


Cluster API Project

[Zoom meeting link](#) (passcode: 77777)


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

Related Documents:

- [AWS Cluster API Implementation Office Hours](#)
- [Azure Cluster API Implementation Office Hours](#)



**For edit access to this document, join
the [sig cluster lifecycle mailing list](#).**



Cluster API Project Meeting Times:

- 10am PT on Wednesdays, weekly (1pm Eastern, 6pm UK time, 1700UTC)

Mission

- [sig-cluster-lifecycle \(SCL\)](#)
- [Cluster API](#)

Meeting Etiquette

- Use the [Raise hand](#) feature of zoom (click participants, then raise hand) to continue current topic by expanding on a point, disagreeing, responding, etc.

- Edit the agenda (this document) to add a new, unrelated topic.
- The moderator for each meeting should help keep us on track by encouraging participants to add new topics to the agenda as we go.

Open Proposals

This section lists all active, open proposals. When a proposal is either moved to a PR or withdrawn, please remove it from this list. The last bullet point links to all open proposal PRs.

- [Bootstrap Reporting](#)
- [Management Cluster Operator](#)
- [KCP Scale-in](#)
- [Load Balancer Provider](#)

Day XX Month - 10am Pacific (template)

Recording TBD

Attending

-

Agenda

- PSAs
- Discussion Topics
Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.
- [name] <<add your items here>>

Wed 9 Dec - 10am Pacific

[Recording](#)

Attending

- Prakash R. - Indi Contributor
- Hardik Dodiya - SAP
- Andy Goldstein, Sagar Muchhal, Sedef Savas, Naadir Jeewa, Fabrizio Pandini, Warren Fernandes, Shyam Sankaran, Nader Ziada, Travis Hall - VMware
- Cecile Robert-Michon, Zach Wachtel, Matt Boersma, James Sturtevant, Mike Kostersitz, Craig Peters, David Justice - Microsoft

- Jan Tilles - Ericsson
- Marcel Müller - Giant Swarm
- Bryan Boreham - Weaveworks
- Amine Hilaly - AWS
- elmiko - Red Hat
- Vincent Batts - Kinvolk
- Dan Finneran - Equinix Metal
- Arvinderpal Wander - AT&T

Agenda

- PSAs
 - [ncdc] personal update
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.
- [marcelmue] Enable running multiple versions of CAPI controllers in parallel [#4004](#)
 - Background: Main concern is that we want to start adopting the upstream controllers, and we have been using the upstream CRDs with our own controllers. We've realised that even though the contract may be stable, but behavioural changes may result from changes in the controller code, which have made us hesitant when changing controllers, which we now do on a case by case basis. There's no mechanism to control which version of a controller is controlling a particular CRD. From the Slack conversation, the discussion is that we run a different management controller and lift and shift resources from one cluster to another. Have therefore made the proposal to run multiple controllers in the same management cluster by reconciling only resources with a particular label.
 - [vince] the core problem is behavioural changes within a version, which we should avoid going towards v1beta1. We need to draw up a plan of action around testing. Whilst the technical solution could be done, but it is an anti-pattern. If you have multiple versions of controllers there may still be API clashes. Could add a selector on controllers, but it wouldn't be a supported path - we need to be better at ensuring behaviour. VMware has attempted to run multiple versions of controllers and it's caused problems for us and are now moving away from it.
 - [marcelmue] Agree that it's an anti-pattern and that it's not easy but we still see it as a necessity. However, not hellbent on the implementation detail, but we have learnt that we do need this level of graduality. These behavioural changes are hard to test for so have erred on the side of over-cautiousness. More than happy to work on the implementation, though understand it wouldn't be super supported but it would help us move towards adopting the upstream controllers.
 - [cecile] What's the blocker from migrating between mgmt clusters
 - [marcelmue] We run 50-60 clusters per management clusters, and it's not unusual for them to have a great variance in k8s versions.

- [andy] Write on the issue some of these blockers, and then we can tease out what to do.
- [sedef] Cutting a v0.4.0-alpha.0 release for the providers that are also upgrading their APIs to v1alpha4.
-

Wed 2 Dec - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham - Weaveworks
- Sagar Muchhal, Andy Goldstein, Naadir Jeewa, Fabrizio Pandini, Nader Ziada, Warren Fernandes, Shyam Sankaran - VMware
- Jason DeTiberus - Equinix Metal
- Vincent Batts - Kinvolk
- David Justice, James Sturtevant, Matt Boersma, Cecile Robert-Michon, Zach Wachtel - Microsoft
- Jan Tilles, Furkat Gofurov - Ericsson
- Danil Grigorev, Christian Glombek, Aravindh Puthiyaparambil, Sebastian Soto - Red Hat
- Hardik Dodiya - SAP
- Chris Hein - Apple
- Arvinderpal Wander - AT&T

Agenda

- PSAs
 - KCP upgrade race condition fixed in <https://github.com/kubernetes-sigs/cluster-api/pull/3964> → 0.3.12
 - Affects v0.3.11 (also affects previous versions but less likely to occur before v0.3.11)
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.
- [vince] Relaxing our backporting guidelines for the current stable release ([PR 3960](#))
- [vbatts] ignition/secrets ([#3761](#), [#3437](#))
 - Few ways to support alternates to cloud-init, like Ignition, with concerns around secrets management.
 - What are the next steps with regards to a PR to include ignition
 - [andy] set up a group and times to chat

- <https://docs.google.com/document/d/1dMUUP7zy4tOhBznaSvqVtH20bmnS8fyUgSn1fU6lZQI/edit>
- [vince] reach out to each infra provider to see what they can offer in terms of expertise. For success getting the google doc around and then the process would be to move to a PR.
- [aravindh] [\[WIP\] Add ability to provision Windows VMs](#)
 - In Openshift, we have an operator that waits for a Windows Machine and then provisions the machine. We need to configure the VM with SSH capabilities. The CAEP suggests using cloudbase-init, and then the providers provision the VM with userdata. This implies the cloudbase-init is pre-baked. Does the customer need to do it, or will Microsoft publish images?
 - [james] The expectation is that users run image builder. Some images will be published on the marketplace.
 - [Christian] we do things on AWS and Azure similarly, the bootstrap data that is written to disk, and on AWS it executes automatically.
 - [david] we can help folks set up images and get them published.
 - [christian] we'd rather not have customers have to bake images, exposing a field for unattend.xml config on the provider spec would be sufficient
 - [cecile] issue with the pr is that it's not clear to users how they can use it and that there's a requirement on an external tool
 - [daniil] we can help out on the proposal about how to consume this data. Right now, we provide the bootstrapping data in a secret that gets to custom userdata.
 - [vince] to clarify, customers should be using image builder to build images. Red Hat need to configure from bare images, so makes sense to look at moving the bootstrapping to the appropriate openshift bootstrap provider.
 - [vince] please open issues around broken documentation.
- [detiber] Thanks to cpanato and other contributors there has been quite a few fixes and updates made to cluster-api-provider-gcp, and we are looking to cut a v0.3.0 release (targeting Cluster API v1alpha3) tomorrow
 - Please reach out if you want to get involved
- [srm09] Please take a look at the [GH issue](#) around CAPI secret type. There is a decision that needs to be made about how to move forward with this and all community input would be appreciated.

Wed 25 November - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham - Weaveworks

- Matt Boersma - Microsoft
- Jason DeTiberus - Equinix Metal
- Fabrizio Pandini - VMware
- Jan Tilles - Ericsson
- John H Terpstra - Dell EMC
- Hardik Dodiya - SAP

Agenda

- PSAs
 - [Fabrizio] v0.3.11 is out !!
 - MHC supports external remediation!
 - KCP remediates unhealthy machines!
 - KCP adds conditions on the CP machines reflecting static pods and etcd status.
 - Questions:
 - [Jan] is there documentation? Yes, the KCP proposal is updated; the MHC configuration does not change (just use the right selector).
 - [Jan] is it possible to use external remediation for CP nodes created by KCP? Yes, KCP takes actions only if MHC sets the ownerRemediation condition.
 - [Jan] Please review <https://github.com/kubernetes-sigs/cluster-api/pull/3857>
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

Wed 18 November - cancelled

This session is cancelled due to Kubecon NA

Wed 11 November - 10am Pacific

Recording TBD

Attending

- Cecile Robert-Michon, Zachary Wachtel - Microsoft
- Jason DeTiberus - Equinix Metal

- Fabrizio Pandini, Sagar Muchhal, Andy Goldstein, Vince Prignano, Warren Fernandes, Nader Ziada, Sedef Savas, Shyam Sankaran, Naadir Jeewa - VMware
- Irvi Aini
- Yuvraj - Nirmata
- Jan Tilles - Ericsson
- Joel Speed, elmiko - Red Hat
- Arvinderpal Wander - AT&T
- Jonathan Bryant
- Chris Hein - Apple

Agenda

- PSAs
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

- [cecile] v0.3.11 release code freeze
 - Decision: End of day tomorrow (November 12th)
 - Release: RC tomorrow afternoon, final tag on Wed 18th morning
 - v0.3.12 for bug fixes only.
- [wfernandes] [Mgmt Cluster Operator CAEP](#) PR
 - Reminder to review the PR, has been open for a while
 - [vince] let's give it another 10 days to give time everyone to review
- [name] <<add your items here>>

Wed 04 November - 10am Pacific

This session is cancelled due to world events.

2020/11/2 - Load balancer meeting

Recording

- [Jan] Can we run the load balancer in the control plane. We don't want it to use a whole entire external set of servers. We are currently create the workload clusters, we delete the management clusters.
- [Andy] Is the Service in this the one of Core v1.Service. If so, for VMC, because it's not normal EC2 instances, we can't put the VMs directly in the vSphere CPI or AWS in tree CPI.

- [Naadir] I think the renamed AWS load balancers project may support this but using a separate CRD
- [Yassine] There may be a chicken and egg situation here if we're using CPI
 - [Moshe] It's a new controller in the management service .
 - [Yassine] If we create the API
- [Jason] There's a limitation that there's only one CPI type allowed for v1.Service
- [Moshe] Believe some CPIs have a filter. Could also have a default implementation, but then there could be 3rd party implementations not using MachineService.
- [Andy] We won't be able to support this today, but if you do have a CPI that is only in tree, then this isn't going to work. There's a question around Service v2 should we use this. We can do everything that is proposed here and not rely on the services API and progress faster, but there's benefits on leveraging the CPI.
 - [Moshe] We wouldn't get rid of the host/port endpoint combo.
 - [Andy] Let's pretend there's only in-tree, then we can't really spin anything up.

Wed 28 October - 10am Pacific

[Recording](#)

Attending

- Ben Moss, Shyam Sankaran, Andy Goldstein, Sagar Muchhal, Nader Ziada, Jason Scarano, Yassine Tijani, Fabrizio Pandini - VMware
- Jan Tilles - Ericsson
- Carlos Panato
- Cecile Robert-Michon, Matt Boersma, Zach Wachtel, David Justice - Microsoft
- Michael Gugino, elmiko, Joel Speed - Red Hat
- Jason DeTiberus - Equinix Metal
- John Northrup - Ripple
- Yuvraj - Nirmata
- Nicole Yson, Dane Thorsen - Newrelic
- Bryan Boreham - Weaveworks
- Xiao - Salesforce
- Arvinderpal Wander - AT&T
- Vincent Batts - Kinvolk
- Hardik Dodiya - SAP
- Prakash Ramachandra - Dell

Agenda

- PSAs
 - [Fabrizio] wrapping up work on v0.3.11
 - 7 PR to go (6 ~ ready), release date still TBD
 - Main themes
 - External remediations
 - Better visibility (machine node conditions, KCP conditions)
 - KCP remediation
 - [Fabrizio] CAEPs
 - [Management cluster operator](#)
 - [KCP scale in](#)

- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [Ben Moss] Permission to create kubemark provider repo under kubernetes-sigs
 - Current repo is <https://github.com/benmoss/cluster-api-provider-kubemark/>
 - [ncdc] API conversion issues
 - <https://github.com/kubernetes-sigs/cluster-api-provider-aws/pull/2074>
 - [kubernetes-sigs/cluster-api#3877](https://github.com/kubernetes-sigs/cluster-api/pull/3877)
 -
 - [name] <<add your items here>>

Wed 21 October - 10am Pacific

Recording TBD

Attending

- Vincent Batts - Kinvolk
- Sagar Muchhal, Andy Goldstein , Gab Satchi, Shyam Sankaran, Warren Fernandes, Sedef Savas, Yassine Tijani , Nader Ziada, Ben Moss, Fabrizio Pandini - VMware
- Bryan Boreham - Weaveworks
- Erwin van Eyk - Platform9
- Yuvraj - Nirmata
- Joe Julian - D2iQ
- Jan Tilles - Ericsson
- James Sturtevant, Zach Wachtel, David Justice, Cecile Robert-Michon, Matt Boersma - Microsoft

- Joel Speed, Danil Grigorev - Red Hat
- Chris Hein - Apple
- Carlos Panato
- John Northrup - Ripple
- Jason DeT huiberus - Equinix Metal
- Andrew DeMaria - Cloudflare
- David Watson
- Xiao - Salesforce
- Hardik Dodiya - SAP
- Jun Zhou - Spectro Cloud

Agenda

- PSAs
 - v1alpha4 is already carrying lots of breaking changes, we should require documentation changes as PRs get merged ([example migration doc](#))
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [Fabrizio] KCP remediation Demo
 - [Fabrizio] CAPI release informing tests are failing CAPG/CAPA. Carlos volunteered to take a look, Thanks!.
 - <https://github.com/kubernetes/kubernetes/issues/95730>
 - <https://github.com/kubernetes/kubernetes/issues/95729>
 - Some project is doing periodic E2E test failure triage, should we do the same for CAPI as well?
 - Google Group for Cluster API alerts: <https://groups.google.com/g/kubernetes-sig-cluster-lifecycle-cluster-api-alerts>
 - [detiber] Scheduling kickoff meeting for Load Balancer Provider working group
 - <https://github.com/kubernetes-sigs/cluster-api/issues/1250>
 - <https://doodle.com/poll/mct22fq43wga8qzp>
 - [jan] Could we open a scale-in PR to the current KCP proposal?
 - https://docs.google.com/document/d/1HWHTL6g45HrrBppGWgDi_1VkW_RMA0Ue3IRBONcqKFKE/edit#heading=h.2lrpo279kcej
 -

Wed 14 Oct - 10am Pacific (template)

[Recording](#)

Attending

- Bryan Boreham - Weaveworks
- Dan Finneran, Jason DeTiberus - Packet
- Jack Francis, Matt Boersma, James Sturtevant, Cecile Robert-Michon, David Justice - Microsoft
- Sagar Muchhal, Andy Goldstein, Vince Prignano, Fabrizio Pandini, Warren Fernandes - Nader Ziada, Sedef Savas, Gab Satchi, Jason Scarano, Ben Moss, Shyam Sankaran - VMware
- Joel Speed, elmiko - Red Hat
- Vincent Batts - Kinvolk
- Jonathan Bryant
- Prakash Ramchandran - Dell
- Yuvraj - Nirmata
- Jan Tilles - Ericsson
- Chris Hein - Apple
- Hardik Dodiya - SAP
- Daniel Lipovetsky - D2IQ
- Jun Zhou - Spectro Cloud

Agenda

- PSAs
 - Feedback needed for Kubernetes and dependencies version upgrades (or required) <https://github.com/kubernetes-sigs/cluster-api/issues/3783>
 - [ncdc] Don't think there are changes that require a newer version of Kubernetes
 - [cecile] One of the concerns was managed K8s weren't up to 1.19
 - [Joel] Do we need to stick to the +/- 2 version compatibility for the client libraries?
 - [elmiko] Can we still operate on 1.17 & 1.18
 - [vince] Generically, given the v1alpha4 release dates, v1.19 will be the -2. If we have a very large maximum skew, there'd be a significant maintenance burden in e2e testing.
 - [ncdc] I don't think we necessarily have to say you must run on v1.19 or else at this point.
 - [cecile] the minimum version could be different right now prior to actual v1alpha4 release, which we can do later along the release cycle.
 - [james] is this a documentation issue, or also an API that can be discovered programmatically such that we can block a controller upgrade because it's the wrong version of k8s

- [vince] in terms of features being used, they may work on older versions, it's more of a suggestion about a support statement on what versions we have tested.
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [christopherhein] New Provider Repo: CAPN
 - <https://sigs.k8s.io/cluster-api-provider-nested>
 - Had an initial kickoff meeting yesterday, anyone looking to help now we're polling for a weekly timeslot <https://doodle.com/poll/bsfcb9b9a76pkfr4>
 - [wfernandes/fpandini] [Management Cluster Operator CAEP](#) update.
 - Looking to move the CAEP to a PR next week. More feedback/comments welcome.
 - Didn't see major blockers, the one big change was single controller multitenancy support for clusterctl move. See the doc for explanation.
 - [james] Windows CAEP
 - Retry support in 1.19 and OsType field on infra machine: <https://github.com/kubernetes-sigs/cluster-api/pull/3616#pullrequestreview-503885559>
 - [fabrizio] retry was added in a rush to fix v1alpha3 and some of it was moved directly into kubeadm. Have asked naadir to test. If we see failures, we should raise issues in kubeadm instead of hiding it inside the script.
 - [naadir] will remove retry from CAPA latest 1.17.x test jobs to start getting signal on kubeadm without experimental retry join
 - [cecile] will do same for CAPZ
 - [vince] which versions got the kubeadm fixes?
 - [lubomir] was added to 1.19 and backported to 1.17 onwards. You can try concretely removing it for all supported versions.
 - [vince] yes & no, because users may want to deploy a previous patch version.
 - [naadir] think it's ok to keep for v1alpha3 and remove for v1alpha4 (and in our support matrix we don't support the prior patch releases)
 - Image-builder pr ready for review: <https://github.com/kubernetes-sigs/image-builder/pull/382>
 - [jackfrancis] clusterctl as kubectrl plugin
 - <https://github.com/kubernetes-sigs/cluster-api/issues/3533>
 - Stepping back from this
 - Might be putting the cart before the horse
 - Need to continue thinking about the problem space, especially across multiple providers

- [jan] add option for scale-in for KCP
 - <https://github.com/kubernetes-sigs/cluster-api/issues/3512>
 - https://docs.google.com/document/d/1HWHTL6g45HrrBppGWgDi_1VkW_RMA0Ue3lRBONcqKFKE/edit#heading=h.2lrpo279kcej
 - Please comment
- [fabrizio] brainstorming doc for kubeadm library (see [issues/2316](#)), feedback are welcome
- [ncdc] Call for proposal/implementation help for v1alpha4
 - Reiterate calls from previous meetings - if you're interested in shaping a feature request, there's an [open roadmap PR](#), please take a look. If you want to get involved, there's people who can mentor and support you, so please reach out.
 - [cecile] Please DM me or other maintainers if you want to talk
- [bmoss] Kubemark provider
 - <https://github.com/benmoss/cluster-api-provider-kubemark>
 - Kubemark is a fake kubelet, doesn't need VMs or anything. Useful for scale testing.
 - [Demo](#)
 - [Autoscaler testing](#) - still very much WIP

Wed 07 Oct - 10am Pacific

[Recording](#)

Attending

- John H Terpstra, Prakash Ramchandra (Dell EMC)
- Matt Boersma, Cecile Robert-Michon, Mike Kostersitz, Zach Wachtel -- Microsoft
- David McKay, Jason DeTiberus (Equinix Metal)
- Chris Hein - Apple
- Michael Gugino, Joel Speed, Michael McCune, Alex Demicev - Red Hat
- John Northrup - Ripple
- Bryan Boreham - Weaveworks
- Vince Prignano, Naadir Jeewa, Warren Fernandes, Timmy Carr, Fabrizio Pandini, Ben Moss, Jason Scarano, Nader Ziada - VMware
- Erwin van Eyk - Platform9
- Nicole Yson - NewRelic
- Arvinderpal Wander - AT&T
- Hardik Dodiya - SAP

