Kubernetes Multi-cluster SIG Meeting Notes/Agenda 2025

This doc: https://tinyurl.com/sig-multicluster-notes

Zoom: https://zoom.us./my/k8s.mc

Recordings: youtube

Meeting details: We meet every Tuesday at 9:30AM Pacific Time / 18:30 CET (<u>convert to nuyour timezone</u>) and bi-weekly Thursday at 6:30AM UTC (<u>convert to your timezone</u>). Join <u>kubernetes-sig-multicluster@googlegroups.com</u> to get a calendar invite.

Archived notes from previous quarters are <u>here</u>. **Federation WG** meeting notes <u>here</u>.

How to sign up

Please add your topic with your [handle] to the next meeting date. If there isn't an entry for the next date yet, it just means no one has volunteered, so go ahead and create a heading 2 weeks after the last date and add your topic there. All are welcome to lead a discussion or present.

If we have too many topics for a week, or if presenters are unable to attend we will shift any spillover to the following meeting.

Backlog - date tbd

- Every call: scrub
 - https://github.com/pulls?q=is%3Aopen+label%3Asig%2Fmulticluster+org%3A kubernetes-sigs+org%3Akubernetes+
 - https://github.com/kubernetes-sigs/about-api/issues
 - o Pull requests · kubernetes-sigs/about-api · GitHub
 - o Pull requests · kubernetes-sigs/cluster-inventory-api · GitHub
 - Pull requests · kubernetes-sigs/mcs-api · GitHub
 - o Pull requests · kubernetes-sigs/multicluster-runtime · GitHub
 - Issues · kubernetes-sigs/sig-multicluster-site · GitHub
 - Pull requests · kubernetes-sigs/sig-multicluster-site · GitHub
 - o Issues · kubernetes-sigs/work-api · GitHub
 - https://github.com/kubernetes-sigs/work-api/pulls

2025-11-11 [AMER/EMEA]

Cancelled (KubeCon)

2025-11-06 [APAC]

Topics

Reducing the call to 30 min and moving it back 30 min

2025-11-04 [AMER/EMEA]

Scrub

- KEP 1645: add traffic distribution and internal traffic policies fields: https://github.com/kubernetes/enhancements/pull/5588 /
 https://github.com/kubernetes-sigs/mcs-api/pull/131
 - Looking for a lgtm
- KEP-5313: Placement Decision API for multicluster scheduling
- KEP-5339: add additional cluster-specific auth info field to the cluster profile object
 - Affiliated cluster profile example API <u>https://github.com/kubernetes-sigs/cluster-inventory-api/pull/27/</u>

2025-10-28 [AMER/EMEA]

Scrub

https://github.com/kubernetes-sigs/cluster-inventory-api/pull/23

2025-10-23 [APAC]

Scrub

• KEP-5339: add additional cluster-specific auth info field to the cluster profile object

2025-10-21 [AMER]

Scrub

- KEP-5313: Placement Decision API for multicluster scheduling
 - Lengthy discussion about the genericity of the API. Users will still need Karmada, Kubefleet etc.; this isn't intended to replace those entirely. But any project wanting to integrate with ArgoCD, MultiKeue, and without its own scheduling component, can rely on a standard integration. This KEP doesn't allow interchanging schedulers. What this KEP enables is that ArgoCD, MultiKeue etc. don't need specific integration plugins to work with other projects.
 - In the KEP diagram, there are three operation arrows: "writes", "reads from", and "performs actions on". This KEP allows the "read from" operation to be generic; the other two still need workload specific implementations.
 - Schedulers support the standard manifests (StatefulSet etc.), the chairs' concern is with application-specific constructs. But the KEP isn't related to that: it only pertains to standard constructs; its goal is to allow a single implementation to be used in all rollout tools, instead of one "reads from" implementation per tool.
 - The placement request is just the key. The KEP only promises that a scheduler will produce a decision with a key, and the workload will also have that key. For rollout tools, it answers the question: "which decision should I apply to this workload I want to deploy?"
- crd: add label with CRD version (merged)
- KEP-5339: add additional cluster-specific auth info field to the cluster profile object
- rename the credential provider field
- KEP 1645: add traffic distribution and internal traffic policies fields / apis: conformance: add traffic distribution and internal traffic policies fields

2025-10-14 [AMER]

Scrub

- https://github.com/kubernetes/enhancements/pull/5314
- https://github.com/kubernetes-sigs/mcs-api/pull/116
- https://github.com/kubernetes/enhancements/pull/5559
- https://github.com/kubernetes-sigs/cluster-inventory-api/pull/23

 https://github.com/kubernetes/enhancements/pull/5588 / https://github.com/kubernetes-sigs/mcs-api/pull/131

Topics

- While discussing IP families reconciliation in Cilium we realized that it is a very similar problem as ports
 - Like the union of port could means that contacting one port might be reaching only some subset of the pods (only those that have this port)
 - o Does the behavior above make sense/should it be changed??
 - For instance we could keep the union but raise a conflict if it's not an exact match everywhere to help users rollout some version using an additional ports but still make it a conflict to signal that there's something to be fixed ~
 - Discussion about <u>ports</u> and union/intersection
 - union is fine with smart tracking implementations (like Cilium port behavior described as only routing to subset)
 - intersection is simpler with dumb implementations (but may drop/conflict if any mismatch/superset)
 - union with a dumb implementation not tracking eligible backends is bad
 UX risks attempting to route traffic to a backends which can't support
 the intended ipFamily/port
 - Probably possible to write conformance tests for this e.g. if there's a mismatch of exported ports/ipFamilies, ensure routing to that succeeds (reaches the subset and not a dead end)
 - Discussion about a status condition that can inform the <u>user</u> that there is a conflict for ports and then reference that in the ipFamilies section so implementors know that it is similar
 - Could be a ServiceExport status condition
 - We like this because it's already a thing and as far as I can tell the target audience (users) would look at this
 - Message should(?) say what behavior the implementation is choosing (if it's a smart tracking union or not, if its intersection or not in the case of ipFamilies)
 - Could be something else (?)

2025-10-07 [AMER]

Scrub

- https://github.com/kubernetes/enhancements/pull/5314
- https://github.com/kubernetes/enhancements/pull/5264 //
 https://github.com/kubernetes-sigs/mcs-api/pull/104
- https://github.com/kubernetes-sigs/mcs-api/pull/116

- https://github.com/kubernetes/enhancements/pull/5559
- https://github.com/kubernetes-sigs/cluster-inventory-api/pull/23

Topics

- [lauralorenz] Quick update on multi cluster monitoring and observability survey project (doc)
 - o Total Number of questions (Qualitative + Quantitative) = 21
 - Total Number of Respondents for survey = 150
 - Average number of responses for all questions = Total number of questions/
 Overall number of responses 73%
 - SIG Contribex supported us with a social media push to Aug 11 September
 22

2025-09-30 [AMER]

Scrub

- https://github.com/kubernetes/enhancements/pull/5314
- https://github.com/kubernetes/enhancements/pull/5264 / https://github.com/kubernetes-sigs/mcs-api/pull/104
- https://github.com/kubernetes-sigs/mcs-api/pull/116
- https://github.com/kubernetes/enhancements/pull/5559

Topics

- [ryan] https://github.com/kubernetes-sigs/cluster-inventory-api/pull/23
- [corentone] ClusterProfile credentials:
 - o extensions passed from CP to CredProvider
 - Naming
- [arthur] proposing to add trafficDistribution & internalTrafficPolicies support for ServiceImport before potentially promoting to v1beta1 and then introducing more complex multi cluster specific load balancing bits
 - https://github.com/kubernetes/enhancements/pull/5588 / https://github.com/kubernetes-sigs/mcs-api/pull/131

•

2025-09-25 [APAC/EMEA]

ClusterProfile credentials

2025-09-23 (cancelled)

2025-09-16

Scrub

- https://github.com/kubernetes/enhancements/pull/5314
- https://github.com/kubernetes/enhancements/pull/5255
- https://github.com/kubernetes-sigs/mcs-api/issues/106
- https://github.com/kubernetes/enhancements/pull/5264
- https://github.com/kubernetes-sigs/mcs-api/pull/116

Topics

- MCS cluster selection and traffic distribution
- [ryan] https://github.com/kubernetes-sigs/cluster-inventory-api/pull/23
- [ryan] PlacementDecision KEP

2025-09-11 [APAC/EMEA]

Topics

- [zhiying] https://github.com/kubernetes-sigs/mcs-api/issues/106 MCS API change
 - Should we make the breaking changes for the IP to get the API right before moving to v1beta1?
- [zhiying] Looking for feedbacks/comments about <u>MCS cluster selection and traffic distribution</u>, especially on the API design

2025-09-09 [AMER]

Scrub

• https://github.com/kubernetes-sigs/mcs-api/pull/124

- [mikeshng] <u>PlacementDecision KEP</u>
 - Updates relating to removing Placement ref and adding in decision-key label

- Discussion on how that affects decoupling
- Looking for an example flow through the diagram (including preferentially at least some of the assumptions about the arrow into the "Scheduler/PlacementController" side)
- Discussion of some use cases people have now where the placementdecision knows more about whether a cluster is draining or not to make better placement decisions
 - This is considered 'Consumer Feedback' in the KEP which is technically out of scope right now

0

2025-09-02 [AMER]

Scrub

Topics

- [lauralorenz] https://github.com/kubernetes/enhancements/pull/5314
 PlacementDecision API
 - o [mike] the KEP shouldn't refer specifically to Placement in the example
 - [laura] how do we implement the "simple" part ("basic workload distribution can be achieved ...")
 - [mike] the labels were key, "multicluster.x-k8s.io/placement" is what the consumer is supposed to look up
 - [skitt] it's difficult to decouple the scheduler and consumer if there is no defined, shared Kind that sits where the placementRef currently is (the consumer needs to know what the thing is that the scheduler wants it to schedule on the specified clusters)

2025-08-28 [APAC/EMEA]

Scrub

•

- [qiujian] Duplicate API definition in about-api: https://github.com/kubernetes-sigs/about-api/pull/31
 - [skitt] As suggested by quijian, we'll get PR 30 merged first, with some changes – in particular, keep the pkg directory for consistency with other projects (mcs-api, work-api), and the CRDs; then qiujian will rework 31 on top, restoring the generated client

2025-08-26 [AMER]

Scrub

- MCS-API related (mainly putting this here to advertise those for review, especially lead/approval ones, nothing changed since last time though)
 - ServiceExport conditions updates: https://github.com/kubernetes/enhancements/pull/5438
 - The corresponding PR on the MCS-API repo has been merged (thanks Stephen!)
 - Merged, thanks Stephen :D
 - Adding IPFamilies field to ServiceImport
 https://github.com/kubernetes/enhancements/pull/5264 /
 https://github.com/kubernetes-sigs/mcs-api/pull/104
 - + this one needs Mike Morris approval/(final?) review in addition to leads
- [arthur] Looking for feedbacks/comments about MCS cluster selection and traffic distribution in case you are not aware of this (not in a rush though, we have time to circle back on this)

Topics

- Adding CRD schema version to MCS-API CRD to facilitate a go program installing the CRD (and being able to check if we are not downgrading them there): https://github.com/kubernetes-sigs/mcs-api/pull/116
 - o Should this be part of this repo release/same version or be decorrelated?
 - Stephen prefers that this should be decorrelated/not part of the release "process"
 - [mike] this could be a bit strange as we'd potentially end up "jumping" several versions in each release
- [Jingming] Support additional endpoint query format endpoint.service.namespace.pod|svc.clusterid.zone: https://github.com/coredns/multicluster/pull/28
 - Next step is to create a PR to the KEP and see if implementors agree on adding this
- [skitt] SIG-MC involvement in the proposed Al Gateway WG (
 [PUBLIC] Kubernetes Al Gateway WG Proposal)
- [shane] (i'm going to be late, probably ~15-20m) follow up on the Al Gateway WG
 - https://github.com/kubernetes/community/pull/8521

2025-08-19 [AMER]

Scrub

•

- [arthur] Proposal / first designs about MCS cluster selection and traffic distribution
 - Discussed API proposals and alternatives (CEL?), connection to <u>well known</u> <u>property ladder PR</u>, connecting with SIG-Network

2025-08-14 [APAC]

Scrub

- [lauralorenz] https://github.com/kubernetes/enhancements/pull/5314
- [arthur] MCS-API related (mainly reminder to review those)
 - Adding IPFamilies field to ServiceImport
 https://github.com/kubernetes/enhancements/pull/5264 /
 https://github.com/kubernetes-sigs/mcs-api/pull/104
 - ServiceExport conditions updates:
 https://github.com/kubernetes-sigs/mcs-api/pull/112 /
 https://github.com/kubernetes/enhancements/pull/5438
- [zhiying] How to resolve conflict serviceExport: https://github.com/kubernetes/enhancements/pull/5463

2025-08-12 [AMER]

Scrub

- [lauralorenz] https://github.com/kubernetes/enhancements/pull/5314
- [arthur] MCS-API related (mainly reminder to review those)
 - Adding IPFamilies field to ServiceImport
 https://github.com/kubernetes/enhancements/pull/5264 /
 https://github.com/kubernetes-sigs/mcs-api/pull/104
 - ServiceExport conditions updates:
 https://github.com/kubernetes-sigs/mcs-api/pull/112 /
 https://github.com/kubernetes/enhancements/pull/5438
 - Now only with reason added, no type modified!
- [owen] https://github.com/kubernetes/enhancements/pull/5457
- [shane] WG Al Gateway Proposal
 - [PUBLIC] Kubernetes Al Gateway WG Proposal
 - Follow-up from last time:
 - what are the expectations of stakeholder SIGs?
 - Could be as low as being cc'd in proposal PRs
 - Alternatively, could expand the goals to include one or more multicluster goals and consider the stakeholder position to be one of more proposals and the group is open to this
 - Prominent use case is an egress use case I have an application that does inference but it's not in my cluster. Could be to a non-k8s cloud but could be a multicluster use case

- more clarity on definition of done?
 - Perception that WGs can run for a very long time and have a very big scope, but this group is trying to avoid that. They have a timebox of 1 year or less
 - Main goals are about several proposals shopped to other SIGs
- Next steps for SIG-MC: discuss over the next ~2 weeks including with leads about these expectations and who/how much peoples are interested in participating

2025-08-05

Scrub

- Adding CRD schema version to MCS-API CRD to facilitate go install (and being able to check if we are not downgrading them there):
 - https://github.com/kubernetes-sigs/mcs-api/pull/116
- One-line change to ClusterProfile KEP: <u>https://github.com/kubernetes/enhancements/pull/5457</u>

Topics

- Changing the ServiceImport API (merging to root vs spec/status): https://github.com/kubernetes-sigs/mcs-api/issues/106
 - Let's try to decide what we change here!
 - Cf the ServiceExport condition change we discussed two weeks ago + https://kubernetes.slack.com/archives/C09R1PJR3/p1753300742058739
- Adding IPFamilies field to ServiceImport https://github.com/kubernetes/enhancements/pull/5264 /

https://github.com/kubernetes-sigs/mcs-api/pull/104

- Submariner is now interested in addition to Cilium!
- Mike has reviewed and generally in favor, some nuance/details to work out.
- [Keith/Mike M] WG Al Gateway
 - Would SIG-Multicluster be interested in being a stakeholder SIG?
 - shane.utt@konghq.com what would expectations be from stakeholder SIGs? Meeting participation from leads, proposal review, something else? Definition of "done" for WG?
- [laura] add terminology doc?
 - https://github.com/kubernetes-sigs/sig-multicluster-site/pull/41
 - https://github.com/kubernetes/community/pull/8536
 - Context from https://github.com/kubernetes/community/pull/8210

О

2025-07-31 (APAC-friendly)

Scrub

Topics

- [mjudeikis] Kube-bind demo
 - Clarification that the goal is to host a KEP specifying the handshake and "service" provided, and kube-bind is an experimental implementation to explore what needs to be in the KEP
 - o Questions about multi-tenancy and the cluster-scoped nature of CRDs
 - The core of the demo is CRDs, but the KEP also needs to handle related objects (for example secrets)

2025-07-29

Scrub

- Adding ServiceImport conditions: https://github.com/kubernetes-api/pull/113
 / https://github.com/kubernetes/enhancements/pull/5439
 - Reminder to review/approve this
- Changing the ServiceImport API (merging to root vs spec/status): https://github.com/kubernetes-sigs/mcs-api/issues/106
 - Let's try to decide what we change here!
 - [mike] suggest deferring this to allow time for the Al Gateways topics
 - Cf the ServiceExport condition change we discussed last week + https://kubernetes.slack.com/archives/C09R1PJR3/p1753300742058739

- [robscott] Multi-Cluster + Gateway Inference Extension
 - [SIG-NETWORK] Multi-Cluster Inference Gateways
 - Want to avoid ClusterIP routing and DNS because routing is L7 with an extension callout (ClusterIP bypasses that)
 - [arthur] If using EndpointSlice, could integrate with MCS?
 - o Is there a version of MCS with a configurable backend?
 - Generalize ServiceExport/ServiceImport? Syncing?
 - Small amount of support for MC Gateway API currently
 - https://gateway-api.sigs.k8s.io/geps/gep-1748/
 - Shared ownership?

- LLMd in SIG Network
- Envoy has AI Gateway project with local routing and out-of-cluster routing (support MCS and GIE) — https://aigateway.envoyproxy.io/ (Tetrate-backed)
- Is the goal to have a common API or an implementation?
 - Common API but would be great to rely on common plumbing (for example Submariner); there are already implementations
- [skitt / shaneutt] ☐ [PUBLIC] Kubernetes Al Gateway WG Proposal and https://groups.google.com/a/kubernetes.io/g/dev/c/XC-8qAyk8W0 — possible combined use-cases for multi-cluster Al gateways
 - Wants to investigate issues closer to the gateway than the LLM, such as semantic routing
 - o Discussions in GIE
 - Enough scope for a WG? Many people interested
 - Open to having MC goals
 - [rob] MC implementation of GIE in-scope for GIE, what other MC scenarios are envisaged? InferencePool wouldn't be in scope for the WG
 - [laura] Rob motivated using MC in an inference environment (GPU scarcity),
 is that the main MC deployment motivation? [shane] yes
 - [shane] Another use-case is the egress use-case, not sure whether that has an MC story
 - [mikem] What egress? [shane] Envoy Gateway, egress option to send to third-party SaaS AI providers
 - [laura] Many MC deployments grew up that way, or for data locality or regulatory requirements, and this remains in scope for AI and many people in SIG-MC have relevant experience
 - [ryan] Production inference always ends up MC. What's the end goal of the WG? API, set of APIs, implementations? Also confused between AI GIE and WG.
 - [shane] Deliverables: definitions, creating a space to collaborate to see if there's strong consensus and producing proposals for other projects (including GIE). WG because larger scope, needed to pull in other people.
 - [rob] Fair representation, both Rob and Shane are maintainers of Gateway API; GIE wants to make it as easy as possible to route to LLMs inside clusters, whereas historically AI usage has involved external providers. Room for both GIE and the WG to co-exist. The WG is just for proposals; those will go elsewhere, the WG can't own APIs.
 - [shane] The WG has exit criteria. If SIG-MC wants to join in, the commitment is capped, hopefully around a year.
 - o [ryan] Where should I go? The WG? Gateway API?
 - [shane] Both, unless you only want InferencePool. Now looking for individual
 +1s, and if the SIG has related goals in mind, SIG sponsorship
- [skitt] CRD version labels (https://github.com/kubernetes-sigs/mcs-api/pull/116) is there anything similar in other projects?
- [embik] Cluster Inventory API support in multicluster-runtime has been merged: https://github.com/kubernetes-sigs/multicluster-runtime/pull/48

 [corentone] ClusterProfile Pluggable Credentials KEP has been merged! https://github.com/kubernetes/enhancements/pull/5338 thanks Stephen for the final approval!

