



KubeCon



CloudNativeCon

India 2024



Fantastic KEPs & Where to Find them —
A Guide to Kubernetes Enhancements!

Who We Are



Mario Jason Braganza
@jasonbraganza



Priyanka Saggu
@psaggu

Agenda

Agenda

- Introduction: what is a KEP?
- KEPs come into being!
- Demo



Hocce Hemisphaerium Firmamenti Sobiesciam Borealem
Diversis novis Asterismis, et quàm plurimis novis Stellis
exornatum atq; auctum, Invietsissimo IOHANNI III.
Poloniarum Regi etc. Dni Sui Clementissimò, ad immor-
talem Gloriã ac perpetuam Eius Memoriam, ob Heroicus
Suae Reg. Mitis Virtutes, et ob Res sub eo Hemisphaeris
fortiter, in maximum Commodum, et Incrementum
Christianae Religionis gestas, ac felicissime peractas,
Humillimo ac subiectissimo animo offert atq; consecrat
Ioh: Hevelius, Auctor, Observator, ac
delineator.

In lucem prodit Anno Æræ Christianæ 1680
Auctoris Sumptibus et Curâ,
Pictore verò Andr. Stech, et
Sculptore Carol de la Haye
Gedani.
Magnitudines Stellarum.

Fantastic Beasts!

Johannes Hevelius, Public domain, via Wikimedia Commons

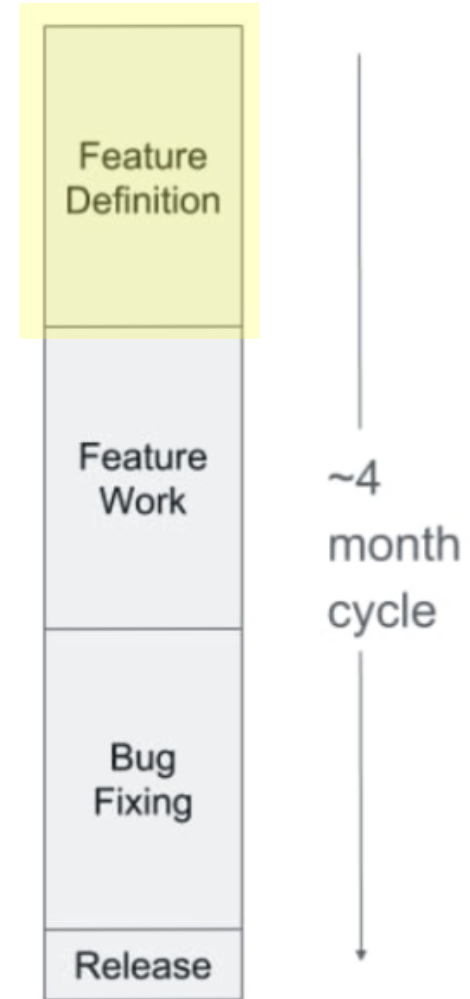


Fantastic KEEPS!

- We, Priyanka and Jason, have been involved in Kubernetes release cycles for years, taking on roles from shadow to Release Lead and Emeritus Advisor. While we've worked with Kubernetes Enhancement Proposals (KEPs) throughout this time, it wasn't until the 1.29 release, while prepping for press interviews, that we really dug into them—and wow, did we learn a lot!
- KEPs aren't just feature proposals; they give you a peek into the entire lifecycle of a feature—designs, alternatives, testing, and final code. They're a treasure trove for understanding Kubernetes, whether you're new or a seasoned contributor.
- We'll walk you through some real KEPs and share our approach to reading and making sense of them.
- If you're curious about Kubernetes and want to explore its features more deeply, join us. There's so much to discover, and we're excited to share what we've learned with you!

- Feature Discussion (per SIG): ongoing
 - Feature Freeze: week ~4


 - Release Branch Creation: week ~11
 - Code Slush / Freeze: week ~12
- ...bugs, testing, bugs, fixing, bugs, ...iterating
- End Code Freeze: week ~15
 - Release: week ~16



Not Just Feature Proposals

- Designs


- Alternatives



- Final Code

- Testing

Treasure Trove



Gus. Jarnet

Toutant

Walk Through a KEP



**Want to know more?
Hop in! And Let's Go!**



Benefits to the Eco-system!

- One of the biggest benefits of understanding Kubernetes Enhancement Proposals (KEPs) is that you get to see the entire lifecycle of a major change in the Kubernetes project.
- Whether it's a new feature, a deprecation, or a significant update, every impactful change in Kubernetes starts with a KEP. These documents track the progress of the enhancement across multiple releases, so you can follow along, understand the rationale behind decisions, and see how it all comes together in the code.
- This process isn't just for Kubernetes; it's a model that many other projects follow too. For example, Python has PEPs (Python Enhancement Proposals), and Go has its Design Proposals. By getting comfortable with KEPs, you're also building a skill set that's transferable to other major open-source projects, especially those under CNCF.
- Plus, by diving into KEPs, you can contribute more effectively. You'll know how to provide valuable feedback, suggest improvements, or even become an active part of the implementation. It's a fantastic way to get more involved in the community and make a real impact.

See the Lifecycle of a Major Change



KEP Process

Feature



Deprecation

Update

Any impactful change

Not just Kubernetes

dj



GO

Dive into Keps!

Contribute More
Effectively

Provide Feedback

Suggest Improvements

Active Part of the
Implementation



Make an Impact!

Let's Start at the Very Beginning

Note to Handout Readers:

**From here onwards, the video explains the text in the screenshots below.
Or please find detailed links to the source at the bottom of most slides**

Just WHAT is a KEP?

Guess What?

KEP-0000

That's documented in a KEP too!

 **wojtek-t** Migrate KEP process template to the new template 353137c · 3 years ago 🕒 History

Preview Code Blame 343 lines (276 loc) · 16.1 KB Raw 📄 ⬇️ ✎ ⋮

Kubernetes Enhancement Proposal Process

Table of Contents

- [Summary](#)
- [Motivation](#)
- [Stewardship](#)
- [Reference-Level Explanation](#)
 - [What Type of Work Should Be Tracked by a KEP](#)
 - [KEP Template](#)
 - [KEP Metadata](#)
 - [KEP Workflow](#)
 - [Git and GitHub Implementation](#)
 - [KEP Editor Role](#)
 - [Important Metrics](#)

master enhancements / keps / sig-architecture / 0000-kep-process / README.md ↑ Top

Preview Code Blame 343 lines (276 loc) · 16.1 KB Raw Copy Download Edit Menu

Summary

A standardized development process for Kubernetes is proposed, in order to:

- provide a common structure for proposing changes to Kubernetes
- ensure that the motivation for a change is clear
- allow for the enumeration of stability milestones and stability graduation criteria
- persist project information in a Version Control System (VCS) for future Kuber纳uts
- support the creation of *high-value, user-facing* information such as:
 - an overall project development roadmap
 - motivation for impactful user-facing changes
- reserve GitHub issues for tracking work in flight, instead of creating "umbrella" issues
- ensure community participants can successfully drive changes to completion across one or more releases while stakeholders are adequately represented throughout the process

This process is supported by a unit of work called a Kubernetes Enhancement Proposal, or KEP. A KEP attempts to combine aspects of

- a feature, and effort-tracking document
- a product requirements document
- a design document

into one file, which is created incrementally in collaboration with one or more Special Interest Groups (SIGs).

master enhancements / keps / sig-architecture / 0000-kep-process / README.md ↑ Top

Preview Code Blame 343 lines (276 loc) · 16.1 KB Raw Copy Download Edit Menu

Summary

A standardized development process for Kubernetes is proposed, in order to:

- provide a common structure for proposing changes to Kubernetes
- ensure that the motivation for a change is clear
- allow for the enumeration of stability milestones and stability graduation criteria
- persist project information in a Version Control System (VCS) for future Kuberernauts
- support the creation of *high-value, user-facing* information such as:
 - an overall project development roadmap
 - motivation for impactful user-facing changes
- reserve GitHub issues for tracking work in flight, instead of creating "umbrella" issues
- ensure community participants can successfully drive changes to completion across one or more releases while stakeholders are adequately represented throughout the process

This process is supported by a unit of work called a Kubernetes Enhancement Proposal, or KEP. A KEP attempts to combine aspects of

- a feature, and effort-tracking document
- a product requirements document
- a design document

into one file, which is created incrementally in collaboration with one or more Special Interest Groups (SIGs).

This process is supported by a unit of work called a **Kubernetes Enhancement Proposal**, or **KEP**. A KEP attempts to combine aspects of

- a feature, and effort-tracking document
- a product requirements document
- a design document

into one file, which is created incrementally in collaboration with one or more Special Interest Groups (SIGs).