Agenda

- PSAs
 - [vince] [Roadmap update for v1alpha4](#)
 - [vince] [Release guidelines updates](#)
 - Main branch is now open for breaking changes
 - Bug fixes and some other important PRs might be backported
 - Mention in PR if it needs backporting
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [wfernandes/fpandini] [Management Cluster Operator CAEP](#) is ready for review!
Related [RFE github issue 3427](#).
 - [Fabrizio] Review doc offline, if people want a meeting to discuss, we can arrange it.
 - [Vince] FYI, it's a way to enable us to move away from using CLI for provider management to CRDs
 - [christopherhein] Process/reqs for a new repo for a new cluster api provider
Related <https://github.com/kubernetes/org/issues/2247>
 - Working on a project called virtual cluster, based out of the multi-tenancy WG, and it's being reworked to fit CAPI. WG are not technically supposed to own code, so need a home.
 - [vince] Have assigned to Timothy St. Clair and Lubomir as SIG Cluster Lifecycle leads to approve.
 - [naadir] First of a few bootstrapping related documents:
 - [Secure node registration](#)
 - Fixes longstanding security hole around hijacking node identity
 - Should allow node labels to be securely applied at creation ([Issue #493](#))
 - [Fabrizio] Could we use a customised bootstrap token given that each token is linked to a specific Machine object
 - [Naadir] We can't do this because we don't know the node name ahead of time, e.g. AWS cloud provider requires PrivateDNSName (which is also part of the attestation), and that only happens post-boot.
 - [Michael Gugino] OpenShift already uses this today, and we use TLS bootstrapping with a separate signer called Cluster Machine Approver.
 - <https://github.com/openshift/cluster-machine-approver>
 - [Jed] How does this work outside of the 3 public cloud providers?
 - [Naadir] Proposing we do a TPM based implementation that covers vcenter

- [elmiko] autoscaler scale from zero, quick question about the taint issue. *once more with feeling!*
 - related, <https://github.com/kubernetes-sigs/cluster-api/pull/2530#issuecomment-705065202>
 - Agreed first iteration would include CPU/RAM/GPU requirements, and we wouldn't do taints, but Stephen Harris raised an issue that AWS requires taints in order to do ASGs? Asked if it's ok to provide an empty list of taints.
 - Are there any issues blocking not having taints in the initial implementation
 - [Hardik] Will need to resolved eventually
 - [Elmiko] We don't want to leave it there, but there's a broader discussion to be had about how Cluster API will handle taints. We'll have to be very clear in the documentation what doesn't work out of the box.
 - [Prakash] The scheduler in the autoscaler will match against node groups that get scaled.
 - [Elmiko] This can lead to poor scheduling if the taints aren't applied.
 - [Jason] There's an issue around this regarding the existing autoscaler implementation, we can treat this separately.
 - [Elmiko] Have an assumption based on the OpenShift implementation, but tend to agree.
- [bryan] Distributed Tracing [issue/doc](#)
 - Done a POC (see prior meeting) for Cluster API.
 - Bits of a proposal need to go into other repos (e.g. controller-runtime)
 - PTAL and we'll try to move it forward.
 - Have branches against 0.5.x and 0.6.x CR, you can take the linked branches and build it.
- [detiber] Will be presenting the [PoC work I've been doing around a minimal embedded api server that supports CRDs](#) at the [SIG API Machinery call today](#)
 - Could potentially use for bootstrapping workflows without needing (e.g. kind) clusters.
 - [Michael Gugino] Really like this model
 - [Jason] We're building PXE workflows for machines etc... , and want to use K8s models for this.
 - [Warren] Does this potentially replace envtest?
 - [Jason] Not sure yet, lots of overlapping concerns trying to solve with this.
- [name] <<add your items here>>

Wed 30 Sept - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham, Scott Rigby - Weaveworks
- Yuvraj - Nirmata
- Joel Speed, Michael 'elmiko' McCune - Red Hat
- Xander Grzywinski - Apple
- David Justice, Matt Boersma, James Sturtevant, Jack Francis, Zach Wachtel, Cecile Robert-Michon - Microsoft
- Andy Goldstein, Naadir Jeewa, Warren Fernandes, Gab Satchi, Ben Moss, Sedef Savas, Fabrizio Pandini, Timmy Carr, Sagar Muchhal, Nader Ziada, Jason Scarano - VMware
- Jason DeTiberus - Packet
- Jun Zhou - Spectro Cloud
- Jan Tilles - Ericsson
- John H Terpstra - Dell EMC
- John Northrup - Ripple
- Hardik Dodiya - SAP
- Xiao - Salesforce
- Daniel Lipovetsky - D2IQ
- Arvinderpal Wander - AT&T

Agenda

- PSAs
 - [v0.3.10](#) on track to be released tomorrow
 - [jan] Was wondering that ncdc left some comments on my external remediation PR, is there time to get work in tomorrow morning EEST.
 - [vince] should be fine, but we can also do a quick follow up
 - [cecile] we have been doing e2e tests with CAPZ, can we have RC can we have another RC before cutting final given the size of changes
 - [vince] Yes, we should also document getting signal from 1 or more providers prior to release
- v1alpha4 planning
 - Will hold a session to re-prioritise the roadmap
 - [Fabrizio] Are we going to have releases of 0.3 after 0.3.10
 - [Vince] For bug fixes, and probably have monthly releases

- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

- [david j] Encryption / protection of Kubeadm bootstrap data (certs, secrets, etc...)
 - Azure provider takes the cloudinit userdata and consumes it directly. The issue is that this userdata is accessible in the Azure portal which is not advisable. We only want the folks that have access to the host to get access to that data. Would like to have encrypted userdata, for example the CAPA provider uses multi-part cloud-init together with AWS Secrets Manager. Would be nice to have a more uniform way to handle this. We could do what the AWS provider is doing, but is this still going to work with Ignition? Can we provide a more uniform interface for this?
 - [Naadir] We are hitting Ignition issue immediately. Want to definitely get it sorted for v1alpha4.
 - [Vince] We have discussed in the past
 - Action Item: Naadir to publish doc on problem statement
- [james s] Windows proposal and OSType information on Inframachine https://github.com/kubernetes-sigs/cluster-api/pull/3616#discussion_r495132309
 - We are pretty close to completing discussion on the CAEP.
 - One of the new things that has emerged that experimentalRetryJoin was very useful for making Windows joins more resilient, but CABPK right now has a Linux specific script. Cecile put forward the idea that the OS type is exposed in the InfraMachine to address this. It is likely we will need OS type for other use cases.
 - [Fabrizio] Right now, kubeadm as a library is a discussion as requirements. I and Naadir will capture all the CAPI based requirements and add to the [kubeadm issue](#).
 - [vince] Potentially have two contract changes if we don't bubble up OSType to Machine object and require bootstrap providers to introspect the infrastructure machine.
 - [James] Is the contract changes part of this KEP or separate?
 - [Vince] Update the contracts in this doc so it's all in once place.
 - [Ben] Would using a cross-platform binary make this whole problem go away.
 -
- [warren] Management cluster operator status update. [Current CAEP](#).
 - Wrt to multi-tenancy, see <https://github.com/kubernetes-sigs/cluster-api/issues/3042>
 - [Fabrizio] CAPA and CAPZ are adding support for handling multi-tenancy within a single instance, which
 - [James De Felice] Right now we have multiple instances of CAPI for each tenant, are we changing this so all controllers are global, or only infra ones?

- [Fabrizio] We would have tenancy arranged by namespace, with credentials linked to certain namespaces.
- [James De Felice] So there will be a management operator that will manage CAPI and infrastructure providers. Everything will be global and managed by the operator. We have a concern around blast radius, and are not excited about creating lots of management clusters.
- [Vince] Philosophically, controllers should never act on global resources, so will be looking at this usage.
- [vince] the namespace flags won't be disappearing just yet, just that clusterctl won't support it
- [name] <<add your items here>>

Wed 23 Sept - 10am inPacific

[Recording](#)

Attending

- Bryan Boreham, Richard Case - Weaveworks
- Joel Speed, Michael 'elmiko' McCune - Red Hat
- Matt Boersma, Cecile Robert-Michon, Zach Wachtel, Jack Francis, James Sturtevant, David Justice - Microsoft
- Warren Fernandes, Andy Goldstein, Vince Prignano, Ben Moss, Fabrizio Pandini, Gab Satchi, Sedef Savas, Timmy Carr, Naadir Jeewa - VMware
- Justin SB - Google
- Mytu Nguyen - New Relic
- Jason DeTiberus - Packet
- John Northrup - Ripple
- Xiao - Salesforce
- Hardik Dodiya - SAP
- Joseph Davis - SUSE
- Prakash Ramchandran - Dell
- Jan Tilles - Ericsson
- Arvinderpal Wander - AT&T

Agenda

- PSAs
 - [vince] v0.3.10-rc.0 has been released
 - [vince] v1alpha4 planning Sept 30th? Please bring items for the [roadmap](#) that you're interested in proposing or you'd like to work on.

- Proposals and reviews in October
- Starting October 1st, we'll start merging some breaking changes to the main branch, to prepare for v1alpha4 API types
- Tentative date ~Q1 2021
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [elmiko] autoscaler scale from zero, quick update
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2530>
 - [fabrizio] [KCP remediation proposal](#) (amends KCP proposal), PTAL
 - TL;dr the remediation will follow the path of deleting a failing node and replacing it with a new one
 - [james] Windows vhd in image-builder:
 - <https://github.com/kubernetes-sigs/image-builder/pull/382>

Wed 16 Sept - 10am Pacific

[Recording](#)

Attending

- Vince Prignano, Naadir Jeewa, Sedef Savas, Warren Fernandes, Andy Goldstein. Sagar Muchhal, Fabrizio Pandini, Nader Ziada, Timmy Carr - VMware
- elmiko - Red Hat
- Leigh Capili, Bryan Boreham, Mark Emeis - Weaveworks
- Cecile Robert-Michon, Matt Boersma, David Schott, Jack Francis, Zach Wachtel, Dinesh Govindasamy, Madhan Mookkandy - Microsoft
- Joseph Davis - SUSE
- Jason DeTiberus - Packet
- John H Terpstra - Dell EMC
- Hardik Dodiya - SAP
- Xiao - Salesforce

Agenda

- PSAs
 - [Fabrizio] /test pull-cluster-api-e2e-full to run on demand the full CAPI e2e suite
 - [Vince] v0.3.10 is slated for the end of month. There are a few PRs including memory leak fixes and changes to the cluster tracker. If you are seeing high memory use, this should fix it. Plan to release a beta at EOW.
 - [ncdc] Can we have it as "RC" instead of beta.

- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

- [cecile/jdef] Adding a sentinel file to the bootstrap provider contract to facilitate detecting bootstrap failure/success via infrastructure providers (<https://kubernetes.slack.com/archives/C8TSNPY4T/p1599744069010700>)
 - When investigating bootstrap failure reporting, we concluded:
 - Using a “daemonset” type thing is too late in the process to be useful, and therefore solutions would be infrastructure specific.
 - We looked at addressing this in CAPZ, and came to the conclusion that we could use a signal that bootstrap is complete, and it's up to the bootstrap provider to determine that, but could be a sentinel file that is touched when bootstrapping is complete, and the infrastructure provider can go and look at that sentinel.
 - Are there objections to this, for v1alpha4. Would change the bootstrap provider contract.
 - [Yassine] Going to be interesting for providers that don't have a way to access machines, would that require SSH access?
 - [Cecile] The bootstrap provider wouldn't be writing the file, but be giving instruction to bootstrap, e.g. for CABK, append the sentinel file write to the cloudinit runcmd.
 - [ncdc] I'm hearing this as a contract change for the bootstrap providers, and infrastructure providers could optionally read the sentinel file.
 - [MarkE] Could we have it in the cluster as some sort of eventing procedure.
 - [ncdc] There's no universal way for code running in Cluster API to reach into a VM to figure out what's happening. And that's no universal secure way to communicate back with the management cluster to say that they're done.
 - [Cecile] We also can't assume kubeadm in the infrastructure provider because kubeadm may not always be used.
 - [Vince] We could add fallbacks today if the file is not found.
 - [Cecile] However, we won't know the reason why the file isn't there.
 - [Vince] We can use a trap exit in the shell script we use for the experimentalRetryJoin
 - [Cecile] This maybe an implementation detail that's not valid for all bootstrap providers.
 - [zach wachtel] Open sourcing of a new CAPI Infrastructure Provider. For AzureStackHCI. <https://github.com/microsoft/cluster-api-provider-azurestackhci>
 - We're open sourcing this provider now, and would like to help out.

- [David Schott] HCI - hyper converged infrastructure
 - [We've been working on the bootstrapping procedure for Windows with SIG Windows.](#)
- [David Schott] How do we get the provider listed in the CAPI book?
 - [Vince] Open a PR to the book
- [elmiko] revisit autoscaler scale from zero proposal
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2530>
 - Seeing an uptick in requests from autoscaler community for this. Need more eyes.
 - Where we are stuck right now is where are we going to annotate information about CPU, Memory, GPU. Andy noted that it is present in the kubeadmconfig, but there's pushback from having to dig so deep.
 - [ncdc] Need to revisit this, but thought we had discussed adding the annotation to the appropriate resource after at least one instance exists and figure out its characteristics and save that information back onto the machineset and then the autoscaler can use that. Case where this doesn't work with is if you start with zero.
 - [elmiko] infrastructure providers could potentially know what this information is and could add this annotation back.
 - [ncdc] If we think about it from a contract perspective for infrastructure providers.
 - [hardik] would be convenient to have the information on machinedeployment/machineset. I see in autoscaler the providers have scripts to save all this information, because it can't be looked up at runtime.
 - [elmiko] that is exactly correct, but CAP* providers should not look into autoscaler code, but the providers should put this information in the appropriate place.
 - [vince] should this be annotations or fields in spec/status, since annotations are hard to version, esp. when we think about v1alpha4. And is it ok for v1alpha4?
 - [elmiko] Annotations are easy, so that's how we did it OpenShift, but agnostic about how its done, and would be fine for v1alpha4.
 - [elmiko] To summarise: Add to CRDs for v1alpha4.
- [james] Windows CAEP PR:
 - <https://github.com/kubernetes-sigs/cluster-api/pull/3616>
 - A few comments were on the Google doc, and have been addressed in PR. Since we have an infrastructure provider with actual Windows support, would like comments from them.
 - Have also been working on image builder scripts to build Windows images.
 - [cecile] Do we want to set up a lazy consensus for this PR?
 - [james] That'd be a way to move things forward.

- [Bryan] Demo of tracing prototype for Cluster API
 - [Prior art](#) - KEP 34
 - Added instrumentation code to controllers that sends data to Jaeger
 - Log the various k8s API server operations that each controller does, and this passes the context through to the next thing in the process.
 - Did this as I was personally uncertain about what things were happening. Believe it could be useful for answering performance or “why is it not working” questions.
 - [James De Felice] What was done in the controllers to enable this
 - [Bryan] Used Jaeger, but could be anything in that vein. In each controller, kick it off and tell Jaeger who’s talking. In every reconcile, we pick up a tracing span from an annotation on the object from the kep in prior art.
 - Also wrapping the runtime client, and adding trace IDs on the object.
 - The same code exists in clusterctl too
 - [Bryan] ideally this would go into controller runtime libraries, and would like support for this.
 - [Lubomir] Did you capture the overhead of the tracing itself?
 - [Bryan] Formally no, if you don’t turn it on, it’s going to do some extra metadata search which is discarded string manipulation, then if you do turn on tracing, then the spans are sent as a UDP packet to a local collector. So the impact is non-zero, but given the amount of json marshalling/unmarshalling going on, it’s probably negligible.
 - [Leigh Capili] And presumably sampling is not particularly useful as this isn’t a high-load web server.
- [Fabrizio] Update on management cluster operator
 - Kick off meeting on Monday, notes are below in this document, [recording](#).
 - Main decision:
 - In scope: cert-manager (optional), providers + config required to install providers
 - Out of scope: cluster-templates & move
 - Next steps
 - Formalize goals
 - Start rallying on an API draft
 - [Cecile] Was CRS considered in scope?
 - [Fabrizio] Considered out of scope because it’s part of the cluster template.
 - [Cecile] Not all templates have CRS’, so require users to apply them first, so would be good to have managed by clusterctl.
 - [Fabrizio]

- [Sedef] KCP remediation
 - https://docs.google.com/document/d/1hJza3X-XbVV_yczB5N6vXbl_97D0bOVQ0OwGovcnth0/edit#heading=h.elgr7djhgier
- [name] <<add your items here>>

Sep 14 Sept - 10am Pacific - clusterctl Management Cluster Operator Notes

Recording

CAEP Document -

<https://docs.google.com/document/d/1ZsusF5c9pYxseuaKxTpctI5aUDqzI0sdCW4xxDbLm3k/edit#>

Action Items

[fabrizio] Will come up with the initial round of API and share with community.

- Motivation
 - Clusterctl is used by 50% of survey respondents. This is to provide another way for installing and upgrading the management cluster.
 - The second use case is the git ops workflow for managing the management cluster.
 - Support team of admins.
 - [robert] infrastructure promotion from different environments such as dev -> staging -> production. As part of the git ops workflow.
 - [charles] they got the management cluster installation in git ops already via helm charts. Struggled with move and upgrades regarding management cluster. Would like for these two steps to have higher priority so it would enable them to move completely to git ops workflow.
 - [fabrizio] Good point. We can address this point in the goals.
 - [vince] The first version of the management cluster operator wouldn't include move because of simplicity and also because move could become backup/restore. Prioritizing the move should be a future goal and should be included in the CAEP.
 - [Robert] If we do move later, then the upgrade would be hard because the objects wouldn't be owned properly.
- Goals
 - [fabrizio] what objects/CRDs are required that should be stored as part of the management cluster?

- E.g. Cert-manager.
Do we want the lifecycle of the cert-manager be managed by this operator?
- [robert] cert-manager already has a way to install. We just document supported versions and let the customer install cert-manager. Also reduces scope of this CAEP.
- [fabrizio] think we should provide same experience as we do today. That is, If the customer already has cert-manager then we use that, else we install it.
- Provider Instance
 - Fetch configuration
 - [fabrizio] This is the core expectation of the operator. Repository configuration. If this stays outside of the cluster, it will be hard to do upgrades. Suggest we include repository information inside the provider.
 - [charles] there could be other ways to provide repository info- spinnaker, argoCD.
 - [fabrizio] component yamls are currently fetched by clusterctl and managed by clusterctl. The operator should continue to keep this layer of abstraction. If we expose this information, we are removing a big part of clusterctl. It is required to manage this in higher level.
Customer should manage the yaml in their gitops workflow but have clusterctl be able to fetch the components.
 - [robert] authentication for these sources of yamls. Would the operator need to know and have access to these toolings (argoCD, spinnaker, weaveworks, Jenkins).
 - [charles] agree to robert's point. There needs to be a clear demarcation for how the changes are made to the yamls and get into the cluster. That way we can be flexible to customer's needs (example, banking customers).
 - [vivek s.] delivery mechanism is custom and doesn't follow gitops workflow.
 - [robert] template for a provider as a CRD is a good way. Maybe it just includes the yaml that needs to be applied within the cluster.
 - [charles] Some customers are a nogo for gitops workflow.
 - [fabrizio] Design will not force any specific usage. *Review the issue attached in the CAEP.* ProviderSpec is a simple abstraction. When user applies this spec, then the operator should take charge of installing and upgrading. We can provide a way to "fetch" abstraction so that there is flexibility within the operator.
 - [richard] agree.

