



KubeCon



CloudNativeCon

Europe 2019

Koping with Change

What kops Learned Adopting
etcd3, cluster-api and CRDs
Justin Santa Barbara & Mike Splain

Introductions



KubeCon



CloudNativeCon

Europe 2019

justinsb@

kops maintainer

Started the project with a few others

Now @google

mikesplain@

kops maintainer

“New Wave” brings new ideas

Now @sonos

kops release - philosophy



KubeCon



CloudNativeCon

Europe 2019

kops releases **after** kubernetes releases

A kops .0 release is believed to be safe for production

Make sure kops works with new k8s!

Make sure k8s release stable: likely .2 or .3

Make sure the ecosystem is working (CNI providers etc)

kops release - reality



KubeCon

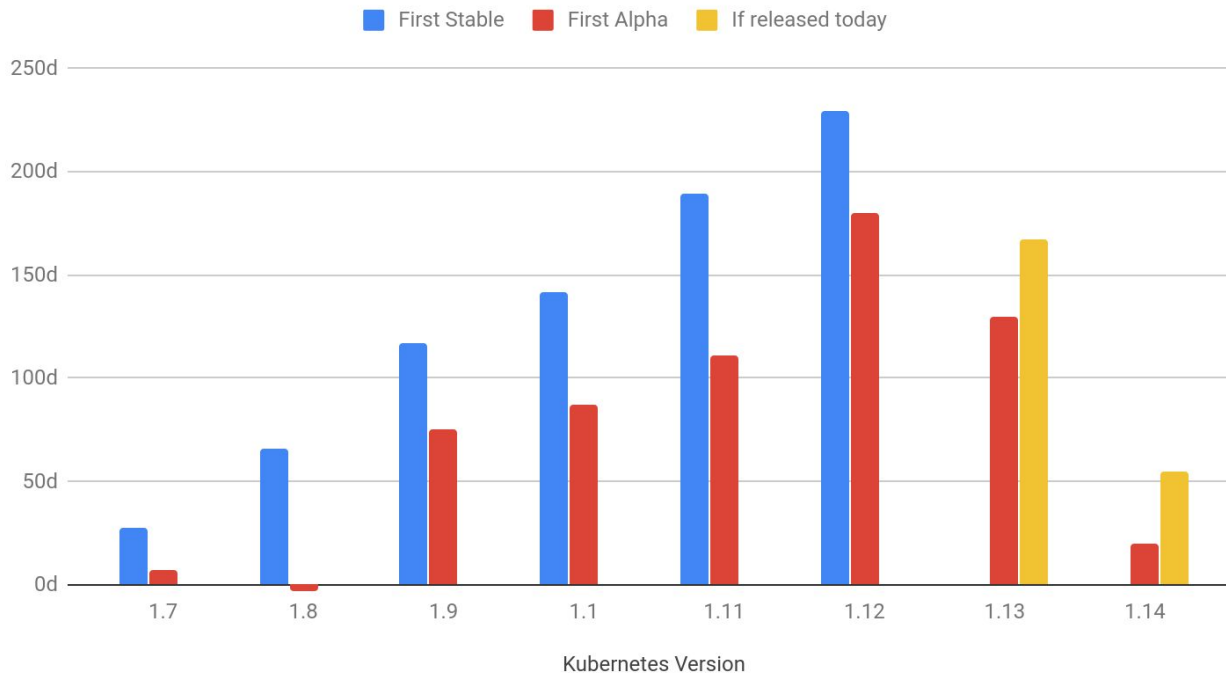


CloudNativeCon

Europe 2019

Thanks to: @peterr

Kops Release Delays



kops release - reality



KubeCon



CloudNativeCon

Europe 2019

Falling further and further behind

Hard for end-users who wanted to use the latest version

Head of line blocking: e.g. k8s 1.12 met release criteria before k8s 1.11 did

We also felt that we were always struggling to keep up

kops release - the fast-forward



KubeCon



CloudNativeCon

Europe 2019

Revised philosophy:

Users like that .0 is production-ready

But ... release alphas & betas earlier to give users a choice

The fast-forward:

Don't put a lot of features into kops 1.12, 1.13, 1.14.