2025-07-22

Scrub

- MCS-API related PRs/issues
 - ServiceExport conditions overhaul: <u>https://github.com/kubernetes-sigs/mcs-api/pull/112</u> / <u>https://github.com/kubernetes/enhancements/pull/5438</u>
 - Adding ServiceImport conditions:
 https://github.com/kubernetes-sigs/mcs-api/pull/113 /
 https://github.com/kubernetes/enhancements/pull/5439
 - Adding IPFamilies field to ServiceImport https://github.com/kubernetes/enhancements/pull/5264
 - Changing the ServiceImport API (merging to root vs spec/status): https://github.com/kubernetes-sigs/mcs-api/issues/106 (this is mainly collecting what implementation is doing)
- [ostrain] CP credentials KEP (github.com/kubernetes/enhancements/pull/5338)

•

Topics

2025-07-17 (APAC-friendly)

- [zhiying] whether to relax the conflict condition on mcs api.
 - conversation on PR: <u>Add conformance test to verify resolution of conflicting</u> <u>service types.</u>
 - Related PR link, <u>KEP 1645: relax the ServiceExportConflict requirements</u>
 - The impact of a conflicting service on existing imported services isn't explicit in the KEP; it's possible to understand the conflict resolution process as implying that ServiceImports are constructed based on the resolution order, but it would be useful to make that explicit in the KEP (Zhiying will submit a PR)
- [ryan] Property Ladder (https://github.com/kubernetes/enhancements/pull/5255)

- [ryan] kube-bind
 - Experimental project
 - Why host in SIG rather than donate to CNCF?
 - Not ready for incubation?
 - Not company-backed
 - Goal is to agree on common behaviour (KEP), not necessarily a common implementation — but there is a goal/need to have a trusted implementation

2025-07-15

Scrub

- KEP-5313 PlacementDecision API
- [ostrain] CP credentials KEP (github.com/kubernetes/enhancements/pull/5338)
 - Still need review from stephen & jeremy
 - Owen is taking over from Corentin as the driver for this PR

- kube-bind discussion started on the mailing list
- Open discussion
 - About why MCS project does not have a centralized published controller especially in light of multicluster-runtime project now
 - Some chat about MCS API integration with istio as a follow up from discussing projects that have a centralized implementation and those that do not
 - Relevant past convo on the istio tie ins with MCS API: https://kubernetes.slack.com/archives/C09R1PJR3/p174679825873259 9?thread_ts=1746790157.822419&cid=C09R1PJR3
 - o Late scrub adds/reminders about
 - Various MCS PRs from last week's scrub
 - Property ladder
 - Ryan to resolve comments in PR, everything not about voting can have a new proposal based on the discussion on the PR
 - More complex is the question of voting entities which is the
 most inquired about by Stephen. Idea that was talked about
 the most was by finding a way to consider a voting entity to be
 an MCS implementation. Some conversation about what is an
 implementation (producer or consumer, how complete) and if
 that can be gamed (aka kubectl apply counts). Stephen
 wants us to follow up on how Gateway API does this since we
 believe they also function in a per-implementation voting
 system

- Alex Scammon inquired about overlap with the batch working group (https://bit.ly/cncf-bsi-wq)
 - The credentials KEP above
 - Orchestration and scheduling

2025-07-08

Scrub

KEP-5313 PlacementDecision API

- [arthur] add ServiceImport conditions
 - https://github.com/kubernetes/enhancements/pull/5439
 - https://github.com/kubernetes-sigs/mcs-api/pull/113
 - o For any status/errors need to be reported at import time
 - First use case is IPFamily/IP protocol incompatibility (i.e.: importing an IPv6 only service in an IPv4 only cluster)
 - Also FYI the existing ServiceExport conditions PRs:
 - https://github.com/kubernetes/enhancements/pull/5438
 - https://github.com/kubernetes-sigs/mcs-api/pull/112
 - Please review the PRs if you are into/interested by MCS
- [arthur] Adding IPFamilies vs removing other fields that can be theoretically inferred from other clusters
 - Continue discussion from previous meeting / related to https://github.com/kubernetes/enhancements/pull/5264
 - supporting "schema" -> https://link.excalidraw.com/l/19CgeSQLzX3/78feRl3QNGK
 - If we were to follow the argument of not following the Service fields and not adding IPFamilies for that reason
 - we would need to consider removing a bunch of fields currently in ServiceImport (see the excalidraw above)
 - -> at least Cilium and Submariner wouldn't be super happy of removing all those fields from the ServiceImport API
- [mjudeikis,embik] Kube-bind discussion follow-up from 07-01
 - There were some follow-up questions. Anything needs addressing?
 - <u>Slides</u>, page: https://github.com/kube-bind/kube-bind
 - [mike] would like more clarity on who expected implementers of a standard would be

- [mike/john] some projects exist with seemingly similar scope such as https://kro.run/, this feels like a worthwhile use case and project, but i'm unsure if it would be appropriate to try to position as a "standard"
 - kro doesn't appear to deal with multiple clusters
- [john] does kube-bind operate on its own or does it rely on a multicluster provider?
 - kube-bind handles this itself
 - The proposal isn't intended to become a "blessed" API to share objects across multi-clusters
- [arthur] do you want to standardise the API, or host the project in SIG-MC?
 - The first step is to see if SIG-MC would like to host the project
 - kube-bind is two things: an API spec and a reference client and backend
 - There is interest in different backend implementations
- [MJ] if many providers have implementations, how do you trust all agents?
 You might want a single provider which talks to multiple backends. You also want to know you're giving minimal access to every provider; if a standard agent manages RBAC in a consistent way, that's useful from a trust perspective
- o [John] S3 type of format for Kube APIs.
 - I do think it might be important to separate out the distribution of changes vs. distributon of trust (be it push or pull)

2025-07-03 (APAC-friendly)

Scrub

•

- [Hongcai Ren@RainbowMango] Discussion the next move of the <u>standard concept</u> definition
 - Discussed in the previous APAC-friendly call and the June 24 call
 - Hongcai thought there would be a formal survey
 - Main question now is whether there's consensus to run a second survey as suggested in
 https://github.com/kubernetes/community/pull/8210#issuecomment-30123045
 - Ryan is not sure we need another survey; the previous survey mostly covered regular attendees to the SIG call. How can we reach a more diverse audience?

- [Hongcai] Issues with first survey: not "formal" enough (votes in a GH discussion), options not representative enough
- [Hongcai] We might need to define more than one term control plane, host cluster, member cluster
- [Hongcai] Most important issue is to get broader participation
- Is a survey effective if we want to define multiple terms
- Hongcai has a plan (new slide in the existing presentation
 Concept survey from Karmada)
- o [jqiu] What is the goal of the survey?
- [Hongcai] Term definitions for use in the SIG, not to force projects to change terminology
- [jqiu] Define concepts and map terms depending on context. Will we mention multiple terminologies?
- [Hongcai] Define a single term for similar concepts: for example hub cluster in OCM, host cluster in Karmada, decide which is used for the SIG.
- [Chen] A terminology file would be helpful, but this feels like a chicken/egg situation. First identify common concepts, and associate terms, or the other way round? We need to disentangle them.
- o [skitt] Concepts defined by the KEPs, we want to choose terms.
- o [Hongcai] Concepts aren't consistent either.
- [MJ] Importance of acceptance criteria how many votes are needed to make the vote "binding"? Want to avoid getting into this situation again if we run a second survey.
- [zhiyinglin] MCS API, export conditions https://github.com/kubernetes/enhancements/pull/5438
 - o This will break existing code, implementers need to change code
 - Would like to hear other people's feedback on the PR
 - https://github.com/kubernetes/enhancements/pull/5439 will have an impact on the parallel spec/status merge discussion
- [Ryan] how can we merge discussions from the US/EMEA and EMEA/APAC calls?
- [Chen] on the credentials front, some questions, will start the discussion on the merged PR (jqiu mentions that the PR was perhaps merged accidentally)
- [Ryan] next step on the credentials front is to integrate for example in ArgoCD; need to look at the multicluster-runtime PR (https://github.com/kubernetes-sigs/multicluster-runtime/pull/48, will use the mcs-api library)

Action items

- skitt to discuss the survey with Josh Berkus
- Hongcai to review https://github.com/kubernetes/community/pull/8210 from a concepts angle
- skitt to ask Gateway API how they handle discussions across calls

2025-07-01

Scrub

- Decision to merge for https://github.com/kubernetes/enhancements/pull/5436?
 - Enables conformance updates in https://github.com/kubernetes-sigs/mcs-api/pull/111
 - Issues around not wanting to require admins to look at all clusters to find a solution
 - Not having the behaviour at all might be better than having an inconsistent behaviour
 - Can be seen as an RBAC consideration.
- [Corentone missing meeting but please address] Decision to merge github.com/kubernetes/enhancements/pull/5338
 - O Please approve and submit?
 - No blocker? Everything has been discussed.
 - o still waiting for jeremy/stephen to review. no new replies
- Owner files updates?
 - <u>KEP2149: Add the well known property ladder</u> still waiting for jeremy/stephen to review. no new replies
 - o KEP4322: add some clusterProfile property blocked by property ladder
 - KEP-5313 PlacementDecision API
 - waiting for chairs to review

- [mjudeikis,sttts,embik] <u>Kube-bind</u> (<u>https://kube-bind.io/</u>) introduction, demo & discussion
 - <u>Slides</u>, page: https://github.com/kube-bind/kube-bind
 - [mike] would like more clarity on who expected implementers of a standard would be
 - [mike/john] some projects exist with seemingly similar scope such as https://kro.run/, this feels like a worthwhile use case and project, but i'm unsure if it would be appropriate to try to position as a "standard"
- [mike/arthur/tpantellis] standardizing ServiceExport condition types and reasons
 - https://github.com/kubernetes/enhancements/pull/5438
 - https://github.com/kubernetes-sigs/mcs-api/pull/112
 - See prior art in https://github.com/kubernetes-sigs/gateway-api/blob/main/apis/v1/gateway_ty
 pes.qo

- Defining ConditionType and ConditionReason constants can be very helpful for conformance tests, plus see docs on positive-polarity (like `{ type: Ready, status: True }`) vs negative-polarity (like where `{ type: Conflicted, status: False }` is the *good* case) conditions.
- [arthur] add ServiceImport conditions
 - https://github.com/kubernetes/enhancements/pull/5439
 - o https://github.com/kubernetes-sigs/mcs-api/pull/113
 - For any status/errors need to be reported at import time
 - First use case is IPFamily/IP protocol incompatibility (i.e.: importing an IPv6 only service in an IPv4 only cluster)
 - -> out of time, bumped to next week
- [arthur] Adding IPFamilies vs removing other fields that can be theoretically inferred from other clusters
 - Continue discussion from previous meeting / related to https://github.com/kubernetes/enhancements/pull/5264
 - suporting "schema" -> https://link.excalidraw.com/l/19CgeSQLzX3/78feRl3QNGK
 - [mike] got preoccupied with status proposals above, didn't get a chance to follow up on this, would appreciate deferring to next week
 - -> out of time, bumped to next week

2025-06-24

Scrub

- need Chair's call (or missing elements) on CP credentials KEP <u>github.com/kubernetes/enhancements/pull/5338</u>
 - Tiny prototype in the PR in the cluster inventory repo how controller would look like when implementing the KEP
- UX research project: https://forms.gle/9hSd5okypD771Nod8
 - o Email for sig-mc email is in draft
 - o [josh] can connect with contributor comms for more assistance
- need Chair's call (or missing elements) on HubCluster definition https://github.com/kubernetes/community/pull/8210

- [arthur,mike morris] discussing about the IP Families PR and mike argument for not adding this field: https://github.com/kubernetes/enhancements/pull/5264
 - Let's keep talking about it on the PR and at the next meeting if necessary as well

- See the comments specifically backtracking from https://github.com/kubernetes/enhancements/pull/5264#issuecommen t-3001129468
- We recovered what has been discussed lately and the summary of what this is originally about
- Important point is that the Cillium architecture does not have as centrallized a controller as other implementations so it needs this information to be signalled at this point
- There is going to be some more engagement with more Cilium people from Microsoft to better understand the architectural situation
- Open question: If we go totally the other direction on this and decide not to put this here, do we also remove SessionAffinity?!
- [Tom P] re: comment on https://github.com/kubernetes-sigs/mcs-api/pull/111#discussion r2162096079
 - Related to conflict conditions on ServiceExport
 - This PR is to add conflict resolution tests for service types (headless vs clusterip)
 - The comment is about propagating the conflict condition back to all clusters who have exported the service, which assumes access to all clusters by the controller that is responsible for this
 - Submariner is not able to do this at this step, or more precisely, does not have access to all the clusters to update the service condition everywhere
 - Could be done in a complicated data forwarding way with annotations etc
 - Questioning whether propagating this up to every ServiceExport is useful
 - [mikemorris] easier to do this in a centralized controller model
 - Option: If we were to handle this, is a technical avenue in the cluster information stanza (right now its only requires its name)?
 - Option: Can we loosen the statement?
 - o [jeremyot] context from the past on conflict resolution
 - Wanted to make it "easy to tell" if exports exporting the same service disagree, under the baseline that outside of a rolling upgrade or something, steady state should be that they are non conflicting
 - Disagreement shouldn't "break" it because we don't want to break rollouts or bad config changes
 - Guarantee the derived service state matched at least one actual configuration that existed somewhere NOT mix and match
 - For example, while "oldest" is an ugly solution in a distributed system (tlmE iS mEaNinGleSs) but met the above bars
 - [mikemorris] Gateway API has a similar well-defined conflict resolution strategy that includes "oldest" to avoid breaking existing config

https://gateway-api.sigs.k8s.io/guides/api-design/?h=conflict#c

onflicts (but doesn't necessarily have the same distributed-system problems as it assumes a single k8s API server as source of time)

 Next step: Tom is opening a PR about softening the language on how many clusters MUST vs SHOULD have the conflict in their ServiceExport

2025-06-19 (APAC-friendly)

- [Hongcai Ren@RainbowMango] Share the <u>survey results from Karmada community</u> regarding to <u>standard concept definition</u>. (<u>Survey result</u>)
 - Glossary (short list of terms) + more detailed documentation of concepts (with connections, for example the references to ClusterProfile)
 - Hongcai prefers "host cluster" / "member cluster" terminology (from Kubefed)
 - "Host cluster" is the Kubernetes cluster where running the control plane components of a multi-cluster system.
 - Terminology is context-dependent:
 - API hosting ClusterProfile?
 - "Host" cluster for end-users will generally refer to "workload" clusters, the opposite of "host" or "hub" cluster
 - [Hongcai] Suggest asking implementer projects
 - o [jqiu] Define the concept in ClusterProfile?
 - [skitt] Review the existing KEPs to see how much clarification / updating would be needed following an agreement on the terminology
 - Naming is hard
 - [jqiu] It might be misunderstood the document is not intended to force projects to use the same term, but to provide a shared understanding of concepts
 - [Hongcai] Hoping that projects would all use the same term; the documentation isn't a MUST but for Karmada, would like to follow the standard from SIG
 - [MJ] Maybe keep a list of how projects refer to concepts; having the same name also doesn't mean we have the same implementation
 - [MJ] Useful to have project-specific terminology matching the project implementation, for example when implementing multicluster-runtime providers
 - Next steps: run a formal survey, Hongcai will raise this with Corentin

2025-06-17

Scrub

- [corentone can't attend 06/17; please record decision] please add
 Survey questions for feedback to the scrub
 - Meeting notes
 - Incorporated feedback from the doc in the Google forms implementation which is at https://forms.gle/9hSd5okypD771Nod8
 - We are ready to send this out according to the <u>proposal doc</u>, starting with this announcement in the sig meeting (please fill out the google form! Send to your coworkers! Send to your friends! Send to your family!!!!! lol)
- [corentone can't attend 06/17; please record decision] need Chair's call (or missing elements) on HubCluster definition https://github.com/kubernetes/community/pull/8210
 - [Hongcai Ren] can't attend either, hope to share the survey result from Karmada community at next APAC-friendly meeting.
 - Meeting notes
 - we want to get the analysis from the survey results before moving on
- [corentone can't attend 06/17; please record decision]
 - Meeting notes
 - Some minor updates needed to pass tests
 - Planning on the exec model (as opposed to a plugin or other library) to avoid, the vision was to avoid recompiling as implied by a plugin model
 - Open question: how can we avoid any recompiling? Does this actually solve this? Linked to an example implementation (https://github.com/kubernetes-sigs/cluster-inventory-api/pull/1 7) that showcases this
 - Discussed in meeting the example implementation and about how the shown library provides both a plugin and an exec model (where an exec model is that the function call is responsible for finding a separately compiled executable). And that what we really get is less that each implementer will need to recompile to use this library once, but that we will not need them to recompile every time an implementer makes a new plugin binary (upgrade or just new to the ecosystem)
 - Open question: representation of certificates? Latest representation is actually serializable unlike others proposed before, plz update KEP with this

Topics

- [lauralorenz] k8s has to move off of Slack, downgrades to Free tier features (after 10 years!) this Friday June 20, during which you are responsible for uploaded files and private message history see <u>announcement</u> here and <u>readme here</u>, <u>github</u> <u>discussion for changes</u>. Back it up!
- [skitt] Reviewers
 - We have them! Jeremy/Stephen announced names from the volunteer pool and will email them directly and set up OWNERS files updates for them
- [lauralorenz] times up for new people's insights on cluster property ladder https://github.com/kubernetes/enhancements/pull/5255. Discussion topics for today:
 - Deprecation / removing property process (kubelet n-3 energy?)
 - Naming conflicts (first come first served?)
 - Conformance obligations / low hanging fruit
 - Versioning?
 - Voting <u>proposal</u> is based on cloud-event model
 - > Who are eligible to vote: Every entity whose member collectively kept good attendance (attend the meeting at least 4 out of the last 5 times) have one vote.

2/3 of all eligible voters vote yes pass the proposal The cool off period is 3 months No veto power

 [7/1 update from lauralorenz] Discussion points that were in comment form in this agenda have now been moved into the PR directly

2025-06-10

Scrub

- KEPs:
 - KEP-1645: define dual stack policies and fields
 - [Mike Morris] Currently working on replies. TLDR: understands the use case, api change may not be necessary, should change the language of the KEP at least. Stephen kitt has a similar implementation in submariner where they handle it.
 - [Arthur] Not exactly like that in submariner. Reminds KEP: Proposing adding an IPFamily in ServiceImport / Most likely useful for at least implementation using derived services
 - [Tom] No need for IPFamily fields. doesn't affect submariner. Happy with it being optional
 - TLDR: needs more discussions

- KEP2149: Add the well known property ladder still waiting for jeremy/stephen to review. no new replies
- o KEP4322: add some clusterProfile property blocked by property ladder
- o KEP-5313 PlacementDecision API
 - waiting for chairs to review
- KEP 5339 ClusterProfile Credentials external providers #5338
 - got a lot a reviews thank you! needs chairs/committee for review/approval
- MCS API changes: does implementers use ServiceImport spec as desired state?

Topics:

- [corentone] what's up with the committee? [bsalamat +1]
- [corentone] Go in a little bit more details of https://github.com/kubernetes/enhancements/pull/5338; Need committee for this one. I think it could go in!

2025-06-05 (APAC-friendly)

- Round of Introduction.
 - Stephen introduced himself
 - Hongcai, the maintainer of Karmada, expressed the gratitude to have an APAC time friendly meeting.
 - Mangirdas Judeiks ("MJ") works on multicluster-runtime and also likes the new APAC time.
 - August is a PM from OCM and he enjoys the meeting time too.
 - Zhiying is from Microsoft
 - o Qiu Jian from Red Hat
 - Vishal from Red Hat (Submariner)
 - Zhujian is a maintainer of OCM
 - Qing Hao is a maintainer of OCM also, interested in placement API
 - Ryan from Microsoft, based on the West coast! Already unofficial meetings with Qiu Jian etc. so excited to have an official replacement
- [Zhiying Lin] call for inputs/ more data points on MCS serviceImport api changes on issue 106
 - o [Hongcai Ren] I can take a look late and leave my comment on that issue
- [Zhiying Lin] the conformance test: service type [issue 67] has no progress for a while. I'm willing to help.
 - Stephen agree that zhiying can help
- [Ryan/Qiujian] <u>KEP 5339 ClusterProfile Credentials external providers #5338</u> updates

- Qiujian introduced the new clusterProfile plugin interface
- Stephen explained that the exec is to match kubeconfig exec mode
- MJ will take a look at the implementation to see if it fits with the multicluster controller runtime
- Determine the constraints of exec v. plugin, the cost of supporting both which shouldn't be that high

2025-06-03

Scrub

- KEPs:
 - o KEP-1645: define dual stack policies and fields
 - KEP4322: add some clusterProfile property
 - KEP2149: Add the well known property ladder aim for June 17 for decision, Stephen Kitt and Jeremy to review, Laura to send an email to the mailing list to draw attention
 - KEP-5313 PlacementDecision API
 - KEP 5339 ClusterProfile Credentials external providers #5338
 - MCS API changes: does implementers use ServiceImport spec as desired state?
 - This is the discussion issue with a concise explanation of the current question that is blocking https://github.com/kubernetes-sigs/mcs-api/pull/85, need input from Submariner and Google

- [bsalamat] SIG leads to announce their picks for the review committee.
- [embik] Quick multicluster-runtime update
 - Alive and kicking
 - Minor updates and improvements; branch open for controller-runtime 0.21 version compatibility
 - Provider implementations (multicluster runtime has an idea of a "cluster provider" with a go implementation the cluster providers can implement to give provider specific access logic) coming up downstream, looking for more!
 - KCP has a 0.1 release and considered usable
 - Hottest newest news is Gardener
 - [corentone] can multicluster-runtime peeps take a look at https://github.com/kubernetes/enhancements/pull/5338 and see if it fits with their life?
 - Also cluster profile at some point
- Survey questions for feedback, still looking for feedback, even just a happy emoji:)

- Backstory of entire effort:
 - Proposal for Engineering-Focused Multi cluster Monitoring and Observa...