- [vivek s.] Not tied to gitops but it needs to be pulled from somewhere. Correct?
- [naadir] clarify the gitops definition vs. fetching the components yaml.
- [charles] helm chart for the management components? Operator can be notified that a new version of chart is available. But still couldn't proceed regarding upgrade. If the operator could handle the upgrade via a https/helm repository/artifactory/harbor that they could point to internally and fetch the upgrade components yaml
- Image Overrides
 - [fabrizio] When we fetch currently, clusterctl does image overrides as another configuration. This falls under the same category as above.
- Flag variables
 - Should be included as part of MVP.
- All these information should be recorded when clusterctl init. And can be used later for upgrade.
- Management Group
 - [fabrizio] Operation on a group of containers. Like, upgrade from v1.x to v1.y.
- Cluster templates
 - [fabrizio] Personally, keep out of scope of the management cluster operator since we already have many ways to manage the workload cluster templates.
- [yassine] challenge/concern is the API we need to surface for this will need to ensure it is declarative and not embed any operation within it.
- [fabrizio] yup...agree modeling this API in a declarative way will take some work.

Sep 9 Sept - 10am Pacific

[Recording](#)

Attending

- Jan Tilles - Ericsson
- Ben Moss, Fabrizio Pandini, Nader Ziada, Andy Goldstein, Warren Fernandes, Gab Satchi, Naadir Jeewa, Sedef Savas - VMware
- Joel Speed, Nir - Red Hat
- Cecile Robert-Michon, Jack Francis, Matt Boersma, David Justice, James Sturtevant - Microsoft

- Jason DeTiberus - Packet
- Bryan Boreham, Mark Emeis, Scott Rigby - Weaveworks
- Mytu Nguyen, Andrew Rudoj, Nicole Yson - New Relic
- Hardik Dodiya - SAP
- Joseph Davis - SUSE
- John Northrup - Ripple
- Erwin van Eyk - Platform9
- John H Terpstra, Prakash Ramchandra - Dell EMC
- Arvinderpal Wander - AT&T
- Scott Rosenberg - TeraSky
- Daniel Lipovetsky - D2IQ

Agenda

- PSAs
 - [vince] KubeCon NA Maintainer track
 - Reach out if you're interested in presenting a session on Cluster API
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [jan] External Remediation Demo
 - Design proposal: <https://github.com/kubernetes-sigs/cluster-api/pull/3190>
<https://github.com/kubernetes-sigs/cluster-api/pull/3606>
 - In bare metal environments, we may not want the default behaviour of deleting a machine, as this will cause a long deprovisioning and provisioning of the bare-metal host.
 - [fabrizio] Kicking off the work on the management cluster operator by defining goals/non goals
 - Proposing some slots next week [Doodle](#)
 - Deadline for the Doodle Sept 11th
 - [fabrizio] SCL [survey results](#)

Wed 2 Sept - 10am Pacific

[Recording](#)

Attending

- Cecile Robert-Michon, Matt Boersma, James Sturtevant, Jack Francis - Microsoft
- Vince Prignano, Fabrizio Pandini, Andy Goldstein, Ben Moss, Warren Fernandes, Yassine Tijani, Sa - VMware
- Joseph Davis - SUSE
- Michael Gugino, elmiko - Red Hat
- Arvinderpal Wander - AT&T
- Jun Zhou - Spectro Cloud
- Bryan Boreham - Weaveworks
- Darryl Darling - IBSS Corporation
- Hardik Dodiya - SAP
- Scott Rigby
- Jason DeTiberus - Packet
- David Watson

Agenda

- PSAs
 - [v0.3.9](#) has been released
 - Problem with upgrading to 1.19.0 currently
 - kubeadm fix: <https://github.com/kubernetes/kubernetes/pull/94398>. Will need to be backported to 1.19.1 (ETA n)
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [james] reminder to look at [CAPI Windows Proposal](#) and comment. Plan to open the PR for the proposal by Sept 8, 2020
 - [cecile] marking breaking changes with :warning:
 - Related: <https://github.com/kubernetes-sigs/cluster-api/pull/3579>
 - [fabrizio] [#3580](#) extend E2E test framework to support a pluggable LogCollector
 - [dlipovetsky] Thanks @fabriziopandini for making CAPD easier (<https://github.com/kubernetes-sigs/cluster-api/pull/3514>)
 - [dlipovetsky] *Limited* in-place upgrades, e.g., kubelet, CRI patch version.
 - Motivation
 - Avoid drain, preserve workloads. (Alternative: [Live Pod Migration](#))
 - Faster patching, e.g., for CVE
 - [bboreham] observes that CAPI documentation defines Machines to be immutable, so you first need to change that rule
 - [mgugino] Thought about updates w/o rebooting in OpenShift.
 - Would need to abstract kubelet version, etc to allow MachineDeployment to work

- PodDisruptionBudget is a solution to preserving workloads (at least temporarily)
- In-place upgrades would add a lot of complexity
- [awander] Want to update certain host configuration (e.g. sysctl) in place. Working on a proposal.
 - [mgugino] See the MachineConfig Operator in OpenShift
- [vprignano] There are some ways to do this outside of CAPI. Doesn't have to be part of the project's support matrix. Time to document these ways?
 - Kubeadm operator effort should support these use cases
 - [dlipovetsky] Happy to write up ideas
- [mgugino] Let's support these use cases "in the community," even if they are not supported in the CAPI code base.
 - [andy] This would be nice. We don't have a single place where the CAPI community can collaborate on use cases.
- [jackfrancis] Updating a component on the node makes it deviate from the "canonical" desired state in the Machine resource
- [zawachte - chat] +1 to this point. Lets say for example I wanted to create an operator that installed helm charts on each target cluster managed by clusterapi. I wouldn't expect to upstream to cluster api, but its certainly a useful operator built ontop of the capi primitives. So we should have a community to talk about these type of tools.
- [chris - chat] yeah but as a consumer of clusterapi you need day 2 functionality, which includes config updates and upgrades.
- [andy] timeline for opening up main branch to breaking changes for v1alpha4
 - [cecile] we should at least wait a week after 0.3.9 release
 - [vince] Proposals first I'd say
 - [jason] +1 to vince, I would like to see proposals and a public planning of what we are targeting for the release before opening
 - [vince] come up with a calendar, might take 6 weeks for all proposals to be approved

Wed 26 Aug - 10am Pacific

[Recording](#)

Attending

- David Watson
- Michael Gugino, elmiko - Red Hat
- Jason DeTiberus - Packet
- Fabrizio Pandini - VMware

- Cecile Robert-Michon, Kalya Subramanian, Matt Boersma, Mark Rossetti, David Justice - Microsoft
- Naadir Jeewa, Sedef Savas, Warren Fernandes, Ben Moss, Andy Goldstein, Vince Prignano, Fabrizio Pandini - VMware,
- Chris Hein, Xander Grzywinski - Apple
- Jun Zhou - Spectro Cloud
- John H Terpstra (Dell EMC)
- Bryan Boreham - Weaveworks
- Arvinderpal Wander - AT&T
- Hardik Dodiya, Prashanth - SAP
- Andrew Rudoj - New Relic
- Daniel Lipovetsky - D2IQ

Agenda

- PSAs
 - [vince] v0.3.9 release soon-ish
 - Was planning to have a release on Aug 31, can punt some things to 0.3.10 as there's a few things in there. Ping if there's things you definitely want in the release.
 - Start filing RFEs for 0.4. We may backport bug fixes to 0.3.x.
 - [ncdc] No intention to backport features to 0.3.x
 - [fabrizio] Working on conditions for the delete workflow in CAPI, CAPD, CAPV <https://github.com/kubernetes-sigs/cluster-api/pull/3527>
 - If you can also add these to your infrastructure providers, please do so.
 - [vince] can we move this 0.3.10?
 - [fabrizio] Current behaviour can cause confusion around state.
 - [Vince] given that infrastructure providers need to adopt this, we can put it in 0.3.9 so infra providers can start implementing.
 - [Cecile] Earlier than later is preferable so there's time before v1alpha4.
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [Kalya Subramanian] Windows support for CAPI
 - https://docs.google.com/document/d/14evDI_3RgEFfchmgPzNw6Ib1vlttNHb9333UnUJ734/edit
 - Intention to file as a CAEP, please comment on Google doc.
 - Crux of the problem is that Windows does not have privileged containers at the present time, but is being worked on.
 - As a stop gap, will use a named pipe to proxy commands onto the host itself (wins.exe), together with CloudBase. Intention is to have a proposal extensible for all infrastructure providers.

- [Michael Michael] Please review this doc and take the journey along with us.
- [Fabrizio Pandini] CAPI & GitOps, [the clusterctl operator](#)
 - My take as of this moment is that CAPI can be operated in a GitOps model for workload clusters, but less so for the management cluster. My view is to produce an operator using the clusterctl API to manage the providers in a declarative approach.
 - [Bryan Boreham] We do a lot of CAPI installs using GitOps and completely ignore clusterctl because it doesn't fit the model, so am interested in what you come up with. We do "bend the rules" for pragmatism, so should bring in Richard Case and other colleagues.
 - [Ben] It's confusing to me talking of a "clusterctl operator", it seems like it's about sharing configurations of a cluster that it's stored in CRDs. Seems weird to build an operator to drive a CLI for driving API.
 - [Fabrizio] One of the other limitations is that many of the settings are envsubst'd, if we can move these configuration into the cluster, the better.
 - [Vince] Would like a step back, that the operator is mainly about managing the lifecycle of the providers rather than cluster lifecycles.
 - [Jack] To Bryan: If you're not using clusterctl, are you using the SDK?
 - [Bryan] A thousand flowers have bloomed, people have done things with helm charts. We started with v1alpha1, and wrote scripts and tools which are mostly open source. One of the reason I am hedging is that all current solutions have issues.
 - [Jack] Is there anything clusterctl have that isn't represented in the K8s API.
 - [Bryan] Pivot primarily. We take the API as the standard, and gitops as the spirit.
 - [Andrew Rudoi] At New Relic, we are a big consumer of `clusterctl move`. We have a strong desire to not have to call `clusterctl move` in a script, for use cases involving environment teardown. The other thing is flavours, which clusterctl has a concept of. We have worked around the idea of flavours with our own operators. For example, we have many AWS accounts, and we know for a fact the subnets we want to use for any cluster, we have to put in defaulting webhooks to always set it.
 - [Fabrizio] Should start a working group / doc to explore the ideas
 - AI: Start WG & doc.
- [Arvinder] `clusterctl rollout` proposal: <https://github.com/kubernetes-sigs/cluster-api/issues/3439>
 - Tried to base it on how kubectrl rollout with the long-term plan for kubectrl to support it directly.

- Some elements need discussion about Cluster API Core. MachineSet fits well with the kubectl rollout model, but less so with kubeadmcontrolplane.
- [Fabrizio] Need to figure out if it makes sense to tackle MachineDeployment and KCP at the same time, or do MD first then KCP.
- [Arvinder] Doc only tackles MDs, but KCP is a bigger task.
- [Cecile] Initial question is that clusterctl is more to operate the management cluster, and we've not cross that boundary that clusterctl does day 2 ops of the workload cluster and would seem to break that principle. In which case, kubectl rollout would be far more preferable.
- [Arvinder] Had a look at kubectl scale, and it didn't seem as simple, and would require upstream changes to make this work. E.g., if you look at kubectl restart, which takes a MachineDeployment, if you look at the code it has a switch statement on types, which means reaching into the code base to support CAPI types. The abstraction for scale is supportable by CRDs, but the one for rollout is not.
- [Jason] If we're worried about overloading clusterctl, then a more tightly scoped krew/kubectl plugin may provide another avenue.
 - refs:
 - <https://kubernetes.io/docs/tasks/extend-kubectl/kubectl-plugins/>
 - <https://github.com/kubernetes-sigs/krew>
- [Arvinder] Is there a reason that clusterctl doesn't do day 2 ops?
 - [Vince] Generally to integrate better with the community with tooling they're already familiar with.
- [David Justice] kubectl would seem a nice landing spot for Day 2 operations. Having that single anchor point for the community is advantageous.
- [Vince] Probably need a larger discussion about the boundary setting.
- AI: Jason to open issue for kubectl plugin
- [cecile] CAPI book branch (<https://kubernetes.slack.com/archives/C8TSNPY4T/p1597914437003400>)
 - Frequent breakage of the docs when the docs against the main branch run ahead of a release. Split on this, since we can have the docs not ready in time for a release, the current process allows it to be kept up to date.
 - [Vince] We can have a release branch that is fast-forwarded to the next release. Which would cover the use case of docs for the next release, but would not work for a main branch in 0.4, where we would need to backport docs to the previous release.
 - [ben] More complicated way would be have docs for multiple versions at the same time.
 - [Vince] we do have `<release-as-subdomain>.cluster-api.sigs.k8s.io`

- [cecile] Do we want a book branch, that will allow us to fix docs separately.
 - [vince] adds significant effort.
 - [cecile] could add a disclaimer about it being based on the latest.
- AI: Vince & Cecile Document change to process
- [cecile] clusterctl version support / back compat with cluster api minor releases (https://github.com/kubernetes-sigs/cluster-api/pull/3484#discussion_r474752889)
 - [cecile] The 3 options we have our:
 - Any version of CAPI is supported by any clusterctl, and needs to be tested.
 - Recommendation to use the same version and warn of potential issues.
 - Strictly block, e.g. clusterctl for v1alpha4 cannot be used with v1alpha3
 - [Vince] Having cross-version compatibility may place significant constraints on development of clusterctl operator for example.
 - [Andrew Rudoj] Should clusterctl should not at least support upgrade from v1alpha3 to v1alpha4. We were bitten significantly at NR because of this.
 - [ncdc] Upgrade MUST work. We don't want clusterctl for v1alpha4 to emit v1alpha3 clusters and vice versa. The change to use envsubst in templates is an example of something we should strive to avoid in a patch release of clusterctl. If we do something like this again, then it should be treated as a regression.
 - [Warren] To clarify, if upgrade should work, if there's a v1alpha3 cluster, is it the v1alpha4 clusterctl responsible for upgrading.

Wed 19 Aug - 10am Pacific

Meeting cancelled due to KubeCon EU 2020 Virtual conference.

Wed 12 Aug - 10am Pacific

[Recording](#)

Attending

- Nir, Michael Gugino, elmiko - Red Hat
- Jason DeTiberus - Packet
- Cecile Robert-Michon, David Justice, Jack Francis, Matt Boersma - Microsoft
- Naadir Jeewa, Warren Fernandes, Nader Ziada, Fabrizio Pandini - VMware
- Joseph Davis - SUSE
- Chris Hein, Xander Grzywinski - Apple
- Prakash Ramachandra - Dell
- Andrew Rudoj, Mytu Nguyen, Nicole Yson - New Relic
- Richard Case, Bryan Boreham - Weaveworks
- John H Terpstra - Dell EMC
- Hardik Dodiya - SAP

Agenda

- PSAs
 - 0.3.8 released
 - 0.3.9 planned for end of August
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [Nir] External remediation proposal
 - [📖 External Remediation Proposal · Issue #3190 · kubernetes-sigs/cluster-api · GitHub](#)
 - [Prakash] Seems complex to me: Is this coming from a bare metal perspective or MaaS (such as Ionic). Am OK with KCP being excluded.
 - [Cecile] Worth clarifying non-goals if this isn't clear.
 - [Maël] question about the KCP upgrade behaviour in case of failures
 - Would like confirmation that whenever there's a KCP upgrade and there's a misconfiguration or something that prevents the upgrade, is it supported that you can revert the configuration. What is the expected behaviour: that the node will come up with the new config, or not. Have seen the nodes reuse the old kubeadm config.
 - [Fabrizio] Expected behaviour is that KCP will roll out new nodes, and if you define a new state mid rollout, that KCP should use the latest spec. If you see strange behaviour, file an issue.
 - [Ben Moss] relevant issue re KCP kubeadm configuration <https://github.com/kubernetes-sigs/cluster-api/issues/2083>
 - [Warren] metadata yaml in provider releases [See issue 3418](#) for context.
 - [Provider Contract for clusterctl](#).

- We are relying on embedded metadata to check on contracts. Can we have a metadata.yaml as part of release artifacts, which would mean we wouldn't run into 3418 as often. Right now, users have to make sure they're using the latest clusterctl.
- [Cecile] For CAPZ, we have no objection to adding metadata.yaml, wasn't aware that was intended.
 - [Naadir in chat] Same for CAPA
 - [Jason in chat] +1
 - [Warren] It is in the clusterctl provider contract
- [Cecile] Open up issues in each provider for the metadata file
- [Fabrizio] Up til now, most providers worked because the same contract was followed, but if you have a larger change, such as a 0.x release, it may not work. Case to be made to include some sort of defaulting, but maybe better to enforce providers stating which contract they're using.
- [name] <<add your items here>>

Wed 5 Aug - 10am Pacific

[Recording](#)

Attending

- Michael Gugino - Red Hat
- Matt Boersma - Microsoft
- Fabrizio Pandini, Gab Satchi, Naadir Jeewa , Nader Ziada, Ben Moss, Andy Goldstein, Warren Fernandes, Sedef Savas - VMware
- Chris Hein - Apple
- Joseph Davis - SUSE
- Spencer Smith - Talos Systems
- Bryan Boreham - Weaveworks
- Hardik Dodiya - SAP
- Andrew Rudoi, Nicole Yson - New Relic
- John H Terpstra (Dell EMC)
- Prakash Ramchandran (Dell)
- Daniel Lipovetsky - D2IQ

Agenda

- PSAs
 - [v0.3.8 has been released](#)
 - [v0.4.0](#) planning reminder

- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

- [vince] SIG Survey for KubeCon EU 2020
<https://docs.google.com/document/d/1OokYTB0k61yWIoDzNnTDNvDfuoUswgyO8P2wzfCY4WY/edit#heading=h.4rulggug5155>
 - [Naadir] How do we publicise this (the recordings are already in)
 - [Bryan] Can contact the moderator for the session, and they can get AV to add in a broadcast slide/message.
 - [lubomir] the Link is hardcoded and should be present in the VOD for the SIG CL intro session. The general announcement with the link was sent to the SIG ML already. VODs cannot be amended.
 - [John Terpstra] Added some qs related to bare metal
- [Fabrizio] Should we define a deadline for the proposal about [testing guideline and conventions for Cluster API](#) ?
 - There are already PRs coming in that implement the proposal.
 - [vince] ginkgo maybe too opinionated, the proposal works towards using go test and envtest exclusively.
- [michaelgugino] - Deletion Lifecycle hooks:
<https://github.com/kubernetes-sigs/cluster-api/pull/3132>
 - [ncdc] have lgtm'd, the implementation is fairly clear and uncontroversial, with a lot of the debate happening around the use cases in the proposal, have resolved a lot of the conversations as a result.
 - Merged!
- [ncdc] EOL for v1alpha3 / opening main branch for breaking changes
 - Proposing: August 31 for 0.3.9, with 0.3.10 for bug fixes by October 1
 - Main branch open from September 1
 - [fabrizio] Should cert-manager upgrade be targeted for 0.3.9 and clusterctl change
 - [Vince] Current cert-manager version being used has CVEs in the base image, and the clusterctl change should be good to enable Talos.
- [ncdc/ben] Review KCP Remediation proposal
 - https://docs.google.com/document/d/1hJza3X-XbVV_yczB5N6vXbl_97D0bOVQ0OwGovcnth0/edit

Wed 29 July - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham, Richard Case - Weaveworks
- Cecile Robert-Michon, Matt Boersma, Ace Eldeib, David Justice, Jack Francis - Microsoft
- Andy Goldstein, Vince Prignano, Jason DeTiberus, Nader Ziada, Warren Fernandes, Sedef Savas - VMware
- Joel Speed, Michael elmiko McCune, Michael Gugino - Red Hat
- Jun Zhou - Spectro Cloud
- Mytu Nguyen, Andrew Rudoi - New Relic
- Joseph Davis - SUSE
- Daniel Lipovetsky - D2IQ
- Chris Hein - Apple

Agenda

- PSAs
 - New agenda template format
 - [vince] [v0.3.8 to be released](#) by tomorrow, or EOD Friday at latest
- Discussion Topics

Use this section for demos, topics you'd like to discuss, issues or pull requests that might need more attention, or any generic questions.