(Lots) of features will land in 1.15

Cut the branches and start the alpha & beta train rolling!

kops release - the fast-forward



KubeCon



CloudNativeCon

Europe 2019

Lemons 🍋:

Slightly confusing releases on github

We have 1.12.1, 1.13 betas, 1.14 alphas etc

So many cherry-picks!

Lemonade:

Landing significant features in master branch

Feels less stressful and release branches feel more stable

(Maybe we will keep the master branch further in the future?)

kops release - the fast-forward



KubeCon



CloudNativeCon

Europe 2019

	A	B	C	D	E	F	G	H	I	J	K
1			1.11	1.12	1.13	1.14					
37	6748	3e81a3f47f Merge pull request #6748 from justinsb/release_process_push_over_ssh	-	-				docs			
38	6753	d1fc04d59e Merge pull request #6753 from granular-ryanbonham/fix_6751	-	#6755				LaunchTemplate fixes			
39	6762	a8e140db0d Merge pull request #6762 from bksteiny/iss6754	-	#6765				Fix AWS VPC IAM permissions			
40	6357	1ff341450b Merge pull request #6357 from sp-jose Luis-ledesma/kubelet/registry_qps	-	#6923				registry-qps kubelet flag		CP AFTER MERGE #6918	
41	6759	9661c57d2d2 Merge pull request #6759 from Vlatombe/patch-1	-	-				docs			
42	6706	b91db4f360 Merge pull request #6706 from granular-ryanbonham/apiserver_cpurequest	-	#6912				cpuRequest for API Server			
43	6772	3f8241ea1b Merge pull request #6772 from justinsb/release_1_14_0_alpha_1	-	-	-	-		docs			
44	6781	713ae16dcd Merge pull request #6781 from justinsb/channels_for_113_and_114	-	-	-	-		meta			
45	6780	2e6c19157f Merge pull request #6780 from justinsb/relnotes_1_12	-	-	-	-		docs			
46	6776	0b3db247ec Merge pull request #6776 from MathieuMailhos/fix_6770	-	#6919	#6793	#6812		Fix: set provisioned iops in terraform templates			
47	6768	ad2d9546b3 Merge pull request #6768 from adrianlyjak/pin-nvidia/6767	-	-	-	-		examples			
48	6769	ca7563d719 Merge pull request #6769 from alrs/pkg-dropped-errors	-	#6911	#6805	#6813		Fix: don't ignore error			
49	6789	fff5dd733 Merge pull request #6789 from xichengliudui/update-go-const-part1	-	?	-	-		Code cleanup			
50	6791	d517e0c58e Merge pull request #6791 from xichengliudui/update-go-const-part3	-	?	-	-		Code cleanup			
51	6766	f5dd4d8ece Merge pull request #6766 from mikesplain/rules_go_0.18.2	-	?	#6806	#6794		Updates bazel			
52	6774	c8163379c7 Merge pull request #6774 from bksteiny/etcd3-doc	-	-	-	-		docs			
53	6787	24edaf8bf1 Merge pull request #6787 from gambol99/http2_param	-	#6913	#6807	#6814		Maps http2 parameters			
54	6420	4d7df29ca5 Merge pull request #6420 from justinsb/rename_urls	-	?	#6816	#6980		Code refactor			
55	6802	5ba001f353 Merge pull request #6802 from xichengliudui/update-docker-README	-	-	-	-		docs			
56	6798	732121d226 Merge pull request #6798 from justinsb/golang_1_11_5	-	-	#6808	-		Update to go 1.11.5 for 1.13			
57	6799	f94d499920 Merge pull request #6799 from justinsb/golang_1_12_1	-	-	-	#6815		Update to go 1.12.1 for 1.14			
58	6809	e497f33104 Merge pull request #6809 from xichengliudui/update-go-const-part5	-	?				Code cleanup			
59	6810	bea86cb4d9 Merge pull request #6810 from xichengliudui/update-go-const-part6	-	?				Code cleanup			
60	6278	00782fd90 Merge pull request #6278 from elithrar/docs/state-store-easier	-	-	-	-		docs			
61	6820	a2ee29c75 Merge pull request #6820 from wangxy518/patch-5	-	-	-	-		docs			
62	6833	ab2413135 Merge pull request #6833 from chrisz100/master	-	-	-	-		channels			
63	6823	3208373d4 Merge pull request #6823 from appvia/canal-3.6.1	-	#6926	#6838	#6839		updates canal, allows setting mtu			
64	6836	0efdd22b8 Merge pull request #6836 from mikesplain/update_readme_chart	-	-	-	-		docs			
65	6840	0709200c3 Merge pull request #6840 from rifelpet/aws-iam-authenticator-api-typo	-	#6922	#6843	#6842		Fix typo in aws-iam-authenticator image field name			
66	6841	8b72574b1 Merge pull request #6841 from mikesplain/remove_verify_bazel_from_travis	-	-	Merged	#6846		bazel fix causing problems			
67	6837	78f5b2c06 Merge pull request #6837 from mikesplain/add_t3a_family	-	#6905	#6838	#6839		Adds t3a family			
68	6835	85293df09 Merge pull request #6835 from wxdao/master	-	#6914	#6929	#6930		AWS ap-east-1			