Scrub

- KEPs:
 - KEP-1645: define dual stack policies and fields
 - KEP4322: add some clusterProfile property
 - KEP2149: Add the well known property ladder aim for June 17 for decision, Stephen Kitt and Jeremy to review, Laura to send an email to the mailing list to draw attention
 - KEP-5313 PlacementDecision API
 - o KEP 5339 ClusterProfile Credentials external providers #5338

- [Pavanipriya] Presenting Engineering Focused multi cluster monitoring and observability User research and survey doc.
 - Proposal for Engineering-Focused Multi cluster Monitoring and Observability S…

 Provide feedback on the survey questions:
 - https://docs.google.com/document/d/1C_ph6inl8gGg8wWPTTLgllPWV7HpWovb44M_3AFfeS0/edit?usp=sharing
- [arthur] ServiceImport status conditions for dual stack handling?
 - For instance when you want to get a ipv6 ServiceImport in an ipv4 cluster
 - Excalidraw schema used to illustrate: https://link.excalidraw.com/l/19CgeSQLzX3/9JldD285kOO
 - Submariner have a use case to signal it to user that import didn't worked out
 - Right now the IP is ignored at dns level if it can not be imported
 - [Arthur] Will create a PR to add ServiceImport status and discuss there opinions about it
- [arthur] Presenting on MCS API change to be done before v1beta1
 - Presentation: MCS API Change
 - Question about why we want to move everything to root
 - ServiceImport is designed to be entirely written by a controller
 - Suggestion to collapse to root instead of spec/status pattern initially proposed during <u>July 23rd, 2024 SIG-Multicluster meeting</u>, citing similarity to resources like <u>EndpointSlice</u> generally written by controllers instead of human request-response desired state reconciliation
 - Azure moved many of the fields to Status

- Azure allows manual creation of ServiceImport because clusters can be asymmetric
- [skitt] Result of the APAC-friendly poll: Thursdays at 6:30 UTC bi-weekly starting next week?
- [bsalamat] Review Committee selection: We have 11 volunteers from 6 companies. SIG leads please make the call.
- [corentone sorry was late] intro to KEP creds please review KEP 5339 -ClusterProfile Credentials external providers #5338 !!!

Scrub

- https://github.com/kubernetes-sigs/mcs-api
 - o apis/v1alpha2: move ServiceImport spec fields
 - Alpha → Beta Graduation (needs the previous
- KEPs
 - o KEP-1645: define dual stack policies and fields
 - KEP4322: add some clusterProfile property
 - KEP2149: Add the well known property ladder
 - o KEP-5313 PlacementDecision API

- [bsalamat] 7 people have shown interest in participating in the review committee.
 Given the number of volunteers so far, we will probably go with a 5 person committee. Please fill out this form by EoD May 20, 2025 if you are interested.
- [arthur] followup discussion for https://github.com/kubernetes/enhancements/pull/4715 after last week scrub
- [corentone] [slides] thoughts about scheduling; Would like to suggest our Placement CRD as a standard. (we can always have an output-only first then standardize the spec)
- [arthur] cool news regarding CoreDNS/MCS-API
- [corentone] Update on management project waiting for Karmada
- [corentone] Credentials; will submit a KEP, but looking for design partners:) There's a couple of us in a slack thread, but would love to take this publicly. See discussion: https://kubernetes.slack.com/archives/C09R1PJR3/p1747786073248159

Scrub

- https://github.com/kubernetes-sigs/about-api
 - Updating ClusterProperty API to v1betav1
 - Clean up duplicate CRD files
 - Should we consider this initiative dead?
 - KEP2149: Add the well known property ladder
- https://github.com/kubernetes-sigs/mcs-api
 - o apis/v1alpha2: move ServiceImport spec fields
 - Alpha → Beta Graduation (needs the previous item)
- https://github.com/kubernetes-sigs/work-api
- KEPs:
 - KEP 1645: add a derived service annotation on ServiceImport
 - o KEP-1645: define dual stack policies and fields
 - o KEP4322: add some clusterProfile property
 - KEP2149: Add the well known property ladder

Topics

- [Mike Spreitzer] KubeStellar
- [mikeshng,ryanzhang-oss] PlacementDecision API proposal and demo. <u>Community</u> doc here, presentation slides here
- [Babak Salamat] Given that we haven't seen any objections to the idea of having a 4-5 member committee for reviewing PRs, let's move forward with identifying the committee members. I would suggest making an online form where volunteers can submit. Volunteers commit to spend 4 hours weekly to review PRs. Once the volunteers are identified, SIG leads will choose 4 or 5 people among them. They will choose people from different companies if there are enough volunteers.
 - [hongcai] +1 again on this idea. But we might don't need to setup the committee, we can nominate some <u>Tech Leads according to sig-governance</u>.
- [lauralorenz] APAC meeting updates on <u>slack thread</u>, tl;dr it will be an additional meeting and Stephen will be the bridge lead. He has opened <u>a poll</u> for APAC interested peoples to give time requests, and also asked in the slack thread about days of week (proposing later in the week)

If you want to be in the reviewing committee of the SIG, please <u>fill out this</u> <u>form</u>.

- [corentone] Short Updates and follow up

 - [corentone] Hub cluster position doc to be approved soon: https://github.com/kubernetes/community/pull/8210
 - o [lauralorenz] APAC: no update
 - o MCS: no update
- [Babak Salamat] Let's organize a 4-5 person committee of volunteers from different companies who can spend at least 2 hours every week to review PRs. Once the committee give their thumbs up to PRs/KEPs, SIG Leads approve them (if they don't have any serious objections).
 - [Mike Morris] I can dedicate time to this starting next week
 - [corentone] Would people be okay if we had the following pillars? We can have a 5min update at each sig-mc meeting to sync on whats happening before our deep dives. See Als in:
 - **■** [PUBLIC] Sig-Multicluster Pillar approach for more engagement
 - o [liqian] +1 on this idea to make faster progress in related projects
 - [hongcai] +1 on this idea. Would like to get more involvement from karmada, kubefleet, OCM, etc.
 - Al: Create a structure to review and approve PRs more promptly
- [raffaele spazzoli] present helium a different approach to multicluster https://github.com/raffaelespazzoli/helium
 - Issues with traditional hub/spoke model (shared state in the hub)
 - Hub is a SPOF
 - Multicluster operators are hard to write (watching multiple control planes)
 - Architecture isn't scalable beyond a few hundred clusters (single partition)
 - Distributed control plane for shared state
 - Fully distributed database stretched across the clusterset, with Kine
 - Shared API domain alongside the local API server (local etcd), shimmed through the local API server
 - Operators can access the shared state but still only access a single API endpoint
 - If needed, clusters can be added without stretching the database
- [corentone] Announce officially ClusterProfile being supported by GKE; official Doc

2025-04-29

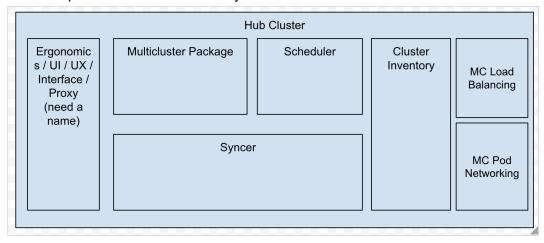
 [lauralorenz] APAC meeting - see also <u>slack thread</u> [also moved to backlog bc leads were not present]

- [ryan, arthur] the MCS v1alpha2 PR (and slack discussion https://kubernetes.slack.com/archives/C09R1PJR3/p1745342913008259)

 https://github.com/kubernetes-sigs/mcs-api/pull/85
 - [PUBLIC] spec, status, root and MCS api VIPs
- [arthur] dual stack MCS-API
 - Proposing the following change to the spec/CRD https://github.com/kubernetes/enhancements/pull/5264
- [corentone] sig organization: pillar streams proposal
 - [PUBLIC] Sig-Multicluster Pillar approach for more engagement
 - No new meeting but reserved time for status update in each pillar at beginning of meeting
 - goal is to encourage each flow to independently move and still get the opportunity to show progress and call out for needs (reviews, feedback, code)
- [corentone] was starting notes around the credentials plugin
 - [PUBLIC] Plugin for ClusterProfile Credentials Notes @ostrain to join me on Google side to help with this.
- [corentone] Hub Cluster PR; what to do with it? I can close it if its not needed need Stephen or Jeremy to weight in.
 - https://github.com/kubernetes/community/pull/8210

2025-04-22

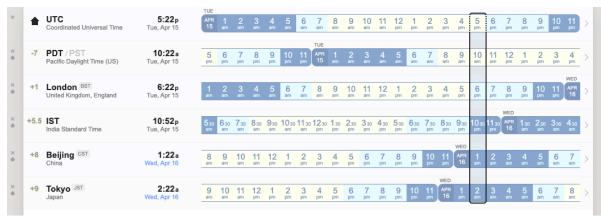
- [lauralorenz] MCO demo □ Kubernetes Multi-Cluster Orchestra Demo
- [bsalamat] Let's set a deadline of KubeCon NA (Nov'25) to address 2-3 major items in our API:
 - a. having Auth in ClusterProfile, so that the MC controllers can access workload clusters with the information in ClusterProfile.
 - b. an MC placement strategy and policy API. This could be a baseline that we can expand later.
 - c. (optional) cluster auto-provisioning API. This one is not a fundamental requirement, but I think it is an important piece of the puzzle to have a practical multi-cluster story in a cloud environment.



- [ryan] the about API property ladder PR https://github.com/kubernetes/enhancements/pull/5255
- [ryan, arthur] the MCS v1alpha2 PR https://github.com/kubernetes-sigs/mcs-api/pull/85

2025-04-15

- Introductions:
 - Shweta Vohra from Booking.com
- [ryan] the updated property KEP https://github.com/kubernetes/enhancements/pull/5185
 - o Inconsistencies around versions of the About API, docs shouldn't be v1
 - o Proposal to split the PR into two:
 - The property ladder first
 - Then properties, following the ladder
 - Behaviour tied to About API the property ladder should apply there too and it might be better to define it there
 - Are all About API properties supposed to be made available in ClusterProfile?
 Should we define a meta property specifying whether a given About API property is expected in ClusterProfile too? See the label explanation in the current proposal
 - Require concrete implementations for each property before a proposal?
 - ClusterProfile shouldn't need About API, the information it contains may be populated without local information
- [babak] How can we move faster on KEPs etc.?
 - Doubled calls
 - Should we create working groups?
 - Credentials
 - Sync
 - Scheduling
 - UI/UX
 - Challenge: APIs move faster when implementations are using them → we need to get implementations to participate in the SIG
 - For example not much feedback from implementers of MCS API (see also skepticism in Istio but no feedback to the SIG)
 - Feedback loop working well in Gateway API
 - Karmada and OCM skew to APAC so having an APAC-friendly call would help
 - Not necessarily meetings, explore other ways to move faster
 - o Pinging each other, discussing PRs etc. in Slack
- [ryan] the workAPI https://github.com/kubernetes-sigs/work-api
- [ryan] placementDecision
 https://open-cluster-management.io/docs/concepts/content-placement/placement/#placementdecisions
 - o [skitt] ping Josh Packer re OCM contributing to the two APIs above
- [mike] trying to pick up v1alpha2 again for MCS API, blockers or work remaining on https://github.com/kubernetes-sigs/mcs-api/pull/85?
- [mike] scheduling APAC-friendly timezone meetings on alternating weeks now that we're moving to a weekly cadence?



o Source for above:

0

https://www.worldtimebuddy.com/?pl=1&lid=100,8,2643743,30,1816670,1850147&h=100&hf=1

FWIW SIG-Release did this time (screenshot below) for APAC this last cycle:



 [corentone] Let's merge or close hub cluster definition? https://github.com/kubernetes/community/pull/8210

2025-04-08

Meeting host: lauralorenz@google.com

- [ryan] the updated property KEP https://github.com/kubernetes/enhancements/pull/5185
- [Arthur] WDYT about MCS-API in a scenario where some clusters are dual stack and some are single stack?
 - Is it fine that clients may only reach a subset of backing pods depending on what IP stack they use?
 - Should we do anything about dual stack? For instance:
 - (strict) Try to reconcile the intersection of ip family among exported
 Service and on Import erroring out if we cannot get at least one of the exported family in our local cluster
 - (loose) Always attempt to get a dual stack service (similar as PreferDualStack) and accept that some client may not be able to reach every backing pods

- Should we amend the KEP or this should be exclusively implementation specific?
- These should be mostly implementation specific but the KEP might need a few tweaks to suggest to take IpFamilies on Service as a hint on what to do

2025-03-18

- Clay Baenziger Would love to walk through his team's infrastructure and how it relates to SIG-MultiCluster projects; perhaps we've had some useful thoughts (see the talk at <u>Cloud Native Rejekts</u>)
- [Arthur] how should we continue the discussion on figuring out what to do with the "trafficDistribution" field (with sig-network?)
 - https://kubernetes.io/docs/concepts/services-networking/service/#traffic-distrib

 ution
 - Will email sig-network+sig-multicluster if sig-network has some opinions on that and restart the discussion in either sig-network or sig-multicluster meeting
- [corentone] present an idea for credentials issuance of ClusterProfile (still an early draft of an idea but want to get feedback and we can work on a KEP or edit on the ClusterProfile KEP
 - □ [PUBLIC] ClusterProfile Credentials via Plugin
- following were not discussed.
 - [ryan] the updated property KEP https://github.com/kubernetes/enhancements/pull/5185
 - [corentone] Let's merge or close hub cluster definition? https://github.com/kubernetes/community/pull/8210

2025-03-04

- [skitt] Onboarding https://groups.google.com/g/kubernetes-sig-multicluster/c/3BfLtXzoJ1A/m/0_MChcxuAQAJ
- [ryan] KEP draft https://github.com/kubernetes/enhancements/pull/5185
 - o A separate PR will be about how to add the last refresh time for each property
 - [NEEDS LINK] upstream k8s KEP discussion of grouping prefixes for existing well-known topology keys
- [ryan/corentone] cross-cluster controller work with ClusterProfile

2025-02-18

- [sttts] From first glimpse, same topic as [max]'s backlog item:
 - [max] discuss cross-cluster controllers/reconciliation as a follow-up to https://kubernetes.slack.com/archives/C09R1PJR3/p1663758534872249?thr ead ts=-1663758534.872249&cid=C09R1PJR3
 - AI: follow up on controller-runtime/operator SDK work and Work API as an example and anything else about multicluster (and multiregion?) control planev
- We have been working on a multi-cluster proposal https://github.com/kubernetes-sigs/controller-runtime/pull/2746 for quite some time.
 https://github.com/kubernetes-sigs/controller-runtime, pull for controller-runtime, implementation OUTSIDE of controller-runtime became feasible.
 https://github.com/kubernetes-sigs/controller-runtime, pull for controller-runtime became feasible.
- We have prototyped this as
 https://github.com/multicluster-runtime/multicluster-runtime and are now looking for a formal home, potentially in sig-multicluster (where IMO it belongs), and with that an sig.k8s.io/multicluster-runtime import path.
 - Also: being owned by a sig allows a number of people to contribute, without going through a complicated OpenSource company approval process.
 - o tl/dr: Multi cluster controllers with controller-runtime
 - o **no fork, no go mod replace**: clean extension to upstream controller-runtime.
 - universal: kind. cluster-api. vcluster through cluster-api. Gardener (tbd). kcp (WIP). BYO. Cluster providers make the controller-runtime multi-cluster aware.
 - seamless: add multi-cluster support without compromising on single-cluster.
 Run in either mode without code changes to the reconcilers.
- Next steps:
 - 1. finding a formal home
 - 2. continue prototyping by adding (experimental) mutli-cluster support to some big projects. Interesting candidates (help wanted!):
 - a. Crossplane (we have started with this one)
 - b. kueue
 - c. kro
 - d. cert-manager
 - e. Argo
 - 3. At some point in the future: reevaluate whether multicluster-runtime stays or is subsumed by controller-runtime taking over. But this process needs real-world experience (= the goal of multicluster-runtime).
- Feedback/questions?
 - Excitement from Jeremy and Corentin, broad agreement this feels like it could be a good fit for SIG-MC
 - Jeremy to coordinate with @skitt on onboarding process to get a repo under kubernetes-sigs

2025-02-04

- [corentone] https://github.com/kubernetes/community/pull/8210/files update
- [skitt] Releasing a new version of the mcs-api repo
- [mike] Followup on port conflict rules clarification should ServiceImport actually use the *intersection* of service ports declared on exported services rather than the *union* as currently specified in the KEP?
 - https://github.com/kubernetes/enhancements/pull/4887#pullrequestreview-24 94246716
 - o https://github.com/kubernetes/enhancements/pull/4887
 - This would likely merit a followup to https://github.com/kubernetes-sigs/mcs-api/pull/78 to clarify expected behavior more clearly in conformance tests.
 - This should be decided before we release 0.2.0
 - ACTION ITEM: requesting eyes on this for review, could likely be considered to be a breaking change, although it is only applicable to edge cases where backends are dissimilar and not the common case where backends are identical

2025-01-21

- [Arthur] FYI sig-multi-network is bootstrapping a new initiative to have a commonly maintained generic endpointslice reconciler
 - Cilium already have a fork of the kubernetes endpointslice reconciler for EndpointSlice syncing in a multi cluster environment
 - Multi network minutes: Multi-Network community sync Notes
- [corentone] go over https://github.com/kubernetes/community/pull/8210/files for final edit
- [ryan] Cluster Properties updated according to feedback

2025-01-07

- [Arthur] Adding targetPort to ServiceImport to fit derived service implementation (or more generally anything that reads from the ServiceImport to do LB/something similar to what a kube-proxy + other kcm controller would do on a Service)
 - o https://github.com/kubernetes-sigs/mcs-api/pull/93
 - Target port was not considered at the time to limit the API exposed to the user

- Will try to pass it down without adding this field to ServiceImport and trying to check how hard it would be this way
- [corentone] management cluster's name poll; results and next steps.[bryan] sizeable number of prs on mcs-site ready for reviews (could be resolved by https://github.com/kubernetes-sigs/sig-multicluster-site/pull/24)

2024-12-10

- [Arthur] Looking for reviews for:
 - Labels/Annotations export:

https://github.com/kubernetes/enhancements/pull/4922

- And https://github.com/kubernetes-sigs/mcs-api/pull/84
- Cilium feature freeze was last friday would need this very very soon (before the release branch out) to possibly have this in Cilium 人
- V1alpha2: https://github.com/kubernetes-sigs/mcs-api/pull/85
 - This PR continue Mike work by fixing the generation but still storing v1alpha1 before coredns support v1alpha2
 - Should decide if we want to do this and if so the sooner the better as it needs coredns support for next steps (and winter break period is soon)
- [Bryan Oliver] Short Docs/Site PR https://github.com/kubernetes-sigs/sig-multicluster-site/pull/27
- [corentone] Talk about standardizing the definition of Management Cluster. (See DRAFT)
- [lauralorenz] in absentia (I need to miss the 12/10 meeting) but wanted to mention Nick Eberts and I (+ anyone else interested!) are planning on hosting the bye week "users" / "use case" meetings discussed last meeting in January after the holidays are over. This will be a scoped series of broad invite meetings with seed questions to gather use cases from multicluster users, and everyone's help socializing the project within their teams / customers / etc will be appreciated
- [mikemorris] assuming we're cancelling the December 24th meeting?

2024-11-26

- Dependency update PRs and v1beta1 tag for ClusterProperty API
 - o Pull requests · kubernetes-sigs/about-api · GitHub
- [Ryan]Discuss an "office hour" option: based on the feedback from KubeCon NA
 - Clay Baenziger (welcome!) mentions using entity relationship diagrams to represent CRDs and help visualise them for users
 - Bryan Oliver mentions a couple of talks he gave, <u>https://www.youtube.com/watch?v=e4vHy7Y-PFc&t=3s</u> and <u>https://www.youtube.com/watch?v=rkteV6Mzifs&t=1653s</u>

- [Arthur] Are you aware of existing projects using/publishing to the about API?
 - Context: thinking about integrating about API in Cilium network policy (aka adding cluster labels and authorize network traffic based on that)
 - Constraint: those labels would need to not change often, we could selectively include/exclude properties though
 - [Laura] proposal from that time: ☐ Proposal: Adoption of ClusterID
 - [Mike NG] Talked about OCM and cluster claim that integrate/plan to integrate with about api:
 - https://open-cluster-management.io/docs/concepts/clusterclaim/

0

- [Arthur] Let's discuss about MCS-API v1alpha2
 - Related PR: https://github.com/kubernetes-sigs/mcs-api/pull/52
 - o Consensus seems to be moving towards a conversion webhook
 - o It's within the SIG's remit to host a conversion webhook
 - Can the SIG ship a webhook image? Perhaps we can just ship the webhook code and leave it up to downstreams to integrate it
 - Consensus to write the conversion code in mcs-api repo and let implementation integrate to build the binary, the image and the deployment YAML
 - [Arthur] Will look into implementing this and chatting with Mike to add this to his PR
- [PUBLIC] test plan updates for MCS api for conformance tests

2024-10-29

- [Arthur] Followup discussion on labels/annotations sync for MCS-API (https://github.com/kubernetes/enhancements/pull/4922)
 - Offline/while discussing on the PR we discovered that:
 - For implementation that doesn't do derived service, this not super useful right now
 - only for service import selector but controller using those is non existent ATM
 - Labels and annotations on the service directly was painful to manage because some tooling add some labels and annotations
 - kubectl add last applied annotation
 - argood add a label to mark the application
 - Last call for reviews
 - Arthur and Mike will work to get the CRD updated assuming the PR is approved
- [Keith] MCS and unroutable endpoints (using gateways)
- [corentone] Talk about the new version of ryan's PR: https://github.com/kubernetes/enhancements/pull/4778

2024-10-15

- [Arthur Outhenin] Planning to add label sync in Cilium's MCS-API
 - For potential integration with ServiceImport selector and syncing those to the derived Service as well
 - Should this be in the KEP too?
 - [mikemorris] Like
 https://kubernetes.io/blog/2023/11/28/gateway-api-ga/#gateway-infrastructure
 -labels?
 - [jeremyot] if we sync metadata it would make sense to sync both labels and annotations
 - [mikemorris] need to be able to specify which labels and annotations to sync
 (a field on the ServiceExport?)
 - [mikemorris] will reference Gateway API design document describing this feature
 - o [Arthur] will open a PR to the KEP for discussion
 - Comparable functionality in Gateway API for populating labels and annotations on generated resource from spec.infrastructure stanzas https://gateway-api.sigs.k8s.io/reference/spec/#gateway.networking.k8s.io/v1.
 GatewayInfrastructure
 - Some mention of motivation for this including syncing <u>Topology Aware</u> Routing annotations, but likely other use cases too
- [Arthur Outhenin] Should the new traffic distribution field be added in the MCS-API KEP / ServiceImport CRD:

https://kubernetes.io/docs/reference/networking/virtual-ips/#traffic-distribution?