The purpose of the KEP process is to reduce the amount of "tribal knowledge" in our community. By moving decisions from a smattering of mailing lists, video calls and hallway conversations into a well tracked artifact, this process aims to enhance communication and discoverability.

A KEP is broken into sections which can be merged into source control incrementally in order to support an iterative development process. An important goal of the KEP process is ensuring that the process for submitting the content contained in [design proposals](#) is both clear and efficient. The KEP process is intended to create high-quality, uniform design and implementation documents for SIGs to deliberate.

Stewardship

The following **DACI model** identifies the responsible parties for KEPs.

Workstream	Driver	Approver	Contributor	Informed
KEP Process Stewardship			SIG Leadership	Community
Enhancement Delivery	Enhancement Owner	SIG leadership (SIG Chairs + TLs)	Enhancement Implementer(s) (may overlap with Driver)	Community

Stewardship

The following DACI model identifies the responsible parties for KEPs.

Workstream	Driver	Approver	Contributor	Informed
KEP Process Stewardship			SIG Leadership	Community
Enhancement Delivery	Enhancement Owner	SIG leadership (SIG Chairs + TLs)	Enhancement Implementer(s) (may overlap with Driver)	Community

Stewardship

The following DACI model identifies the responsible parties for KEPs.

Workstream	Driver	Approver	Contributor	Informed
KEP Process Stewardship			SIG Leadership	Community
Enhancement Delivery	Enhancement Owner	SIG leadership (SIG Chairs + TLs)	Enhancement Implementer(s) (may overlap with Driver)	Community

Stewardship

The following DACI model identifies the responsible parties for KEPs.

Workstream	Driver	Approver	Contributor	Informed
KEP Process Stewardship			SIG Leadership	Community
Enhancement Delivery	Enhancement Owner	SIG leadership (SIG Chairs + TLs)	Enhancement Implementer(s) (may overlap with Driver)	Community

Stewardship

The following DACI model identifies the responsible parties for KEPs.

Workstream	Driver	Approver	Contributor	Informed
KEP Process Stewardship			SIG Leadership	Community
Enhancement Delivery	Enhancement Owner	SIG leadership (SIG Chairs + TLs)	Enhancement Implementer(s) (may overlap with Driver)	Community

What Type of Work Should Be Tracked by a KEP

What Type of Work Should Be Tracked by a KEP

The definition of what constitutes an "enhancement" is a foundational concern for the Kubernetes project. Roughly any Kubernetes user or operator facing enhancement should follow the KEP process. If an enhancement would be described in either written or verbal communication to anyone besides the KEP author or developer, then consider creating a KEP.

What Type of Work Should Be Tracked by a KEP

Similarly, any technical effort (refactoring, major architectural change) that will impact a large section of the development community should also be communicated widely. The KEP process is suited for this even if it will have zero impact on the typical user or operator.

What Type of Work Should Be Tracked by a KEP

KEPs will also be used to drive large changes that will cut across all parts of the project. These KEPs will be owned by SIG-architecture and should be seen as a way to communicate the most fundamental aspects of what Kubernetes is.

What Type of Work Should Be Tracked by a KEP

As the local bodies of governance, SIGs should have broad latitude in describing what constitutes an enhancement that should be tracked through the KEP process. SIGs may find it more helpful to enumerate what *does not* require a KEP, than what does. SIGs also have the freedom to customize the KEP template according to their SIG-specific concerns. For example, the KEP template used to track API changes will likely have different subsections than the template for proposing governance changes. However, as changes start impacting other SIGs or the larger developer community outside of a SIG, the KEP process should be used to coordinate and communicate.

KEP Template

The template for a KEP is precisely defined [here](#).



liggitt Update KEP template skew questions 8ef33fe · last year History

Name	Last commit message	Last commit date
..		
README.md	Update KEP template skew questions	last year
kep.yaml	Cleanup prr-approvers field from kep.yaml	2 years ago

README.md

```
KEP-NNNN: Your short, descriptive title
```

- [Release Signoff Checklist](#)
- [Summary](#)
- [Motivation](#)



master

enhancements / keps / NNNN-kep-template /

↑ Top

KEP-NNNN: Your short, descriptive title

- [Release Signoff Checklist](#)
- [Summary](#)
- [Motivation](#)
 - [Goals](#)
 - [Non-Goals](#)
- [Proposal](#)
 - [User Stories \(Optional\)](#)
 - [Story 1](#)
 - [Story 2](#)
 - [Notes/Constraints/Caveats \(Optional\)](#)
 - [Risks and Mitigations](#)
- [Design Details](#)
 - [Test Plan](#)
 - [Prerequisite testing updates](#)
 - [Unit tests](#)
 - [Integration tests](#)
 - [e2e tests](#)
 - [Graduation Criteria](#)



master

enhancements / keps / NNNN-kep-template /

↑ Top

- [Design Details](#)
 - [Test Plan](#)
 - [Prerequisite testing updates](#)
 - [Unit tests](#)
 - [Integration tests](#)
 - [e2e tests](#)
 - [Graduation Criteria](#)
 - [Upgrade / Downgrade Strategy](#)
 - [Version Skew Strategy](#)
- [Production Readiness Review Questionnaire](#)
 - [Feature Enablement and Rollback](#)
 - [Rollout, Upgrade and Rollback Planning](#)
 - [Monitoring Requirements](#)
 - [Dependencies](#)
 - [Scalability](#)
 - [Troubleshooting](#)
- [Implementation History](#)
- [Drawbacks](#)
- [Alternatives](#)
- [Infrastructure Needed \(Optional\)](#)

master enhancements / keps / NNNN-kep-template / README.md ↑ Top

Preview Code Blame 807 lines (618 loc) · 29.9 KB Raw Copy Download Edit View

```
60 -->
61 ✓ # KEPN: Your short, descriptive title
62
63 <!--
64 This is the title of your KEP. Keep it short, simple, and descriptive. A good
65 title can help communicate what the KEP is and should be considered as part of
66 any review.
67 -->
68
69 <!--
70 A table of contents is helpful for quickly jumping to sections of a KEP and for
71 highlighting any additional information provided beyond the standard KEP
72 template.
73
74 Ensure the TOC is wrapped with
75 <code>&lt;!-- toc --&rt;&lt;!-- /toc --&rt;</code>
76 tags, and then generate with `hack/update-toc.sh`.
77 -->
78
79 <!-- toc -->
80 - [Release Signoff Checklist](#release-signoff-checklist)
81 - [Summary](#summary)
82 - [Motivation](#motivation)
83   - [Goals](#goals)
84   - [Non-Goals](#non-goals)
85 - [Proposal](#proposal)
86   - [User Stories (Optional)](#user-stories-optional)
```

🔑 master | enhancements / keps / NNNN-kep-template / README.md | ↑ Top

Preview | **Code** | Blame | 807 lines (618 loc) · 29.9 KB | Raw | 📄 | ⬇️ | ✎ | ▾ | 🗨️

```
61 # KEP: Your short, descriptive title
248 ## Details
256
257 ▾ ### Test Plan
258
259 <!--
260 **Note:** *Not required until targeted at a release.*
261 The goal is to ensure that we don't accept enhancements with inadequate testing.
262
263 All code is expected to have adequate tests (eventually with coverage
264 expectations). Please adhere to the [Kubernetes testing guidelines][testing-guidelines]
265 when drafting this test plan.
266
267 [testing-guidelines]: https://git.k8s.io/community/contributors/devel/sig-testing/testing.md
268 -->
269
270 [ ] I/we understand the owners of the involved components may require updates to
271 existing tests to make this code solid enough prior to committing the changes necessary
272 to implement this enhancement.
273
274 ▾ ##### Prerequisite testing updates
275
276 <!--
277 Based on reviewers feedback describe what additional tests need to be added prior
278 implementing this enhancement to ensure the enhancements have also solid foundations.
279 -->
280
```


KEP Metadata

There is a place in each KEP for a **YAML document that has standard metadata**. This will be used to support tooling around filtering and display. It is also critical to clearly communicate the status of a KEP.

