checkpointing and HA

- Leigh is interested in helping with HA
 - The only reason for us to do checkpointing is to bootstrap
 - One idea is to use envoy to bootstrap kubelet to join master
 - Work here with both SIG cluster lifecycle and auth
 - Goal is to enable self-hosting by default
 - Updating the feature proposal (KEP) is needed (Is this it @tim <https://github.com/kubernetes/kubernetes/issues/489>)
- Self hosting is higher priority because the cost of upgrades without it is high
 - Certificate bootstrap problem is a dependency (e.g. checkpointing)
 - HA can be done via ansible and following Jamie's doc
- [Leigh] etcd TLS on init/upgrade
 - <https://github.com/kubernetes/kubernetes/pull/57415>
- [Robert] Review the [kubeadm to GA](#) document
 - Can we GA without HA (contentious)?2f h1-report