Dealing with External Change:

The Move to etcd3

kops & etcd3



KubeCon



CloudNativeCon

Europe 2019

Moving from etcd2 -> etcd3

Lemons 🍋

etcd2 and etcd3 are completely different systems
etcd2 does not support migration to etcd3

As a kubernetes community, we probably should have pushed back

As a kops developer, I probably should have just sent a PR to do migration

A big source of the increasing time delay

Not really something that is user-visible

kops & etcd3



KubeCon



CloudNativeCon

Europe 2019

Lemonade

- Through the worst of it (🙄) and can start adding “real features” again
- Developed etcd-manager, merging with etcdadm to form a community project
- Making sure that etcd management is available to all, no longer “just kops”
- Etcd-manager adds etcd backups
- Etcd-manager allows for resizing etcd clusters
- We turned on full etcd-TLS
- We isolated the etcd cluster from end-user applications (security)

API aggregation -> CRDs



KubeCon



CloudNativeCon

Europe 2019

Backing the wrong horse:

API Aggregation -> CRDs

API aggregation -> CRDs



KubeCon



CloudNativeCon

Europe 2019

Systemic change: Aggregated API server -> CRDs

Kops was an early adopter of aggregated apiservers

But ... we never got it to a releasable state ... we vendored a lot of k/k libraries, it never offered a stable API, was a ton of work to rebase

Aggregated apiservers not being improved; toolkit effectively deprecated

Big push on CRDs to bring them up to parity; CRDs are the future

API Aggregation harder for end users to install

Tight coupling with our docs system etc

API aggregation -> CRDs



KubeCon



CloudNativeCon

Europe 2019

Moving to kubebuilder & CRDs

Less & simpler code; a lot less dependencies

controller-runtime is a more stable API, and is intended to be user-facing

ahmet's [doc generation tooling](#) much easier than full api machinery

Much easier to install CRDs on a k8s cluster (no need for an etcd, less invasive permissions)



Getting things too right:

Cluster-API

Proposal: NodeSets for API-driven node management #16199

 **Closed** justinsb wants to merge 1 commit into `kubernetes:master` from `justinsb:nodesets` 

 Conversation 41

 Commits 1

 Checks 0

 Files changed 1



justinsb commented on Oct 24, 2015

Member



This proposal suggests implementing (as an experimental API extension) NodeSets, which will allow management of Nodes through the API.



bgrant0607 commented on Mar 31, 2016

Member



I agree with [@roberthailey](#). This is outside the scope of the core.