 **wojtek-t** Cleanup prr-approvers field from kep.yaml fb9265b · 2 years ago [History](#)

Code **Blame** 49 lines (43 loc) · 1.22 KB Raw 📄 📄 ✎ ⌵ ⌵

```
1 title: KEP Template
2 keps-number: NNNN
3 authors:
4   - "@jane.doe"
5 owning-sig: sig-xyz
6 participating-sigs:
7   - sig-aaa
8   - sig-bbb
9 status: provisional|implementable|implemented|deferred|rejected|withdrawn|replaced
10 creation-date: yyyy-mm-dd
11 reviewers:
12   - TBD
13   - "@alice.doe"
14 approvers:
15   - TBD
16   - "@oscar.doe"
17
18 see-also:
19   - "/k8s/sig-aaa/1234-we-heard-you-like-keps"
20   - "/k8s/sig-bbb/2345-everyone-gets-a-kep"
21 replaces:
22   - "/k8s/sig-ccc/3456-replaced-kep"
23
```

} *Driver*

 **wojtek-t** Cleanup prr-approvers field from kep.yaml fb9265b · 2 years ago 🕒 History

Code | **Blame** | 49 lines (43 loc) · 1.22 KB Raw 📄 ⬇️ ✎ ⌵ 🗨️

```
1 title: KEP Template
2 keps-number: NNNN
3 authors:
4   - "@jane.doe"
5 owning-sig: sig-xyz
6 participating-sigs:
7   - sig-aaa
8   - sig-bbb
9 status: provisional|implementable|implemented|deferred|rejected|withdrawn|replaced
10 creation-date: yyyy-mm-dd
11 reviewers:
12   - TBD
13   - "@alice.doe"
14 approvers:
15   - TBD
16   - "@oscar.doe"
17
18 see-also:
19   - "/k8s/sig-aaa/1234-we-heard-you-like-keps"
20   - "/k8s/sig-bbb/2345-everyone-gets-a-kep"
21 replaces:
22   - "/k8s/sig-ccc/3456-replaced-kep"
23
```

} Approver

 **wojtek-t** Cleanup prr-approvers field from kep.yaml fb9265b · 2 years ago 🕒 History

Code Blame 49 lines (43 loc) · 1.22 KB Raw 📄 ⬇️ ✎ ⌵ ⌂

```
1 title: KEP Template
2 keps-number: NNNN
3 authors:
4   - "@jane.doe"
5   owning-sig: sig-xyz
6   participating-sigs:
7     - sig-aaa
8     - sig-bbb
9   status: provisional|implementable|implemented|deferred|rejected|withdrawn|replaced
10  creation-date: yyyy-mm-dd
11  reviewers:
12    - TBD
13    - "@alice.doe"
14  approvers:
15    - TBD
16    - "@oscar.doe"
17
18  see-also:
19    - "/keps/sig-aaa/1234-we-heard-you-like-keps"
20    - "/keps/sig-bbb/2345-everyone-gets-a-kep"
21  replaces:
22    - "/keps/sig-ccc/3456-replaced-kep"
23
```





master

enhancements / keps / NNNN-kep-template / kep.yaml

Go to file



wojtek-t Cleanup prr-approvers field from kep.yaml

fb9265b · 2 years ago History

Code

Blame

49 lines (43 loc) · 1.22 KB

Raw



```
1 title: KEP Template
2 kep-number: NNNN
3 authors:
4   - "@jane.doe"
5   owning-sig: sig-xyz
6   participating-sigs:
7     - sig-aaa
8     - sig-bbb
9   status: provisional
10  creation-date: yyyy-mm-dd
11  reviewers:
12    - TBD
13    - "@alice.doe"
14  approvers:
15    - TBD
16    - "@oscar.doe"
17
18  see-also:
19    - "/keps/sig-aaa/1234-we-heard-you-like-keps"
20    - "/keps/sig-bbb/2345-everyone-gets-a-kep"
21  replaces:
22    - "/keps/sig-ccc/3456-replaced-kep"
23
```

Informed

👤 wojtek-t Cleanup prr-approvers field from kep.yaml | fb9265b · 2 years ago | 🕒 History

Code | Blame | 49 lines (4,123 bytes) · 1.22 KB | Raw | 📄 | 📥 | ✎ | ⌵ | 🗑️

```
1 title: KEP Template
2 kep-number: NNNN
3 authors:
4   - "@jane.doe"
5 owning-sig: sig-xyz
6 participating-sigs:
7   - sig-aaa
8   - sig-bbb
9 status: provisional|implementable|implemented|deferred|replaced
10 creation-date: yyyy-mm-dd
11 reviewers:
12   - TBD
13   - "@alice.doe"
14 approvers:
15   - TBD
16   - "@alice.doe"
17
18 see also:
19   - "/keps/sig-aaa/123456-headers-like-keps"
20   - "/keps/sig-bbb/123456-headers-like-keps"
21 replaces:
22   - "/keps/sig-ccc/3456-replaced-kep"
23
```

Informed

==

Community!

🔍 master | enhancements / keps / sig-architecture / 0000-kep-process / README.md | ↑ Top

Preview | Code | Blame | 343 lines (276 loc) · 16.1 KB | Raw | 📄 | ⬇️ | ✎ | ⌵ | ☰

KEP Metadata

There is a place in each KEP for a YAML document that has standard metadata. This will be used to support tooling around filtering and display. It is also critical to clearly communicate the status of a KEP.

While this defines the metadata schema for now, these things tend to evolve. The KEP template is the authoritative definition of things like the metadata schema.

Metadata items:

- **title** Required
 - The title of the KEP in plain language. The title will also be used in the KEP filename. See the template for instructions and details.
- **status** Required
 - The current state of the KEP.
 - Must be one of `provisional`, `implementable`, `implemented`, `deferred`, `rejected`, `withdrawn`, or `replaced`.
- **authors** Required
 - A list of authors of the KEP. This is simply the GitHub ID. In the future we may enhance this to include other types of identification.
- **owning-sig** Required
 - The SIG most closely associated with this KEP. If code or other artifacts will result from this KEP, then it is expected that this SIG will take responsibility for the bulk of those artifacts.
 - Sigs are listed as `sig-abc-def`, where the name matches the directory entry in the `kubernetes/community` repo.

master enhancements / keps / sig-architecture / 0000-kep-process / README.md ↑ Top

Preview Code Blame 343 lines (276 loc) · 16.1 KB Raw Copy Download Edit Menu

- **participating-sigs** Optional
 - A list of SIGs that are involved or impacted by this KEP.
 - A special value of `kubernetes-wide` will indicate that this KEP has impact across the entire project.
- **reviewers** Required
 - Reviewer(s) chosen after triage, according to the proposal process.
 - If not yet chosen, replace with `TBD`.
 - Same name/contact scheme as `authors`.
 - Reviewers should be a distinct set from authors.
- **approvers** Required
 - Approver(s) chosen after triage, according to the proposal process.
 - Approver(s) are drawn from the impacted SIGs. It is up to the individual SIGs to determine how they pick approvers for KEPs impacting them. The approvers are speaking for the SIG in the process of approving this KEP. The SIGs in question can modify this list as necessary.
 - The approvers are the individuals who decide when to move this KEP to the `implementable` state.
 - Approvers should be a distinct set from authors.
 - If not yet chosen, replace with `TBD`.
 - Same name/contact scheme as `authors`.
- **editor** Required
 - Someone to keep things moving forward.
 - If not yet chosen, replace with `TBD`.

master enhancements / keps / sig-architecture / 0000-kep-process / README.md ↑ Top

Preview Code Blame 343 lines (276 loc) · 16.1 KB Raw Copy Download Edit

- **creation-date** Required
 - The date that the KEP was first submitted in a PR.
 - In the form `yyyy-mm-dd` .
 - While this info will also be in source control, it is helpful to have the set of KEP files stand on their own.
- **last-updated** Optional
 - The date that the KEP was last changed significantly.
 - In the form `yyyy-mm-dd` .
- **see-also** Optional
 - A list of other KEPs that are relevant to this KEP.
 - In the form `KEP-123` .
- **replaces** Optional
 - A list of KEPs that this KEP replaces. Those KEPs should list this KEP in their `superseded-by` .
 - In the form `KEP-123` .
- **superseded-by**
 - A list of KEPs that supersede this KEP. Use of this should be paired with this KEP moving into the `Replaced` status.
 - In the form `KEP-123` .

KEP Workflow

A KEP has the following states:

Where to Find them!

Git and GitHub Implementation

KEPs are checked into the Enhancements repo under the `/keps` directory. KEPs in SIG specific subdirectories have limited impact outside of the SIG and can leverage SIG-specific OWNERS files.

New KEPs can be checked in with a file name in the form of `draft-YYYYMMDD-my-title.md`. As significant work is done on the KEP, the authors can assign a KEP number. No other changes should be put in that PR so that it can be approved quickly and minimize merge conflicts. The KEP number can also be done as part of the initial submission if the PR is likely to be uncontested and merged quickly.

Git and GitHub Implementation

KEPs are checked into the Enhancements repo under the `/keps` directory. `KEPs in SIG specific subdirectories` have limited impact outside of the SIG and can leverage SIG-specific OWNERS files.

New KEPs can be checked in with a file name in the form of `draft-YYYYMMDD-my-title.md`. As significant work is done on the KEP, the authors can assign a KEP number. No other changes should be put in that PR so that it can be approved quickly and minimize merge conflicts. The KEP number can also be done as part of the initial submission if the PR is likely to be uncontested and merged quickly.