 - [cecile] KCP machine creation in parallel
 - First cp node is bootstrapped, but could others be done in parallel?
 - Depending on version of kubeadm there is a parallel join bug that can come up. Also want to ensure consistency in handling etcd quorum.
 - What is the motivation for creating in parallel? They are a relatively stable set with not much churn so operations to add new would be infrequent.
 - Reducing create cluster time, could be 10-15 minutes to get control back
 - more details on Parallel join at <https://github.com/kubernetes/kubeadm/issues/1646>
 - And <https://github.com/kubernetes/kubeadm/issues/2005>
 - Quorum could be biggest potential issue. If a node is unable to join for some reason it would be more complex to handle failures.
 - Also not sure we have signal on concurrent joins, see <https://github.com/kubernetes/kubeadm/issues/1661>
 - May be worth looking at timing and understand what can be done in parallel. Biggest time is bringing up infrastructure.
 - [jack] observation that etcd leader election changes impact time overall. During cluster creation events want to take scheduling operations offline until etcd leader election done. Could potentially add 20 minutes to time.

- Do you see leader elections happen for first few or more than 2? Can't confirm if seeing for higher numbers. 2 member etcd count is the "stable state"
- Took one at a time approach because expected the number of replicas is relatively small for most use cases
- [wfernandes] projects that use envtest for unit tests examples
- [ncdc] Reminder about v1alpha4 call for feature requests
 - [vince] Backlog grooming Friday 8AM PT, zoom going to be posted on Slack. Goal is to go over v0.3.x and v0.4.0 milestones and revisit / retriage issues.
- [name] <<add your items here>>

Wed 22 July - 10am Pacific

[Recording](#)

Attending

- Vince Prignano, Andy Goldstein, Yassine Tijani , Nader Ziada, Warren Fernandes, Ben Moss, Jason DeTiberus — VMware
- Bryan Boreham - Weaveworks
- Andrew Rudoj - New Relic
- Michael elmiko McCune, Michael Gugino - Red Hat
- Xuan Gong - Salesforce
- Krishna Posa - AT&T
- Daniel Lipovetsky - D2IQ
- Hardik Dodiya - SAP
- Erwin van Eyk - Platform9
- Prakash Ramchandran - Dell
- David Justice, Cecile Robert-Michon - Microsoft

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] Preparing for [v0.3.8](#)
 - KCP's ClusterConfiguration wasn't allowed to be nil, causing a rollout <https://github.com/kubernetes-sigs/cluster-api/pull/3356/files>
 - Workaround: set `clusterConfiguration: {}` under KCP's spec
 - [ClusterResourceSetBindings weren't always getting deleted when deleting a Cluster](#)

- Demos/POCs
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [ncdc/vince] v1alpha4 planning
 - <<add your items here>>
 - [prakash] Questions about merging
 - <https://github.com/kubernetes-sigs/cluster-api-provider-openstack/pull/594>
 - Someone can pair with prakash (reach out on slack)
 - Also, the #cluster-api-openstack slack channel is the right place to ask about PRs on the openstack provider
- New issue triage

Wed 15 July - 10am Pacific

Recording TBD

Attending

- Bryan Boreham, Richard Case - Weaveworks
- Michael Gugino, Joel Speed - Red Hat
- Jason Tarasovic - Paylt
- Cecile Robert-Michon, Matt Boersma, Ace Eldeib - Microsoft
- Xuan Gong - Salesforce
- Warren Fernandes, Jason Scarano, Andy Goldstein, Nader Ziada - VMware
- Andrew Sauber - Linode
- Jun Zhou - Spectro Cloud
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] [v0.3.7 has been released](#)
 - [vince] v1alpha4 planning in ~2-3 weeks, please open issues for features you'd like to see/tackle
 - Cecile has been promoted to maintainer
- Demos/POCs
- RFCs / discussion on future proposals [time-box to 5m]
 - Autoscaler scale from 0 <https://github.com/kubernetes-sigs/cluster-api/pull/2530>
 - Infrastructure Capabilities on Machines <https://github.com/kubernetes-sigs/cluster-api/pull/2927>

- Machine deletion phase hooks
<https://github.com/kubernetes-sigs/cluster-api/pull/3132>
- External Remediation <https://github.com/kubernetes-sigs/cluster-api/pull/3190>
- Discussion Topics
 - [Fabrizio] Working on [testing guideline and conventions for Cluster API](#) as a follow up of recent problems on tests (flakes, fake client limitations, see [3287](#)). Looking forward to community feedback on the direction we should take.
 - [Andy] [KubeadmControlPlane spec mutability](#) - request for comment
 - <<add your items here>>
- New issue triage

Wed 8 July - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham - Weaveworks
- Naadir Jeewa, Andy Goldstein, Ben Moss, Nader Ziada, Fabrizio Pandini, Warren Fernandes - VMware
- David Justice, Matt Boersma, Cecile Robert-Michon - Microsoft
- Joel Speed, Michael elmiko McCune, Michael Gugino - Red Hat
- Andrew Sauber - Linode
- Chris Hein - Apple
- Jason Tarasovic - Paylt
- Daniel Lipovetsky - D2IQ
- Andrew Rudoi - New Relic
- Joseph Davis - SUSE
- Hardik Dodiya - SAP

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [naadir/cecile] If there are new features you want to highlight for your infrastructure provider for the Kubecon deep dive on August 20, please contact Naadir or Cecile in Kubernetes slack, or email [jeewan\[at\]vmware.com](mailto:jeewan@vmware.com).
 - [vince] v0.3.7-rc.0 has been released, rc.1 tomorrow, actual release by EOW / early next week.
 - See <https://cluster-api.sigs.k8s.io/roadmap.html> for a rough list of changes (there is lot more in this release, notes will be compiled later).
- Demos/POCs

- General questions / provider implementer questions
 - [wfernandes] [Allow a user specifiable node draining timeout](#). Why do we wait forever when evicting pods but only 20s when deleting pods directly during node draining?
 - [detiber] right now, the timeout is the only form of interaction with PDBs, we have to make sure this is respected. Cluster autoscaler does do some introspection of the workloads
 - [michael] you could create an external controller to implement a different behaviour via annotations
 - [ncdc] External controller is useful, but would also be nice to not need it
 - [thorsen] if the default is changed, that would need advance notice, and if it was would need configuration in any case. Otherwise in the process of debugging a possible drain node behaviour.
 - [gugino] +1 on not changing default
 - [wfernandes] happy to leave the default but make it configurable.
 - [cecile] Adopting ClusterResourceSets
 - How should infra providers go about adopting it in terms of timelines
 - [sedef] Open issues are not blockers. CRS is available in the latest 0.3.7 alpha.
 - Is an experimental template, so shouldn't be in the default template.
 - [vince] CRS could be defaulted to true
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [cecile] Status on [Cluster API Bootstrap Reporting](#) / v1alpha4 roadmap?
 - [naadir] had paused for the 0.3.7 timeframe but happy to revisit asap. Seems pretty difficult to do in an infrastructure agnostic way.
 - [cecile] was holding off pending this CAEP, but if infrastructure agnosticism is difficult, will proceed with implementation in CAPZ.
 - [vince] Upgrade cert-manager
 - <https://github.com/kubernetes-sigs/cluster-api/issues/2240>
 - [James] Want to make sure we take an approach to make sure cert-manager is version agnostic. We do put an onus on API stability, so hopefully can address this.
 - <<add your items here>>
- New issue triage

Wed 1 July - 10am Pacific

[Recording](#)

Attending

- Fabrizio Pandini, Naadir Jeewa, Andy Goldstein, Vince Prignano, Warren Fernandes, Dan Finneran, Jason Scarano, Jason DeTiberus, Ben Moss - VMware
- Michael elmiko McCune, Joel Speed, Michael Gugino, Alex Demicev - Red Hat
- Bryan Boreham - Weaveworks
- Andrew Rudoj, Dane Thorsen - New Relic
- David Watson
- Spencer Smith - Talos Systems
- Joseph Davis - SUSE
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ
- Hardik Dodiya - SAP

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - v0.3.7-beta.0 going to be released today
<https://cluster-api.sigs.k8s.io/roadmap.html>
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - michaelgugino: Machine Deletion Lifecycle Hooks
 - CAEP: <https://github.com/kubernetes-sigs/cluster-api/pull/3132>
 - Potential Implementation:
<https://github.com/kubernetes-sigs/cluster-api/pull/3273>
 - Michaelgugino: Node Maintenance Lease KEP
 - <https://github.com/kubernetes/enhancements/pull/1411>
 - [wfernandes] Template defaults with envsubst
<https://github.com/kubernetes-sigs/cluster-api/pull/3270>
 - [ncdc] Request for feature requests
 - [rudoj] Flexibility in MachinePools
 - How will they support different tiers of configuration? E.g. ASGs in AWS provide a few different levels of configuration options.
 - Cluster-autoscaler understands all the individual cloud provider autoscaling implementations.
- New issue triage

Wed 24 June - 10am Pacific

[Recording](#)

Attending

- Ben Moss, Jason Scarano, Andy Goldstein, Vince Prignano , Nader Ziada, Naadir Jeewa, Warren Fernandes, Fabrizio Pandini - VMware
- Bryan Boreham, Mark Emeis - Weaveworks
- Michael elmiko McCune, Joel Speed - Red Hat
- Jason Tarasovic, Paylt
- Prakash Ramachandra, Dell
- Seth Pellegrino – New Relic
- Hardik Dodiya - SAP
- Prankul
- Daniel Lipovetsky - D2IQ
- Cecile Robert-Michon, James Sturtevant - Microsoft
- Dhawal Patel, Nordstrom

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - Currently blocked on building images for new k8s versions (WIP fix [here](#))
<https://github.com/kubernetes/kubernetes/issues/92242>
 - [vince] Timeline for v0.3.7
 - [ncdc] [New reviewing guidelines](#)
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [vince] Roadmap <https://cluster-api.sigs.k8s.io/roadmap.html>
 - [prankul] Add clusterctl command
<https://github.com/kubernetes-sigs/cluster-api/issues/2542>
 - [dhawal patel] container linux / flatcar support
 - [detiber]
 - Would require changes in [Image builder repo](#)
 - [PR here](#) for adding initial flat car support to image builder
 - <<add your items here>>
- New issue triage

Wed 17 June - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham - Weaveworks
- Nir, Joel Speed, Michael elmiko McCune - Red Hat
- Jason Scarano, Fabrizio Pandini, Naadir Jeewa, Yassine Tijani, Warren Fernandes, Nader Ziada, Robert Van Voorhees - VMware
- Matt Boersma, David Justice - Microsoft
- Jason Tarasovic, PayIt
- Jun Zhou - Spectro Cloud
- Joseph Davis - SUSE
- Seth Pellegrino – New Relic
- Daniel Lipovetsky - D2iQ

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - Proposed updates to CONTRIBUTING [#3191](#)
 - Proposed addition of REVIEWING.md [#3194](#)
 - Proposed updates to project roadmap [#3192](#)
 - Auto-approvals are being removed for PRs made by maintainers
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
 - External Remediation PR [#3190](#) [Nir]
- Discussion Topics
 - [Jason Tarasovic] Roadmap / vision for workload cluster upgrades
 - [Vince] In v1alpha3, there's two areas of upgrade: MachineDeployments, that were already supported since v1alpha2, and KubeadmControlPlane which now supports the upgrade of control planes internally, for both single and HA control planes.
 - [Jason T] Are there any plans to provide helpers for the common features such as Kubernetes version.
 - [detiber] In the past, there was discussion about changing the version in one place and have it flow through to all downstream resources, need a proposal on how this would work to go forwards.

- [vince] there were concerns previously about ordering, desires for all deployments to be upgraded at once or not.
- [fabrizio] There is documentation that can be improved, and additionally could look at the e2e tests as to how upgrades proceed.
- [Joel] MachineHealthCheck currently checks Machine across clusters
 - Not yet an issue, but was doing testing and seems that MHC isn't checking the cluster field and is returning machines from all clusters.
 - [detiber] is this where the label might match across multiple clusters
 - [joel] Yes, you would end up with machines from multiple clusters matching.
 - [vince] (in chat) We should force the label matching for the cluster. +1 from Fabrizio.
 - Tis a bug.
- New issue triage

Wed 10 June - 10am Pacific

[Recording](#)

Attending

- Seth Pellegrino, Andrew Rudoj - New Relic
- Vince Prignano, Fabrizio Pandini, Andy Goldstein, Yassine Tijani, Nader Ziada, Ben Moss, Warren Fernandes, Sedef Savas - VMware
- Joseph Davis - SUSE
- Justin SB - Google
- Michael Gugino, Michael elmiko McCune, Danil Grigorev, Alex Demicev - Red Hat
- Cecile Robert-Michon, Ria Bhatia - Microsoft
- Gianluca Arbezzano - Packet
- Jun Zhou - Spectro Cloud
- Max Rink - DTAG
- Jan Tilles - Ericsson
- Mark Emeis, Bryan Boreham - Weaveworks
- Mael Valais - Ori Industries
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ Hi!
- David Watson

Agenda

- Welcome/Introduction for new attendees

- PSAs
 - [vince] v0.3.7 ~2-3 weeks time frame
- Demos/POCs
 - [elmiko] cluster autoscaler support for management/workload cluster separation
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [naadir] Running CPI([Cloud Provider Interface](#)) on the management cluster
 - <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere/issues/924>
 - How are providers handling external cloud provider set up today?
 - [cecile] we are using post-apply, and are planning to use cluster resource set.
 - [vince] Should we do this in this series or later?
 - [cecile] Not urgent, we have some time given external CCM ([Cloud Controller Manager](#)) is mandatory in 1.21
 - AI: Get some lines on the roadmap doc for 0.4
 - [elmiko] capi-provider-provider or federated capi and other silliness
 - Came out of discussion of cluster autoscaler, could have a provider that could manage other capi management clusters. Have there been any other discussions in CAPI?
 - [vince] there's been off the cuff discussions on various items that
 - [elmiko] If we can reuse some of the abstractions and building blocks, it will make it easier for other projects to consume like the way autoscaler does.
 - [justinsb] before CRDs we had API aggregation (which came with a lot of baggage), but did allow us to represent the objects in both places. Personally prefer the kubeconfig approach.
 - [naadir] Hashing & Upgrades (<https://github.com/kubernetes-sigs/cluster-api/issues/3170>)
 - we use hashes to determine if capi(?) has changed and needs an upgrade
 - a change in 0.3.0 inadvertently changed the hash
 - additive changes are now causing this to break
 - this is a PSA to bring more attention to the changes
 - [spellegrino] in k/k hashes are not being used to determine deployment upgrade, it does a logic-aware comparison of the templates. this is a direction we should push towards, more semantic comparisons.
 - [mgugino] since we are still in alpha, and are making many changes, it seems like we will oscillate a lot in this space.
 - [ruboi] seeing that any version starting with 0.3.2 were changing hashes frequently and re-creating the same machines with different hashes
 - [vince] we should revisit the hashing function

- seems like consensus is that we should explore an alternative to hashing
 - [ace] what was the reasoning behind using hashes as opposed to more semantic style detection?
 - seems maybe historical based on other examples(?)
 - [spellegrino] not every change should cause a controller change
 - [vince] this is a blocker for 0.3.7
 - [vince] also need more testing around this, specifically targeting rolling upgrades and the hashing methods(or whatever replaces them)
 - [vince] let's a group together on slack or zoom to discuss further
 - [ruboi] whatever we do, please make it backward compatible
 - [gianluca] We (Packet) would like to move [cluster-api-provider-packet](#) to github.com/kubernetes-sigs
 - [vince] in favor, there is a process in place to get this accepted
 - [justin] don't think it would count as a separate project, so it should just take a *rubber stamp* to get it added
 - [lubomir] community file called sigs.yaml, should be added in here
 - <https://github.com/kubernetes/community/blob/master/sigs.yaml#L979-L994>
 - [justin] might be nice to add to test grid as well, but not required
 - [justin] next step, propose at the next sig cluster lifecycle meeting and/or on mailing list
 - [elmiko] autoscaling from/to zero ref: <https://github.com/kubernetes/autoscaler/issues/3150>
 - Requires touching each provider to expose additional requirements, so the autoscaler knows what kind of machine to remake when coming back up from 0, and some additional logic to expose taints.
 - this provider specific information is also used by the autoscaler for predictive simulations about scaling operations.
- New issue triage

Wed 27 May - 10am Pacific

[Recording](#)

Attending

- Vince Prignano, Jason Scarano, Naadir Jeewa, Andy Goldstein, Fabrizio Pandini, Robert Van Voorhees, Jason DeTiberus, Ben Moss, Yassine Tijani, Warren Fernandes, Justin Seely, Nader Ziada, Sedef Savas — VMware
- Andrew Rudoi - New Relic
- David Justice, Cecile Robert-Michon, Matt Boersma - Microsoft

- Mael Valais - Ori Industries
- Joel Speed, Alex Demicev - Red Hat
- Bryan Boreham - Weaveworks
- Justin SB - Google
- Daniel Lipovetsky - D2IQ
- Joseph Davis - SUSE

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - KCP adoption functionality has been merged ([PR](#)) :tada:
 - Conditions proposal has merged, implementation has started
 - MachineHealthCheck condition-based proposal amendment is almost ready ([PR](#))
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [ncdc] [CAPG flakes due to machine/project name length](#)
 - [yassine] Non-remediation use cases in MHC
 - In CAPV, we have states that do not require automatic remediation by MHC and require operator action. E.g. datastore is full and the operator needs to clean some disk space. We would like MHC to not delete the machine when the machine fails in this case.
 - [vince] do you expect MHC to re-evaluate after a certain time.
 - [joel] Have seen a similar issue with spot instances where spot instance is too low and AWS won't give a new instance.
 - [naadir] maybe a UX would be to end up with a condition on the machineset to show the problem to the end user and provide a hint to not continue remediation until the condition is resolved.
 - [joel] maybe you could hot loop, but with backoff.
 - [Sedef] ClusterResourceSet [CAEP](#)
 - If there are no further comments, let's proceed to merge
 - <<add your items here>>
- New issue triage

Wed 20 May - 10am Pacific

[Recording](#)

Attending

- Andrew Rudoi - New Relic
- Jason Scarano, Andy Goldstein, Naadir Jeewa, Nader Ziada, Fabrizio Pandini, Warren Fernandes, Sedef Savas, Gab Satchi, Yassine Tijani, Tim St. Clair, Ben Moss - VMware
- Nir, Michael Gugino, Michael Elmiko McCune, Alex Demicev, Danil Grigorev, Joel Speed - Red Hat
- Joseph Davis - SUSE
- Cecile Robert-Michon, David Justice, Matt Boersma - Microsoft
- Jun Zhou - Spectro Cloud
- Bryan Boreham - Weaveworks
- Gianluca Arbezano - Packet
- Daniel Lipovetsky - D2IQ
- Jacob Blain Christen - Rancher Labs
- Hardik Dodiya - SAP
- Prakash Ramchandra - Dell
- Moshe Immerman - Flanksource

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [v0.3.6 has been released](#)
 - Breaking Change: MachineHealthCheck spec.selector field cannot be empty.
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
 - [Naadir] [Bootstrap reporting/failure detection](#)
 - Use cases wanted
 - Meeting on Friday at 19:00 UTC (link in doc above)
 - [Tim] What do we have now?
 - [Naadir] Only the happy case, and otherwise MHC can timeout a machine and have it deleted and recreated.
 - [Rudoi] Right now we have to shell on the machine and find out what happened. Really want a way to find out Terminal conditions where whatever the MHC does is not going to resolve the problem.