- Something to sneak in for alpha2 perhaps?
- [mikemorris] was discussing exactly this situation internally with current plans to collapse spec/status, we would effectively lose the option to add any fields like this that we expect to be human-authored if entire resource is expected to always be machine-written
- Consensus is to postpone after v1beta1 but continue discussing in the meantime
- Discussed whether user story of exported service owner configuring this field on the underlying Service and having that imply routing configuration preference in remote clusters makes sense
- [Ryan Zhang] Quick update on the authorization through clusterProfile API.
 - Talked to SIG-Auth, the conclusion was to remove the secret from the API (to avoid any risk of leaks); API Machinery agreed
 - Going back to the drawing board, would only keep the principal, not the credentials, with an audience field to represent who can talk to whom
 - Users would need something else for credentials
 - [corentone] would the only alternative be federated credentials? OIDC?

- Didn't go into implementation details
- Desire is to bridge the gap between all different federation-related tools that already have their own way to store credentials for different clusters (Kueue, ArgoCD, Istio etc.)
- Could perhaps arrange for the Token API to give tokens which are valid on the leaf cluster

2024-10-01

- [Arthur Outhenin] Does anyone else than Cilium use/plan to use the multicluster coredns plugin (<u>GitHub - coredns/multicluster: CoreDNS plugin implementing K8s multi-cluster services DNS spec.</u>)?
 - Context: trying to simplify this process:
 https://github.com/cilium/cilium/blob/main/Documentation/network/clustermes
 h/mcsapi.rst#installing-coredns-multicluster
 - Created by AWS/EKS but not very active ATM
 - Will try to check with CoreDNS maintainers if it can be merged in tree/compiled in by default
- [Arthur Outhenin] Looking for reviews/opinions on this MCS-API KEP PR
 https://github.com/kubernetes/enhancements/pull/4672 about conditions/adding infos if the local service is involved in the conflict (last meeting topic)
- [corentone] discuss https://github.com/kubernetes/enhancements/pull/4778
 - Rotating secrets should be possible using existing features of the Secret API
 - More usage guidance (similar to the MCS controller descrpition in the main API KEP) would be useful
 - o Ask SIG Auth for their opinion, once we have a consensual proposal

2024-09-17

- [Arthur Outhenin] Cllium doesn't currently fully respect the OwnerReference on EndpointSlice
 - Small context schema: https://link.excalidraw.com/l/19CgeSQLzX3/9H0fOHTTPgn
 - [Several] Agreement that the essential requirement (if that) is that EndpointSlices are deleted along with their ServiceImports, and that owner references (direct or indirect) are only the suggested mechanism. The conformance test should only test the deletion, not the implementation.
- [Arthur Outhenin] ServiceExport conflict condition
 - Follow up from a thread on slack: https://kubernetes.slack.com/archives/C09R1PJR3/p1725449448075499
 - Small Context schema: https://link.excalidraw.com/l/19CgeSQLzX3/XBtR5mzBmo

- [Jeremy] In MCS services are peers, so any conflict means the overall service isn't healthy and reflects a misconfiguration. However some working state is better than outage.
- o [Arthur] Will look into describing more information to add in the conditions.
- [Ryan Zhang] Go over the presentation of CloudProfile API integration with Kueue setup in KubeCon HK: <u>Connecting the Dots: Towards a Unified Multi-Cluster Al/ML Experience Qing Hao, Red Hat & Chen Yu (youtube.com)</u>
 - Add clusterProfile docs by zhiying-lin · Pull Request #22 · kubernetes-sigs/sig-multicluster-site (github.com)
 - https://github.com/michaelawyu/fleet/blob/kubecon/demo/apis/clusterinventor y/v1alpha1/clusterprofile_types.go
- [Mike Morris] how to handle breaking change for v1alpha2 ServiceImport in https://github.com/kubernetes-sigs/mcs-api/pull/52?
 - Remove v1alpha1 directory?
 - Safe to set v1alpha2 as storage version?
 - What about storage version for ServiceExport which is staying identical, and could have both v1alpha1 and v1alpha2 versions (and where we likely want to avoid breaking end users if possible, because this the resource we expect them to actually author, as opposed to ServiceImport where a controller should be able to write the new v1alpha2 version)
 - Upgrade notes for implementations and/or end users?
- [Jon H] PSA: MCS API repo now has presubmit tests running on PRs (build/test/etc.). Please let me know if you run into any issues.
 - Issue I just discovered: <u>pull-mcs-api-verify job reports success but is actually failing · Issue #74 · kubernetes-sigs/mcs-api (github.com)</u>

2024-09-03

- [mikeshng] Quick KubeCon China recap
 - CNCF YouTube Playlist
 - o SIG-Multicluster Intro and Deep Dive Presentation Vid [ENG] PDF
 - Developing a Standard Multi-Cluster Inventory Vid [CHN]
- [jackfrancis] MCS API E2E status update

2024-08-20

- [skitt et al] E2E and conformance tests
 - E2E as an API change verification tool, especially in preparation for v1alpha2
 - Currently on MCS API but probably of interest for About, ClusterProperties etc.
 - [PUBLIC] test plan updates for MCS api for conformance tests
 - https://github.com/kubernetes-sigs/mcs-api/pull/61 shows an example of the current conformance report

- [Jon H] PR to implement the last e2e test described by the KEP: add test 3: connectivity to all endpoints of service imported by multiple clusters by nojnhuh · Pull Request #62 · kubernetes-sigs/mcs-api (github.com)
- [tpantelis] Questions about the spec:
 - Are ServiceImports intended to be aggregate objects? I.e. the MCS controller creates an aggregated ServiceImport common across clusters. Laura confirms.
 - ClusterSetIP: clusterset-wide or per-cluster, so what happens with aggregated ServiceImports?
 - The array of IPs was intended for dualstack

2024-08-06

- [jackfrancis] MCS API E2E tests next steps?
 - Agreement on collapsing all fields to root, NOT moving to status Mike M to update PR
 - Mike M to update KEP clarifying ServiceImport should only be created by a controller, not user-initiated request-response pattern, ping Jeremy OT for review
 - Get review from Karamada and other stakeholders who haven't been participating in this conversation before merging, then tag v1alpha2, wait minimum soak period, time for implementations to update to a v1alpha2 release, then should be clear for v1beta1
 - Jack & John will continue in parallel with e2e tests in coordination with Stephen and Laura and Tom
- [ryanzhang] <u>KEP-4322: Add credentials definitions by ryanzhang-oss · Pull Request</u> #4778 · kubernetes/enhancements (github.com)

2024-07-23

- [ryanzhang] Continue discussion on adding credentials to the clusterProfile API
 - o ClusterProfile credential Google Docs
- [mikemorris] would appreciate reviews on <u>adding an alternatives considered section</u> to the KEP-4322 Cluster Inventory API
- [mikemorris] next steps on https://github.com/kubernetes-sigs/mcs-api/pull/52 for rethinking ServiceImport `spec` fields?
 - Previously we discovered that Google's MCS implementation may just have a public docs issue and not have actually moved the spec fields to status - can we get confirmation one way or the other?

- Laura planning to look into this
- Jeremy mentioned wanting an update to the MCS API KEP clarifying who writes ServiceImport CRDs to a cluster - an automated or manual (human) controller, but not a human writing it as a request and expecting a response?
 - This (and collapsing the spec/status stanzas entirely) would likely foreclose a cross-ClusterSet (non-sameness) manual usage of ServiceImport are we okay with that?
 - No objections from implementations present in meeting (Azure Fleet, Submariner, Google)
- Feedback from Karamada needed, who has a contact and can reach out?
 - Ryan to reach out with Mike after PR is updated
- [laura] What does Istio MCS API impl do? Consume ServiceImports only, or act as controller to create ServiceImport from ServiceExports?
 - Mike to investigate
- Consensus on final direction moving fields to status, or eliminating spec/status stanza entirely to be more like Endpoints?
 - Mike to update PR to eliminate spec/status, collapse all fields to root

2024-07-09

- [ryanzhang] Continue discussion on adding credentials to the clusterProfile API
 - o <u>ClusterProfile credential Google Docs</u>
- [corentone] Continue discussion on Multicluster Controllers API to leverage ClusterProfile
 - [Public] Multicluster Feature API
- [jackfrancis] MCS API graduation status update
 - E2E test progress
 - https://github.com/kubernetes-sigs/mcs-api/pull/51
 - https://github.com/kubernetes-sigs/mcs-api/pull/53

2024-06-25

- [mikemorris] ClusterProfile KEP
 - https://github.com/kubernetes/enhancements/pull/4690 summarizing past namespace-scope vs cluster-scope discussion in "alternatives considered"
- [mikemorris] move ServiceImport spec fields to status
 - o https://github.com/kubernetes-sigs/mcs-api/pull/52
 - Lack of clarity on whether GKE actually moved spec fields to status or if https://cloud.google.com/kubernetes-engine/docs/how-to/multi-cluster-services/ s#consuming_cross-cluster_services is a docs bug, needs further investigation.
 - Discussed that maybe simply removing the spec/status stanza would be preferable if intent is never for this to follow a request/response pattern? This

may limit future extensibility to cross-clusterset imports/exports, but maybe that's okay as guidance in the past has been to avoid that pattern.

- Discussed a KEP clarification PR [mike: can't recall specific language we wanted to update]
- Other pending functionality request here is tenancy, to exclude same-named namespaces in different clusters in a clusterset from being able to export services - see https://github.com/kubernetes/community/pull/6748
- o **TODO**: Need input/feedback from Karmada
- [ryanzhang] Continue discussion on adding credentials to the clusterProfile API
 - ClusterProfile credential Google Docs
 - [corentone] slides to help with discussion and intro a "Feature API" (To be
 Named) ☐ [Public] Multicluster Feature API
 - o [mike] ran out of time, bumping to next meeting
- [jackfrancis] MCS API Graduation status update (unless we already discuss this as part of Mike's agenda item above :))
 - E2E test progress
 - https://github.com/kubernetes-sigs/mcs-api/pull/51
 - https://github.com/kubernetes-sigs/mcs-api/pull/53
 - API changes (we will probably have already discussed this, see Mike's note above)
 - https://github.com/kubernetes-sigs/mcs-api/pull/52

2024-06-11

- [Arthur Outhenin] Could we standardize the annotation "multicluster.kubernetes.io/derived-service" in the KEP if the MCS API implementation is using derived/shadow services
 - Context: To facilitate the support of ServiceImport in Cllium GatewayAPI implementation, I used this to "swap" the ServiceImport by the actual derived service inside the internal Cilium logic
 - [mikemorris] Gateway API using similar label for derived resources in GEP-1762
 - gateway.networking.k8s.io/gateway-name: my-gateway (where my-gateway is the name of the Gateway resource)
 - [Arthur] no one was strongly against, will send a PR so that people can comment there
 - PR is here: https://github.com/kubernetes/enhancements/pull/4715
- [ryanzhang] Looking for sign-off on KEP PR
 https://github.com/kubernetes/enhancements/pull/4620
 which gives more formal definition of what a cluster inventory is
- [ryanzhang] Continue discussion on [PUBLIC] ClusterInventory API access/credentials Google Docs

2024-05-28

- [ryanzhang-oss] ClusterProfile KEP
 - o https://github.com/kubernetes/enhancements/pull/4620
- [munnerz] Presenting Kubernetes apiserver proxy to introduce 'workspaces' (a collection of namespaces in a single cluster), with workspace-scoped watches & authZ (<u>slides</u>, <u>demo</u>)
- [Jon Huhn] Looking for feedback on MCS PR to unblock e2e tests and beta graduation: <u>update go, deps, and codegen by nojnhuh · Pull Request #49 · kubernetes-sigs/mcs-api (github.com)</u>

2024-05-14

- [mikeshng, ryanzhang-oss] InventoryCluster => ClusterProfile rename KEP PR.
 Define ClusterProfile scope KEP PR.
- [ryanzhang-oss] multicluster sig volunteer
 - [mikemorris] improve the conformance tests to drive implementation engagement
- [jackfrancis] MCS API graduation status
 - https://github.com/kubernetes-sigs/mcs-api/issues/48

- [Arthur Outhenin] Wondering if for MCS API it would make sense to add precedence for conflict resolution (headlessness/session affinity) to the local service if a ServiceExport exist

 - Tim Hockin said that this was proposed during the initial design but that it was decided to go with the oldest ServiceExport so that the conflict is resolved the same way everywhere for predictability
 - There was also a discussion about having a rule for the majority as well instead of oldest
 - No strong opinion from anyone on the call
 - [Arthur] Will try to write something into the alternative section that this was not considered to have strong consistency across all imported services
- [corentone, mikeshng] Discussion / Decision? <u>Cluster Profile API: cluster scoped or namespace scoped</u> New section: "One or multiple inventories?"

- Need to explicitly call out the relationship between clusterSet and a ClusterProfileList (aka ClusterInventory);
 - 3 options: (A) Inventory 1:1 Set; (B) Inventory N:1 Set; (C) Inventory1:N ClusterSet
 - A+B are fairly similar; guarantee that an inventory has sameness (preferred by participants to anchor definition of clusterset as an API – to confirm in KEP comments)
- ClusterSet<>Inventory decision to inform Cluster v Namespace
- Al mikemorris@ to send a PR to update the KEP
 - sig-mc chairs to approve to move forward
- o CRD cluster vs namespace scope will follow KEP

- [Arthur Outhenin] Adding condition to MCS API ServiceImport status (would be useful for the Cilium implementation)
 - [mikemorris] +1 in favor of this, i have a use case in mind for conditions on ServiceImport status for cross-ClusterSet non-sameness export/import pairs (see third bullet point in https://github.com/kubernetes-sigs/mcs-api/issues/48#issuecomment-204372
 7874 for more detail)
 - As per discussion in the meeting, the conflict is more something that the producer exporting the service should be resolving so even if the implementation deals with it at import time it's nicer to expose it on the export object
- [mikemorris] planning to propose breaking change moving at least IPs (possibly Ports and Type fields too) on ServiceImport from `spec` to `status` as discussed in https://github.com/kubernetes-sigs/mcs-api/issues/48#issuecomment-2043727874
 - Would this require a KEP or (as this doesn't use the "core" API group) just a PR to https://github.com/kubernetes-sigs/mcs-api with a v1alpha2 bump?
 - PR on MCS API repo is good place to start for async review/feedback
 - Other changes from <u>past voluntary API review</u> we may want to consider at this time?
 - [tim] alternate universe where users manually import a service? Would need to rethink this a bit
 - Original design was for MCS controller to own/write ServiceImport resources
 - Status cannot be written at resource creation time
 - [ryanzhang] we may not have IPs at object creation time, need to wait for controller to reconcile
 - [tim] Controller does need to set IPs on Service in spec, we just do this synchronously.

- [tim] if we were inventing Service today, would put IPs in status. At least in Service, need to allow user to request a specific IP. Original intent was to be consistent with Service for ServiceImport design.
- [jeremy] IPs on spec allows single-write instead of double-write if it were to be on status
 - [clayton coleman] status being creatable needs an official position?
 - [tim] dropped requirement
 - [clayton] not sure if we dealt with all the implications
 - [tim] It's not like this is a high-traffic object, it's less about optimization than simplicity
- [jeremy] is this more an internal resource like EndpointSlice where spec/status isn't relevant? Do we ever expect users to create a ServiceImport manually?
 - [mikemorris] potentially yes, enough that i wouldn't want to completely collapse/eliminate spec and status stanzas - not currently in scope for MCS API, but for cross-ClusterSet export/import pairs, a controller would not know in which namespace it should automatically create a ServiceImport without "sameness" parity between export and import locality
- [corentone] follow up on ClusterInventory credentials discussion and push/pull, exploring different patterns:

 [PUBLIC] Multicluster Push vs Pull: permissions
- [ryanzhang,mikeshng] finalise <u>cluserInventory name(clusterProfile)</u> and scope
 [Public] API scope of Cluster Inventory API
- [clayton] We are proposing a new working group, WG-Serving, to make accelerated inference better on Kubernetes and also improve serving workloads in general
 - Baking Kubernetes great for accelerated workloads: a serving working ...
 - Likely there are multi-cluster scenarios that are relevant, such as ensuring workloads can use capacity effectively across regions or reliability needs across clusters
 - [mikemorris] I'm not clear on if/how this differs from the WG Device Management proposal scope?
 - https://groups.google.com/a/kubernetes.io/g/dev/c/YWXGXe07A5w
 - ■ WG Device Management Agenda and Notes
- [jackfrancis] MCS API alpha-beta graduation
 - Just discussed some concrete things in this meeting that could facilitate this
 - [jeremy] if we do a breaking spec -> status change that could introduce some friction, need to do this before moving to beta

- [skitt] KubeCon follow-up on ClusterInventory:
 - [PUBLIC] ClusterInventory API access/credentials

- [jackfrancis and Jon Huhn] ClusterInventory reference implementation demo: https://github.com/Jont828/inventory-cluster, https://github.com/nojnhuh/cluster-inventory-capi
- [mikeshng] Namespace vs ClusterScope
 [Public] API scope of Cluster Inventory API
- ClusterInventory naming decision
 - Tentative agreement to move forward with ClusterProfile for CRD name?
 - Discussed how we could still keep "Cluster Inventory API" as project name (similar to About API with ClusterProperty CRDs as example)

2024-03-12 (special session)

- [lauralorenz] [[PUBLIC] ClusterInventory API access/credentials
 - Push/pull model is the data structure the same? Fields need to be optional at least. Some concern that the data structures will be different
 - Think about potential architectures, then have a concrete conversation about having credentials
 - Good enough to have admin credentials, and then SAs could be made per particular use case, but important to have at least the cluster credentials for controllers
 - Is endpoint & ca-cert enough for controllers; per controller credentials?
 - → Is there even a use case for putting cluster inventories somewhere untrusted/or that can't have admin credentials?
 - o If we allow providers to middleman the credentials, do we make integrations on top of this too brittle?
- [jackfrancis] ClusterInventory reference implementations (Cluster API, Fleet, e.g.)

2024-03-05

- [mikeshng,mikemorris] <u>KEP-4322: Cluster Inventory API</u> rename API to something else? Feedback from the community: ClusterInventory implies multiple things, while InventoryCluster feels a bit awkward and doesn't align with the usual k8s naming pattern.
- [lauralorenz] slight follow up on "push"/credentials thoughts from last time (super WIP slides, open for working session next week at this time)

2024-02-20

- [lauralorenz] want to talk about if clusterinventory should abstract anything related to cluster credentials. There is a use case in multi kueue
 - It had been talked about in the community doc before but held until Phase 2
 - There is a meta question about whether this is in scope for this API at all

- There is a point that the API to be 'generically useful' especially in small integration stories (as opposed to consumers who might get the same data from Cluster API, which comes with more stuff/requires blessed providers), it needs something like this
- Scalability convo tracking into design details tabled for next time
- o "Push" subteam will be thinkin' on this

2024-02-06

- [Stephen Levine] <u>sync-controller</u> <u>slack context</u> Teleport Cloud's approach to multi-cluster with CRD management plane
- [mikeshng] KEP-4322: Cluster Inventory API rename survey results. See ref.