House of the Kubernetes Enhancements:

git.k8s.io/enhancements/keps

enhancements Public

Edit Pins Watch 376 Fork 1.5k Star 3.5k

master 11 Branches 0 Tags Go to file Code

k8s-ci-robot Merge pull request #4984 from soltys... 82825fa · 2 days ago 8,783 Commits		
.githubhooks	.githubhooks/pre-commit: Fix shebang and si...	3 years ago
.github	.github: Add dependabot config	2 years ago
api	Update kep parser to get rid of prr-approv...	2 years ago
cmd	make test: stop hitting github for unit tests	3 years ago
docs	Update API review project links	last year
hack	spell: ignore YAML files	6 months ago
internal	Update to newer version of go-mod-outda...	2 years ago
keps	Merge pull request #4984 from soltysh/ku...	2 days ago



About

Enhancements tracking repo for Kubernetes

- kubernetes enhancements
- kubernetes-releases k8s-sig-architecture

- Readme
- Apache-2.0 license
- Code of conduct
- Security policy
- Activity
- Custom properties
- 3.5k stars
- 376 watching
- 1.5k forks



master

enhancements / keps /

Go to file



Add file



k8s-ci-robot Merge pull request #4887 from MrFreezeex/kep1645-port-rules

178cf35 · 1 minute ago History


Name	Last commit message	Last commit date
..		
NNNN-kep-template	Update KEP template skew questions	last year
prod-readiness	Merge pull request #4833 from sanposhiho/async-preemption	2 months ago
provider-aws	KEP-630: Fix typo in README.md	last year
sig-api-machinery	Clarify JSON Patch support in mutating admission policy	3 weeks ago
sig-apps	Update KEP-3335 kep.yaml to status: implemented	last month
sig-architecture	Merge pull request #4794 from pohly/compatibility-versions-linter	3 weeks ago
sig-auth	updating as per review	2 weeks ago
sig-autoscaling	Merge pull request #4389 from CatherineF-dev/patch-1	9 months ago
sig-cli	KEP-4706: update its status to rejected, after sig-arch discussion	5 days ago
sig-cloud-provider	[KEP-2395] Book keeping for cleaning up removal of in-tree cloud prov...	7 months ago
sig-cluster-lifecycle	Modify KEP based on implemented PR	last week

k8s-ci-robot Merge pull request #4887 from MrFreezeex/kep1645-port-rules 178cf35 · 1 minute ago History

Name	Last commit message	Last commit date
..		
NNNN-kep-template	Update KEP template skew questions	last year
prod-readiness	Merge pull request #4833 from sanposhiho/async-preemption	2 months ago
provider-aws	KEP-630: Fix typo in README.md	last year
sig-api-machinery	JSON Patch support for merging admission policy	3 weeks ago
sig-apps	Update KEP-1645: port rules	last month
sig-architecture	Merge pull request #4794 from pohly/compatibility-versions-linter	3 weeks ago
sig-auth	updating as per review	2 weeks ago
sig-autoscaling	Merge pull request #4389 from CatherineF-dev/patch-1	9 months ago
sig-cli	KEP-4706: update its status to rejected, after sig-arch discussion	5 days ago
sig-cloud-provider	[KEP-2395] Book keeping for cleaning up removal of in-tree cloud prov...	7 months ago
sig-cluster-lifecycle	Modify KEP based on implemented PR	last week

Fun Fact!



 **k8s-ci-robot** Merge pull request #4794 from pohly/compatibility-versions-linter 🗨️ ✓ 290206d · 3 weeks ago 🕒 History

Name	Last commit message	Last commit date
..		
0000-kep-process	keps: add latest-milestone and stage to implemented KEPs	3 years ago
1143-node-role-labels	regen tables of contents	11 months ago
1194-prod-readiness	keps: add latest-milestone and stage to implemented KEPs	3 years ago
1333-conformance-without-beta	Make PRR metadata required (#2672)	3 years ago
1618-conformance-profiles	Add section on roles related to certification	4 years ago
1635-prevent-permabeta	keps: add latest-milestone and stage to implemented KEPs	3 years ago
1659-standard-topology-labels	KEP 1659/2335: Correct typo from 'Kuberntes' to 'Kubernetes'	7 months ago
2527-clarify-status-observations-vs-rbac	Merge pull request #3121 from thockin/2527-clarify-status-observation...	5 months ago
3136-beta-apis-off-by-default	Mark KEP 3136 as implemented	2 years ago
3935-oldest-node-newest-control-plane	Mark 3935 as implemented	last year



Search this site...

Resources

Community Calendar

Services and Requests

Kubernetes 1.32 Release Information

Default Branch Migration

Enhancements

Kubernetes Enhancement Proposals (KEPs)

List of Kubernetes enhancements

These data come from the [kubernetes/enhancements](#) repository on GitHub.

To see the original Enhancement issue, click on the KEP number. To see a summary of what the KEP does, click on the Title.

Note

If you notice an empty cell in the table, the KEP may not have updated information. Please raise an [issue in the kubernetes/enhancements](#) repository.

KEP List

Show 10 entries

KEP Number	Title	SIG	Author	Created	Last Updated	Milestone
95	Graduate CustomResourceDefinitions to GA	API Machinery	@jpbetz , @roycai , @sttts	2018-04-15	2018-04-24	
365	Paginated API Lists	API Machinery	@smarterclayton , @wojtek-t	2017-08-29		alpha: v1.8 beta: v1.9 stable: v1.29
492	Graduate Admission Webhooks to GA	API Machinery	@mbohlool	2019-01-27	2019-02-04	
555	Anniv	API	@anelice , @ieffree	2018-03-28	2021-09-21	alpha: v1.14

www.kubernetes.dev/resources/keps/



Search this site...

Resources

Community Calendar

Services and Requests

Kubernetes 1.32 Release Information

Default Branch Migration

Enhancements

Kubernetes Enhancement Proposals (KEPs)

List of Kubernetes enhancements

These data come from the [kubernetes/enhancements](#) repository on GitHub.

To see the original Enhancement issue, click on the KEP number. To see a summary of what the KEP does, click on the Title.

Note

If you notice an empty cell in the table, the KEP may not have updated information. Please raise an [issue in the kubernetes/enhancements](#) repository.

KEP List

Show 10 entries

Search:

KEP Number	Title	SIG	Author	Created	Last Updated	Milestone
3960	Pod lifecycle sleep action	Node	@AxeZhan	2023-04-22		alpha: v1.29 beta: v1.30 stable: v1.32
4818	Allow zero value for Sleep Action of PreStop Hook	Node	@sreeram-venkitesh	2024-09-16		alpha: v1.32

Showing 1 to 2 of 2 entries (filtered from 525 total entries)

Previous Next

<end> 0000-kep-process </end>

Kubernetes

Feature (enhancement) stages!

Or ...

KEP stages!

Feature stages

A feature can be in *Alpha, Beta* or *GA* stage.

Feature stages

An *Alpha* feature means:

- Disabled by default.
- Might be buggy. Enabling the feature may expose bugs.
- Support for feature may be dropped at any time without notice.
- The API may change in incompatible ways in a later software release without notice.
- Recommended for use only in short-lived testing clusters, due to increased risk of bugs and lack of long-term support.

Feature stages

An *Alpha* feature means:

- Disabled by default.
- Might be buggy. Enabling the feature may expose bugs.
- Support for feature may be dropped at any time without notice.
- The API may change in incompatible ways in a later software release without notice.

--feature-gates = ... , YourAlphaFeatureGate=true

Recommended for use only in short-lived testing clusters, due to increased risk of bugs and lack of long-term support.

Feature stages

A *Beta* feature means:

- Usually **enabled by default**. Beta API groups are **disabled by default**.
- The feature is well tested. Enabling the feature is considered safe.
- Support for the overall feature will not be dropped, though details may change.

Feature stages

A *Beta* feature means:

- Usually enabled by default. Beta API groups are disabled by default.
- The feature is well tested. Enabling the feature is considered safe.
- Support for the overall feature will not be dropped, though details may change.

Feature stages

A *Beta* feature means:

- Usually enabled by default. Beta API groups are **disabled by default**.
- The feature is well tested. Enabling the feature is considered safe.
- Support for the overall feature will not be dropped, though details may change.

***[https://kubernetes.io/docs/reference/using-api/
#enabling-or-disabling](https://kubernetes.io/docs/reference/using-api/#enabling-or-disabling)*** 🖱️

Feature stages

A *General Availability (GA)* feature is also referred to as a *stable* feature. It means:

- The feature is always enabled; you cannot disable it.
- The corresponding feature gate is no longer needed.
- Stable versions of features will appear in released software for many subsequent versions.

Feature stages

A *General Availability* (GA) feature is also referred to as a *stable* feature. It means:

- The feature is always enabled; you cannot disable it.
- The corresponding feature gate is no longer needed.
- Stable versions of features will appear in released software for many subsequent versions.