- [ncdc] [ClusterResourceSet](#) (was: PostApply) request for use cases
 - Not meant to replace cluster addons, more of a bridge. Key use case is installing CNI.
 - Originally, we were going to use label selectors on a single clusterresourceset, and the controller would apply that to every matching cluster. However, downstream we found we could drive this “per-cluster” instead of via a label selector. Do people have a need for this capability or are people thinking of specifying it at a per-cluster level. Feedback required on UX.
 - [Daniel] For the label selector approach, does that mean there is a single clusterresourceset that applies to some set of clusters, and then what happens when it gets removed, do you no longer know what the source is?
 - [ncdc] We propose a method by which to track what happens. The design makes sense from an API perspective, but not necessarily from a CLI or UI perspective.
 - [Rudoj] The label selector pattern is used heavily at New Relic, e.g in ArgoCD, we match on environment label. Think it applies quite well as a pattern, but have questions around lifecycle management, does the CNI get reapplied?
 - [ncdc] The proposal right now is for a one-time apply, but with continuous reconciliation in a future enhancement.
 - [tim] Am torn, traditionally liked labels, but one thing I don’t like is about non-fungible properties of a cluster, but labels are. A label selector implies someone can switch a cluster from “dev” to “prod”. If labels aren’t to be used, then a way for multiple clusters to reference common resources.
 - [yassine] analogous to the service endpoint, whereas you need to keep track of references. If we’re also acting on a per-cluster basis, then how is it operated at scale.
 - [ncdc] on the CLI perspective, we were looking at “clusterctl config cluster” - would that create one CRS per cluster, which isn’t efficient but separates concerns. We may tweak the proposal to make it easier to do a per cluster set up.
- Discussion Topics
 - [Naadir] [Draining Service Type=LoadBalancer](#)
 - [michael gugino]
 - [yassine] was similar to when finalizers weren’t available for service type=loadbalancer,
 - [tim] does enforced/uniform tagging of cloud provider resources help?
 - [Naadir] Not necessarily, especially if people create additional things and don’t tag them

- [david] if you couldn't delete the vpc because there are dependent resources, would you continue to try to reconcile?
 - [Naadir] Do that now, more of a UX discussion
 - [david justice] Azure sees people go in and add resources that weren't part of CAPI install that create dependencies. May need message to indicate it will try but stop at some point
- [michael] Wasn't there a KEP about having the cloud provider interfaces in Kubernetes tag everything?
- [moshe] move toward declarative management of cloud provider resources + a custom finalizer that CAPI can add to these resources to handle deletion (don't let people manage things out of band)
 - [detiber] How do we design this given the split between management cluster and workload clusters (e.g. the workload clusters have Service type=LoadBalancer, the management cluster has Clusters/Machines)
- [rudoij] Is the expectation that a CAPA-created VPC is only for Cluster API things? Not sure that it's in the scope of Cluster API to clean up other cloud resources beyond the ones it manages directly.
- [tim] Cluster API clusters that use CPI - the CPI should tag everything it creates with a consistent tag or tags (makes cleanup easy). The broader story of cleaning up other resources is probably a policy decision. We may have some situations that require user intervention.
 - [naadir] CPI does currently tag the resources it creates. CAPA just doesn't have code to handle all of them (NLB, ALB, for example).
 - [tim] We should just be able to leverage finalizers on these resources
 - [detiber] But the resources are created in the workload clusters. We'd have to issue deletes against all these resources in the workload clusters.
 - [tim] There should be a way to tell the CPI to do cleanup. This seems like CCM problem IMO.
- [Nir] [CAEP External Remediation](#) and [PR #3056](#)
 - AI: Set up dedicated meeting
- [michaelgugino] [Machine Lifecycle Hooks](#)
 - AI: Open up PR
 - [vince] maybe scope it down to pre-delete for now.
 - [michael] pre-drain and pre-delete are related enough that we can address two use cases simultaneously.
- [wfernandes] [Clusterctl Extensible Template](#) CAEP merging in this week :)
- [Gianluca] Say hello and an unplanned presentation (thanks Vince :)) about [Packet and its ClusterAPI](#)
 - How do we get this SIG adopted?

- [Jason] Model is that there are contributors across companies, but maybe different for a provider that is developed and maintained for and by a provider. Would go up to SIG Cluster Lifecycle, then a request for repositories to be made, requisite integrations with prow etc...
- [Tim] Modus operandi for vendors is to get several approvers. The best way is to use the Kubernetes [standard project template](#). Repo which includes all the logistics around OWNERS files.
- Main repository guidelines: <https://github.com/kubernetes/community/blob/master/github-management/kubernetes-repositories.md>
- Shoutout to Fabrizio for helping out
- [joadavis] How hard is it to plug in a provider project outside of k8s-sigs (from a separate repo)?
 - [ncdc] Can be added to <https://cluster-api.sigs.k8s.io/clusterctl/provider-contract.html>. Additionally been discussion around documentation and how it could be pulled in, tbd.
- New issue triage

Wed 13 May - 10am Pacific

[Recording](#)

Attending

- Jan Tilles (Ericsson)
- Joel Speed, Michael 'elmiko' McCune, Nir (Red Hat)
- David Justice, Cecile Robert-Michon, Ace Eldeib, Matt Boersma (Microsoft)
- Seth Pellegrino – New Relic
- Jason Scarano, Ben Moss, Vince Prignano, Fabrizio Pandini, Warren Fernandes, Andy Goldstein, Tim St. Clair, Gab Satchi, Dan Finneran, Nader Ziada- VMware
- Bryan Boreham - Weaveworks
- Joseph Davis - SUSE
- Spencer Smith - Talos Systems
- Jacob Blain Christen - Rancher Labs
- Claudia Beresford, Mael Valais, Rabieh Fashwall - Ori Industries
- Moshe Immerman - Flanksource
- Hardik Dodiya - SAP
- Prakash Ramachandra (Dell)
- Verónica López- DigitalOcean

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] Cutting v0.2.11 (v1alpha2) release tomorrow
 - [Relevant bug found in conversions](#)
 - [vince] Getting ready for [v0.3.6 milestone](#)
 - [cecile] CAPZ [v0.4.3](#) released with experimental support for MachinePools
 - [tim] will there be broader adoption?
 - yes
 - [fabrizio] how can this work with the health check?
 - [vince] It won't work directly, but MHC could be modified to work with MachinePools, or we could add a new MachinePoolHealthCheck type of solution
 -
- Demos/POCs
- General questions / provider implementer questions
- RFCs / discussion on future proposals [time-box to 5m]
- Discussion Topics
 - [Conditions CAEP update](#) (Fabrizio)
 - Proposal updated with uses cases and examples
 - We are ~aligned with a [Kubernetes KEP for standardising conditions](#); the proposal introduces an additional for improving observability on CAPI long running task
 - Big effort for on defining guidelines (thanks @vince) i.e.
 - How to define condition semantic/polarity:
"Status=True" → good → the component is ready for serving application workloads
 - How to implement conditions in the context of the Cluster API hierarchy of objects and its extensibility points.
 - There will be a new util.conditions package that simplifies implementation for providers / enforce compliance with the guidelines.
 - No breaking changes/support for incremental adoption across provider
 - MHC proposal needs another update (Ben Moss)
 - Opened a PR [here](#)
 - If you have edit permissions on a Machine, but not delete, the ability to add the unhealthy annotation effectively gives you delete permission
 - The new PR replaces the annotation with conditions, to avoid the permissions escalation
 - [CAEP Extensible Template Processing for clusterctl](#) (Warren F.)
 - Sync between CAPV and CAPM3 about IPAM work (Maël K.)
 - [CAPV issue](#)

- [Yassine] More inclined to say this is a more on-prem provider specific set of use cases. Not sure we want to do it at the CAPI level.
 - [Naadir] If there's common enough use cases across the on-prem providers, then would be worth co-operating on a common effort, even if not in the top-level Cluster API.
 - [CAEP External Remediation](#) (Joel Speed on-behalf of Andrew Beekhof)
 - Looking nearly ready, could solve MHC annotation issue raised recently
 - [Kubeadm proposal for customising pod manifests via patch files](#) (Naadir Jeewa on behalf of Lubomir Ivanov)
 - If you need to make custom changes to etcd, api server beyond volume mounts, this will be the mechanism.
 - [Andrew Sauber] We do patch CoreDNS after the fact. Is there a summary of why JSON is preferred over kustomize?
 - [Fabrizio] Yep the kubeadm KEP describe reasons behind the switch to json patch
 - Provider multitenancy (Naadir Jeewa)
 - Allowing providers to control resources across different accounts, [starting with AWS](#)
 - Security model is to have accounts as cluster scoped resources, with a declaration of allowed namespaces, similar to [Service API evolution](#).
 - CAPBK files from secret (Ace Eldeib)
 - <https://github.com/kubernetes-sigs/cluster-api/pull/3038>
 - [andy] we've started lazy consensus on some of the CAEPs
- New issue triage

Wed 6 May - 10am Pacific

[Recording](#)

Attending

- Cecile Robert-Michon, Ria Bhatia, David Justice, Ace Eldeib - Microsoft
- Michael Gugino, Joel Speed, Michael elmiko McCune - Red Hat
- Naadir Jeewa, Andy Goldstein, Warren Fernandes, Yassine Tijani, Nader Ziada, Sedef Savas, Ben Moss, Gab Satchi - VMware
- Bryan Boreham - Weaveworks
- Joseph Davis - SUSE
- Jacob Blain Christen - Rancher Labs
- Maël Valais, Claudia Beresford, Harry Lawrence, Michael Kashin - Ori Industries
- Sudeep Batra - Ericsson
- Daniel Lipovetsky - D2IQ

- Erwin van Eyk - Platform9
- David Watson
- Andrew DeMaria - Cloudflare
- Yuvraj - Nirmata
- Hardik Dodiya - SAP

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] [v0.3.5 has been released with bug fixes and support for Kubernetes v1.18](#)
 - [vince] Bug in MachineDeployment controller with mixed v1alpha2 / v1alpha3 objects ([PR](#))
- Demos/POCs
- General questions / provider implementer questions
 - [Maël K] What is the reason behind making *KubeadmCommands and files field of KCP Spec immutable ? (related to <https://github.com/kubernetes-sigs/cluster-api/issues/3014>)
 - [ncdc] Are open to making fields mutable, but rollbacks are difficult to orchestrate. Is copied to the Machine object because that's the object the core controllers use.
 - [Bryan B] Why is providerID field required on both [Machine](#) and on [infrastructure-specific Machine](#) objects?
 - [ncdc] Ultimately need to link a node to a machine.
 - [bryan b] find it interesting that it looks in both places
 - [ncdc] the machine controller is the only one that looks in both places for the purpose of copying. Anything else that is doing things with nodeRef is doing it with the Machine object. Could evaluate whether we need to continue doing it.
 - [bryan b] to me, the infrastructuremachine is where it should be.
 - [vince] it's optional on the Machine object, it was more for other controllers that need to do it. There is a log message for the error, but contractually not required.
 - [Maël Valais] capd-manager:v0.3.0 image is not present on [gcr.io](#)? The v0.2.1 is there, did we forget to push the v0.3.0?
 - Provider was previously in its own repo ([cluster-api-provider-docker](#)) and is now in the [cluster-api](#) repo
 - [d lipovetsky] the experience has changed: <https://cluster-api.sigs.k8s.io/clusterctl/developers.html>. The image isn't being published right now, follow up with issue and ping me (@dlipovetsky on slack).
 - [Mael Valais, 21 days later] Created [issue 3101](#)
- RFCs / discussion on future proposals [time-box to 5m]

- [cecile] [Mounting etcd data disk proposal](#)
 - Ability to mount disk requires changes to kubeadmbootstrap to support relevant parts of cloud-init wrt to filesystem creation and mounts.
 - There is a bit of a difficulty knowing which disk the etcd ends up on, but shouldn't affect the changes to the bootstrap config.
- [fabrizio] [Conditions \(PR\)](#)
- [naadir] [CAPA controller multi-tenancy](#) (maybe relevant to other providers)
- [sethp] Infrastructure linking through CAPI ([related CAPA issue](#))
 - Want to allow in the AWS case the ability to change tags and security groups without destroying machines. Requires ability to figure out where the machine as far as versioning is in relation to the template.
 - [vince] Worth producing a proposal
- Discussion Topics
 - [ria] Doc support from the cncf ([Issue](#))
 - We want users to land in the CAPI book
 - Intent to allow providers to write docs in a structure that can roll up into the book
 - CNCF can provide support in the form of doc writers
 - [elmiko] Experiences from the past is that we can produce templates in the main repo with instructions on how to copy them.
 - [ria] we want this to look very similar across providers.
 - Markdown can live in provider repos, with a job to pull them in to the docs
 - [vince] how to handle alpha version changes to docs? Some may go in at different times.
 - May need to tag docs for matching to releases
 - [ria] Happy birthday, Cecile!!
 - <<add your items here>>
- New issue triage

Wed 29 April - 10am Pacific

[Recording](#)

Attending

- Andrew Rudoj, Seth Pellegrino - New Relic
- Bryan Boreham, Cornelia Davis - Weaveworks
- Tim St. Clair, Andy Goldstein, Vince Prignano, Gab Satchi, Nader Zini, Yassine Tijani, Warren Fernandes, Sedef Savas, Naadir Jeewa, Jason DeTiberus - VMware
- Prakash Ramachandran - Dell

- Cecile Robert-Michon, Craig Peters, Ria Bhatia, Matt Boersma, David Justice, Kei Yoshikoshi - Microsoft
- Joseph Davis - SUSE
- Joel Speed, Michael Gugino, Michael elmiko McCune - Red Hat
- Jacob Blain Christen - Rancher Labs
- Daniel Lipovetsky - D2IQ
- Shuting Zhao - Nirmata
- Verónica López- DigitalOcean
- David Watson
- Hardik Dodiya - SAP

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [v0.3.4 has been released!](#)
 - But there is a new [clusterctl init bug](#) with a fix in flight
 - [New proposal guidelines have been merged](#)
 - [michaelgugino] What is an experiment and why would I use it?
 - [vince] It's based after Kubernetes feature gates, so we can add new features that don't have any API or compatibility guarantees. You could add new controllers, new CRDs that live in a separate API group. These are disabled by default, and we ultimately want to promote them (if it makes sense) - either into Cluster API "proper" (non-experimental), or to external places.
 - [Kubernetes 1.18 support is underway](#) (v0.3.5)
 - [Fabrizio] New presubmit job pull-cluster-api-e2e
 - It tests the new quick-start workflow (with clusterctl), pivoting to self hosted, KCP & MD upgrades
 - It is based on a set of new helpers recently added in the test framework
 - Providers can use the new helpers or directly the entire specs, which are designed to run with any infrastructure provider
 - The pull-cluster-api-capd-e2e is going to be deleted
 - See the new job in testgrid [here](#)
 - [Cecile] What is the state of including tests for providers?
 - [Fabrizio] Is possible to enable in the configuration file, but of the opinion that in the end those tests should be in the provider repos.
 - [Yassine] What's the timeline for deprecation?
 - [Fabrizio] The tests are broadly compatible, but don't think we can delete those bits of the test framework that will go unused until next 0.x release.
- Demos/POCs
 - [elmiko] cluster-autoscaler with capd provider implementation

- General questions / provider implementer questions
 - [Shuting] Does CAPI want to support in place upgrades?
 - [daniel] depends on what you want. Can write a infra provider that will use the same machine with same os, but to cluster api it looks like a cluster api Machine has been deleted.
 - Implementing bare metal scenario, limited set of machines so in-place desired.
 - Usual assumption is to drain before an upgrade
 - [bryan] One place where the question “drain before upgrade” is debated: <https://github.com/kubernetes/website/issues/12326>
 - General policy for cluster-api is immutable
 - This has been an ask several times in the past
 - Sync up on slack channel, start writing a proposal around in-place upgrades and requirements
 - we should have user stories related as well
- RFCs / discussion on future proposals [time-box to 5m]
 - [michaelgugino]: Machine Lifecycle Hooks
 - https://docs.google.com/document/d/1pQLxOYA9r95jnrUZPimE4oajYjIPMcC9SPX2O8_nmRE/edit#
 - Reboot remediation support
 - https://docs.google.com/document/d/1fOdOzQgH6L8PWDPDHTfd_kUchSHUSsRNJMGfjkAvahc/edit
 - MHC External remediation
 - <https://docs.google.com/document/d/1V59vCguy5doNBL05lsRTIfon7AhUZ2oHfvhw1T4to6s/edit>
 - [Fabrizio] Conditions
 - https://docs.google.com/document/d/1eWTDcNp8p-m-RheaX4WJmmb0YGDSm_Ruj1GjtVZ_9Jk/edit
 - Goal of the first iteration to define Condition type so we can start implementing the first set of conditions and gather user feedbacks
 - There is a similar [KEP](#) ongoing in Kubernetes, trying to keep alignment but this is still WIP
 - Main difference so far is the introduction of a new field to provide actionable feedback on what is happening while a condition is not yet true (give visibility in long-running actions).
 - v0.3.x compatible (only field addition, no change of contracts)
 - Clusterctl extensible template processing
 - <https://docs.google.com/document/d/1GBYR9Dm0igw7FaDDwIM7OldhB5TMuixHswq7TAO1bQg/edit>
 - Post apply
 - <https://docs.google.com/document/d/1IWLGN66roMjXL49gKO6Dhwe7yzCnvgrCtjR9mu4rvPc/edit>
 - [michaelgugino]: Node maintenance lease

- <https://github.com/kubernetes/enhancements/pull/1411>
-
- Discussion Topics
-
- New issue triage

Wed 22 April - 10am Pacific

[Recording](#)

Attending

- Joel Speed, Michael elmiko McCune, Nir, Michael Gugino- Red Hat
- Bryan Boreham, Mark Emeis - Weaveworks
- Manik Taneja - Spot.io
- Jan Tilles, Sudeep Batra, Feruzjon Muyassarov - Ericsson
- Moshe Immerman - Flanksource
- Jason Scarano, Naadir Jeewa, Jason DeTiberus, Andy Goldstein, Vince Prignano, Gab Satchi, Fabrizio Pandini, Sidharth Surana, Nader Ziada, Sedef Savas, Yassine Tijani, Warren Fernandes - VMware
- Daniel Lipovetsky - D2IQ
- Seth Pellegrino - New Relic
- Maël Valais, Michael Kashin - Ori Industries
- Justin SB - Google
- Jacob Blain Christen - Rancher Labs
- Arjun Hemrajani - Razorpay
- David Justice - Microsoft
- Joseph Davis - SUSE
- Prakash Ramchandran (Dell)
- Nate Franzen - NetApp
- Hardik Dodiya - SAP
- David Watson

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] Revisiting issue-proposal and focusing on the CAEP process ([PR](#))
 - The gist is that there are 20+ issue proposals with various levels of information.