2024-01-23

- [mikeshng] KEP-4322: Cluster Inventory API additional feedback and discussions. Based on the comment, rename to something else? "Cluster" might be too common and used by ClusterAPI already. "ClusterProfile" or something else or leave it?
 - NodeStatus is too vague and might cause confusion to consumers and users.
 - Cluster name uniqueness should show examples instead of providing options. One example should cover how to handle multiple cluster managers. ie use prefix.
 - Rename ClusterInventory API to something else. Create a poll to gather feedback (WIP).
 - ClusterAPI integration future goal. Not just for collaboration but ensure functionalities are not overlapping.

2023-11-28

[Mike Helmick] intro, describing Roblox approach to multi-cluster

2023-11-14

- [ArangoGutierrez,mikeshng] KEP-4322: Cluster Inventory API
 - Jeremy: how to deal with different definition unit measurements due to different hardware etc.
 - Carlos: how to aggregate cluster inventory into groups to differentiate for different usage (ARMs vs GPU etc) might solve the above(?)
 - Carlos: cluster name option-2 add to spec.clustermanager.group.name maybe?
 - Jeremy: If we really want to create a standardize format for uniqueness? Or set a requirement that is unique then it will make it easier. Maybe it's too early for uniqueness standadization.
 - Carlos: Let's relax the rules now, and ask for feedback so in the future if needed, we can harden the rule.

- Stephen: display name vs unique name Alot of tools have generated/unique names already. Once API is set it's hard to undo so it's better to have a relax requirement to start with.
- Carlos: Igtm can be resolved. Let Qiu Jian know.

2023-10-31

- [ArangoGutierrez] NFD ClusterFeature API How can NFD help to discover features for ClusterInventory
- [qiujian16,mikeshng,ryanzhang-oss] PR: KEP-4322: Cluster Inventory API welcome any comments, suggestions and additional use cases.

2023-10-17

[mcosbuc] As a follow-up to the round of introductions in the last call, I wanted to bring up the MongoDB operator and the multi-cluster related challenges we're facing, and the SIG topics we're following: MongoDB operator multi-cluster challenges.

2023-09-05

- [mikeshng,ryanzhang-oss,jnpacker] <u>ClusterInventory API proposal</u> addressing feedback from the previous meeting. Clarify and categorize "Healthy" and detailed providers implementation plans.
 - Healthy breakdown. Multiple "components" allowed
 - ControlPlaneHealthy apiserver/readyz, controller-manager/healthz, scheduler/healthz, etcd/healthz
 - AllNodesHealthy when all nodes are reporting ready
 - Potential Future "EnoughNodesHealthy"? some threshold in the number of ready nodes.
 - Multi-provider registration, two potential ideas
 - Non-goal this would effectively preclude future expansion to include this and de-conflicting even for multiple providers of the same type is not quaranteed
 - Cluster-scoped registration with no vanity names pets
 - metadata.Name is always generated (never specified)
 - Spec.ClusterManager similar to ApplyOptions.FieldManager. This must match for every entity expected to update status. Could be a multi-cluster provider name for instance: fleet-manager/instance-1.
 - open-cluster-management/instance-2
 - Specific Label (x-k8s.io/cluster-manager) must match the spec.ClusterManager. This allows for listing, "show me all my instances".

- Add human readable alias in spec or annotation
- This approach
 - o avoids conflicting creates between providers
 - Different providers could be permission partitioned by a clever admission plugin that matches on spec.ClusterManager.
 - Allow more memory efficient list/watch by provider
 - Extensions on the platform can have visibility across multiple providers
 - o Multiple providers won't fight on name
 - Guarantees that the cluster names won't be pets

2023-08-22

 [jeremyot] About API promotion https://github.com/kubernetes/enhancements/pull/4156

2023-08-08

• [qiujian16,deads2k,mikeshng,ryanzhang-oss, dtzar] <u>ClusterInventory API proposal</u> addressing feedback from the previous meeting.

2023-07-11

 [qiujian16, deads2k, mikeshng, RainbowMango, ryanzhang-oss] <u>ClusterInventory</u> <u>API proposal</u>

2023-05-30

- [keithmattix] Enhancements process timeline
 - Interest in standardizing GEP type out-of-tree enhancement proposals keith to float with SIG-Arch
- [shane] follow-up on SIG MC machinery, MCS, Gateway API implementations, e.t.c.
 - Shane has asked more folks from SIG-Net interested in MC to join our meetings.
- [andander@redhat.com] KubeStellar Mutlicluster Configuration Management for Edge, Multi-Cloud, and Hybrid Cloud - https://kubestellar.io - would like an opportunity to introduce our work https://docs.kubestellar.io/release-0.2/Getting-Started/guickstart/
- Matthew Clarke The Spotify deployment infrastructure team has a multi-cluster library client we're interested in open sourcing, we're just thinking about the best place to open source this

 [costin@google.com] need for API for local/multicluster semantics for custom domains.

2023-05-16

- [lauralorenz] office hours / kubecon debrief
 - SIG-MC related talks:
 - Two Houses, Both Alike in Dignity: Gateway API and MCS API (panel) https://sched.co/1Hydh
 - SIG-Multicluster Intro and Deep Dive (maintainer talk) https://sched.co/1HvTl
 - People who use / know about MCS API consider it standard
 - End user interest less at the API problem space as they are interested in:
 - Networking, of which the MCS API assumes is resolved
 - Even if this is not in the "charter" of SIG MC, it is the first thing
 a lot of people run into if they are not already on a single
 platform that has networking solved
 - Shared services
 - Making a central vault available across multiple clusters
 - Sensitive workload but that you still want to access from a variety of different clusters, MCS makes it easy because the consuming clusters can get DNS to it without having other access to the workload
 - Replication stuff
 - Investigating service meshes
 - Tenancy management with clusters as a side effect of that; how do I manage multiple distinct teams
 - Still an opportunity for broadcasting the MCS API more to people who need a way to connect clusters especially the tenancy management part
 - Get a sense that many people are looking for a multicluster solution, part about it is getting the word out, part of it is articulating it in a user-digestible form, part of it is about addressing some of the surrounding problems / interfaces users run into
 - o GEP-1748: Gateway API Interaction with Multi-Cluster Services

Do we need "more" / "the" upstream implementation of MCS API / other SIG MC projects?

- It's a segue into the next topic.......
- [shaneutt / davemartin] MCS, ingress, in-cluster, e.t.c.
 - Idea is that if there was more machinery / pluggable implementation then the convergence of related APIs and projects becomes more realistic. It would give implementers something to go "towards" to converge, "step up"
 - o Related: SIG-Network posting and

2023-04-04

- FYI only: ready to approve PR for "one dot" sublabel-disambiguated DNS naming of Headless service pods: https://github.com/kubernetes/enhancements/pull/3918
 - After this, About API beta bump PR will be rebased on this and be ready as all the graduation blockers and PRR review is addressed

2023-03-21

- [robscott] GEP-1748: Gateway API interaction with Multi-Cluster Services
- [lauralorenz] [[PUBLIC] MCS API: do we need to define EndpointSlice behavior
 - [mikemorris] the "could be gateways" bit would be nice to clarify explicitly, it feels like a bit implicit/allowed/by omission currently, refs https://github.com/kubernetes-sigs/gateway-api/pull/1843#discussion_r11426 47596
- [lauralorenz] cluster ID kep cleanups:
 - Allowing clusterid to have "one dot" as a subdomain for extra disambiguation (ie location disambiguation); ready for comments or lgtm: https://github.com/kubernetes/enhancements/pull/3918
 - Beta bump: waiting to rebase on PR covering the disambiguation of pod DNS, placeholder (with some language improvements) is in https://github.com/kubernetes/enhancements/pull/3652/files
 - PRR bump: asked for a PRR reviewer and filled out beta requirements in https://github.com/kubernetes/enhancements/pull/3917 if you like reading PRR reviews and want to give a chill Igtm on that too feel free to
- [krzykwas] PR for location-disambiguated DNS naming of Headless service pods: https://github.com/kubernetes/enhancements/pull/3918

2023-03-07

- [lauralorenz on behalf of robscott]
 - [SIG-NETWORK] Gateway API Interaction with Multi-Cluster Services
- [Krzysztof/lauralorenz] location disambiguation in multicluster DNS proposal
 [PUBLIC] location disambiguation in multicluster DNS
- [srampal] Go over the <u>multi-cluster networking slides from last time</u> with more time for discussion this time
- [mikemorris] more context/"why" for locality subsets for MCS DNS for headless services (as discussed in 1/24 meeting), how it relates to cluster subset targeting, is there a doc/proposal for this somewhere yet?
 - Answered in second agenda item!

2023-02-21

- [mikeshng, mikemorris, deads2k, jnpacker] gathers interests around topics such as cluster registry, multi-cluster control plane, multi-cluster controller, centralized broker, work-api, etc. Community doc.
 - Notes in doc, but pulling up a few:
 - Discussed (a) cluster registry and (b) multicluster control plane as related, possibly dependent projects
 - (c) Work API with its own open items, but also as a reference of (a) and (b) that would conform to a more general purpose standards-driven variation of them
 - Mutlicluster controllers has a direct action item to connect with what is already going on in controller runtime / operator SDK on this point
- [srampal] Some multicluster service networking model/ framework topics (network topologies, policy)
 - [mikemorris] GAMMA is looking at AdminNetworkPolicy as possible future point of AuthZ convergence (but no concrete work towards this yet)
 - [mikemorris] multi-cluster N/S load balancing with Gateway API? Perhaps make Gateway API routing ClusterID-aware for use in these applications?
- [lauralorenz] quick update: conformance tests
 - [PUBLIC] test plan updates for MCS api for conformance tests
 - https://github.com/kubernetes-sigs/mcs-api/tree/master/conformance
- [lauralorenz] quick update: sig mc website
 - Moved to https://github.com/kubernetes-sigs/sig-multicluster-site
 - Issues to work on are now in that repo

2023-02-07

Agenda postponed to 2023-02-21

2023-01-24

- [lauralorenz] sig mc website project update
 - pending official repo
 - Latest preview is at https://lauralorenz.github.io/sig-multicluster-site-proposal/
 - o current TODOs to clean it up before it hits the prime time are here
- [lauralorenz] conformance testing suite update
 - o Overall plan/intention described in
 - [PUBLIC] test plan updates for MCS api for conformance tests

- skitt@ and lauralorenz@ working on POCs for next meeting. Feel free to contact if you have some ideas or interest in this project
- [lauralorenz] MCS questions Re: voluntary api review
 - ServiceImport.Status.Clusters is this in use by any implementations now? I see its use potentially for endpoint ttl. I find the status fields not well defined in the spec
 - Possibly overspecified here? Under emphasized
 - Use conditions instead even though that will diversify across vendors but since the implementations have a lot of room on ServiceImport
 - Question of what end users need (ex ServiceExport.Status.Conditions) versus the consumer
 - [tom] Submariner uses the clusters field for the id of the cluster that is exporting that slice of the serviceimport, technically could embed an annotation
 - [mike] the non-aggregated ServiceImports with a single remote cluster listed in status actually is quite similar to Consul's "cluster peering" (which does not assume sameness) implementation displaying PeerName for imported services and might be interesting for modelling a "not-sameness" use of ServiceImport in the future
 - Per-cluster IPs vs clusterset-wide VIP?
 - o GKE does per-cluster VIP
 - Al: consider clarifying the use of per producing cluster metadata in the ServiceImport (right now we keep "all" of that in the EndpointSlice but it seems there is more use for it)
 - Could it be the ServiceImport → ServiceImportSlice!!
 Maybe not required
 - Room for region in headless multicluster DNS and generally on region as a disambiguating feature in multicluster
 - [jeremyot] location makes sense as something of this tier not necessarily region (zone, region, data center, rack, etc)
 - Trust boundaries across regions?
 - [mike morris] very interested in network topology this is attempting to model
 - As a sub-question, i thought cluster name was not supported in MCS DNS? Or is that different for headless?
 - Different for headless Laura :)
 - Some overlap with topology-aware routing in SIG-Network, multicluster doesn't meaningfully participate in that currently
 - Should user have routing configuration control over crossing more significant boundaries?
 - [stephen] scalability boundaries in data shared across a clusterset?
 - Options:
 - Do nothing

- Make another DNS slot for "location"
- Arbitrary number of sub-labels? Kinda scary!
- Sameness for sure still applies
- o [mike ng] after KubeCon US, talk of cross-cluster, multi-cluster controllers?
 - Community interest for sure, nothing concrete like a one-pager yet, needs an active champion to start a focused discussion

2022-11-29

- [lauralorenz/npintaux] update on sig mc contribex
 - Proposal for sig mc website: [FUBLIC] sig mc website proposal
 - Repo request (if the above is ok):
 https://github.com/kubernetes/org/issues/3841
 - o [npintaux] show demo site
- [skitt] Conformance tests
 - https://github.com/kubernetes-sigs/mcs-api/issues/23: ServiceExport v.
 ServiceImport (see also https://github.com/kubernetes-sigs/mcs-api/pull/25)
 - What behaviours to test (aka completing <u>https://github.com/kubernetes-sigs/mcs-api/issues/14</u>)
 - Submariner test suite: https://github.com/submariner-io/lighthouse/tree/devel/test/e2e/discovery
 - Let's not mutate the existing e2e test and instead focus on developing an actual conformance suite
 - Support testing across more than two clusters with appropriate tests on three or more clusters, plus checking on large numbers of clusters; also test the single-cluster case
 - First priority is still to update the e2e tests so that they actually match the KEP
 - The tests should make sense if you've only read the spec (not seen an implementation); every test should point to the section of the spec that it's testing
 - o Test plan in the spec: Kubernetes Multi-Cluster Services
 - Finer-grained tests so that implementers can see what is missing (and so the SIG can see what parts of the spec implementations aren't satisfying)
 - o Talk about the contents of the tests offline

2022-11-15

- [lauralorenz] MCS API e2e/CI deep dive and running conformance tests
 - Showed a demo on running the e2e test suite against:
 - 1. Kind clusters with the demo implementation of an MCS controller that is in the MCS API repo
 - 2. Any other clusters that have your MCS controller implementation configured against them

- Just switch your KUBECONFIG1 and KUBECONFIG2 environment variables to be the ones holding your two clusters' config
- API implementers: please use the above directions to try out the e2e tests against your implementation and let us know if you see issues with them as conformance tests!
- Feedback today on e2etests as conformance tests:
 - Using a different api group / parameterizing the GVR
 - Could use the unstructured client
 - Could make the GVR paramterized entirely –
 - Version of kubernetes
 - Valid to check the e2e against the version
 - Thinking that 1.21 is best for e2e conformance because EndpointSlice v1 is from there and it is significantly mature at this point
 - If there are parts of the implementation of the e2e tests that are an issue depending on the version of the control plane
 - MCS API minimum version required in the spec
- [pmorie/jeremyot] Kubefed
 - o 6 weeks from today will cut a tombstone commit for the repo
 - Reminder this is not deletion, just archival

2022-11-01

- [lauralorenz] quick kubecon recap / debrief
 - Clusterset to clusterset
 - Esp re: self service
 - MCS and GAMMA initiative / service mesh integration
 - [jeremyot] Coordination
 - Config and workloads into clusters
 - Enforcement policy
 - How to deal with mismatching k8s versions or CRDs being different (and how does this affect namespace sameness)
 - What does a multi-region control plane look like is it "actually" a multi-region cluster
 - Contribex
 - Tutorials for how to get started / lack of documentation
 - Not enough other topics besides networking

2022-10-28

• SIG Multicluster Intro & Deep Dive was streamed at Kubecon, video, slides

2022-10-24 (<u>special meeting</u> @ kubecon contributors' summit 11:15am ET / 15:15 UTC)

Add your topics!

- multicluster controllers and operators
 - [makkes, from slack] The only project in that space I have found is https://github.com/admiraltyio/multicluster-controller and that hasn't received any updates for 2 years.
 - [haosdent] https://github.com/karmada-io/karmada
 Our company use this to operate multiple clusters
 - [lauralorenz, from slack] there was some very brief discussion about this last year in sig-mc about what we should do generally [1] and a cameo by howardjohn around then too about leader election in k8s generally [2] but AFAIK there is no recent work to develop a framework or centralized standard yet, just the individual examples (besides kubefed and aws cloudmap already mentioned, i'd mention <u>submariner.io</u>, all service meshes generally, and one off multicluster operators like the k8ssandra operator [3]).
 - [1] https://docs.google.com/presentation/d/10wuprHoal4bu9KAS5fnSwcet rSiEsnTnnwhBx8CyTTQ/edit?usp=sharing
 - [2] https://youtu.be/xEN5XEqPbW8
 - [3] https://dok.community/blog/developing-a-multi-cluster-operator-with-the-e-k8ssandra-operator/
- Multicluster workload placement
 - o Work API
 - Having a <u>virtual office hour on Wednesday</u> as part of Kubecon
- multicluster statefulsets
 - StatefulSetSlice KEP
 - ^ in service of supporting migration from cluster to cluster, POC built on MCS was demo'd at SIG-MC on 7/26: demo on youtube
- MCS for multi network
 - Multi network for single cluster is in progress by SIG-Network
 - Regular sync: Multi-Network community sync Notes
 - Current work is a KEP draft, solidifying API schema:
 - KEP: Multi-Network Requirements
- Multi cluster network policy
 - [thockin, comments on this prior proposal] Another dimension to explore is relying on something like ServiceAccount (which could be more easily carried as metadata on the wire) to enable MCNP. I don't think we've done enough to explore that.
- [mikemorris] extending ServiceImport to different identity spaces? ("outside the clusterset")

- [??] Multicluster as an anti-goal; namespace isolation limits leading to multicluster but then they proliferate
 - o [jeremyot] cluster is important when you need to know about them
- MCS API could use contributions for e2e tests to get to beta
- Want to move past the current projects and start next wave

SIG-MC contributor experience mindmeld: what would make you contribute to SIG MC?

- Need a deep dive about the MC e2e / CI setup to get folks involved in those issues and on running those tests on alternate cloud providers
- Not immediately obvious as an end user on how to do multicluster things, getting a good test environment using multicluster features like MCS or a mesh
- Hard as an outside to tell what are parallel explorations / competing in ideas vs implementing agreed upon directions of the SIG +1 +1
 - Example of Gateway API on what a vendor who conforms, what would they have, and as an end user how to start
 - [jeremyot] particularly called out that workload placement is in the parallel exploration mode
- Non-networking / non-service mesh topics not clear

Coordination breakout notes TODO: Jeremy to add

Networking breakout notes

- GAMMA subproject / Gateway API
- Istio and multi cluster services whats up with that
 - [john] Idea originally was assuming sameness everywhere and was super automatic, now trying to reconcile with MCS and whether to migrate the sameness everywhere UX as opt in, with tooling to recreate it, or what
 - [mikemorris] overlap here with Consul; sameness everywhere is what doesn't work in larger organizations
 - [mikemorris] Interest in having teams opting into MCS
 - [bowei] MCS has the export concept and the stitch together metadata with import; main capability is expressing exports symmetrically. does ingress / gateway work for asymmetric
 - [lauralorenz] Exports can be asymmetric but then there is still a risk without another policy engine in between that a non-allowed cluster could export into
 - The conceit of MCS is that its simple because it makes a lot of assumptions i.e. namespace sameness
- Example problem: Ingress that is integrated with Clouds so they can have geographically dispersed DNS so they can go to different clusters
 - At least something that they can unify on that then end user companies can implement against
- Clusterset to clusterset, thinner trust line but can keep the clusterset as the atomic
 unit

- O Where should the clusterset end?
- [mikemorris] one option is "this group of humans is administering these clusters"
 - But if there isn't a central administration team / if there is a self service model then its difficult to use something like MC Admin Network Policy as the engine to enfore this
- Where is it different for workload placement vs MCS

•

2022-10-18

- [lauralorenz] Kubecon NA next week
 - Contributor summit brainstorm on Monday (<u>event</u>, register for contributor summit before Thursday <u>here</u>)
 - Still scheduling a hybrid meet & greet Tuesday during the Contributor Summit unconference
 - SIG-MC maintainer track update and Q&A is Friday (event)
 - Meet & Greets also being hosted on the Friday lunch time but don't think SIG-MC will host one then

2022-09-20

- [pmorie] follow-up: archival of kubefed
- [@pwschuurman] Cross cluster StatefulSet migration, <u>KEP-3335</u>

2022-08-09

• [pmorie] follow-up: discussion on archiving kubefed

2022-07-26

- [@pwschuurman] Cross cluster StatefulSet migration
 - o <u>Slides</u>
- [lauralorenz] id.k8s.io rename survey results
 - Least hated is `cluster.clusterset.k8s.io`, provisionally working with that for next round of API review and will publish data to the mailing list
- [lauralorenz] FYI Gateway API GAMMA <u>submeeting</u> will be discussing Gateway API for east/west traffic today at 3pm PT, early exploration <u>doc link is here</u>

2022-07-12

• [Zach Zhu(@zqzten)] willing to help maintain KubeFed, I can share our use case and some internal enhancements to everyone interested, and have a discussion on the future maintenance of the project

- [lauralorenz] About API api review discussion point: id.k8s.io vs mcs-id.k8s.io for ultimate mega clarity of what this is used for and future-proofing id.k8s.io ClusterProperty (more background)
 - Naming brainstorm details in <u>mailing list message here</u>

2022-06-28

 [lauralorenz] – I can't attend 6/28 so this may get bumped unless someone else can lead discussion: Follow up on https://kubernetes.slack.com/archives/C09R1PJR3/p1655184999186449

2022-06-14

- [Hongcai Ren(@RainbowMango)] Request <u>review the proposal</u> for adopting ClusterID from <u>Karmada project</u>.
- [sanskar] questions about sig-mc generally
 - What is the charter / role of SIG-MC especially in light of other third-party solutions in the space
 - o How much do we work with other SIGs

2022-05-31

• [Hongcai Ren(@RainbowMango)] presentation about how <u>Karmada</u> sync resource status. Continue to discuss <u>this work-api proposal</u>.