THEORY ENDS!

***Let's see an actual
Kubernetes Enhancement Proposal***

*from its inception ...
to its fruition!*

pod LifecycleHandler should support a sleep option #114465

Edit New issue

Closed thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026



thockin commented on Dec 13, 2022 Member ...

Many people use `exec` to run `sleep` in their container. This works, but it's kludgy and requires them to have a sleep binary in the image. We should offer sleep as a first-class thing.

Assignees ⚙️

No one—[assign yourself](#)

Labels ⚙️

- area/api
- help wanted
- kind/feature
- needs-triage
- priority/backlog

pod LifecycleHandler should support a sleep option #114465

Edit New issue

✓ Closed thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026



thockin commented on Dec 13, 2022 Member ...

Many people use `exec` to run `sleep` in their container. This works, but it's kludgy and requires them to have a sleep binary in the image. We should offer sleep as a first-class thing.

Assignees ⚙️

No one—[assign yourself](#)

Labels ⚙️

- area/api
- help wanted
- kind/feature
- needs-triage
- priority/backlog



Pods-lifecycle-events-yaml-before-sleep-handler.yaml

```
apiVersion: v1
kind: Pod
metadata:
  name: lifecycle-demo
spec:
  containers:
  - name: lifecycle-demo-container
    image: myimage
    lifecycle:
      preStop:
        exec:
          command: ["/bin/sh", "-c", "while true; do sleep 10; done"]
```



Pods-lifecycle-events-yaml-with-sleep-handler.yaml

```
apiVersion: v1
kind: Pod
metadata:
  name: lifecycle-demo
spec:
  containers:
  - name: lifecycle-demo-container
    image: myimage
    lifecycle:
      preStop:
        sleep:
          seconds: 10
```

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

aojea commented on Dec 16, 2022

Member



It's kind of a hack, but it's what people are doing today and it works.

this sounds like a house of bugs for sig-node and sig-network, is not only the impact on the pod lifecycle that is worrisky, because there still a lot more work to have it completely right:

[#113606](#)

[#113717](#)

, is also the impact on things like probes, per example, there were important regressions that went unnoticed during several releases

more user stories from the community ...

pod LifecycleHandler should support a sleep option #114465

Edit New issue

Closed thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026



thockin commented on Dec 13, 2022 Member

Many people use `exec` to run `sleep` in their container. This works, but it's kludgy and requires them to have a sleep binary in the image. We should offer sleep as a first-class thing.

Assignees ⚙️
No one—[assign yourself](#)

Labels ⚙️
`area/api` `help wanted` `kind/feature`
`needs-triage` `priority/backlog`



✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

AxeZhan commented on Mar 29, 2023

Member

...

I'd like to try to solve this issue.

This looks challenging, so I'm a little afraid, if I drop out, someone else is welcome to take over.

But I got some questions as I never wrote `KEP` before.

1. How to determine the 4 digits in front of `kep` 😊
2. Does this need a `feature gate`? I read the [document](#) and it says we can simply add a new field. Do you mean we should add a feature gate, and make this feature default false in alpha?
3. For which sig should I put the `KEP`? sig/node?

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

AxeZhan commented on Mar 29, 2023

Member

...

I'd like to try to solve this issue.

This looks challenging, so I'm a little afraid, if I drop out, someone else is welcome to take over.

But I got some questions as I never wrote `KEP` before.

1. How to determine the 4 digits in front of `kep` 😊
2. Does this need a `feature gate`? I read the [document](#) and it says we can simply add a new field. Do you mean we should add a feature gate, and make this feature default false in alpha?
3. For which sig should I put the `KEP`? sig/node.

*People with time, bandwidth,
and expertise/willing to learn,
to get the work going.*

<https://github.com/kubernetes/kubernetes/issues/114465>

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

AxeZhan commented on Mar 29, 2023

Member

...

I'd like to try to solve this issue.

This looks challenging, so I'm a little afraid, if I drop out, someone else is welcome to take over.

But I got some questions as I never wrote `KEP` before.

1. How to determine the 4 digits in front of `kep` 😊
2. Does this need a `feature gate`? I read the [document](#) and it says we can simply add a new field. Do you mean we should add a feature gate, and make this feature default false in alpha?
3. For which sig should I put the `KEP`? sig/node?

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

How to determine the 4 digits in front of kep

Open an issue - whatever issue number you get is your KEP number.

Does this need a feature gate? I read the document and it says we can simply add a new field. Do you mean we should add a feature gate, and make this feature default false in alpha?

All new APIs need at least one alpha release. We can debate whether it needs a beta release, but alpha is a requirement.

For which sig should I put the KEP? sig/node?

yes

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

AxeZhan commented on Apr 6, 2023

Member



@thockin I wrote a [draft KEP](#), can you take a look and give me some advice first?



pod lifecycle sleep action ☆ 📄 🕒 ☁
File Edit View Tools Help

Request edit access Share

Sleep Action

axe zhan

Overview

This document discusses an action for a container which simply sleeps, to handle new connections from clients that have not yet finished endpoint termination during the pod termination.

Problem

The container and endpoint termination processes are asynchronously initiated when the pod is terminated. Because of this asynchronous nature, if the container process is stopped immediately, the clients might still try to connect to the container after the shutdown, causing connection errors during a rolling update.

For this problem, many people use exec to run sleep in their container. This works, but it's kludge and requires them to have a sleep binary in the image. We should offer sleep as a first-class thing.

This document suggests adding a sleepAction for **LifecycleHandler**, so users will be

Document tabs

- pod lifecycle sleep acti...
- Sleep Action
- Overview
- Problem
- Proposal Summary
- Detailed Design
 - SleepAction
 - LifeCycleHandler
 - handlerRunner.Run
- Feature Gate
- Graduation Criteria
 - V1.28 (alpha)
 - V1.29 (beta)
 - V1.30 (ga)
- Test Plan
 - Unit Tests
 - Integration Tests
 - E2e Tests
- Upgrade/Downgrade Strategy
 - Upgrade
 - Downgrade
- Production Readiness Review ...
 - How can this feature be ena...
 - Can the feature be disabled ...
 - Does enabling the feature c...

The image shows a Google Docs interface for a document titled "pod lifecycle sleep action". The document is shared and has a "Request edit access" button. The main content area displays the title "Sleep Action" by "axe zhan". A left sidebar shows a table of contents with sections like Overview, Problem, Proposal Summary, Detailed Design, Feature Gate, Graduation Criteria, Test Plan, Upgrade/Downgrade Strategy, and Production Readiness Review. The "Overview" section is currently selected and visible in the main content area.

Sleep Action

axe zhan

Overview

This document discusses an action for a container which simply sleeps, to handle new connections from clients that have not yet finished endpoint termination during the pod termination.

Problem

The container and endpoint termination processes are asynchronously initiated when the pod is terminated. Because of this asynchronous nature, if the container process is stopped immediately, the clients might still try to connect to the container after the shutdown, causing connection errors during a rolling update.

For this problem, many people use `exec` to run `sleep` in their container. This works, but it's kludge and requires them to have a `sleep` binary in the image. We should offer `sleep` as a first-class thing.

This document suggests adding a `sleepAction` for **LifecycleHandler**, so users will be

✓ Closed

pod LifecycleHandler should support a `sleep` option #114465

thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026

@thockin I've updated the doc. Added some test plan and Production Readiness Review, PTAL. If it's ok I'll create a KEP based on this doc later some time.

<https://docs.google.com/document/d/1IOXSBlzXQwPddOgqF5LIX4F9TwyATSSUrjLCgU1-Zhs/edit#heading=h.51a06ysnxedv>



thockin commented on Apr 21, 2023 via email 

Member

Author



Go for KEP


 Closed

pod LifecycleHandler should support a `sleep` option #114465
thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026



This was referenced on Apr 22, 2023

Introducing Sleep Action for PreStop Hook kubernetes/
enhancements#3960

 Open

KEP-3960: Introducing Sleep Action for PreStop Hook kubernetes/
enhancements#3961

 Merged



AxeZhan mentioned this issue on Jul 3, 2023

Open

9 of 12 tasks

Introducing Sleep Action for PreStop Hook #3960

AxeZhan opened this issue on Apr 22, 2023 · 81 comments

Enhancement Description

- One-line enhancement description (can be used as a release note): Add Sleep Action for PreStop Hook
- Kubernetes Enhancement Proposal: <https://github.com/kubernetes/enhancements/blob/master/keps/sig-node/3960-pod-lifecycle-sleep-action/>
- Discussion Link:
 - [related issue](#)
 - [google doc](#)
- Primary contact (assignee): [@AxeZhan](#) [@charles-chenzz](#)
- Responsible SIGs: sig-node
- Enhancement target (which target equals to which milestone):
 - Alpha release target (x.y): 1.29