- [vince] Align on code freezes and releases ([issue](#))
- Demos/POCs
- General questions / provider implementer questions
 - Where can we review upcoming roadmap items?
 - <https://cluster-api.sigs.k8s.io/roadmap.html>
 - What's the maturity of existing providers?
 - <https://cluster-api.sigs.k8s.io/reference/providers.html>
 - [seth-p] we have 6 clusters of >100 node CAPA clusters, with the caveat that we use self-managed clusters.
 - [yassine] we have users with CAPV in production. The guess we can get more info via Slack.
 - [cecile] Reach out to the per-provider Slack to learn about specific providers.
 - Will clusterAPI deploy k8s inside k8s or using standalone nodes where components run as systemd services? Or we give consumers options to do both?
 - [vince] CAPI doesn't necessarily have k8s-in-k8s yet, there was an experimental implementation. Cluster API is however pluggable, so would allow the possibility to use a different implementation.
 - [Add Kubernetes provider to pre-defined list · Issue #2813 · kubernetes-sigs/cluster-api](#)
- RFCs / discussion on future proposals [time-box to 10m]
- Discussion Topics
 - [detiber] Adding Gab Satchi as an additional maintainer for cluster-api-provider-gcp: <https://github.com/kubernetes-sigs/cluster-api-provider-gcp/pull/294>
 - [ncdc] If you want to be an additional maintainer for this provider or any other, please reach out on Slack (using the specific or generic #cluster-api) channel.
 - [moshe] Service Health for rolling updates
 - How do you do a rolling update of a cluster whilst maintaining the health of the service underneath. A pod disruption budget solves some of this, but not always. E.g. a stateful service like elasticsearch, pods that may come up as part of nodes being rolled out may be marked by ready, but due to replication, etc... may not be truly ready.
 - [detiber] there is a shared responsibility here, right now we are completely mindful of PDBs, but would be interesting to see how we hook into this without needing to do the things the cluster autoscaler does.
 - [bryan] it's better to hold the ready status on the pod to block the rolling update
 - [moshe] the pod may be ready to serve traffic even if state isn't fully replicated for example
 - [bryan] it might be better to instead expand the definition of 'ready', and this is what weave does

- [vince in chat] document the use cases
- AI: Moshe to file CAEP
- [wfernandes] Clusterctl templating
 - <https://github.com/kubernetes-sigs/cluster-api/issues/2339>
 - Today, clusterctl does naive string templating, which causes issues with customisations. There has been discussion of using ytt. The idea is to define an interface to allow you plugin a templating engine.
 - Airship is using kustomize and doing templating. How would that change?
 - [warren] it should work, because of the pluggable model, and in addition the variable substitution will still continue to work.
 - [andy] kustomize used currently to build out release bundles. This change would make templating pluggable, and may have other options with a contrib folder for those.
- [sedefsavvas] Post-apply experimental feature
<https://github.com/kubernetes-sigs/cluster-api/issues/2395>
- [Feruz] External Machine Remediation
 - <https://docs.google.com/document/d/1V59vCguj5doNBL05lsRTIfon7AhUZ2oHfvhw1T4to6s/edit>
 - Setup separate public meeting to discuss
- New issue triage (new issues with no milestones)
 - Discussion on support for adopting an existing cluster into cluster-api, vs leaving that as an unsupported but allowed process and then possibly documenting/supporting a migration path

Wed 15 April - 10am Pacific

[Recording](#)

Attending

- Vuk Gojnić, Max Rink, Marcel Fest - Deutsche Telekom
- Charles Sibbald, Mark Ramm, Bryan Boreham, Mark Emeis - Weaveworks
- Jan Tilles, Feruzjon Muyassarov - Ericsson
- Caleb Bron - Rancher Labs
- David Justice, Cecile Robert-Michon, Ace Eldeib, Ria Bhatia - Microsoft
- Joel Speed, Michael elmiko McCune, Nir, Michael Gugino - Red Hat
- Ben Moss, Jason Scarano, Yassine Tijani, Andy Goldstein, Vince Prignano, Jason DeTiberus, Warren Fernandes, Fabrizio Pandini , Nader Ziada, Gab Satchi - VMware
- Yuvraj - Nirmata
- Verónica López - DigitalOcean
- Hardik Dodiya - SAP
- Prakash Ramchandran - Dell

- David Watson

Agenda

- Welcome/Introduction for new attendees
- PSAs
- General questions / provider implementer questions
- Demos/POCs
 - [vukg, casibald] Das Schiff - Deutsche Telekom Kubernetes Engine based on CAPI, Helm and GitOps (for off-line follow-up we are available in Kubernetes Slack(GitHub) under Vuk Gojnic(vukg), Charles Sibbald(casibald), Maximilian Rink(MaxRink), Marcel Fest(cellebyte)).
- Discussion Topics
 - [bmoss] Thoughts on adding [Codecov](#) or similar tools to CAPI projects
 - Lots of +1s as long as it's done through Prow and non-blocking
 - [wfernandes] Added this issue to track - <https://github.com/kubernetes-sigs/cluster-api/issues/2918>
 - [jan] External machine remediation. <https://github.com/kubernetes-sigs/cluster-api/issues/2846>
 -
- New issue triage

Wed 8 April - 10am Pacific

Recording TBD

Attending

- Daniel Lipovetsky - D2IQ
- Cecile Robert-Michon, Ace Eldeib - Microsoft
- Joel Speed, Michael elmiko McCune, Michael Gugino - Red Hat
- Mark Emeis, Bryan Boreham - Weaveworks
- Jason DeTiberus, Vince Prignano, Naadir Jeewa, Sedef Savas, Gab Satchi, Nader Ziada - VMware
- Jun Zhou - Spectro Cloud
- Prakash Ramchandran (pramchan) Dell
- Hardik Dodiya - SAP
- Joseph Davis - SUSE

Agenda

- Welcome/Introduction for new attendees

- PSAs
 - Roadmap updates / focus on v1alpha3+
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2882/files>
 - Want to point out the e2e spreadsheet linked above about what tests we want to have. Please take a look if you are a contributor or infrastructure provider to help with robustness.
 -
 - [daniel] v1alpha3 blog
 - <https://github.com/kubernetes/website/pull/20140>
 - Preview:
 - <https://deploy-preview-20140--kubernetes-io-master-staging.netlify.com/blog/>
 - I wrote the names of the contributors for each feature from memory. If anyone has been missed, ping in Slack.
- General questions / provider implementer questions
- Demos/POCs
- Discussion Topics
 - [pramchan] What is the cluster-api responsibility regarding ability to update kubernetes components configuration of a running cluster. i.e. adding arguments or removing from the api server configuration. I.e. CAPBK?
 - [vince] CAPI aims to manage immutable infrastructure. It's not necessarily the role of CAPI to manage running infrastructure. CAPI today can't do an in-place upgrade on a running machine from say 1.16.2 to 1. "Replace" rather than "upgrade"
 - XRef [Kubeadm operator KEP POC](#)
 - [pramchan] Will v1alpha4 be compatible with v1alpha3?
 - [vince] we reserve the right to make breaking changes, we document these in the book.
 - [pramchan] What about semver guarantees?
 - [vince] CAPI follows upstream Kubernetes guarantees
 - <https://kubernetes.io/docs/concepts/overview/kubernetes-api/#api-versioning>
 - [elmiko] jin relation to the notion of compatibility, breaking changes, and semver. just a warning not to read too much into the ideas expressed at semver.org in relation to the version numbers we are using.
 - [joadavis] There's a bunch of environment variables that OpenStack requires. Not sure what's being asked for and no default values or examples.
 - [vince] Open an issue to update the quickstart to explain what these are, but additionally look through the yaml, and the environment variables should be visible as "\${ FOO}"
 - for reference,
 - <https://kubernetes.slack.com/archives/C8TSNPY4T/p1586295111284200>

-
- New issue triage

Wed 1 April - 10am Pacific

[Recording](#)

Attending

-
- Michael elmiko McCune, Joel Speed, Michael Gugino - Red Hat
- Justin SB - Google
- Joseph Davis, Danny Sauer - SUSE
- Seth Pellegrino – New Relic
- Verónica López- DigitalOcean
- Sudeep Batra - Ericsson
- Cecile Robert-Michon, David Justice - Microsoft
- Fabrizio Pandini, Jason DeTiberus, Ben Moss, Dan Finneran, Naadir Jeewa, Vince Prignano, Nader Ziada, Gab Satchi, Yassine Tijani, Sedef Savas, Warren Fernandes - VMware
- Hardik Dodiya - SAP
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ
- Nate Franzen - NetApp
- Bryan Boreham - Weaveworks

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [cecile/vince] Cluster API Deep Dive today at 1:20pst as part of Virtual Rejekts
 - Youtube livestream: <https://youtu.be/cZEcdTOwV-A>
 - [elmiko] autoscaler updates for cluster-api
 - readme has landed, [pull request](#)
 - cluster-api label added, [quick filter link](#)
 - next steps are focused on cleaning up code and adding more tests (unit and e2e)
 - [Joel] Spot instance proposal PR has been added
 - No further comments for a couple of weeks on gdoc
 - Please review/add further comments
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2817>

- Copied from the google doc, fixed up comments from Jason
- Let me know if the reviewers list should be updated
- General questions / provider implementer questions
 - [Cecile] clusterctl quickstart instructions move faster than infra providers on adopting new versions of capi.
 - Say we release 0.3.4 next week, the quickstart also updates to 0.3.4. Don't expect a breaking change in a patch version, but that may break a downstream infrastructure provider.
 - [Vince] In general, a CAPI release shouldn't break everyone, and we'd have to revoke it if it did so.
 - [Cecile] Are there ways to find out it is broken before users tell us?
 - [Vince] Are planning on doing this for e2e
 - [fabrizio] currently adding clusterctl to the e2e framework such that a downstream provider could execute a clusterctl e2e
- Demos/POCs
- Discussion Topics
 - [ncdc] clusterctl go library changes
 - Would like to propose to keep API stable, but for the rest of the code in cmd/clusterctl, would like to change function signatures and datatypes
 - [tim st clair] might be acceptable at this early stage of clusterctl, but will eventually need to define a contract. Given what has happened with client-go 1.18, it's clear people will not be happy with this change.
 - [david justice] would it make sense to separate the modules - what's the api and the rest in /cmd/clusterctl?
 - [vince] have gone in the other direction in the past to simplify release
 - [elmiko] improving the state of quickstart and developer docs
 - Emerged out of discussions with new users and debugging sessions working through the docker deployer. I and another user have come up with install instructions for Fedora and Ubuntu respectively. Would it make sense to take out Docker and put that in the development section, which mirrors comments from people involved in CAPI development.
 - [docs/book: Update CAPD quickstart once ships in this repository](#)
 -
 - [ncdc] bunch of discussion above about what to do with the docker provider
 - [Revisit documentation structure](#)
 - [ria] would like to get more of a team together to talk about doc structure.
 - [elmiko] +1, probably could easily excise the docker docs into a separate developer docs.
 - [ncdc] have talked in 1 or 2 meetings about what to do with the docs and where CAPD fits in. Know Daniel Lipovetsky

had issues with downranking the importance of the Docker provider given not everyone has budget to use Amazon, Azure, Google etc...Happy to have a link along the lines of "i don't have access to a provider, see this"

- +1 from Joseph Davis for having/keeping a way for people to consume without an external provider
- +1 to moving out from Warren
- [cecile] Clusterctl github API limits issue
 - [Testing of clusterctl exhausts Github API limits · Issue #2450 · kubernetes-sigs/cluster-api](#)
 - People attempting to use clusterctl frequently hit the issue
 - [ncdc] Which API is being hit?
 - [fabrizio] the rate limit is very low
 - [ncdc] could we cache some of the data locally?
 - [vince] we have done some work in the docs publishing tool, so can look at this.
- [cecile] Expand test framework to include k8s upstream testing <https://github.com/kubernetes-sigs/cluster-api/issues/2826>
 - There's lots of details on the above, so take a look
- [vince] KCPv2 robustness and incremental improvements <https://github.com/kubernetes-sigs/cluster-api/issues/2753>
 - If you have ideas on how to make KCP better take a look at the above
 - Starting off with unit tests and e2e tests, then refactoring
- [prakash] CAPI Control Plan follow up from last week
 - <https://etherpad.openstack.org/p/airship-virtual-meetup-2020>
 - Refer CAPBK <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/control-plane.html>
 - Would like review of consumption of v1alpha3 in Airship via [templating of the KubeadmControlPlane types](#). Should we consume the default template and use kustomize vs. copying it and then completely customise it.
 - Make a template from scratch, or use the provided with the release. Question about when to customize for Metal3 and add configuration - before handing to cluster-api or after.
 - [Naadir] The contract in Cluster API is the CRD, not necessarily the default templates that are in the release which are really for quickstart instructions. [For Cluster API Provider vSphere. we generate our own templates programmatically.](#)
 - [Tim] Packet is developing a provider based on [Tinkerbell](#)
- [Bryan] Should put a password on the Zoom meeting. [Other tips](#)
 - The trolls are about

- [tim] Is a balance, there is a host password, but if we lock it down too much, people can't join. We follow the guidance from contrib-x, but will follow up incase there's updated guidance.
- [jason] may be worth it to re-enable the waiting room by default.
 - Would have a co-host watching who's entering the room
- [lubomir] it's hard to tell if whoever is joining is malicious or not
- Possible to disable screen sharing and co-hosts.
- [Yassine] cloud-init progression
 - When working on CAPV, we saw issues we couldn't see why cloud-init failed without logging on
 - <https://github.com/kubernetes-sigs/cluster-api/issues/2554>
 - CAPV community meeting - people thought it was strange that we don't allow two-way communications between workload cluster and management cluster
 - May not be suitable for all deployments, but maybe it could be opt-in?
 - [Jason] Would limit bootstrapping, because you would already need a kubernetes clusters.
- <<add your items here>>
- New issue triage

Wed 25 March - 10am Pacific

[Recording](#)

Attending

- Bryan Boreham, Richard Case - Weaveworks
- Jason Scarano, Jason DeTiberus, Gab Satchi, Warren Fernandes, Andy Goldstein, Yassine Tijani - VMware
- Joel Speed, Michael elmiko McCune, Michael Gugino - Red Hat
- Prakash Ramchandra - Dell
- Jason Tarasovic - Paylt
- Maël Valais, Matt Shimwell - Ori Industries
- Hardik Dodiya - SAP
- Cecile Robert-Michon, James Sturtevant - Microsoft
- Pawel Bek - Codahead
- Joseph Davis - SUSE

Agenda

- Welcome/Introduction for new attendees

- PSAs
 - Preparation for [v0.3.3](#) release (due possibly friday 27th)
 - Kubeadm retry joins: <https://github.com/kubernetes-sigs/cluster-api/pull/2763>
 - Bug in MD <https://github.com/kubernetes-sigs/cluster-api/pull/2740>
 - KCP scale down bug fix <https://github.com/kubernetes-sigs/cluster-api/pull/2768>
 - CRD contract label cannot have commas <https://github.com/kubernetes-sigs/cluster-api/issues/2775>
- Demos/POCs
- Discussion topics
 - [detiber] Followup on idea of using nested staging/production image repositories: <https://github.com/kubernetes/k8s.io/pull/658#pullrequestreview-374651405>
 - Jason to draft a proposal
 - [prakash] Ask from Airship team for CAPI Control Plane roadmap on March 31 st
 - AT&T Airship team is interesting in learning more about cluster-api and control plane issues
 - <https://www.airshipit.org/>
 - there is an airshipctl that works with cluster-api, they are looking to find out where the gaps are between it and clusterctl
 - they would like to learn more about cluster-api technically, this would be better if we (cluster-api) had some requirements or technical questions we could prepare to answer
 - this would be best started via an email conversation and then we can escalate to meetings. sounds like there was an email, but perhaps it went to the wrong person.
 - Vince and Prakash to sync offline
 - [andy] revisiting the topic of reducing/removing provider office hours
 - any questions around this? or any specific provider questions?
 - Prakash's question fits this nicely ;)
 - we need to make this a top level item on the agenda, namely "Provider questions"
 - [tim] kind of a PSA that tinkerbell has finally been opensourced for those looking at bare metal solutions. <https://github.com/ackeam/tinkerbell>
 - <Add agenda items here>
- New issue triage

Wed 18 March - 10am Pacific

[Recording](#)

Attending

- Michael Gugino, Michael elmiko McCune, Joel Speed - Red Hat
- Jason Scarano, Andy Goldstein, Vince Prignano, Chuck Ha, Yassine Tijani, Gab Satchi, Fabrizio Pandini, Dan Finneran, Noa Amran, Jason DeTiberus - VMware
- Seth Pellegrino – New Relic
- Cecile Robert-Michon - Microsoft
- Bryan Boreham - Weaveworks
- John Eldridge - NTTData
- Maël Valais, Nick Sopuch - Ori Industries
- Marko Mudrinić - Loodse
- David Watson
- Daniel Lipovetsky - D2IQ
- Jason Tarasovic - Paylt

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - Cluster-API Provider for DigitalOcean has moved to [v1alpha2](#)
 - Cluster-API Provider for DigitalOcean is looking for new maintainers
 - Previous maintainers have decided to step down because they don't have the capacity to maintain the project anymore
 - It has been proposed to promote cpanato and prksu to maintainers as they have been active contributors and are interested to maintain the project ([#145](#))
 - The PR is on hold until 3/18 EOD for lazy consensus
 - We are looking for new contributors and maintainers!
 - <https://groups.google.com/d/msg/kubernetes-sig-cluster-lifecycle/iKvpeKAJmNo/gMtU0TuLAWAJ>
 - Preparing for [v0.3.2](#) release tomorrow
 - KCP bug fixes for self-managed clusters
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2700>
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2696>
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2695>
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2711>
 - This will ensure that the health check takes into account that the number of expected etcd nodes can be different than the number of control planes that are currently up, allowing for better reentrance. It'll also greatly reduce the number of etcd leadership changes we make during an upgrade, which went from N (where N is the number of replicas) to 1.

- KCP bug — no timeout was set on the remote client
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2708>
- KCP bug — CoreDNS upgrade image should be OCI compliant
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2699>
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2697>
 -
- [dlipovetsky] V1alpha3 [blog post draft](#).
 - The first version is done. Please provide feedback using comments/suggestions as soon as you can. The blog must be submitted by next Monday.
 - Have we overlooked any important features
 - If you worked on a feature that's covered, are there any technical errors? Are you properly credited?
 - Anything else.
- capz [v0.4.0](#) release (v1alpha3) is now available
- Demos/POCs
- Discussion topics
 - [detiber] wg-k8s-infra is wondering our thoughts on nested staging repos for related projects:
 - <https://github.com/kubernetes/k8s.io/pull/658#pullrequestreview-374651405>
 - [Joel/Michael] Spot instances proposal
 - Please add comments to https://docs.google.com/document/d/1naxBVVII_O-u6TchvQyZFblakrwU9qAzYD4akyV68nQ/edit#
 - Hoping to solidify into PR soon
 - [Bryan] The “Cluster API Provider Implementers' office hours” meetings seem very sparse - should the format change?
 - xref <https://github.com/kubernetes-sigs/cluster-api/issues/1027>
 - [dlipovetsky] It's become a place for CAPI “newbies” to ask questions. We haven't had that space in the main CAPI meeting for a while.
 - [tstclair] A “getting started” FAQ (addressing more basic questions) would help
 - [tstclair] We're efficient in the main meeting. Is there room for a Q&A section following the regular discussion/issue triage
 - [cecile] The main meeting can be intimidating. Let's have a dedicated time to ask basic questions.
 - **Decision:** cancel the separate provider implementers' office hours and roll them into this meeting. Set up a dedicated slot in the agenda for Q&A.
 - Removing meetings from community page: <https://github.com/kubernetes/community/pull/4636>
 - [tim] How was the CAPZ conversion to v1a3? Have people started using it?
 - [cecile] Way easier than v1a2 to v1a3

- Not sure if anyone has tried it yet
- <Add agenda items here>
- New issue triage

Wed 11 March - 10am Pacific

[Recording](#)

Attending

- Daniel Lipovetsky - D2IQ
- Bryan Boreham - Weaveworks
- Michael elmiko McCune, Joel Speed, Michael Gugino - Red Hat
- Cecile Robert-Michon, David Justice, James Sturtevant - Microsoft
- Yassine Tijani, Andy Goldstein, Jason DeTiberus, Gab Satchi, Nader Ziada, Fabrizio Pandini, Chuck Ha, Noa Amran - VMware
- Hardik Dodiya - SAP
- Justin SB - Google
- Paweł Bek - Codahead
- Nate Franzen - NetApp

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [V0.3.0!!!](#)
 - [dlipovetsky] v1alpha3 [blog post draft](#). Still in progress, feedback is very welcome!
- Demos/POCs
 - [joelspeed] Machine Health Check Demo
- Discussion topics
 - [detiber] Last minute FailureDomain support tweaks
 - PSA to anyone upgrading on Infra Provider to v1alpha3
 - [andy] Status check for infrastructure providers
 - CAPA
 - [detiber] release expected later today or sometimes tomorrow
 - CAPV
 - [yassine] release expected tomorrow
 - CAPZ
 - [cecile] ~week timeframe, discussing it at community meeting
 - CAPMetalKube (CAPM3)
 - [mael] release is out