2022-05-17 (cancelled due to KubeCon EU)

2022-05-03

- [mikeshng] Work API resource status sync enhancement <u>proposal</u> presentation and reference implementation demo. <u>Slides</u>. Requested review for this <u>proposal</u> as well.
- [lauralorenz / ishmeetm] FYI About API roadshow to SIG-net and Cluster API this and next week
- [lauralorenz] API beta blockers status (now a doc!)
 - E2E test coverage Q for MCS API (bookmark)
 - DNS label/subdomain follow up (bookmark)

- [lauralorenz] demo of About API CRD
 - Shoutout again to IshmeetMeta@ for her excellent work on this!!!!!!
- [lauralorenz] next steps on About API (<u>slides</u>) especially re: CEL validation and controller implementation

2022-03-22

- [lauralorenz in absentia] I can't make today but FYI:
 - Please comment / edit on
 - [PUBLIC] k8s processes for multicluster support brainstorm
 - Ready to flip the bit on ClusterID KEP to `implementable` in PR#3251
 - Added a graduation criteria re: CRD CEL validation I will be asking the group about in the future
 - Didn't talk to network policy subproject yet about WG-multicluster-networking

2022-03-08

- [lauralorenz] KEP updates
 - ClusterID: Still negotiating on the API review (see <u>my Unicode</u> <u>manifesto/screed</u> if you want your eyes to burn) but close!, shout out to IshmeetMeta@ helping with the CRD implementation
 - Multicluster DNS: "just nits" from SIG-Network are in so waiting for lgtm, material changes from prior review cycles are summarised in this comment
- [lauralorenz] want to get thoughts on including anything multicluster in the official k8s production readiness review questionnaire now or in the future
 - o [lauralorenz] Make a brainstorming doc and email out
 - Update 3/21:
 - [PUBLIC] k8s processes for multicluster support brainstorm
- [lauralorenz] want to get a temperature on a WG between SIG-network and SIG-mc, with a larger number of related proposals recently it has come up in a few conversations I've had individually with people but not sure if it's a better system than we have now
 - Multi network
 - Network policy
 - [lauralorenz] Will go to network policy subproject meeting and see if that's the place to be for this
- [aymen] Any development roadmap for KubeFed?
 - Haven't heard anything recent, some projects were in flight 6 months ago
 - [pmorie] May need to archive since has been waiting for beta 2+ years
 - [liqian] If this project isn't the SIG-MC recommendation for this problem going forward, do we need to reinvigorate/reassess/provide a new rec?
- [liqian] What is the community recommendation for multicluster N/S traffic
 - Noticed that GKE supports MCS ServiceImports as a backend for Gateway API
 - [jeremyot] SIG-Network community was already describing the Gateway API as the solution for multicluster N/S

- Observe that there is an interesting matrix of optional upstream add ons multicluster depends on ie Gateway API + MCS; since CRDs are the future tbd on that problem getting bigger/needing a real solution
- SIG-MC top line problems to solve these days
 - Spreading resources around
 - [mike ng] Currently in progress in the form of Work API with recent PR that needs feedback
 - Multicluster controllers / leader election
 - o ?? may have missed one

2022-02-08

- [donaldh] Present some multicluster use cases and discuss adding capabilities to MCS to support them.
 - Copy of Supporting New Use Cases in MCS
- [jeremyot] APIs, next steps for graduation: Cluster ID and MCS API

2022-01-25

- [sanjeev] introduce early ideas on multi-cluster network policy and invite feedback https://docs.google.com/document/d/1sLq 1CwfdvhmqxslQRdt7-RjLFsZRJkd/edit#
- [andrews] KEP for AdminNetworkPolicy (previously ClusterNetworkPolicy): https://github.com/kubernetes/enhancements/pull/2522

2021-12-14

• [lauralorenz] MCS API beta blockers checkin (very festive slides)

2021-11-30

- [mikeshng] Demo of Work API and discussion about future plans. (repo)
- [lauralorenz] Updates (especially on KEPs)
 - o s/ClusterClaim/ClusterProperty in Cluster ID KEP (PR#3103) merged
 - PRR for Cluster ID KEP (<u>PR#3036</u>) merged
 - Updated grad criteria for MCS API based on convos a while ago (<u>PR#2821</u>) merged
 - Want to boost Jaromir's work on the <u>multicluster DNS plugin for CoreDNS</u>
 - This was a beta graduation blocker for MCS API

2021-11-16

2021-11-02

• [aattuluri] Review and final thoughts on MCS API changes to support multi-network scenarios. Doc.

2021-10-19

- [lauralorenz / robscott] Multicluster Topology Aware Hints, refresher and multicluster brainstorm (<u>slides</u>)
- [Jaromir Vanek] Short demo of multi-cluster DNS plugin for CoreDNS (code)

2021-09-21

- [lauralorenz] Multicluster DNS updates
 - Discussed on PR with SIG-Network, going for not requiring PTR records in multicluster case
 - Also discussed scale concerns due to possible large size of multicluster DNS records
 - Still looking for collaborators for a CoreDNS plugin as a reference implementation for MCS DNS -- if you are interested lmk!
- [lauralorenz] ClusterProperty
 - Last we talked 6/8, there was an effort underway as a spinoff from sig-arch to propose better CRD bootstrapping, which we considered a strong signal to be comfortable deciding ClusterPoperty should be a CRD (see "<u>To CRD or Not</u> <u>To CRD</u>" section in PR)
 - AFAIK that effort has gotten stuck
 - Planning on continuing the path of ClusterProperty should be a CRD even without any plans of improved bootstrapping, but opening the floor for comments on that
- Kubecon Talks are SOON!
 - Oct 13 11:55 am PT <u>SIG-Multicluster Intro & Deep Dive</u>
 - Oct 15 2:30 pm PT <u>Here Be Services: Beyond the Cluster Boundary with</u> <u>Multicluster Services</u>
 - Thank you again to everyone who joined us for our recording!

2021-08-24

 [lauralorenz/skitt] Anyone who wants to join as our virtual audience for our <u>KubeCon</u> talk recording, we are hosting next Tuesday 8/31 at this time

- o [add to cal]
- [lauralorenz] Multicluster DNS updates -- discussing PTR records on the PR with SIG-Network (slides)

2021-07-27

- [howardjohn] Discuss MC coordination
 - https://github.com/kubernetes/enhancements/issues/2835

0

2021-07-13

- [lauralorenz] Quick KEP updates
 - Multicluster DNS: addressed reviewer comments finally from way back from April, and recent comments from Miguel @ submariner.io
 - Next step: request approvals again
 - ClusterID: Catching up backlog of action items starting with <u>updating MCS</u>
 API grad criteria in relation to it back from 5/11 discussion
 - Next step: s/ClusterClaim/ClusterProperty/g, kubebuilder into github.com/sig-mc/about-api, API review request, circle back on CRD bootstrapping
- [lauralorenz] Multi-cluster controller coordination discussion from last time (<u>slides</u>, <u>the</u>
 <u>PR</u> under discussion)
 - Please share with us your use cases; even if they are all of the same pattern
 we want to be able to share this with any larger initiatives going forward. DM
 me, include in the notes, or add to the slides

2021-07-06

- [jeremyot] Multi-cluster controller coordination (e.g. https://github.com/kubernetes/kubernetes/pull/103442)
- [jeremyot] Change to bi-weekly?

2021-06-08 Agenda

- [skitt] ClusterSetIPs
 - How essential are they? (Beyond <u>the guarantees required of Service cluster IPs.</u>)
 - How are clusterset-wide ClusterSetIPs distinguished from clusterset-specific ClusterSetIPs? (On the face of it they shouldn't need to be.)
 - Should kube-proxy end up being responsible for implementing ClusterSetIPs, in the same way it's in charge of implementing virtual IP for Services other than Externalname? (Enforcing implementation, but not in MCS API implementers, in the K8s core.) But there are scalability issues with handling this in kube-proxy (it needs to know about all the endpoints).

- Ultimately we want to avoid *requiring* any functionality in kube-proxy that should be an implementation detail.
- What about aggregated service imports? Don't they end up requiring a common ClusterSetIP across all clusters?
- [lauralorenz] CRD bootstrapping updates/observations (slides)
- [anil attuluri] MCS with Multi-network setup (document)

2021-06-01 Agenda

- [lauralorenz] ClusterIDs do we want to allow them to be subdomains (aka have dots in them) or only strictly valid DNS labels? (this is an alpha blocker)
- [gabe] thinking more about services beyond clustersets: What needs to be the same across implementations? What guarantees do we want to provide to users?
 - "the world beyond my namespaces doesn't matter to me"
 - Proposed invariant: Given a set of clusters, and a set of namespaces, "the MCS invariant" holds for that set of clusters+namespaces when:
 - for any pod in any of those namespaces, the set of ServiceImports (multicluster services) reachable by that pod does not depend on which cluster the pod is deployed to
 - AND for any ServiceExport in any of those namespaces, the availability (reachability) of the exported Service does not depend on which cluster the Service is deployed to
 - Then a "clusterset" is a set of clusters where the MCS invariant holds for *all* namespaces on all the clusters
 - But an implementation may offer the ability to establish the MCS invariant for a user-defined subset of namespaces...
 - feedback from call
 - Q: what if I have ns admin on clusters A, B, C but not on D?
 - is my original use case actually satisfied / durable in this case?
 - what if I have NS admin on some other NS in cluster D?
 - possible out: kubernetes admin prevents discovery/connectivity from that other namesapce to the one I don't own on D?

2021-05-11 Agenda

- [lauralorenz] Discussion points deep-dive/follow up from last time's alpha->beta convo (slides)
- [gabe rosenhouse] Sharing services beyond ClusterSets

2021-05-04 (skipping for Kubecon EU)

2021-04-27 Agenda

- [lauralorenz] MCS API alpha->beta graduation update and discussion (slides)
- [lauralorenz] A close race between `about.k8s.io` and `self.k8s.io` but the former is the winner using non-random tiebreakers for IRV. In any case, I will use this data when going for formal API approval.

2021-04-13 Agenda

- [lauralorenz] DNS updates/discussion especially re: SRV records
 - The guestions I asked SIG-Network 2 weeks ago are on this slide
 - Known use cases of SRV records from SIG-Network
 - Active Directory (?? I think for clients to find an AD component called domain controller?)
 - VOIP, I think specifically <u>SIP protocol</u> (TL;DR: for this SIP server to find its next hop SIP server, and automatically decide what transport protocol e.g. TCP it should connect on)
 - etcd cluster bootstrapping (TL;DR: for an etcd node to find more/new etcd node friends in the cluster and get together)
 - Open question: are the existing cluster.local SRV records sufficient for these cases?
 - [jeremyot] sounds like we have enough use cases that we should keep SRV records in the multicluster DNS proposal
 - Reminder that the PR is open for comment here
- [lauralorenz] Quick call out that the <u>cluster ID apiGroup poll</u> is open until this Friday 4/16; we now have a <u>sigs repo</u> that we will change the name of based on that + API review
 - Reminder that the current PR is available here

2021-03-30 Agenda

- [lauralorenz] Cluster ID / DNS updates
 - Went to SIG-Arch last Thursday, we are proceeding with a sigs repo CRD
 - There is some work going on to provide a way to bootstrap CRDs on cluster start, that we might want to get into in the future
 - There is separate interest in reviving some other centralized projects (like meta.clusterName) for other use cases
 - All the notes from the discussion are on the <u>SIG-Architecture agenda</u>, and <u>here are the slides</u> I had shared with them
 - Going to SIG-Network this Thursday
 - Mainly expecting to ask about SRV records
 - Generally people I've talked to have been on the side of including them to keep parity with the existing spec
 - Pull request is now open with all of our feedback from 2 weeks ago included
- [jeremyot] Annual report: https://github.com/kubernetes/community/pull/5669

2021-03-16 Agenda

- [lauralorenz] Multicluster DNS 101 presentation (<u>slides</u>, <u>doc</u>)
- [skitt] Multicluster service selection (bringing https://github.com/submariner-io/enhancements/pull/2/files upstream) — do we need anything more at the MCS level than "implementations should do the right thing™ by default"?

2021-03-09 Agenda

- [lauralorenz] ClusterID KEP apiGroup name (<u>pull request</u>, <u>specific PR comment link</u>)
- [lauralorenz] Let's talk multicluster DNS (doc is at bit.ly/k8s-multicluster-dns)
- [pmorie] cluster-registry archival

2021-02-23 Agenda

- [pmorie] need to make a decision on cluster-registry repo

2021-02-09 Agenda

 [lauralorenz] Update and discussion on cluster ID KEP for implementable status (<u>slides</u>, <u>new pull request</u>)

2021-02-02 Agenda

- [lauralorenz] Revisit cluster ID release target re: in-tree vs sigs repo
- [lauralorenz] Revisit Cluster ID naming for real sheet
- [pmorie] opened ticket to make Work API repo (kubernetes-sigs/work-api)

2021-01-26 Agenda

- [jeremyot] Cluster ID release timeline <u>1.21</u>?
- [lauralorenz] Revisit Cluster ID naming sheet

2021-01-19 Agenda

- [pmorie] where to put artifacts / proposals related to work API? Doesn't look like it should be a kep.
- [ytsarev] k8qb demo (slides)

2021-01-12 Agenda

- [lauralorenz] ClusterID KEP touch base (slides, pull request direct link)
- [qiujian] work api update work API

2020-12-15 Agenda

[lauralorenz] ClusterID KEP PR revisit (slides, pull request direct link)

2020-12-08 Agenda

- [hectorj2f] Present the ROADMAP for kubefed to become beta.
- [kevin-wangzefeng] Kubernetes federation evolution demo and API proposal
- [mdelder] Share announcement of http://open-cluster-management.io/ community meeting for this Thursday

2020-12-01 Agenda

- [kevin-wangzefeng] kubernetes federation evolution
- [lauralorenz] cluster ID KEP PR

_

2020-11-10 Agenda

[jeremyot] <u>Cluster ID proposal</u> round 2 - KEP?

2020-10-27 Agenda

[jeremyot] <u>Cluster ID proposal</u>

2020-10-20 Agenda

- [jeremyot] Registry / Cluster ID use-cases, potential basic API

2020-10-13 Agenda

- [qiujian] work API design followup
- [hectorj2f] kubefed: handle old features deletion.

2020-10-06 Agenda

- [skitt] Use-case scrub, in particular ClusterSet which we haven't discussed much yet
 - Cluster registration
 - ClusterSet

- [skitt] Is the MCS API definition ready for external use (with the v1alpha1 caveats of course)?
 - The answer is yes, it's been ready since it was pushed to the sigs repo

2020-09-22 Agenda

- [andrewsykim] Discuss exporting Service Type=NodePort
- [vthapar] Demo of MCS API implementation with Submariner+Lighthouse
 - https://docs.google.com/presentation/d/e/2PACX-1vS0e2jBhIA0q
 Lo3STASFXLwKGKefspSHhbJJO2ksNsSq7Z5aluJwcBQMIQE2Mdi2g/pub

2020-09-08 Agenda

- [hermanbanken] https://github.com/Q42/mc-robot/: similar to ServiceExport, uses external sync mechanism (PubSub), exports NodePort services & creates Service with manual Endpoints pointing to NodePort of external cluster. Would love to learn about how this relates to existing tools out there (e.g. replacements) and what we can learn from each other.
 - Example ServiceSync resource:

 https://github.com/Q42/mc-robot/blob/master/deploy/examples/mc.q42.nl_v1_servicesync_cr.y

 aml
 - Example of actual in-cluster MC Robot status:
 https://gist.github.com/hermanbanken/71700d4f5d44a35654b82019c7204b12
- [qiujian] work api design:
 https://docs.google.com/document/d/1cWcdB40pGg3KS1eSyb9Q6SIRvWVI8dEjFp9
 RI0Gk0vg/edit

2020-08-25 Agenda

- [hermanbanken] https://github.com/Q42/mc-robot/: similar to ServiceExport, uses external sync mechanism (PubSub), exports NodePort services & creates Service with manual Endpoints pointing to NodePort of external cluster. Would love to learn about how this relates to existing tools out there (e.g. replacements) and what we can learn from each other. hermanbanken: Can't make it:(
- [gilesheron] Demo of an MCS implementation.
- [jeremyot] MCS API PR 1
- [pmorie] reminders:
 - public use case docs:
 - Cluster registration
 - <u>ClusterSet</u>
 - Work API draft document shared

2020-08-18 Agenda (cancelled)

2020-08-11 Agenda

- [richardmcsong] Loblaw multi-cluster model
- [pmorie] ClusterSet use case document
- [jeremyot] <u>sigs.k8s.io/mcs-api</u> <u>next steps</u> (DNS, Implementations)
 - o [andrewsykim] cluster API based implementation?
 - [andrewsykim] PR to assume clustersetIP is always set on ServiceImport https://github.com/kubernetes-sigs/mcs-api/pull/1
 - [andrewsykim] consideration of exporting Service Type=NodePorts

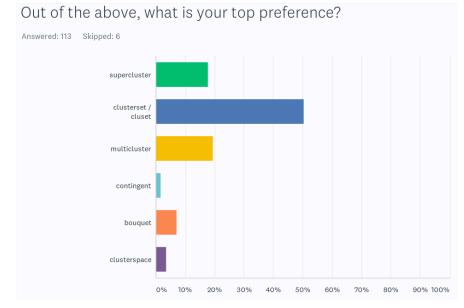
0

2020-08-04 Agenda

- [jberkus] Naming update
 - Survey Results:

*	LOVE IT 🔻	LIKE IT ▼	UNDECIDED ▼	HATE IT ▼	TOTAL ▼	WEIGHTED _ AVERAGE
▼ supercluster	14.66% 17	27.59% 32	21.55% 25	36.21% 42	116	2.21
	35.59% 42	36.44% 43	16.95% 20	11.02% 13	118	2.97
▼ multicluster	18.80% 22	41.03% 48	24.79% 29	15.38% 18	117	2.63
▼ contingent	2.65% 3	10.62% 12	26.55% 30	60.18% 68	113	1.56
▼ clusterspace	2.59% 3	18.10% 21	33.62% 39	45.69% 53	116	1.78
▼ bouquet	6.90% 8	4.31% 5	21.55% 25	67.24% 78	116	1.51

C



- ClusterSet is the choice
- [jeremyot] MCS-API alpha implementation https://github.com/jeremyOT/mcs-api

2020-07-28 Agenda

0

•

2020-07-21 Agenda

- [jeremyot] outcome: <u>Kube-proxy implementation strategy</u>
- [kevin-wangzefeng] Thoughts on Kubefed API

2020-07-14 Agenda

- [jeremyot] MCS status update
 - Kube-proxy implementation strategy
 - Try it out: https://github.com/jeremyOT/mcs-demo

2020-07-07 Agenda

- [jberkus/jeremyot] Naming survey results and next steps (survey round 2)
- [pmorie] https://github.com/kubernetes/org/issues/1499 follow up
- [hectorj2f] https://github.com/kubernetes-sigs/kubefed/pull/1238 and https://github.com/kubernetes-sigs/kubefed/pull/1238 and https://github.com/kubernetes-sigs/kubefed/pull/1238 and https://github.com/kubernetes-sigs/kubefed/pull/1238 and https://github.com/kubernetes-sigs/kubefed/pull/1237
- [kevin-wangzefeng] Thoughts on Kubefed API something like server-side "federate"?

2020-06-23 Agenda

- [jeremyot] Supercluster can we come up with a better name?
 - o Criteria:
 - Descriptive
 - Shorter is better
 - No overlap with related prior work
 - Not overly generic
 - Not negative connotation
- [vllry/jqiu] work API update

2020-06-16 Agenda (Cancelled)

2020-06-09 Agenda

- [jeremyot] Bringing back Cluster ID doc
- [jeremyot] Multicluster Services
 - https://github.com/kubernetes/enhancements/pull/1810
 - o Selective import
- [quinton] I see I'm still listed as a SIG lead.
 https://github.com/kubernetes/community/tree/master/sig-multicluster. Clearly I'm not actually doing that, and should be replaced. How would the SIG like to proceed with this?