 Closed

pod LifecycleHandler should support a `sleep` option #114465
thockin opened this issue on Dec 13, 2022 · 24 comments · Fixed by #119026



This was referenced on Apr 22, 2023

Introducing Sleep Action for PreStop Hook kubernetes/
enhancements#3960

 Open

KEP-3960: Introducing Sleep Action for PreStop Hook kubernetes/
enhancements#3961

 Merged



AxeZhan mentioned this issue on Jul 3, 2023

Open

9 of 12 tasks

Introducing Sleep Action for PreStop Hook #3960

AxeZhan opened this issue on Apr 22, 2023 · 81 comments

✓ Alpha

- ✓ KEP (`k/enhancements`) update PR(s): [🔗 KEP-3960: Introducing Sleep Action for PreStop Hook #3961](#)
- ✓ Code (`k/k`) update PR(s): [🔗 Introducing Sleep Action for PreStop Hook kubernetes#119026](#)
- ✓ Docs (`k/website`) update PR(s): [🔗 Add information about PodLifecycleSleepAction \(KEP3960\) website#43428](#)

📄 master [enhancements](#) / [keps](#) / [sig-node](#) / [3960-pod-lifecycle-sleep-action](#) / [README.md](#) ↑ Top

Preview Code Blame Raw Copy Download Edit

Summary

This KEP proposes the addition of a new sleep action for the PreStop lifecycle hook in Kubernetes, allowing containers to pause for a specified duration before termination. This enhancement aims to provide a more straightforward way to manage graceful shutdowns and improve the overall lifecycle management of containers, and to handle new connections from clients that have not yet finished endpoint termination during the pod termination.

An example:

To use this feature to achieve zero downtime for nginx, we need to deploy a deployment with sleep prestop hook and a service.

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Goals

- Allow containers to perform cleanup or shutdown actions before being terminated, by sleeping for a specified duration in the preStop hook.
- Improve the overall reliability and availability of Kubernetes applications by providing a way for containers to gracefully terminate.

🔗 Non-Goals

- This KEP does not aim to replace other Kubernetes features that can be used to perform cleanup actions, such as init containers or sidecar containers.
- This KEP does not aim to provide a way to pause or delay pod termination indefinitely.

📄 master [enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action](#) / README.md ↑ Top

Preview Code Blame Raw Copy Download Edit ⋮

Proposal

We propose adding a new sleep action for the PreStop hook, which will pause the container for a specified duration before termination. The API changes will include the following:

- Extending the LifecycleHandler object to support a new Sleep field.
- Adding a SleepAction object that includes a Duration field to specify the sleep period in seconds.

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

Preview Code Blame Raw Copy Download Edit

User Stories (Optional)

Story 1

As a Kubernetes user, I want to configure my container to sleep for a specific duration during graceful termination and I want to do it without needing a sleep binary in my image.

Story 2

As a Kubernetes user, I want to configure my nginx service to be able to run with zero downtime. Previously, I use `command: ["/bin/sh", "-c", "sleep 20"]` in prestop hook with exec command to delay the shutdown, and let the Endpoint Controller remove the pod first. But this requires me to have a sleep binary in my image, and I want to do this more conveniently.

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md ↑ Top

master

Preview Code Blame Raw Copy Download Edit

Design Details

Implementation

- Adding a SleepAction object that includes a Duration field to specify the sleep period in seconds.

```
type SleepAction struct {
    // Seconds is the number of seconds to sleep.
    Seconds int32
}
```

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Design Details

- Adding a Sleep field to the LifecycleHandler struct, which represents the duration in seconds that the container should sleep before being terminated during the preStop hook.

```
type LifecycleHandler struct {
    // Sleep represents the duration in seconds that the container should s
    Sleep *SleepAction
}
```

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Design Details

- When Kubernetes executes the preStop hook with sleep action, it'll simply sleep for a specific seconds.

```
func (hr *handlerRunner) Run(ctx context.Context, containerID kubecontainer.ContainerID) error {
    switch {
    case handler.Exec != nil:
    case handler.HTTPGet != nil:
    case handler.Sleep != nil:
        hr.runSleepHandler(ctx, handler.Sleep.Seconds)
        return "", nil
    default:
    }
}
```

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Design Details

```
func (hr *handlerRunner) runSleepHandler(ctx context.Context, seconds int32) {
    c := time.After(time.Duration(seconds) * time.Second)
    select {
    case <-ctx.Done():
        // early termination
        // some logs
        return
    case <-c:
        // sleep expired
        // some logs
        return
    }
}
```

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Test Plan

[x] I/we understand the owners of the involved components may require updates to existing tests to make this code solid enough prior to committing the changes necessary to implement this enhancement.

Prerequisite testing updates

Unit tests

alpha:

- Test that the runSleepHandler function sleeps for the correct duration when given a valid duration value.
- Test that the runSleepHandler function returns without error when given a valid duration value.

enhancements / k8s / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Graduation Criteria

Alpha

- Feature implemented behind a feature flag
- Initial unit/e2e tests completed and enabled
- Documentation is added to demonstrate why this is useful in nginx scenario and how exactly nginx needs to be configured.

📄 master [enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action](#) / README.md ↑ Top

Preview Code Blame Raw Copy Download Edit ⋮

Upgrade / Downgrade Strategy

Upgrade

The previous PreStop behavior will not be broken. Users can continue to use their hooks as it is. To use this enhancement, users need to enable the feature gate, and add sleep action in their prestop hook.

Downgrade

It will silently drop the sleep prestop hook if someone wants to add that to the object (or create object with it). However, existing objects that have it set will not be cleared.

📄 master [enhancements](#) / [k8s](#) / [sig-node](#) / [3960-pod-lifecycle-sleep-action](#) / README.md ↑ Top

Preview | Code | Blame Raw | 📄 | ⬇️ | ✎ | ⌵ | ☰

Production Readiness Review Questionnaire

Feature Enablement and Rollback

How can this feature be enabled / disabled in a live cluster?

- Feature gate (also fill in values in `kep.yaml`)
 - Feature gate name: PodLifecycleSleepAction
 - Components depending on the feature gate: kubelet,kube-apiserver

Does enabling the feature change any default behavior?

No

enhancements / keps / sig-node / 3960-pod-lifecycle-sleep-action / README.md

↑ Top

master

Preview Code Blame

Raw Copy Download Edit

Alternatives

Another way to run `sleep` in a container is to use `exec` command in `preStop` hook like `command: ["/bin/sh", "-c", "sleep 20"]`. However, this requires a sleep binary in the image. We should offer sleep as a first-class thing.

KEP-3960 is merged ...

... now, to code implementation!

Introducing Sleep Action for PreStop Hook #119026

Edit

<> Code

Merged k8s-ci-robot merged 5 commits into [kubernetes:master](#) from [AxeZhan:sleepAction](#) on Oct 17, 2023

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994



AxeZhan commented on Jul 3, 2023 • edited

Member

What type of PR is this?

/kind feature

What this PR does / why we need it:

This is the implementation for [kubernetes/enhancements#3960](#)

Which issue(s) this PR fixes:

Fixes [#114465](#)

Special notes for your reviewer:

Does this PR introduce a user-facing change?

A new sleep action for the PreStop lifecycle hook is added, allowing containers to

Reviewers

- dipan-db
- charles-chenzz
- aojea
- kannon92
- pacoxu
- bart0sh
- caesarxuchao
- gjkim42
- bobbypage
- mrunalp
- thockin

Assignees

Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a



validation

AxeZhan committed on Oct 15, 2023

fa01dfd



implementation

AxeZhan committed on Oct 15, 2023

3a96afd



add a simple e2e test

AxeZhan committed on Oct 15, 2023

332779a



Commits on Oct 16, 2023

codegen

AxeZhan committed on Oct 16, 2023

5c0618b



Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a

Merged

Introducing Sleep Action for PreStop Hook #119026



Changes from 1 commit

File filter

Conversations

Jump to



Review changes

10 pkg/apis/core/types.go

```
2143 // ExecAction describes a "run in container"
2144 // action.
2145 type ExecAction struct {
2146     // Command is the command line to execute
2147     // inside the container, the working directory for the
2148     // command is root ('/') in the
2149     // container's filesystem. The command is simply
2150     // exec'd, it is
2151     // not run inside a shell, so traditional
2152     // shell instructions ('|', etc) won't work. To use
2153     // a shell, you need to explicitly call out
2154     // to that shell.
2155     // +optional
2156     Command []string
2157 }
```

```
2143 // ExecAction describes a "run in container"
2144 // action.
2145 type ExecAction struct {
2146     // Command is the command line to execute
2147     // inside the container, the working directory for the
2148     // command is root ('/') in the
2149     // container's filesystem. The command is simply
2150     // exec'd, it is
2151     // not run inside a shell, so traditional
2152     // shell instructions ('|', etc) won't work. To use
2153     // a shell, you need to explicitly call out
2154     // to that shell.
2155     // +optional
2156     Command []string
2157 }
```

```
2153 + // SleepAction describes a "sleep" action.
2154 + type SleepAction struct {
2155 +     // Seconds is the number of seconds to
2156 +     // sleep.
2157 +     Seconds int64
2158 }
```