- [vince] schedule some code walkthroughs
 - Clusterctl v2
 - Kubeadm control plane
 - <https://www.youtube.com/watch?v=e9Wnhoh0Fy0> [pawel - api code tour]
- [james s] managed cluster support update
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2045>
- New issue triage

Wed 4 March - 10am Pacific

[Recording](#)

Attending

- Andrew Rudoi - New Relic
- michael elmiko mccune, Michael Gugin, Eric Duen, Alberto - Red Hat
- Cecile Robert-Michon, David Justice, James Sturtevant - Microsoft
- Jason DeTiberus, Vince Prignano, Andy Goldstein, Nader Ziada, Warren Fernandes, Gab Satchi, Yassine Tijani, Sedef Savas, Fabrizio Pandini, Noa Amran, Chuck Ha - VMware
- Paweł Bek - Codahead
- Jun Zhou - Spectro Cloud
- Mark Emeis, Bryan Boreham - Weaveworks
- Prakash Ramachandran - Dell
- Daniel Lipovetsky - D2IQ
- Moshe Immerman - Flanksource
- David Watson
- Hardik Dodiya - SAP

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] RC.3 dropping today ~2/3pm PST
 - [vince] On track to release v0.3.0 on March 10th
- Demos/POCs
 - [alberto] Autoscaler demo.
 - <https://github.com/kubernetes/autoscaler/pull/1866>
- Discussion topics
 - [elmiko] recap of autoscaler integration meeting

- [michaelgugino] cluster scale from zero UX discussion proposal: <https://github.com/kubernetes-sigs/cluster-api/pull/2530>
 - [andrew rudoi] should there be some linkage, or usage, of the capabilities feature that gugino is also working on?
- [james s] support for managed clusters
 - <https://github.com/kubernetes-sigs/cluster-api/issues/980>
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2045>
-
-

Tue 3 March - 1400 UTC Autoscaler integration ad-hoc meeting

Attending

- michael elmiko mccune, alberto, joel speed, michael gugino - red hat
- Jason DeTiberus, Naadir Jeewa - VMware

Agenda

- intros
- [alberto] review of autoscaler sig call
 - sync up with maintainers
 - happy with the current pr
 - <https://github.com/kubernetes/autoscaler/pull/1866>
 - red hat has agreed to help support and maintain
 - few details need to finish then pr can merge
- [alberto] want to make sure we are on the same page, and start to form a roadmap
 - how to organize work that will come
 - we need a followup to the autoscaler pr to help fill the details
- [jason] from cluster-api side can get a review on the autoscaler pr (1866), might take a few days given the other work going on
 - [alberto] aware that there are some concerns about the design and details, main goal now is to get a first version merged then we can start to improve what is there
 - [jason] this sounds good, we need a place to start and then start making the necessary changes
 - we have strength in identifying problems, we need to follow up with creating solutions to those problems

- [michael] getting a piece in place allows us to start getting feedback on how it is used and where the problems are.
- [alberto] other actions items
 - currently a few maintainers are being added to the autoscaler project, does anyone else want to be added to that list?
 - once the PR is merged, would be good to do a demo to the CAPI meeting
 - +1s from Jason DeTiberus, Michael McCune and Naadir Jeewa
 - Should we have regular Cluster API / Autoscaler sync meetings, or use the one of the existing autoscaler or cluster api meetings?
 - McCune - Preference would be not to have a separate meeting. Readout to Cluster API meeting as needed would be fine, and can join the autoscaler meeting when needed.
- [mccune] the key question, maybe not to address now is the atomicity question wrt to deletes, but don't know how the autoscaler folk will want to address.
 - [alberto] there have been concerns about it in the past, and not sure what the path forward it. How will the pieces fit in the future (machinesets, machinepools etc...)
 - [jason] even if we did have the ability to specify which machines to scale down in an atomic manner, there will be a significant difference in behaviour between a CAPI managed machineset vs one that isn't, leading to user confusion. If we can specify on the autoscaler side that X implementation is autoscaler compliant such that the autoscaler can trust the behaviour of the provider. A path forward would be some sort of library that could be consumed by both autoscaler and CAPI controllers. We need to then figure out how machinepools fit into this, but possibly not the biggest issue right now, compared to "scale from 0", which we should address first, and then figure out how machinepools fit. Machinepools is being bumped into the experimental group for v1alpha3.
- [gugino] a lot of the provider information (e.g. AWS) is baked into the autoscaler at compile time.
 - [jason] worry about just maintaining tables in code, and having to do releases just for updating tables would not be ideal, and supporting things like GovCloud. If we can do something at run-time, that'd be preferable.
 - [gugino] Some issues around run-time consumption, as AWS billing data, e.g. is IAM restricted and only available on 2 regional endpoints. Could do it via some aspect of configuration, e.g. annotations.
 - [jason] Worry here is that although from a product standpoint, this is fine because it can be baked into a product, from upstream, this creates a footgun for users. Also worry about the use of annotations vs. a part of a spec because it becomes an informal API (i.e. Ingress)
 - [gugino] Ideally 90% of users wouldn't need to touch this. If it's in a spec, we potentially see multiple things having to modify the spec.

- [jason] we see this already in CAPI, e.g. ProviderID and highlights a weakness in the Kubernetes model around status/spec wrt to persistent fields that only controllers care about.
 - [mccune] we need the concrete problems in front of us to figure the way forward/
 - [jason] we are moving forward with API review, and we should perhaps raise the concerns that we have during this process with folk like Liggitt.
 - [mccune] we should bubble up this concern (persistable fields etc...) sooner rather than later as it impinges on the future direction.
 - [jason] +1 particularly as this impacts integration
- [alberto] wrt to reaching beta by end of the year, does this include the providers
 - [jason] primary concern is the core types, but don't have the confidence on the AWS provider, but sure we would find issues during API review. GCP on the other hand has been languishing. Hoping we don't have to go in lockstep with all the providers. We should have a bit of independence there.
- **ACTION ITEMS BELOW**
- get pr merged (see above)
 - [joel] does this contain bug fixes we have done downstream since it was initially added?
 - [alberto] yes, it has been rebased on the internal changes we made
- [gugino] we need to file an issue about what our scale from zero should look like, what fields need to be exposed, etc. **action item for Gugino**
- [jason] make sure we have issues for consistent scale down experience between autoscaler and manual use cases.
- [gugino] when we do docs it would be nice to link to specific autoscaler issues, to help give more details.
 - [jason] we might be able to embed some queries that will help guide users to find the properly labelled issues.
- [gugino] there are several doc areas (missed the specifics) that should be improved during this effort

Wed 26 February - 10am Pacific

[Recording](#)

Attending

- Justin SB - Google
- Michael elmiko McCune, Joel Speed, Michael Gugino, Eric Duen, Alberto - Red Hat
- Chirag Aggarwal - GoJek

- Jason DeTiberus, Vince Prignano, Tim St. Clair, Andy Goldstein, Yassine Tijani, Sedef Savas, Noa Amran, Nader Ziada - VMware
- Prakash Ramchandra - Dell
- Moshe Immerman - Flanksource
- Bryan Boreham, Leigh Capili - Weaveworks
- Cecile Robert-Michon - Microsoft
- Satish Ashok - siva
- Seth Pellegrino – New Relic
- Daniel Lipovetsky - D2IQ

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [vince] RC.2 dropping today ~2/3pm PST
 - [vince] Tentative release date set for v0.3.0 is March 10th
 - [vince] Reviewer list
 - Promote Cecile: <https://github.com/kubernetes-sigs/cluster-api/pull/2376>
 - [detiber] Meet & Greet Breakfast at KubeCon EU in Amsterdam
 - Sign up here: https://docs.google.com/forms/d/1bohKAqjh8LGD_DqebBqfS0B1sQIPc8MSNrbdNg3D3Ac/edit
 - [detiber] Cluster API t-shirts
 - <https://forms.gle/j2R9RyGWsMd2rfcTA>
 - [v1alpha2/v1alpha3 API review with Jordan Liggitt was yesterday](#)
- Demos/POCs
- Discussion topics
 - [joelspeed] Adding support for Spot Instances
 - <https://github.com/kubernetes-sigs/cluster-api/issues/1876>
 - should there be a google doc or straight to a pr/doc?
 - google doc is a good place to start, we are pretty chatty and the doc helps to organize thoughts before the caep
 - what is the process, or who should review, before moving from doc to caep?
 - anyone can review, project maintainers should have the say when things move to caep
 - [vince] are we looking to use spot instance for machine pool? (missed the question)
 - initially managed by machinesets, but later it could be migrated to machinepools
 - it would start in the infrastructure template for the machine
 - there is an implementation for machine pool in azure [vince]

- proposal process is lengthy and reusing the machine pool work might make this easier to ingest initially
 - [joel] is there any research that should be looked at specifically for this effort?
- [Alberto] revive cluster autoscaler effort.
 - Integration via <https://github.com/kubernetes-sigs/cluster-api/issues/2369>
 - Merge <https://github.com/kubernetes/autoscaler/pull/1866>
 - And add >= v1alpha2 support <https://github.com/openshift/kubernetes-autoscaler/pull/118>
 - want to review effort and figure out next steps
 - we (red hat) in a good position to drive this work and maintain going forward
 - does machineAutoscaler CRD make sense? this could help drive the autoscaler effort forward.
 - the pr against k/autoscaler is in the hands of sig-autoscaler, we might need to work closer with them
 - do we need a mini-wg to get this done? [andy]
 - [jason] post v1alpha1 work, does that include support for machine pools as well?
 - it doesn't currently, but it could have an interface for a scalable resource that could plug into machine pools (or other abstractions in the future)
 - [jason] ++ to mini-wg for plotting out next steps, in the past we've had concerns about how the autoscaler interacts with the cluster-api. there are some existing issues around this integration that need to be worked out.
 - [andy] suggest that people who are interested should self-organize and figure out a time to sync up.
 - Interested in mini-wg for autoscaler integration: jdetiber, michaelgugino, elmiko, joelspeed, evalsocket, Alberto
 - [meeting message and invite](#)
- [vince] Feature gates and future (post v1alpha3) work!
 - FG PR: <https://github.com/kubernetes-sigs/cluster-api/pull/2376>
 - important for future work, it will add ability to test PoCs and other feature type experiments
 - works like it does for kubernetes
 - there will be a process for submitting code to experimental folder
 - if you are looking to do new features, this is a good venue for it. we want to put all our experiments here to help promote and vet them.
 - reach out to Vince for more info
 - [jason] gates have solved some issues, but not all issues around types and deployments of experimental crds. this will need some attention.
- [leigh] Cluster Addon work -- nice overlap /w [PostApply proposal by Sedef Savas](#)

- CAPI proposal
[https://groups.google.com/forum/?utm_medium=email&utm_source=foote
r#!msg/kubernetes-sig-cluster-lifecycle/hXu3mXOx5wQ/dU5NXfKcAAAJ](https://groups.google.com/forum/?utm_medium=email&utm_source=footer#!msg/kubernetes-sig-cluster-lifecycle/hXu3mXOx5wQ/dU5NXfKcAAAJ)
- Addons API and example:
<https://github.com/kubernetes-sigs/cluster-addons/tree/master/installer>
Proposes AddonInstallerConfiguration
- Kubeadm KEP (WIP POC in open PR)
<https://github.com/kubernetes/enhancements/pull/1308>
- [leigh] want to start an open conversation about composing and sharing around this
- on cluster addon side there is an installer for these, see links above ^^
- would be nice to start using this api to bundle things across the ecosystem
 - currently needs more use and testing, looking for feedback and issues
- [andy] we don't want to duplicate effort in the addons group
- [tim] good first use/example using cluster apply, eg applying policies for newly created clusters. could use this mechanism as an example of how to deploy these (this == sedef's proposal).
- [leigh] no comments on sedef's proposal, but like the idea, it composes well with similar effort they have considered.
- [prakash] missed a point in earlier topics about bare metal
 - would like to have a cluster type for validation in bare metal situations
 - is there a cluster type called "bare metal"
 - [andy] perhaps in metal kube or an infra provider doing bare metal machines
 - when we create a cluster type, don't we need to source the bare metal machines? and does that fit in this context?
 - [michael] do everything in your power to not call things "bare metal", just name it something else
 - what do we document versus what do we deploy? in document we keep some type called "bare metal" is there a problem with that?
 - [andy] what does "cluster type" mean in this context?
 - think control plane cluster versus a deployed worker(?) cluster
 - there is no distinction like this currently, but there are fields for infrastructure references which could refer to specific provider information [andy]
 - cluster-api just has a reference to the cluster type as defined by the provider, but this doesn't have an impact on the data model internally.
 - infrastructure cluster might be marked differently by the provider, as opposed to a deployed cluster

- [michael] please stop using “bare metal”, there is some provisioned service. it may provide a certain type of machine, but there is a service behind the provisioning process. we should avoid calling specific implementations the “bare metal provider”, metalkube is a provider, so are other types that might provide similar hardware.
- v1alpha3 burndown / triage
 - 3 issues (see issue comments)
 - #2429
 - #2447
 - #2452

Tue 25 February - v1alpha2/v1alpha3 API review session

Attending

- Daniel Lipovetsky - D2IQ
- Jason Scarano, Naadir Jeewa, Jason DeTiberus, Warren Fernandes, Noa Amran, Andy Goldstein, Nader Ziada - VMware
- Michael elmiko McCune, Michael Gugino - Red Hat
- Jordan Liggitt - Google

Notes

[Recording](#)

API Types

Cluster Types

- Added some fields (Paused, ControlPlaneRef, ControlPlaneEndpoint)
- [liggitt] What's the version skew policy?
 - [ncdc] Controllers will use the latest version. No codified policy for “external” clients.
- [ncdc] In v1alpha2, we ran one controller per namespace. In v1alpha3, we rely on webhooks for conversion, and there can only be a single webhook for the cluster.
 - [ncdc] Only one API version can run within a namespace

- [ncdc] We have users on v1alpha2. We'd like to find major show stoppers for upgrading to v1alpha3.
 - For example, storage version already changed in cluster-api master branch. There's a recommendation to roll out the storage version over two releases.
 - [liggitt] That's a recommendation specifically for "native" types. Doesn't apply to CRDs.
 - [ncdc] Could be useful for rollback.
 - [liggitt] All existing users/controllers are v1alpha2. Once v1alpha3 is released, then controllers will use that, and external clients are expected to upgrade soon.
 - [liggitt] Lossy conversion is a potential problem to look for.
- Changed the APIEndpoints slice to ControlPlaneEndpoint scalar.
 - When we do a v1alpha2 status update, check to see if it actually updates spec.controlPlaneEndpoint in etcd (or is it converted on the way out), and if it increments generation
- When moving from controller-based validation to webhook validation, updates could now be blocked on validation failures
- We aren't currently fuzzing/testing 2 to 3 to 2
- Things to check in conversion (v1a2 to v1a3)
 - V1a2 object w/bogus json
 - V1a2 object w/ok json but it doesn't unmarshal into the expected type
 - V1a2 object unmarshals but fails openapi validation
 - V1a2 object unmarshals but fails webhook validation (namespace comparison, etc)
 - Want to avoid bad data getting persisted (e.g. it passes validation pre-conversion)
- Defaulting webhook should check for ClusterName and the label
- ClusterName should be immutable
- Use the `Type` field on the bootstrap Secret
 - Update controller to make sure it's operating on secrets of this type
 - Update bootstrap provider implementer's doc
- Consider excluding metadata from json marshalled data for conversions
- Are webhooks intercepting all versions?
 - no, just v1alpha3
 - consider using matchPolicy to intercept all versions
<https://kubernetes.io/docs/reference/access-authn-authz/extensible-admission-controllers/#matching-requests-matchpolicy>
- How is namespace populated in calls to webhook Default()? From admission attributes or object metadata content
 - <https://github.com/kubernetes/kubernetes/issues/88282>
 - Test submitting an object with no namespace set in the manifest (don't use kubectl b/c kubectl will fill in the namespace before sending it to the apiserver)
 - AI: Check what controller-runtime does when a namespace is in a webhook request attributes, but not in the object

- Examine machine validation re bootstrap configref & datasecretname
- File issue to discuss switching from CRD labels to a separate typed resource (or configmap)
-

A look at particular changes

Wed 19 February - 10am Pacific

[Recording](#)

Attending

- Seth Pellegrino, Andrew Rudoj, Jeremy Voorhis, Sean Kane, John Goldsmith, Karthik Ganguru - New Relic
- Michael elmiko McCune, Joel Speed, Danil Grigorev - Red Hat
- Justin SB - Google
- Jun Zhou - Spectro Cloud
- Jason DeTiberus, Andy Goldstein, Yassine Tijani, Jason Scarano, Sedef Savas, Warren Fernandes, Nader Ziada, Fabrizio Pandini, Deepika Pandhi, Naadir Jeewa, Chuck Ha, Noa Amran - VMware
- Daniel Lipovetsky - D2IQ
- David Watson
- Cecile Robert-Michon, David Justice - Microsoft
- Satish Ashok - Diamanti
- Eric Staples - NetApp

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - v0.3.0-rc.0 with v1alpha3. Another RC will be cut today, with new RCs every Wednesday unless the unlikely event that there's no code and looking for first/second week of march for GA
 - CRDv1 is now required to consume v0.3.0 ([CAPI PR](#))
 - [Kustomize config changes for multi-tenancy support when using webhooks](#)
 - [CABPK v0.1.6](#)
 - If you have a v1alpha2 cluster and you want to upgrade to v1alpha3, all CRDs need "preserveUnknownFields:false". CABPK has a [new release](#) to do this. Infrastructure providers also require this.
 - MachinePool PR merged!

- Demos/POCs
- Discussion topics
 - [rudoij] KubeadmControlPlane KubeadmConfigSpec mutability
 - https://github.com/kubernetes-sigs/cluster-api/pull/2334#discussion_r379606144
 - What is an upgrade really?
 - Change infraRef
 - Change kubernetesVersion
 - Change joinConfiguration?
 - Believe machine replacement is the way to make these changes rather than modifying a live machine
 - [detiber] we originally wanted to keep it out of scope to avoid shooting ourselves in the foot, that said, in the process of implementation, a lot of healthchecks and state checks have been added, which makes it pretty safe to remove the immutability of those parts as we should be able to block.
 - [vince] Can this be done in a .1 release rather than today?
 - [rudoij] sure, it's an otherwise small change. What drives not doing it now?
 - [vince] It's a design change, and would like to spend more time thinking through the implications.
 - [jason] it's a hashing change and that of the validation webhook, but it is trivial in the scope of the larger implementation.
 - [vince] could be placed behind a feature flag
 - [ncdc] either method would be fine
 - [ncdc] v1alpha3 API review meeting update
 - Doodle poll: <https://doodle.com/poll/zbp84fctxwwannu6>
 - Tuesday 25th at 14:00-15:00 EST is the current winner
 - Calendar invite will be sent by EOD today
 - [vince] Code freezing policies
 - <https://github.com/kubernetes-sigs/cluster-api/issues/2372>
 - Would like a timeframe for when feature requests come in

Wed 12 February - 10am Pacific

[Recording](#)

Attending

- Michael elmiko McCune, Michael Gugino, Eric Duen - Red Hat
- Eric Staples - NetApp
- Jason DeTiberus, Andy Goldstein, Naadir Jeewa, Yassine Tijani, Noa Amran, Jason Scarano, Fabrizio Pandini, Warren Fernandes, Chuck Ha, Sedef Savas - VMware
- Justin SB - Google
- Craig Peters, David Justice, Ace Eldeib - Microsoft
- Seth Pellegrino, Andrew Rudoj – New Relic
- Jun Zhou - Spectro Cloud
- Erwin van Eyk - Platform9
- David Watson
- Satish Ashok - Diamanti
- Moshe Immerman - Flanksource
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ

Agenda

- Welcome/Introduction for new attendees
- Demos/POCs
 - [\[Sedef\] Apply resources to cluster post-creation](#)
 - Intended to make it easier to consistently apply the same addons
 - Adds `postApplyAddons` array field to the cluster spec, referencing secrets containing the addon
 - Applies to the workload cluster using the generated secret for the kubeconfig
 - [justinsb] reminds me of the work with the addon operator, though this is for applying to another cluster. Should the addons be in line in the cluster or should they be their own object with their own statuses. There's been a suggestion from openshift folk about distribution using container images vs. git or https. There is however a huge difference with the addon operator work that this is cross-cluster.
 - [andy g] given there is a working group for addons, we should work with that group to ensure there is convergence
 - [justinsb] yes, some of the same questions have arised, such as "how often to apply the yaml"
 - [sedef] what would need to occur during delete?
 - [justinsb] work through references and run delete on each one.
 - [moshe] Feel that his should be external to the cluster controller
 - [david watson] +1, a similar approach was taken at samsung.
 - [daniel] uncomfortable with seeing the field inside the cluster type.