2020-05-26 Agenda

- [hectorj2f] kubefed must be focused on resource federation. So, ServiceDNSRecord controller might not have a lot of future in future kubefed versions. Which are your thoughts?.
- [hectorj2f] Federated resource status: kubefed should be able to report the status of the deployed federated resources in addition to the propagation status. This would improve the troubleshooting of resources deployed across clusters.
- [pmorie] What to do about cluster-registry repo? (issue)
- [jeremyot] <u>Multicluster services KEP</u> (<u>Demo repo</u>)
 - https://github.com/kubernetes/enhancements/pull/1810

2020-05-12 Agenda

- [jeremyot] Multicluster services KEP (Demo repo)
 - Next steps
 - Approvers / reviewers

2020-04-28 Agenda

• [jeremyot] Multicluster services KEP (Demo repo)

2020-04-21 Agenda

Focus topic: [vallery] <u>Multicluster work API</u>

2020-04-14 Agenda

Focus topic: [jeremyot] <u>Multi-Cluster services API</u> (<u>Draft KEP</u>)

2020-04-07 Agenda

- [hectorj2f] Kubefed: proposal to offer a pull-reconciliation based approach and other improvements. [slides]
- [hectorj2f] Kubefed: Should we cut a new release this week?
- [vallery] Multicluster work API
- [jeremyot] follow up on namespace sameness
- [jeremyot] Multi-Cluster services API (Draft KEP) breakout meetings / form a WG?

2020-03-24 Agenda

- [pmorie] how to record "principles" that are not yet API?
- [hectorj2f] Metrics PR needs some reviews
- [jeremyot] Put Namespace Sameness wherever we decide above?
- [jeremyot] Discuss moving http://bit.ly/k8s-mc-svc-api-proposal to github

2020-02-25 Agenda

- [jimmidyson] Kubefed: Scalability testing and optimisation
 - o [hectorj2f] Metrics
 - o Controller runtime
 - Optimizing cluster and sync controllers
 - Access to cloud resources for scale testing?
- [jimmidyson] Kubefed: Restarting releases
- [jeremyot] Multi-Cluster Services API proposal http://bit.ly/k8s-mc-svc-api-proposal

2020-01-14 Agenda

- [jkitchener] Razee demo
 - Go to <u>razee.io</u> for more info. Link to slack channel from there and github repos

- Or reach out to @kitch on kube community slack
- [pmorie] Delete Kubefed WG meetings for now and fold into main SIG multicluster meeting?
- [hectorj2f] D2iQ would be interested in helping to maintain kubefed. I could also share our use case to everyone.
- [hectorj2f] #PR1171: I'd like to know if it could get merged into master even if it won't be released any time soon.

2019-12-17 Agenda

- [skitt] Submariner deployment demo using subctl
- [Josh] Review of <u>survey data</u>
- [jeremyot] Demo of multi-cluster service deployment with EndpointSlice / topology

2019-12-3 Agenda

- [pmorie] follow-up from Kubecon NA 2019
 - Use cases thrown out in meeting:
 - Inventory of clusters
 - API endpoints
 - Attributes
 - Location
 - Type
 - Registration into inventory
 - Registration vs. attestation to capabilities
 - Who can designate capabilities
 - Inventory of some set of resources running in union of clusters
 - Deployments
 - Statefulsets
 - Daemonsets
 - My CRD type here
 - Anything I annotate as being important in an aggregated view
 - Cluster to cluster operations
 - Services that span multiple clusters
 - Services that communicate between clusters (service A runs in cluster 1 and talks to service B in cluster 2)
 - Selection of the services (local vs remote) based on health or other rules.
 - Edge of cluster
 - Ingress into multiple clusters
 - App / workload deployment
 - Example: sharded DB
 - Placement policies for where workloads are deployed
 - Clusters are categorized/identified via a set of attributes

- Deployment of resources are not designated to specific clusters but rather cluster attributes
- Storage
 - Data migration from cluster A to B as workloads move
 - Designate a cluster as a primary/secondary in DR scheme, make sure data is available on secondary
- Service identity management
 - Management of identity in the face of heterogeneous auth provider
 - o Also, rbac
 - Network policy
- Future: Existing tools demo
 - o Chris Kim: Rancher Auth
 - Jake Kitchener: <u>razee.io</u> (multicluster inventory and deployment)
 - Skupper.io (who works on the project and can demo)?
 - Dan Berg can demo but not an expert by any means.
 - o mangelajo: L3 connectivity between clusters with submariner
 - Nice I've heard that submariner and skupper are similar but I don't know enough to speak to the differences.
- [mangelajo] multi cluster services: continuing the talk we had on the kubecon about how multi cluster services (discovery, behaviour, locality, endpoint slices, etc,) could behave in multi cluster environments.
 - Definition: A service that users in cluster A work in the same way as they would work with services running in cluster A, where the service may not be running in cluster A
 - Topology relation: pick a service in my same topological-region
- [pmorie] Tim's Multi-Cluster thoughts doc -https://docs.google.com/document/d/1G1lflukib7Fy_LpLUoHZPhcZ5T-w52D2YT9W1

 465dtY/edit#

2019-10-22 Agenda

[pmorie] Finalize <u>usage survey</u>

2019-08-13 Agenda

- [pmorie] is anyone using the cluster-registry?
- [mangelajo] <u>Layer 3 connectivity between clusters KEP</u>

2019-07-16 Agenda

- [adrianludwin] Discuss issues raised in <u>Cluster ID proposal</u>
 - What is the *immediate* driving usecase for this proposal, and how will this constrain the initial implementation (eg controller vs apiserver builtin)?

- The doc has longer-term usecases but nothing really concrete.
- Initial suggestion: audit logs?
- Bonus points if it could work with a CRD implementation, not an apiserver builtin.
- Should Clusters know their own Aliases?
 - If not, should we even raise the concept in the proposal?
 - If so, how should they be added and removed?
- Should Aliases never, rarely or frequently be removed?
- [Andrew Myhre] Hello from Capital One

2019-06-18 Agenda

- [majopela] Submariner proposal to become part of the multicluster-sig
 - https://groups.google.com/forum/#!topic/kubernetes-sig-multicluster/ly3HFyVz
 B3k
 - AI: miguel/livnat to check availability for presentation in next sig-network or another one-off timeslot

2019-06-04 Agenda

- [adrianludwin] Introduce Cluster ID proposal
 - o Doc here

2019-05-07: cancelled due to lack of agenda

2019-04-30 Agenda

- [pmorie] Paul to collate the pointers and prepare/update the draft for "Cluster Identity" discussion next week.
 - o Doc here
- Collect pointers for kubecon barcelona sig update.

2019-04-23 Agenda

- [pmorie] Cluster identity, personal notes from prior experience, what are the attributes we're looking for?
 - Presented some ideas around the background and some historical evidence of a similar attempt in k8s
 - Other potential stakeholders:
 - SIG Multicluster
 - SIG cluster lifecycle (and Cluster API inside that)

- Istio?
- Networking
- Service Catalog
- Problems this intends to solve [collection of inputs from attendees]
 - [csbell]Any automation for multiple clusters would need a durable coordinate, which cannot be IPs (can change), or certificates (can be rotated). An identifier which sticks forever.
 - Something globally unique id ideally should be sufficient, and not sure if it should also be immutable.
 - [Other questions] Should it be user friendly? Is it solely for automation?
 - [Bowei] The definition should ideally be derived from the use cases.
 - [Quinton] What do we define as a cluster?
 - Say by etcd, its nodes, etc. [but can be totally different over a period of time]
 - [Csbell] We can consider it as a kubernetes endpoint.
 - [Bowei] Can it be defined as a resource akin to other resources in a k8s cluster?
 - [Quinton] Shouldn't it be the instance of etcd (or an associated durable store) that is associated with the cluster?
 - [Adrian] Can differentiating a cluster be viewed as a different cluster if a particular resource created in one cannot be seen into another (say if a given cluster is restored from an etcd, it will still have "that" same resource).
 - In-progress definition could be identify a cluster from its apiserver (as store where a resource could persist).

Use Cases:

- Any automation which involves multiple clusters in need of a durable identifier.
- Uls or any other software which might need to interact with other clusters (kubeconfig inclusive).
- It might be beneficial to keep any other details (for example how to reach the cluster, an IP or endpoint) and keep only the cluster id in scope.
- An identifier needed for auditing and logging.
 - Which can mean some amount of immutability.
- It would be beneficial to have the identifier usable in conjunction with another name (say a subname in the dns name).
- Preventing human errors fail an operation in case an operation needs a unique id and is performed on the same id (or cluster).
- Paul to collate the pointers and prepare/update the draft for discussion next week.

2019-04-09 Agenda

- [Ariel Adam] We will follow up on the multi cluster networking topic we discussed on the 12th of March presenting the detailed problem domain and a proposed solution.
 We are looking for additional companies to join this effort and can enhance or modify things accordingly
- Presented decomposition of problem space:
 - Cross-cluster tunneling
 - Cross-cluster routing
 - Cross-cluster service discovery
 - o Cross-cluster network policy
- Additional problems identified:
 - Cluster identity (csbell)
 - Ingress-related (steve sloka)
 - Software frontend LB between Kube and Non-Kube endpoints (North-->South Traffic)
 - Trust model in the presence of multiple clusters with different operators
 - How to communicate the cluster-locality of different endpoints? Should an application in a pod be able to tell whether it is connecting to a local or remote cluster?

2019-03-12 Proposed Agenda

 [Ariel Adam] Our team at Red Hat have started working on the multi cluster networking problem (tunnels/VPNs, cross cluster routes, cross cluster service discovery, cross cluster network policies). We believe it's important we don't reinvent the wheel and we hope to collaborate with others. We'd like to take a few minutes and present what we are doing.

2019-02-26 Agenda

Agenda/Discussion

• [Itulloch] Asked to present at Kubernetes Meetup in Toronto March 6th and would like to represent the current state of Multicluster as best as possible! I would appreciate any feedback on slides for the presentation. Please let me know if there's anything in particular I should highlight (link for slides:

https://docs.google.com/presentation/d/1FTJPLq87tz2bSnYCKOsj0EbRtAA-CifRRdHBpl0iGlE/edit?usp=sharing --this can be commented on outside of the SIG meeting and if this remains the only agenda item, I don't think we need to hold the meeting just to talk about this)

2019-01-28 Agenda

• [Demo] @adrienjt wants to demo multicluster-scheduler, multicluster-service-account, and multicluster-controller which he recently open sourced.

https://admiralty.io/blog/using-admiralty-s-multicluster-scheduler-to-run-argo-workflows-across-kubernetes-clusters/

2018-12-18 Agenda

- [irfanurrehman/nikhiljindal] Consolidate and revisit the delta between the multicluster ingress API proposed at https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress/pull/212 and the services and ingress APIs implemented at https://github.com/kubernetes-sigs/federation-v2/blob/master/pkg/apis/multiclusterdns/v1alpha1/ingressdnsrecord_types.go
 - [nikhiljindal] Still waiting for reviews on the PR (<u>link</u>). Nothing to report as of now.

https://admiralty.io/blog/using-admiralty-s-multicluster-scheduler-to-run-argo-workflows-acros s-kubernetes-clusters/

2018-12-04 Agenda

Attendees

- Red Hat:
 - → Ivan Font?
 - → Paul Morie?
 - Maru Newby?
- Huawei:
 - → Shashidharatd
 - Irfanurrehman ✔
 - Quinton Hoole ✔
- Google:
 - Christain Bell 🗸
 - Nikhil Jindal

- Rohit Ramkumar
- Pivotal:
 - Gabe Rosenhouse?
- IBM:
 - Alexey Roytman?
- Enactive Networks:
 - o Greg Zuro
- NEC
 - Shubhendu Poothia

Agenda/Discussion

- [csbell, nikhiljindal] Update on MulticlusterIngress
 - Kubemci CLI being used by multiple users:
 https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress/issues/117
 - First stab at API: https://github.com/GoogleCloudPlatform/k8s-multicluster-ingress/pull/212
 - Shared with community for early feedback.
 - There are similar APIs implemented as MultiClusterDNS, serviceDNS and ingressDNS in federation V2. It makes sense to collaborate on the same.

2018-11-06 Agenda

- [irfanurrehman] status updates needed for kubecon China sig-multicluster intro and deep dive. [please update the slides here].
 - Federation Slides ready [folks can give feedback here].
 - Kubemci received from nikhiljindal.
 - o Cluster-registry Jonathan/Paul please update.
- [quinton] <u>csbell@google.com</u> mentioned that he plans to give an update on Multicluster Ingress in the coming months (as of Oct 9).
 - o This week or future??
- [pmorie] Added updates to the ClusterRegistry proposal. https://github.com/kubernetes/cluster-registry/pull/261

2018-10-23 Agenda

- [pmorie] SIG Charter follow-up
 - Initial charter PR has merged

- Should we refine in a follow-up
- [pmorie] Cluster connection API
 - Should we rescope this as "HealthCheckedCluster" ?
 - Consensus: no
 - Al: Paul to define API that content of secret should deserialize into
- [irfan] Collecting details for kubecon sig updates (China and NA).

2018-10-9 Agenda

- [pmorie] SIG Charter PR
- [pmorie] Cluster connection API

Cluster-api vs cluster registry. There seems to be some confusion about the difference - add to faq (quinton)

Federation-v1 vs v2. - add to faq (quinton).

2018-09-25 Agenda

• Cancelled due to lack of any agenda/discussion items.

2018-09-11 Agenda

- Red Hat:
 - o Ivan Font
 - Paul Morie
 - 0
- Huawei:
 - o Shashidharatd
 - o Irfanurrehman
- Google:
 - Christain Bell
 - Ray Colline
 - John Murray
- Pivotal:
 - o Gabe Rosenhouse
- IBM:
 - Alexey Roytman
- GKE Policy Management [John Murray & Ray Colline]
 - Overview of a recently announced tool for managing namespace configs and associated policies across multiple clusters from a central repository.
 - Overview presentation (Demo @ 17:04)

• [Paul] Took an item to update cluster registry to specify a kubeconfig like info mapped into the cluster-registry; yet to update that.

2018-08-28 Agenda

- Red Hat:
 - Ivan Font
 - Paul Morie

0

- Huawei:
 - Shashidharatd
 - Irfanurrehman
 - Quinton
- Google:
 - Jonathan
- [pmorie] New SIG charter PR
 - Keeping it minimal and getting an approval from steering-committee might be the first sensible thing to do.
- [perotinus] Moving clusterregistry. Cluster to beta; minor update to cluster registry.
 - Quinton: How do people use cluster registry without auth-info (which seems to currently be an unused field in cluster-registry).
 - Paul: The field is just to indicate what/how the auth should be done. Personally does not feel that cluster object need to have the auth info field.
 - Quinton: How does cluster registry become useful without the authentication information for a user who discovers the clusters in cluster registry?
 - Jonathan: The initial idea of auth field in cluster registry against a cluster was to provide enough info to the user to figure out how to get this auth info.
 - Quinton: Ideally we should be able to do away with the FederatedCluster resource in federation and enable cluster registry to provide more useful information rather than just the API endpoint of the cluster.
 - Need to have more discussions on this before further moving it to beta.

2018-08-14 Agenda/Meeting Notes

- Red Hat:
 - Ivan Font
 - Lindsey Tulloch
 - Maru Newby
 - Mathusan Selvarajah
- Huawei:
 - o Irfan Ur Rehman

- Shashidharatd
- IBM:
 - Alexey Roytman

Agenda/Notes

- As there was nothing preset on the Agenda and no new points there to discuss either, we did cancel the meeting for this week.
- The next meeting on 28th Aug (Tuesday), will stand cancelled if nothing appears on the agenda until 27th Aug (Monday). In that case the cancellation notice will be sent out on 27th.

2018-07-31 Agenda/Meeting Notes

- Red Hat:
 - o Paul Morie
 - o Lindsey Tulloch
 - Mathusan Selvarajah
- Huawei:

Agenda/Notes

Paul received ping from steering committee to move SIG charter to completion; will
cut a new PR from simplified template and post to mailing list.

2018-07-17 Agenda/Meeting Notes

- Google:
 - Ian Chakeres (ianchak@google.com, ianchakeres@github)
 - Jonathan MacMillan (perotinus@github)
 - Nikhil Jindal (nikhiliindal@google.com, nikhiliindal@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
- Huawei
 - o Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)

- Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Containership
 - Norman Joyner (<u>norman@containership.io</u>, normanjoyner@github)
- ??

Agenda/Notes

- Sig update needed at the community meeting
 - Main update for cluster registry would be CRD based work, some updates relating the namespacing of cluster resource.
 - Quinton took the responsibility to do the update.
- Kubemci update:

0

- Cluster registry update:
 - After moving to CRD, the validation for fields lacks as of now.
 - Cluster resource is namespaced.

0

- Federation V2 update:
 - We have been working on trying to put an alpha release out.
 - A release (v0.0.1) is out, which we call out at alpha.
 - We have some level of feature parity with federation v1 with some discrepancy.
 - There is a substantial change in the semantics of deployment and in federation V2 we use CRD based API Server removing the dependency of a separate etcd.
 - It is possible to federate resources (a simple federation) without writing code, including CRDs.

0

2018-07-03 Meeting Notes

- Google:
 - o lan Chakeres (<u>ianchak@google.com</u>, ianchakeres@github)
 - Jonathan MacMillan (perotinus@github)
 - Nikhil Jindal (<u>nikhiljindal@google.com</u>, nikhiljindal@github)
- Redhat:
 - Maru Newby (marun@redhat.com, marun@github)
 - Ivan Font (ifont@redhat.com, font@github)

- Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
- Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (shashidhara.td@huawei.com, shashidharatd@github)
- Containership
 - Norman Joyner (<u>norman@containership.io</u>, normanjoyner@github)

Agenda/Notes

- Kubemci update:
 - 0
- Cluster registry:
 - C
- Federation V2:
 - С

2018-06-19 Meeting Notes

- Google:
 - o lan Chakeres (ianchakeres@github)
 - Jonathan MacMillan (perotinus@github)
 - Nikhil Jindal (<u>nikhiljindal@google.com</u>, nikhiljindal@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
- Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (shashidhara.td@huawei.com, shashidharatd@qithub)

- Kubemci update:
 - Recently 0.4 was released, with some new features, docs and example. A blog post also up for GCP.
 - Good engagement from the community.
 - Exploring an active controller for multicluster ingress, using cluster-registry.
 - Need to brainstorm ideas wrt to possible controllers.

- Some discussions ongoing with networking sig, to explore possibilities of expanding the k8s ingress API to support multicluster use cases.
- Cluster registry:
 - o The move to CRD was smooth.
 - There hasn't been much of activity post that.
- Federation V2:
 - Transition to kube-builder, having some problems because of the CRD support for subresource is pretty restrictive in 1.10; looking forward to 1.11.
 - Most of the transition to crd is working. There are some issues, which need to be ironed out.
 - Irfan has completed work on scheduling type for deployment and replicaset, giving feature parity with federation V1, and shashi has completed similar feature parity wrt to the cross cluster service discovery.
 - Shashi would be giving out a demo of cross cluster service discovery in tomorrow's federation WG sync.

2018-06-05 Meeting Notes

- · Google:
 - Ian Chakeres (ianchak@google.com, ianchakeres@github)
 - Jonathan MacMillan (perotinus@github)
- Redhat:
 - Maru Newby (marun@redhat.com, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
- Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (shashidhara.td@huawei.com, shashidharatd@github)
- Samsung: Leah Petersen

- For the benefit of new contributors; A brief update planned on the three current sub-projects/efforts ongoing as part of this sig:
 - o Kubemci
 - Cluster-registry
 - Federation V2
- Research use-cases from CERN
- Status update:

- Cluster-registry needed defined ns reserved for specific reasons, proposals for the same are:
 - i. Global cluster resource (kube-multicluster-cluster)
 - ii. Multicluster component controllers and admin stuff (kube-multicluster-system)
- If there is no concern about this, these will be used up for functionality in further releases.
- Quinton: use case for propagating cluster list into individual clusters, federated ns?
- Jonathan: haven't considered, (or allowed) that use case as of now. It will however not very difficult to define namespaces which can be federated.
- Maru: there may be a desire to federate system namespaces in the future, but as a longer term goal. Something to come back to later, it may be the case that the multicluster-system namespace is always most useful as a non-federated NS.