Merged

Introducing Sleep Action for PreStop Hook #119026

Review changes

Changes from 1 commit File filter Conversations Jump to

10 pkg/apis/core/types.go

```
2423 type LifecycleHandler struct {
2424     // Exec specifies the action to take.
2425     // +optional
2426     Exec *ExecAction
2427     // HTTPGet specifies the http request to
2428     // perform.
2429     // +optional
2430     HTTPGet *HTTPGetAction
2431     // Deprecated. TCPSocket is NOT supported
2432     // as a LifecycleHandler and kept
2433     // for the backward compatibility. There
2434     // are no validation of this field and
2435     // lifecycle hooks will fail in runtime
2436     // when tcp handler is specified.
2437     // +optional
2438     TCPSocket *TCPSocketAction
```

```
2429 type LifecycleHandler struct {
2430     // Exec specifies the action to take.
2431     // +optional
2432     Exec *ExecAction
2433     // HTTPGet specifies the http request to
2434     // perform.
2435     // +optional
2436     HTTPGet *HTTPGetAction
2437     // Deprecated. TCPSocket is NOT supported
2438     // as a LifecycleHandler and kept
2439     // for the backward compatibility. There
2440     // are no validation of this field and
2441     // lifecycle hooks will fail in runtime
2442     // when tcp handler is specified.
2443     // +optional
2444     TCPSocket *TCPSocketAction
2445     + // Sleep represents the duration that the
2446     + // container should sleep before being terminated.
2447     + // +featureGate=PodLifecycleSleepAction
2448     + // +optional
2449     + Sleep *SleepAction
```


Merged

Introducing Sleep Action for PreStop Hook #119026

Review changes

Changes from 1 commit File filter Conversations Jump to

```
9 pkg/features/kube_features.go

@@ -636,6 +636,13 @@ const (
    // Adds pod.status.hostIPs and downward API
    PodHostIPs featuregate.Feature =
        "PodHostIPs"

    // owner: @AxeZhan
    // kep: http://kep.k8s.io/3960
    // alpha: v1.29
    // Enables SleepAction in container
    // lifecycle hooks
    PodLifecycleSleepAction featuregate.Feature =
        "PodLifecycleSleepAction"

    // owner: @Huang-Wei
    // kep: https://kep.k8s.io/3521
    // alpha: v1.26

@@ -1063,6 +1070,8 @@ var defaultKubernetesFeatureGates = map[featuregate.Feature]featuregate.FeatureS
```

Merged

Introducing Sleep Action for PreStop Hook #119026

Review changes

Changes from 1 commit File filter Conversations Jump to

9 pkg/features/kube_features.go

```
639 + // owner: @AxeZhan
640 + // kep: http://kep.k8s.io/3960
641 + // alpha: v1.29
642 + //
643 + // Enables SleepAction in container
    lifecycle hooks
644 + PodLifecycleSleepAction featuregate.Feature
    = "PodLifecycleSleepAction"
645 +
639 // owner: @Huang-Wei
640 // kep: https://kep.k8s.io/3521
641 + // alpha: v1.26
646 // owner: @Huang-Wei
647 // kep: https://kep.k8s.io/3521
648 // alpha: v1.26
@@ -1063,6 +1070,8 @@ var defaultKubernetesFeatureGates = map[featuregate.Feature]featuregate.FeatureS
1063
1064 PodHostIPs: {Default: false, PreRelease:
    featuregate.Alpha},
1065
1070
1071 PodHostIPs: {Default: false, PreRelease:
    featuregate.Alpha},
1072
1073 + PodLifecycleSleepAction: {Default: false,
    PreRelease: featuregate.Alpha},
1074 +
```

Merged

Introducing Sleep Action for PreStop Hook #119026

Review changes

Changes from 1 commit File filter Conversations Jump to

10 staging/src/k8s.io/api/core/v1/types.go

```
@@ -2272,6 +2272,12 @@ type ExecAction struct {
2272     Command []string `json:"command,omitempty"
      protobuf:"bytes,1,rep,name=command"`
2273 }
2274
2275 + // SleepAction describes a "sleep" action.
2276 + type SleepAction struct {
2277 +     // Seconds is the number of seconds to
      sleep.
2278 +     Seconds int64 `json:"seconds"
      protobuf:"bytes,1,opt,name=seconds"`
2279 + }
2280 +
2275 // Probe describes a health check to be performed
      against a container to determine whether it is
2276 // alive or ready to receive traffic.
2277 type Probe struct {
2281 // Probe describes a health check to be performed
      against a container to determine whether it is
2282 // alive or ready to receive traffic.
2283 type Probe struct {
@@ -2667,6 +2673,10 @@ type LifecycleHandler struct {
2667     // lifecycle hooks will fail in runtime
      when tcp handler is specified.
2673     // lifecycle hooks will fail in runtime
      when tcp handler is specified.
```

Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59 Commits 5 Checks 0 Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a

validation

AxeZhan committed on Oct 15, 2023

fa01dfd

Merged

Introducing Sleep Action for PreStop Hook #119026

Review changes

Changes from 1 commit File filter Conversations Jump to

73 pkg/apis/core/validation/validation.go

```
2987  
2988 func validateClientIPAffinityConfig(config  
*core.SessionAffinityConfig, fldPath *field.Path)  
field.ErrorList {  
2989     allErrs := field.ErrorList{}  
2990     if config == nil {
```

```
2989  
2990 + func validateSleepAction(sleep *core.SleepAction,  
gracePeriod int64, fldPath *field.Path)  
field.ErrorList {  
2991 +     allErrors := field.ErrorList{}  
2992 +     if sleep.Seconds <= 0 || sleep.Seconds >  
gracePeriod {  
2993 +         invalidStr := fmt.Sprintf("must be  
greater than 0 and less than  
terminationGracePeriodSeconds (%d)", gracePeriod)  
2994 +         allErrors = append(allErrors,  
field.Invalid(fldPath, sleep.Seconds, invalidStr))  
2995 +     }  
2996 +     return allErrors  
2997 + }  
2998 +  
2999 func validateClientIPAffinityConfig(config  
*core.SessionAffinityConfig, fldPath *field.Path)  
field.ErrorList {  
3000     allErrs := field.ErrorList{}  
3001     if config == nil {
```

Merged

Introducing Sleep Action for PreStop Hook #119026



Changes from 1 commit

File filter

Conversations

Jump to



Review changes

73 pkg/apis/core/validation/validation.go



@@ -3128,19 +3139,27 @@ func validateHandler(handler commonHandler, fldPath *field.Path) field.ErrorList

```
3128         allErrors =
append(allErrors, validateGRPCAction(handler.GRPC,
fldPath.Child("grpc")))...)
3129     }
3130 }
```

```
3139         allErrors =
append(allErrors, validateGRPCAction(handler.GRPC,
fldPath.Child("grpc")))...)
3140     }
3141 }
3142 +     if handler.Sleep != nil {
3143 +         if numHandlers > 0 {
3144 +             allErrors =
append(allErrors,
field.Forbidden(fldPath.Child("sleep"), "may not
specify more than 1 handler type"))
3145 +         } else {
3146 +             numHandlers++
3147 +             allErrors =
append(allErrors,
validateSleepAction(handler.Sleep, gracePeriod,
fldPath.Child("sleep")))...)
3148 +         }
3149 +     }
```

Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a



validation

AxeZhan committed on Oct 15, 2023

fa01dfd



implementation

AxeZhan committed on Oct 15, 2023

3a96afd



Merged

Introducing Sleep Action for PreStop Hook #119026



Changes from 1 commit

File filter

Conversations

Jump to



Review changes

23 pkg/kubelet/lifecycle/handlers.go

119

128

```
129 + func (hr *handlerRunner) runSleepHandler(ctx
    context.Context, seconds int64) error {
130 +     if !
        utilfeature.DefaultFeatureGate.Enabled(features.PodL
        ifecycleSleepAction) {
131 +         return nil
132 +     }
133 +     c := time.After(time.Duration(seconds) *
        time.Second)
134 +     select {
135 +     case <-ctx.Done():
136 +         // unexpected termination
137 +         return fmt.Errorf("container
        terminated before sleep hook finished")
138 +     case <-c:
139 +         return nil
140 +     }
141 + }
142 +
```