- [seth] new relic has found it helpful to separate out addons required for the functioning of the cluster (i.e. CNI) from others, and would be a replacement for the addons workflow.
- [andy] for the machine load balancer proposal, we've suggested doing it in the contrib directory for experimentation, how would people feel having this as a separate controller in an experimental dir.
 - 6 +1s received
- [justin] sometimes we say it shouldn't be in core is that we're not sure where it should be, whereas with this we seem more agreed that the functionality is important, but we're less sure about the schema & object structure
- [vince] would like to see this behind a feature flag. Just to note we do have a federation mess at present, so keeping things together would be good.
- Discussion topics
 - [Andy] v1alpha2 to v1alpha3 API review with Jordan Liggitt
 - Proposed dates: Feb 25, 26, or 28, sometime between 8am-6pm EST
 - Looking to avoid gotchas that will break users
 - **Action item: Andy to send out Doodle for times**
 - [elmiko] cluster autoscaler / cluster api integration
 - reference: <https://github.com/kubernetes/enhancements/issues/609>
 - i'm curious about the status and any thoughts people might have
 - [jason] we've gone through this a few times, and keep coming back to, is that although we could get it working against the autoscaler repo, the integration would be different in which they consider scaling logic atomic, they also inspect instances and pick one for deletion. In cluster api, we deal with replicas. The current implementation idea is for autoscaler to mark a machine for deletion, however in an ideal scenario we need to define a better way to interoperation. The logic for picking machines for deletion could be exposed as a library consumable by both, however this affects the design of the autoscaler, and have not had the hard discussions yet.
 - [michael gugino] Me and Andrew McDermott agreed with autoscaler folk that a patch could be merged with the proviso that behavioural differences could be documented. Consensus was to get it merged against v1alpha1, and now we need to figure out priorities.
 - [ncdc] is OpenShift PM interested in getting this merged?
 - [elmiko] Definitely interest in autoscaling, and we'd like to align more closely with upstream.
 - [gugino] we have something working with v1alpha1, but no PM push to specifically merge it upstream

- [ncdc] we are wrapping up v1alpha3 in the next few weeks. Let's try to start planning for v1alpha4. I will lean on Vince, Jason, elmiko and Michael Gugino to work towards a plan.
- Action item: Verify that there is an issue to figure out the autoscaler plan and rally there.
- [vince] Controller runtime release happening today / CRDv1 by EOW in CAPI
 - This will be a breaking change. Documentation will be added to the [v1alpha2 → v1alpha3 section of the CAPI book](#).
 - Management cluster will need to be running on v1.16
- [vince] Cluster API multi-tenancy and webhooks
 - Issue: <https://github.com/kubernetes-sigs/cluster-api/issues/2275>
 - TL;DR: Conversion webhooks can only run in a single place, they're configured in the CRD spec.
 - Causes problems such as "if I deploy a new instance, which one is canonical for the webhooks?"
 - For CAPA, we don't support multi-tenancy within a single deployment
 - [andy] our current release includes a single manager per provider and a single ball of yaml with the deployment, CRDs, RBAC etc... however we can't continue to treat the CRDs and conversion webhooks as a single group because they're cluster-wide singletons, and the most recent deployment will win. We're at the mercy of API machinery in the sense there isn't true CRD multi-tenancy.
 - Questions:
 - [vince] Does anyone run multiple instances of a provider [in the same management cluster]
 - [andy] Have some folks running a management cluster for lots of AWS accounts, and they may want to test a newer version in one of those namespaces, and obviously the CRD changes will apply to all of them and the webhooks change. If there's a bug fix or a new field, then that's one use case, the other is that CAPA uses a single set of credentials per instance.
 - [jun zhou] Single management cluster with multiple copies of CAPA and CAPV (1 namespace per copy per credentials)
 - [Seth] we do run 60 instances, but run self-managed clusters.
 - [Satish] Another option is for the provider to work with multiple credentials
 - [vince] big change and out of scope for initial v1alpha3

- [\[naadir\] CAPA issue for interested peeps](#)
- [jason] we need to keep in mind that some of the use cases are related to managed service providers which mean we need to ensure security wrt access of credentials.
- [justin] Could get stricter about the API vs the manager. We could separate out the webhook controller, and say you must run the highest version of the webhook regardless of the controller revision.
 - [david justice] +1 to breaking out the webhook
 - [ace eldelb] +1 for webhooks, and figure out multitenancy later
- [ncdc] if we had non-breaking API changes and everything is additive, it'd be fine, but the problem arose because of the conversion webhooks
 - [ncdc] Do we need to support multi-tenancy in CAPI v1alpha3?
- Possible solutions:
 - Webhooks in a separate namespace:
 - <https://github.com/kubernetes-sigs/cluster-api/pull/2279>
 - Requires changes to *every* provider config/ directory.
 - Changes are quite invasive and depart from kubebuilder model (more maintenance down the line to keep them up-to-date).
- v1alpha3 burndown
 - KubeadmControlPlane
 - Most disruptive prereqs have merged (etcd health check, refactoring reconciliation loop)
 - Ongoing work for scale up & down
 - Rebasing upgrade branch on top of latest changes
 - Working on control plane health checks
 - clusterctl v2
 - Small changes required for webhook/tenancy issue
 - On track for release
 - Working on e2e tests
 - MachineHealthCheck
 - No updates, but there is a [PR that needs reviews](#)
 - MachinePool
 - Working on addressing code review comments (mostly small items)
 - [PR for review](#)

Wed 5 February - 10am Pacific

[Recording](#)

Attending

- Michael elmiko McCune, Michael Gugino - Red Hat
- Andrew Rudoi - New Relic
- Jason DeTiberus, Tim St. Clair, Chuck Ha, John Harris, Jason Scarano, Andy Goldstein, Warren Fernandes, Noa Amran, Nader Ziada - VMware
- Hardik Dodiya - SAP
- Dax McDonald - Rancher
- Joe Rocklin - Siemens Digital Industries Software
- Cecile Robert-Michon - Microsoft
- Prakash Ramchandran - Dell
- Daniel Lipovetsky - D2IQ

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - [detiber] Meet & Greet Breakfast at KubeCon EU in Amsterdam
 - Sign up here:
https://docs.google.com/forms/d/1bohKAgjh8LGD_DqebBqfS0B1sQIPc8MSNrbdNg3D3Ac/edit
 - [vince] Conversions and some gotchas
 - <https://book.kubebuilder.io/multiversion-tutorial/conversion-concepts.html>
 - Want to ensure a good experience for users upgrading from v1alpha2
 - We originally started with the autogenerated conversions, however:
 - If you add a new field, and you convert from v1alpha3 to v1alpha2 and back again, you lose data. There is a utilconversion package in CAPI that you can use to `storhttps://github.com/kubernetes-sigs/cluster-api/blob/master/docs/book/src/roadmap.mde data in an annotation and allow it to be restored.`
 - However, where there is a change of data, you need to add manual conversions, and you should add conversion test cases that cover these - this should test conversion from v1a3 → v1a2 → v1a3.
 - Have added a fuzzy testing function in the utilconversion package that has successfully found issues in our conversions.

- [Elmiko] How does the fuzzing work?
 - testing of many input fields, those that exist and those that don't exist. (1000s of combinations)
- Highly recommend that you use it. We are currently applying it to CAPA.
- [ncdc] We originally went down the wrong path of returning an error if we couldn't convert a field, and in talking to SIG API Machinery that you can never return an error if conversion has an issue. Returning an error is the equivalent of panicking in terms of severity, hence adding the round-trip conversions.
- [jason] does it make sense to round-trip from v1a2 → v1a3 → v1a2, in that they submit a v1a2 resource to a cluster with v1a3 storage formats and want to convert back?
- [tim] traditionally if you store as a new version, there aren't actual guarantees of data fidelity going to an older version.
 - [vince] you can actually request the older version, and as the CRD versions are registered, they will go through the conversions.
 - [ncdc] will have to see how it goes
- Demos/POCs
 - [chuck] A brief overview of running the end-to-end tests locally for fast signal on complex changes
 - [elmiko] how modest a machine do we need to run this?
 - 6GB ram, 4cpu runs in < 5 minutes
 - [rudoj] are there e2e tests you can run on a laptop for CAPA
 - [chuck] There are existing ones that are run in Prow. Would like to convert them over to the e2e framework
 - [detiber] There is work in progress to port more tests to the master branch. Prow uses a Janitor process to clean up AWS resources.
 - [cecile] CAPZ also has e2e tests
- Discussion topics
 - [vince] Cutting a beta tag this week without a release?
 - No dissent noted
 - [fabrizio] Can we also include CAPD yaml & image?
 - [vince] Not planning on doing a full release w/assets, but will chat offline
 - [vince] CRD contracts and conversions
 - [rudoj] KubeadmControlPlane Machine adoption
 - [https://github.com/kubernetes-sigs/c\[vince\] Conversions and some gotchascluster-api/issues/2214](https://github.com/kubernetes-sigs/c[vince] Conversions and some gotchascluster-api/issues/2214)
 - Clarity on v1alpha2 -> v1alpha3 migration path
 - [rudoj] UX for MachineDeployment updates (e.g. upgrade Kubernetes version)

- Being immutable is a confusing UX for users. Do we really get a benefit from it
- [ncdc] Possibly worth reviewing, but for post v1alpha3
- [detiber] One of the main issues is that we don't have a mechanism for tracking updates, or enabling reverts
- [rudoj] Roadmap for defaulting based on infra-bootstrap pairings
 - Was discussed at the f2f in September. In the decoupling effort of v1alpha2, we lost a bunch of UX, e.g. I have to manually set the cloud provider in the kubeadmconfigs, because they are infrastructure agnostic. We are going to approach this with some webhook logic, that we may be able to upstream at some point.
 - [vince] we have an issue that we had opened for using configmaps for defaulting. Believe it was put on the roadmap.
- [elmiko] Project culture
 - There's a lot of good for newcomer issues, and others with comments and assignments with no work.
 - [tim] [community/grooming.md at master · kubernetes/community](#)
 - [vince] reach out if necessary
- v1alpha3 status check

Wed 29 January - 10am Pacific

[Recording](#)

Attending

- Tim St. Clair, Warren Fernandes, Yassine Tijani, Jason DeTiberus , Nader Ziada, Naadir Jeewa, Chuck Ha , Noa Amran, Jason Scarano- VMware
- Joel Speed - Red Hat
- Marcel Müller - Giant Swarm
- Jun Zhou - SpectroCloud
- Ace Eldeib, Ria Bhatia - Microsoft
- Dax McDonald - Rancher
- Hardik Dodiya - SAP
- Vishal Anarse - Diamanti
- Niketu Parekh, Sudeep Batra - Ericsson / At&t
- Joe Rocklin - Siemens Digital Industries Software
- Daniel Lipovetsky - D2IQ

Agenda

- Welcome/Introduction for new attendees

- PSAs
 - [ncdc] [Draft roadmap](#)
 - Nothing set in stone, and should be constantly revised
 - Please provide feedback
 - Mainly focus on what should be upcoming for v1alpha4
 - [vince] Has been open for 5 days, stay open for 2 more, and maybe merge in
- Demos/POCs
- Discussion topics
 - [marcelmue] labels on metrics services need to be aligned. [#2070](#) got merged but there is no clarity which labels should be on the metrics services / which labelselector should be used. Related PRs: [#2129](#) [2128](#)
 - AI: File issue describing agreed solution
 - [tim] A lot of times we go through what feels like a “pause” cycle after a release, please start drafting your CAEPs for v1alpha4.
 - [ncdc] Some are linked in the roadmap
 - [niketu] what’s the plan for introducing HA for all Cluster-API components? Assume it will be straight fwd to change replicas?
- v1alpha3 issue/PR burndown
- Backlog grooming

Wed 22 January - 10am Pacific

[Recording](#)

Attending

- Tim St. Clair, Warren Fernandes, Nader Ziada, Yassine Tijani, Andy Goldstein, Chuck Ha, Noa Amran, Himanshu Pandey, Jason DeTiberus - VMware
- Sudeep Batra - Ericsson
- Michael Gugino - Red Hat
- Andrew Rudoj - New Relic
- Mark Emeis - Weaveworks
- Jay Pipes - AWS
- Cecile Robert-Michon, Ria Bhatia, David Justice, Cecile Robert-Michon- Microsoft
- David Watson
- Alan Meadows - AT&T
- Daniel Lipovetsky - D2IQ
- Moshe Immerman - Flanksource
- Satish Ashok - Diamanti
- Maël Valais - Ori Industries

Agenda

- Welcome/Introduction for new attendees
- PSAs
- Demos/POCs
- Discussion topics
 - [dlipovetsky] Control plane replicas that are machines can also run non-control plane workloads. How might this affect (a) control plane workloads (b) control plane management implementation? Should this be explicitly unsupported? See issue [#2064](#)
 - [sudeepbatra] Understanding of the E2E Flow or sequence of actions involved in CAPI v1alpha2 compared to v1alpha3 using the controllers (capi-controller, capbpk-controller(kubeadm),infrastructure-provider-controller). Something like the cluster provisioning process in [Controller collaboration](#) but in more detail.
 - [ncdc]
 - We should try to consolidate some of the documentation
 - Some of the docs may also be out of date.
 - Most recent development docs for master are at <https://master.cluster-api.sigs.k8s.io>
- Backlog grooming

Wed 15 January - 10am Pacific

[Recording](#)

Attending

- Justin SB - Google
- Cecile Robert-Michon, Ria Bhatia, Ace Eldeib - Microsoft
- Yassine Tijani, Nader Ziada, Andy Goldstein, Vince Prignano, Naadir Jeewa, Jason Scarano, Andrew Kutz - VMware
- Joel Speed, Eric Duen, Michael “elmiko” McCune, Michael Gugino - Red Hat
- Daniel Lipovetsky - D2IQ
- Cornelius Keller, Marcel Müller - Giant Swarm
- Sudeep Batra - Ericsson
- Moshe Immerman - Flanksource
- Hardik Dodiya - SAP
- Erwin van Eyk - Platform9

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - New releases:
 - CAPI: [v0.2.9](#)
 - CAPA: [v0.4.8](#)
 - NPE fixes, set tags on ebs volumes
 - RC.1 tentative release date, Feb 14th?
 - Will be a stable tag that providers can use to update to v1alpha3 if they are not doing so on master
- Demos/POCs
 - None
- Discussion topics
 - [ace] machinepools update
 - CAPI controller is up for review, currently adding tests. Next step will be the CAPZ side implementation.
 - [Andrew Kutz] Is there guidance to update CAPD tests to account for CAPI changes. Should it be done, if it's possible?
 - [ncdc] For MachinePools, can't see it a way to test it for Docker. CAPD is to ensure code is compiling. Behavioural testing is important, but not sure CAPD can cover everything, where MachinePools relies on cloud provider primitives.
 - [Lubomir] May be able to fabricate something that pretends to be a cloud provider implementation that uses Docker.
 - [Ria Bhatia] Documentation structure
 - Have updated Azure book, and also make it fit with the CAPI Book.
 - Is there a standard for what goes in the CAPI book or the provider
 - Where will users go to, there's a lot of jumping to start off with.
 - [ncdc] Don't have a clear answer. There are other providers not sponsored by SIG CL, so expect for them there will always be some jumping back and forth, but for achitecture and structure, would like to see someone who has a background in technical documentation to help.
 - AI: Ria to create an issue.
 - [Vince] When we created the book, the idea was to do a one pager. There's external links in the docs that serve up docs from the provider repository.
 - AI: Andrew to link up resource for writing docs
 - [cornelius] is KubadmControlPlane ready to serve as a template for other Control plane Providers?

- Description of behavior of a generic control plane provider:
 - <https://master.cluster-api.sigs.k8s.io/architecture/controllers/control-plane.html>
 - [moshe] ControlPlane Health
 - <https://github.com/kubernetes-sigs/cluster-api/blob/master/docs/proposals/20191017-kubeadm-based-control-plane.md#control-plane-healthcheck>
 - There are a list of healthchecks that may be performed. The hard part is the healthchecks
 - Using etcdadm?
 - [Naadir] etcdadm manages etcd differently to kubeadm. Don't want to cross streams at present
 - Use of kubeadm upgrade
 - [Andy] No, we will reimplement
 - [Naadir] kubeadm currently lives in k/k, not easily consumable as library work. There is work in progress in the 1.18/1.19 stream to move kubeadm out of tree, and then we'll sync there.
- Backlog grooming
- [Andrew Kutz]
 - Old issue, want to check status of work on switching Events to Conditions
 - <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere/issues/723>
 - [ncdc] CAPA has a bunch of Events that are likely to be good to convert to Conditions.

Wed 8 January - 10am Pacific

[[this may become a new document]]

[Recording](#)

Attending

- VMware - Andrew Kutz, Warren Fernandes, Tim St. Clair, Naadir Jeewa , Nader Ziada, Noa Amran, Vince Prignano, Chris Saunders, Yassine Tijani, Andy Goldstein, Jason Scarano, Fabrizio Pandini
- Red Hat - Michael Gugino, Michael McCune "elmiko",
- Tilt - Maia McCormick, Dan Miller, Nick Santos
- SAP - Hardik Dodiya, Prashanth
- Flanksource - Moshe Immerman
- D2IQ - Daniel Lipovetsky
-

Agenda

- Welcome/Introduction for new attendees
- PSAs
 - CAPV is also using the CAPI e2e test framework to do all of the pre-submit e2e testing on CAPV master (v1a3). Many thanks to Chuck for all his help!
- Demos/POCs
 - [akutz] A demo of an in-tree, rough implementation of the Machine Load balancer CAEP with CAPV and HAProxy
- Discussion topics
 - [Maia] Tilt Cluster API Performance: pre-loading images
 - Understand that CAPI is using Tilt
 - [ncdc] Chuck Ha did the initial work to get Tilt to work with CAPI, and then added a Tiltfile to the CAPI repo that would work with the federation of CAPI
 - [Maia] The Tiltfile seems kind focused, is there interest in using other clusters
 - [Naadir] Use a long-lived home lab cluster, so have made the kind preloading conditional, and have hadn't had issues since
 - [Chuck] What are your recommendations?
 - [Maia] Using a GKE cluster can prevent laptops dying, initial image seeding may take some time, but live update should be good
 - [Maia] What has the live update experience been like?
 - [ncdc] After some changes, it's about 10 seconds, and that's great
 - [chuck] We do something slightly differently from the documentation, we use a local resource to build the binary locally and push it into the image instead of the source code diff.
 - [akutz] At VMware, we have found Tilt to be useful in forcing people to understand how to deploy all of the time. In addition, previously testing was reliant on Prow pre-submit, now we have a consistent way to run tests in a distributed fashion.
 - [Maia] Notice you are using the start.sh/restart.sh scripts on non-Docker (i.e. containerd). Any issues?
 - [Chuck] No issues
 - [akutz] Good that this aligns with what we expect for machine images
 - [Maia] We will be working on addressing speed ups for CAPI as a test case for other users.
 - [Akutz] Would be great to have integration with <https://asciinema.org/> , to record the output.
 - [Chuck] Could be an extension of the snapshot

- Backlog grooming

Meeting notes from 2019 can be found [here](#).