Cluster API



KubeCon



CloudNativeCon

Europe 2019

Lemons 🍋

It can be even harder when something comes along that is “closer to home”

kops InstanceGroups are very similar to cluster-api MachineDeployments

kops InstanceGroups are (today) richer than MachineDeployments

Cluster API



KubeCon



CloudNativeCon

Europe 2019

We can “shim” InstanceGroups to MachineDeployments

When an InstanceGroup CRD is created, we create an owned MachineDeployment

Lets users choose: use either InstanceGroups or MachineDeployments directly

Still early days for cluster-api, we'll see if wrapping MachineDeployments becomes common

Cluster API



KubeCon



CloudNativeCon

Europe 2019

Lemonade:

cluster-api has a pre-built server-side model

People porting cluster-api to lots of cloud providers

MachineSet is an object we are missing in kops

Better integration with autoscaler etc

Bundles & Operators



KubeCon



CloudNativeCon

Europe 2019

Be the change:

Addon Bundles & Operators

Addons - Bundles & Operators



KubeCon



CloudNativeCon

Europe 2019

Bundles of Addons

The next area - being developed (even more) actively

Currently kops addon manifests & versions are baked into a kops release

Bundles: controlled releases of addons without doing a kops release

Addons - Bundles & Operators



KubeCon



CloudNativeCon

Europe 2019

Addon operators

Controller: Lets us check the health / status of addons long after cluster-installation

Can orchestrate complicated updates

Being developed in sig-cluster-lifecycle addon subproject

kops as components




KubeCon



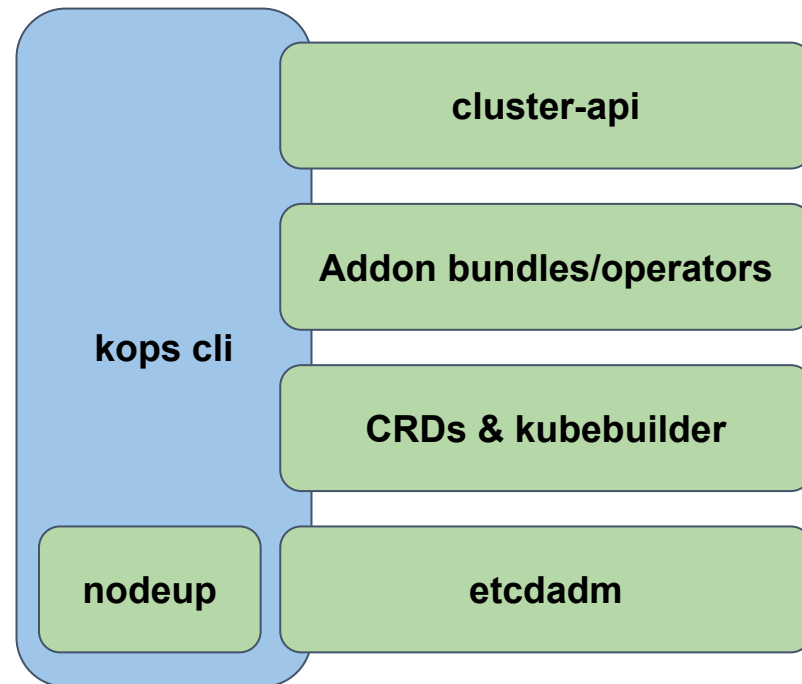
CloudNativeCon

Europe 2019

PAST



FUTURE



What we learned!



KubeCon



CloudNativeCon

Europe 2019

Summary

Kubernetes is changing very rapidly

Often for the better, but with some detours

Some changes are just work

Some changes duplicate something you've been doing - maybe provide a shim

Embrace new contributors - they are more in tune with the project now

Consider “jumping ahead” if you've been falling behind ... doing some “no feature” releases

Thank you / Join us!



KubeCon



CloudNativeCon

Europe 2019

Thank you to all our contributors, in github, slack or elsewhere!

Join us on Slack:

#kops-users

Or for development:

#kops-dev

Office hours:

Odd weeks

Friday 12 noon EDT / 9 AM PDT / 18:00 CEST

Deep dive: Kops

Tomorrow, Thursday, May 23 - 11:55 - 12:30





KubeCon



CloudNativeCon

Europe 2019

Demo



KubeCon



CloudNativeCon

Europe 2019

Demo of k115 branch with cluster-api integration, CRDs and bundles?