2018-05-22 Meeting Notes

Attendees

- Google:
 - o lan Chakeres (<u>ianchak@google.com</u>, ianchakeres@github)
 - Jonathan MacMillan (perotinus@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (lttulloch@redhat.com, onyiny-ang@github)
- Huawei
 - o Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (shashidharatd@github)
- •
- Taylor Carpenter Vulk Coop (CNCF <u>Cross-cloud Cl</u>, CNFs)
- Samsung: Leah Petersen

Agenda/Notes

Kubecon retro:

- Update from Bob killen on scientific use cases relevant to federation and multi-cluster as his take from kubecon EU (volunteered to do so in federation WG sync, was not present in the last multi-cluster sync)
- Kubemci update from Google key people for this item were not present.
- Kubemci, any new federated Ingress API wrt Federated Ingress in federation V2, can some sync happen?
 - Shashi and lan can get in sync relating the same.
- Cluster Registry status update:
 - Moved clusters to CRD's
 - i. Reduced code and tooling burden.
 - Moved to namespaced cluster registries
 - Validation needs to be confirmed.
 - o Plan to move API to beta soon (before end June).
 - i. Might add simple status fields to status.
 - ii. Some auth-related changes to the API to support controllers vs end-user use.
- Federation v2 Status Update:
 - Low level API (mostly marun and ivan):
 - 3 main APIs are complete (template, placement, overrides, per higher level type). Also a push cluster propagator (marun), configured by FederatedTypeConfig.
 - ii. Work on federating CRD's (sync only) marun and ivan -
 - iii. Next steps:
 - 1. Move control plane to CRD's
 - iv. Target date for Federation v2 Alpha: tbd, hopefully end Jun/early July.
 - Higher level APIs
 - i. Scheduled types (deployments and replicasets) alpha impl done irfan.
 - ii. Federated services (shashi).
- New contributors:
 - o CERN
 - i. Quinton to follow up on details with CERN
 - Cisco
 - i. Particular interest in ingress implementation for multicluster.
 - Also, expressed the interest in large scale scalability test (100+ clusters) for federation.
 - Github
 - i. Quinton to follow up.
 - Upbound
 - i. Quinton to follow up.
- SIG meeting videos
 - o Jonathan having trouble uploading videos. Help needed.

2018-05-8 Meeting Notes

Attendees

- Google:
 - o Christian Bell (csbell@google.com, csbell@github)
 - Hievda Ugur (<u>hugur@google.com</u>)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Ian Chakeres (ianchakeres@github)
- Huawei
 - o Irfan Ur Rehman (irfan.rehman@huawei.com, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
 - Paul Morie (<u>pmorie@redhat.com</u>, pmorie@github)
 - Scott Collier (scollier@redhat.com, scollier@github
- Containership
 - Norman Joyner (<u>norman@containership.io</u>, normanjoyner@github)
- Heptio
 - Fabio Yeon (<u>fabio@heptio.com</u>, fabioy@github)
- Cisco
 - Megan O'Keefe (<u>meokeefe@cisco.com</u>, m-okeefe@github)
 - Mike Napolitano (<u>mnapolit@cisco.com</u>, mikeynap@github)

- Kubecon retro:
 - Operator framework
 - Feedback on V2:
 - There were not many questions about the state and future maturity of V1.
 - ii. At Kubecon, we have users and developers.
 - Multi-cluster session: well-attended. People were interested and had many questions.
 - Cluster Administrator level: which one is more important?
 - i. Cluster config

- ii. RBAC API
- Istio may require changes in API. Christian will reach out to them to gather the requirements.

2018-04-24 Meeting Notes

Attendees

- Google:
 - Jonathan MacMillan (<u>dvorak@google.com</u>, perotinus@github)
 - o Christian Bell (csbell@google.com, csbell@github)
 - Hievda Ugur (<u>hugur@google.com</u>)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Nikhil Jindal (nikhiljindal@github)
 - → Matt Delio (mdelio@github)
- Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
 - Paul Morie (<u>pmorie@redhat.com</u>, pmorie@github)
 - Scott Collier (scollier@redhat.com, scollier@github
- Containership
 - Norman Joyner (<u>norman@containership.io</u>, normanjoyner@github)
- Heptio
 - Fabio Yeon (<u>fabio@heptio.com</u>, fabioy@github)

- Resolve open comments for Kubecon EU <u>Slidedeck</u>
 - Quinton will update the mission and the slide after that.
- [pmorie] Updates re: federation-v2:
 - Repo donation rules to kubernetes-sigs: https://github.com/kubernetes/community/pull/2049
 - kubernetes-sigs/federation-v2 is now set up!
 - o A basic integration test and some verification jobs are running in travis now.
 - A set of steps is available to run e2e in minikube.
- Status of <u>SIG charter</u>?
- [Cluster registry update] Should it be cluster scoped or namespaced.

- Concluded that it would be namespaced, but with a special meaning for namespace in clusterregistry.
- The development speed should be much faster as of now.
- Still to conclude as to who would be taking up the work of migrating to CRD.
- [pmorie] any possibility to move this meeting to either Monday or Thursday?
 - o TBD.
- Paul will try to update the sig charter using the <u>guidelines listed by the steering</u> committee.

2018-04-10 Meeting Notes

Attendees

- Google:
 - Jonathan MacMillan (<u>dvorak@google.com</u>, perotinus@github)
 - Christian Bell (<u>csbell@google.com</u>, csbell@github)
 - Hievda Ugur (<u>hugur@google.com</u>)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Nikhil Jindal (nikhiljindal@github)
 - → Matt Delio (mdelio@github)
- Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
 - Paul Morie (pmorie@redhat.com, pmorie@github)
 - Scott Collier (scollier@redhat.com, scollier@github
- Containership
 - Norman Joyner (<u>norman@containership.io</u>, normanjoyner@github)
- Heptio
 - Fabio Yeon (<u>fabio@heptio.com</u>, fabioy@github)

- <u>SIG charter</u> discussion (Paul Morie)
- [Quinton] We should prepare a ~1 year roadmap in the next few weeks. See <u>SIG-PM</u> for details. Suggest carving up as follows:
- We did need to have a roadmap by 4th of april as per the mail communication from sig-pm, we are already overdue.
 - o Federation: Irfan, Paul Morie, Maru

i.

- Cluster Registry: Jonathan, ??
 - A 12 month roadmap for the same would be something around the ecosystem of cluster registry. Cluster registry itself would ideally not be in development that long.
- KubeMCI: Nikhil, Greg
- o ??
- Cluster registry API discussion (Jonathan)
 - Cluster registry API [namespace vs cluster scope](<u>link</u>) (Presented by jonathan)
 - The principle goal is solving customer problems
 - We should schedule a follow up meeting to discuss this further. Assigned to dvorak
 - Figure out who owners, stakeholders, opinion holders are
 - o Paul is focusing on OpenShift.
- Kubecon EU Slidedeck available in draft form. needs input!
 - Homework for everyone: add your input and roadmap slide. The roadmap slide can be shared with the PMs.

2018-03-27 Meeting Notes

Attendees

- Google:
 - Jonathan MacMillan (<u>dvorak@google.com</u>, perotinus@github)
 - Christian Bell (<u>csbell@google.com</u>, <u>csbell@github</u>)
 - Nikhil Jindal (nikhiljindal@github)
 - Mike Rubin (matchstick@github)
 - → Greg Harmon (G-Harmon@github)
 - Matt Delio (mdelio@github)
 - Daniel Nardo (dnardo@github)
 - Nilesh Junnarkar (njunnark@github)

◆ Huawei

- Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
- Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
- Shashidhara TD (shashidhara.td@huawei.com, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>ltulloch@redhat.com</u>, onyiny-ang@github)
 - Paul Morie (<u>pmorie@redhat.com</u>, pmorie@github)
 - Scott Collier (scollier@redhat.com, scollier@github
- VMware:
 - Tim Pepper (<u>tpepper@vmware.com</u>, tpepper@github)

- Pivotal
 - Amey ?
- Other
 - Nikhil Chandrapa

Agenda/Notes:

1. Agenda:

- a. Helm v3 Presentation (Matt Farina)
 - i. Followup discussion with Federation WG seems to be useful.
 - ii. Helm v3 design proposal
- b. SIG-ContribX Q1 Road Show (Tim Pepper)
- c. Fnord workload prototype demo (Ivan Font)
- d. [Quinton] Heads-up regarding <u>SIG charter refresh</u> and <u>SIG repos</u>.
 - i. Volunteers to cut required PR's to refresh SIG charter according to new template? Irfan volunteered. Thanks!
- e. Some issues uploading meeting recordings to YouTube. Any volunteers to debug?
- f. Cluster registry is moving to CRD's.

2018-03-13 Meeting Notes

Attendees

- Google:
 - Christian Bell (<u>csbell@google.com</u>, csbell@github)
 - Nikhil Jindal (nikhiljindal@github)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Matt Delio (mdelio@github)
 - Daniel Nardo (dnardo@github)
 - Nilesh Junnarkar (njunnark@github)

→ Huawei

- Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
- Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
- Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (<u>Itulloch@redhat.com</u>, onyiny-ang@github)
 - Scott Collier (<u>scollier@redhat.com</u>, scollier@github

Agenda/Notes:

2. Agenda:

- a. Maru: Does it make good sense to do away with the complexity of deploying cluster registry the way it is being done right now and use CRDs instead to install cluster registry to an existing k8s cluster.
 - Jonathan/Ivan are evaluating doing the same. Jonathan explained that this was not done earlier, because the CRDs missed some functionality needed by cluster-registry.
- b. Shashi: Shall we make the federated ingress implementation a little generic (which seems quite closely implemented keeping GCE in mind), to something like federated ingress equals only reconciling the ingress objects to member clusters.
 - Quinton mentions that the whole concept of ingress has been under discussion and it might make sense to evaluate reworking on the implementation of the same.
- c. Megan: (Cisco) Beginning to evaluate load balancers outside the clusters and outside the cloud providers. Quinton/Shashi/Megan/Redhat might be interested in talking about the same in a different meeting.

d.

- 3. Cluster Registry updates and/or blockers (if any).
 - a. Updates if any.
- 4. Federation updates and/or blockers (if any).
 - a. Updates if any.
- 5. Federation V2 API WG discussions.
 - a. The meeting notes for these discussions are available here.
- 6. Others:
- 7. Expectations by next meeting.

2018-02-27 Meeting Notes

Attendees

- Google:
 - Christian Bell (<u>csbell@google.com</u>, csbell@github)
 - Nikhil Jindal (nikhiljindal@github)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Matt Delio (mdelio@github)
 - Daniel Nardo (dnardo@github)
 - Nilesh Junnarkar (njunnark@github)

◆—Huawei

- Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
- Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
- Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (ltulloch@redhat.com, onyiny-ang@github)
 - Scott Collier (<u>scollier@redhat.com</u>, scollier@github
- Diamanti
 - Chakri Nelluri (<u>chakri@diamanti.com</u> chakri-nelluri@github)
- Samsung
 - Leah Petersen (<u>I.petersen@samsung.com</u>, leahnp@github)
- Heptio
 - Fabio Yeon (<u>fabio@heptio.com</u>, fabioy@github)
- AWS
 - Arun Gupta (<u>arun.qupta@gmail.com</u>, arun-gupta@github)

Agenda/Notes:

8. Agenda:

a. Short update on the progress on federation API discussions (if needed).

9. Cluster Registry updates and/or blockers (if any).

a. One of the outstanding major issue is if cluster resource should be a namespaced resource, work ongoing to resolve this.

- b. There has been talks to rename the cluster-registry to something more generic, discussions ongoing for the same.
- c. The cluster-registry repo will be split into multiple repos to ease vendoring the cluster-registry api server.

10. Federation updates and/or blockers (if any).

- a. We have been currently working on:
 - i. Kubefed functionality (RBAC, command config, etc..).
 - ii. Update vendoring mechanism and vendor latest code.
 - iii. Reported Issue fixes.

11. Federation V2 API WG discussions.

- a. The meeting notes for these discussions are available here.
- 12. Others:
- 13. Expectations by next meeting.

2018-02-13 Meeting Notes

Attendees

- Google:
 - Christian Bell (<u>csbell@google.com</u>, csbell@github)
 - Nikhil Jindal (nikhiljindal@github)
 - Mike Rubin (matchstick@github)
 - Greg Harmon (G-Harmon@github)
 - Matt Delio (mdelio@github)
 - Daniel Nardo (dnardo@github)
 - Nilesh Junnarkar (njunnark@github)

◆—Huawei

- o Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
- Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
- Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (lttulloch@redhat.com, onyiny-ang@github)
 - o Scott Collier (scollier@redhat.com, scollier@github
- Diamanti
 - Chakri Nelluri (<u>chakri@diamanti.com</u> chakri-nelluri@github)
- Samsung
 - Leah Petersen (<u>I.petersen@samsung.com</u>, leahnp@github)

- Heptio
 - Fabio Yeon (fabio@heptio.com, fabioy@github)
- AWS
 - Arun Gupta (<u>arun.qupta@gmail.com</u>, arun-gupta@github)

Agenda/Notes:

14. Agenda:

- a. Discuss Project Status doc circulated by csbell.
- 15. Cluster Registry updates and/or blockers (if any).
- 16. Federation updates and/or blockers (if any).
 - a. API versions, which specific api version should we expose in federation and what API version should we use to store in federation registry?
 - The actual state of the API does not really match the current GA state of many APIs.
 - ii. One option is that we can make this explicit in the documentation and
 - iii. Another option is to try and wrap the API exposed through federation into some version that is non GA.
 - b. Release and CI update.
 - c. Current status of Issues.

17. Federation V2 API WG discussions.

- a. The meeting notes for these discussions are available here.
- 18. Others:
- 19. Expectations by next meeting.

2018-01-30 Meeting Notes

Attendees

- Google:
 - Christian Bell (<u>csbell@google.com</u>, csbell@github)
 - Nikhil Jindal (nikhiljindal@github)
 - Jonathan (perotinus@github)
 - Mike Rubin (matchstick@github)

- Greg Harmon (G-Harmon@github)
- Matt Delio (mdelio@github)
- Daniel Nardo (dnardo@github)
- Nilesh Junnarkar (njunnark@github)

◆—Huawei

- o Irfan Ur Rehman (irfan.rehman@huawei.com, irfanurrehman@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (Itulloch@redhat.com, onyiny-ang@github)
 - Scott Collier (scollier@redhat.com, scollier@github)
- Diamanti
 - Chakri Nelluri (chakri-nelluri@github)
- Samsung
 - Leah Petersen (<u>I.petersen@samsung.com</u>, leahnp@github)

Agenda/Notes:

20. Agenda:

- a. (Nikhil): Should we communicate more broadly what all different solutions is the sig working on? There seems to be some confusion amongst people outside the sig based on recent slack and email discussions.
 - i. Sig members introduced themselves.
 - ii. Will start an email about scheduling an out of band meeting for providing updates to the sig and open it for Q&A (office hours)
 - iii. Christian will send out a doc with update on what all projects we are working on. It will eventually go as a markdown in community meeting.

21. Cluster Registry updates and/or blockers (if any).

- a. V.0.0.3 release cut out for cluster registry.
- b. Jonathan evaluating a testing plan, to find out what might be best strategy to test out cluster-registry.
- c. Also working out on a better strategy to enable other projects to vendor in cluster registry.

22. Federation updates and/or blockers (if any).

- a. Some issue fixes and minor feature changes have happened in previous few weeks.
- b. Effort ongoing to vendor in the latest k8s.
- c. Some tasks ongoing to complete for federation 1.9.0 release.

23. Federation V2 API WG discussions.

- a. Discussions happening to refine and get consensus of the concepts and the problem space.
- b. The meeting notes for these discussions are available here.

24. Others:

25. Expectations by next meeting.

- a. Csbell will circulate a document to summarise the efforts that are ongoing as part of sig-multicluster.
- b. Complete documentation update to give out the correct status and pointers to federation release and binaries (Irfan).

26. Last two weeks in Multi-cluster Summary (Top 2-4 topics, one sentence, one link)

- a. This will be distributed to: http://lwkd.info/ The goal is to provide a high level update so others can follow what the SIG is doing. More high level than these meeting minutes.
 - i. Cluster registry updates (scollier will summarize)
 - ii. Office hours (need details about when this is going to happen)
 - iii. TBD after Christian summarizes his notes.

2018-01-16 Meeting Notes

Attendees

- Google:
 - Christian Bell (csbell@google.com, csbell@github)
 - Jonathan MacMillan (<u>dvorak@google.com</u>, perotinus@github)
 - Greg Harmon (gharm@google.com, G-Harmon@github)
 - o Gautam Nirodi (qnirodi@github) Istio
 - Rafal Gajdulewicz (<u>gajduler@google.com</u>, rafax@github)

◆—Huawei

- Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
- Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
- Shashidhara TD (<u>shashidhara.td@huawei.com</u>, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Paul Morie (<u>pmorie@redhat.com</u>, pmorie@github)
 - Ivan Font (<u>ifont@redhat.com</u>, font@github)
 - Lindsey Tulloch (lttulloch@redhat.com, onyiny-ang@github)
- Concur
 - Dan Wilson (<u>danw@concur.com</u>, emaildanwilson@github)
- CoreOS
 - Ryan Phillips (<u>ryan.phillips@coreos.com</u> rphillips@github)

- Eric Chiang (<u>eric.chiang@coreos.com</u> ericchiang@github)
- Microsoft
 - Erik St. Martin (erikstmartin@microsoft.com erikstmartin@github)
- Heptio
 - Matt Moyer (<u>moyer@heptio.com</u> mattmoyer@github)
- VMware
 - Tim Pepper (<u>tpepper@vmware.com</u> tpepper@github)
 - Jared Rosoff (<u>irosoff@vmware.com</u> jsr@github)
- Diamanti
 - Chakri Nelluri (chakri@diamanti.com chakri-nelluri@github)
- GitHub
 - Aaron Brown (<u>aaronbbrown@github.com</u>, aaronbbrown@github)

Agenda/Notes:

27. Cluster Registry updates and/or blockers (if any).

- a. V0.0.2 was released.
- b. Items needed to take the API to beta are being discussed.
- c. Milestones are already clearly defined.
- d. Pauls update- cluster registry API is cluster scoped; evaluating to change it to be namespace scoped.

28. Federation updates and/or blockers (if any).

- a. Concluded on v1.9.0 as the version name for this federation release mapping to k8s v1.9.0.
- b. Current federation latest works with k8s 1.9.0+ until k8s latest.
- c. Federation however vendors in a little older version of k8s and k8s/staging libs (is this relevant?).
- d. Is backward compatibility needed by any of the users? Question raised again in last federation WG sync here.
 - Quinton We need to have at least some amount of backward compatibility is needed if we want to support upgrades.
 - ii. Forward compatibility can solve the upgrade scenario to some extent.
- e. Any sharing needed for api related work (or probably in federation WG sync).

i.

f. ...

29. Slot (15 mins) for ISTIO.

- a. Istio team wants to roll out support for k8s multicluster.
- b. Would want to adopt the cluster-registry and the actual usage of this registry (as of now) will be used for k8s clusters only.

- c. Would, if need be use annotations to extend (anything that might be needed) on cluster registry api.
- d. Gautam shared some details about how they want to do cross cluster service discovery.
- e. Design doc for Istio Mesh Expansion (Multi-cluster) https://goo.gl/AQARBs

30. Any relevant Kubecon Updates (pending, is this needed now?):

- a. Christian/Quinton...
- b. Two instances of sig updates happened. Some questions about where we are headed.
- c. Some details/indications about the use cases on k8s docs was one important point asked by community.
- d. Some discussions with oracle happened, they are trying for at least 2 projects for k8s incubation. They have 3 projects which expand or recycle multi-cluster concepts. One of their priorities was to move ahead faster compared to this sig.
- 31. Some multi-cluster use cases have been discussed in federation WG sync which are conducted every wednesday 9:30 PT.
 - a. The notes are available here.

32. Expectations by next meeting

a. ..

2018-01-02 Meeting Notes

Attendees

- Google:Christian Bell (<u>csbell@google.com</u>, csbell@github)
- Greg Harmon (gharm@google.com, G-Harmon@github)

О

- Jonathan MacMillan (<u>dvorak@google.com</u>, perotinus@github)
- → Huawei
 - Irfan Ur Rehman (<u>irfan.rehman@huawei.com</u>, irfanurrehman@github)
 - Quinton Hoole (quinton.hoole@huawei.com, quinton-hoole@github)
 - Shashidhara TD (shashidhara.td@huawei.com, shashidharatd@github)
- Redhat:
 - Maru Newby (<u>marun@redhat.com</u>, marun@github)
 - Paul Morie (pmorie@redhat.com, pmorie@github)
 - Ivan Font (ifont@redhat.com, font@github)
 - Lindsey Tulloch (Itulloch@redhat.com, onyiny-ang@github)

- Concur
 - Dan Wilson (danw@concur.com, emaildanwilson@github)
- CoreOS
 - Ryan Phillips (<u>ryan.phillips@coreos.com</u> rphillips@github)
 - Eric Chiang (eric.chiang@github)
- Microsoft
 - Erik St. Martin (erikstmartin@github)
- Heptio
 - Matt Moyer (<u>moyer@heptio.com</u> mattmoyer@github)

Agenda/Notes:

33. Federation release versioning.

- a. Suggestions are here, need to conclude on an agreeable versioning.
- b. We need to update (put placeholders) redirecting readers reaching the older paths in both documentation and k8s repo Irfan.

34. Cluster Registry Updates

- a. The <u>first release</u> was cut out.
- b. Next step for the cluster registry might be to try aggregating the use of the same with federation as a reference for similar projects.

35. Alternative for kubefed as helm/charts

- a. Is it more maintainable solution?
- b. Cluster registry, already faced a hurdle using helm/charts.
- c. The general consensus is that helm/charts is a better mechanism to maintain the fcp deployment, but only if they can support the ssl related stuff.
- d. More discussion can follow on "issue".

36. Slot (15 mins) for ISTIO.

a.

37. Any relevant Kubecon Updates (for those who did not attend):

a. Cristian/Quinton...

38. Expectations by next meeting

a. ..