120 func (hr *handlerRunner) runHTTPHandler(ctx

143 func (hr *handlerRunner) runHTTPHandler(ctx

Merged

Introducing Sleep Action for PreStop Hook #119026

Changes from 1 commit File filter Conversations Jump to

Review changes

23 pkg/kubelet/lifecycle/handlers.go

@@ -86,6 +87,14 @@ func (hr *handlerRunner) Run(ctx context.Context, containerID kubecontainer.Cont

```
86         klog.V(1).Errors(err, "HTTP
lifecycle hook for Container in Pod failed", "path",
handler.HTTPGet.Path, "containerName",
container.Name, "pod", klog.KObj(pod))
87     }
88     return msg, err
```

```
87         klog.V(1).Errors(err, "HTTP
lifecycle hook for Container in Pod failed", "path",
handler.HTTPGet.Path, "containerName",
container.Name, "pod", klog.KObj(pod))
88     }
89     return msg, err
90 +     case handler.Sleep != nil:
91 +         err := hr.runSleepHandler(ctx,
handler.Sleep.Seconds)
92 +         var msg string
93 +         if err != nil {
94 +             msg = fmt.Sprintf("Sleep
lifecycle hook (%d) for Container %q in Pod %q
failed - error: %v", handler.Sleep.Seconds,
container.Name, format.Pod(pod), err)
95 +             klog.V(1).Errors(err, "Sleep
lifecycle hook for Container in Pod failed",
"sleepSeconds", handler.Sleep.Seconds,
"containerName", container.Name, "pod",
```

Merged

Introducing Sleep Action for PreStop Hook #119026



Changes from 1 commit

File filter

Conversations

Jump to



Review changes

58 pkg/api/pod/util.go

```
599 + func podLifecycleSleepActionInUse(podSpec
    *api.PodSpec) bool {
600 +     if podSpec == nil {
601 +         return false
602 +     }
603 +     var inUse bool
604 +     VisitContainers(podSpec, AllContainers,
    func(c *api.Container, containerType ContainerType)
    bool {
605 +         if c.Lifecycle == nil {
606 +             return true
607 +         }
608 +         if c.Lifecycle.PreStop != nil &&
    c.Lifecycle.PreStop.Sleep != nil {
609 +             inUse = true
610 +             return false
611 +         }
612 +         if c.Lifecycle.PostStart != nil &&
    c.Lifecycle.PostStart.Sleep != nil {
613 +             inUse = true
614 +             return false
```

Merged

Introducing Sleep Action for PreStop Hook #119026



Changes from 1 commit

File filter

Conversations

Jump to



Review changes

58 pkg/api/pod/util.go

```
562 +     if !
utilfeature.DefaultFeatureGate.Enabled(features.PodL
ifecycleSleepAction) && !
podLifecycleSleepActionInUse(oldPodSpec) {
563 +         for i := range podSpec.Containers {
564 +             if
podSpec.Containers[i].Lifecycle == nil {
565 +                 continue
566 +             }
567 +             if
podSpec.Containers[i].Lifecycle.PreStop != nil {
568 +
podSpec.Containers[i].Lifecycle.PreStop.Sleep = nil
569 +             }
570 +             if
podSpec.Containers[i].Lifecycle.PostStart != nil {
571 +
podSpec.Containers[i].Lifecycle.PostStart.Sleep =
nil
572 +             }
573 +         }
```

Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a



validation

AxeZhan committed on Oct 15, 2023

fa01dfd



implementation

AxeZhan committed on Oct 15, 2023

3a96afd



add a simple e2e test

AxeZhan committed on Oct 15, 2023

332779a



Merged

Introducing Sleep Action for PreStop Hook #119026

Changes from 1 commit ▾ File filter ▾ Conversations ▾ Jump to ▾ ⚙ ▾


Review changes ▾

add a simple e2e test

< Prev Next >

master (#119026)

v1.33.0-alpha.0 ... v1.29.0-alpha.3

 **AxeZhan** committed on Oct 15, 2023


commit 332779a3b29c3007d22b33865547401951dcc6d5


49 test/e2e/common/node/lifecycle_hook.go

```
@@ -544,3 +544,52 @@ func getSidecarPodWithHook(name string, image string, lifecycle *v1.Lifecycle) *
544         },
545     }
546 }
547 +
548 + var _ =
549 +     SIGDescribe("[Feature:PodLifecycleSleepAction]",
550 +         func() {
551 +             f := framework.NewDefaultFramework("pod-
552 +                 lifecycle-sleep-action")
553 +             f.NamespacePodSecurityEnforceLevel =
554 +                 admissionapi.LevelBaseline
555 +             var podClient *e2epod.PodClient
```



The Hitchhiker's Guide to Testing Kubernetes - Sreeram Venkitesh, SIG Testing


[Click here to add to My Schedule.](#)


 Tuesday December 10, 2024 10:00am - 10:35am IST


 Convention Centre | Level 4 | Room 404

A vital part of contributing to Kubernetes is testing. Reading existing tests lets you understand how a particular component works. You're also expected to write tests when making changes to the code base. The K8s codebase has different kinds of tests like unit, end to end and conformance tests. Given the scale of K8s, all of these are important to make sure that a feature works as intended and no edge cases come up when all the components are deployed in a production environment.

If you're getting started with contributing to K8s, learning the entire testing infrastructure can be overwhelming. This talk aims to give an introduction of the different types of tests in K8s and how you can setup your dev environment to run them locally. Along with unit tests, the talk also aims to give a demo of setting up tools like kubetest2 and hydrophone to run e2e and conformance tests with a kind cluster. After the talk you'll be ready to write your own tests for K8s and improve the test coverage.

Speakers



Sreeram Venkitesh

Software Engineer, BigBinary

Sreeram Venkitesh is a Software Engineer at BigBinary and is an active contributor to Kubernetes. He is part of the 1.29 release team.

Introducing Sleep Action for PreStop Hook #119026

Edit <> Code

Merged k8s-ci-robot merged 5 commits into `kubernetes:master` from `AxeZhan:sleepAction` on Oct 17, 2023

Preview Give feedback

Conversation 59

Commits 5

Checks 0

Files changed 81

+2,331 -994

Commits on Oct 15, 2023

api changes

AxeZhan committed on Oct 15, 2023

9327f3a



validation

AxeZhan committed on Oct 15, 2023

fa01dfd



implementation

AxeZhan committed on Oct 15, 2023

3a96afd



add a simple e2e test

AxeZhan committed on Oct 15, 2023

332779a



Commits on Oct 16, 2023

codegen

AxeZhan committed on Oct 16, 2023

5c0618b



codegen master (#119026)
v1.33.0-alpha.0 v1.29.0-alpha.3

AxeZhan committed on Oct 16, 2023 commit 5c0618b848d9e0d7fbc2d7a5c41e85735c173fd

18 api/openapi-spec/swagger.json

Load diff
Some generated code...

23 api/openapi-spec/v3/api_v1_openapi.json
2901],
2902 "description": "HTTPGet specifies the
http request to perform."

Auto-generated code!
From your local clone of git.k8s.io/kubernetes –

\$./hack/update-codegen.sh

Demo!

KEP-3960 is implemented ...

next ... documenting new API!

Add information about PodLifecycleSleepAction (KEP3960) #43428

Edit <> Code

Merged k8s-ci-robot merged 1 commit into `kubernetes:dev-1.29` from `AxeZhan:sleepAction` on Oct 20, 2023

Conversation 27 Commits 1 Checks 3 Files changed 2 +6 -1



AxeZhan commented on Oct 11, 2023 • edited Member
PodLifecycleSleepAction updates as part of [kubernetes/enhancements#3960](#)

- Reviewers
- thockin
 - tengqm
 - mengjiao-liu
 - mikedanese
 - reylejano

k8s-ci-robot added the `do-not-merge/work-in-progress` label on Oct 11, 2023

k8s-ci-robot added this to the `1.29` milestone on Oct 11, 2023

- Assignees
- thockin



netlify bot commented on Oct 11, 2023 • edited

Labels

Merged

Add information about PodLifecycleSleepAction (KEP3960) #43428

Changes from all commits File filter Conversations Jump to

0 / 2 files viewed

Review changes

```
content/en/docs/reference/command-line-tools-reference/feature-gates.md

@@ -170,6 +170,7 @@ For a reference to old feature gates that are removed, please refer to
170 | `PodDisruptionConditions` | `true` | Beta | 1.26 |
171 | `PodHostIPs` | `false` | Alpha | 1.28 |
172 | `PodIndexLabel` | `true` | Beta | 1.28 |
173 + | `PodLifecycleSleepAction` | `false` | Alpha | 1.29 |
173 | `PodReadyToStartContainersCondition` | `false` |
Alpha | 1.28 |
174 | `PodSchedulingReadiness` | `false` | Alpha | 1.26
| 1.26 |
175 | `PodSchedulingReadiness` | `true` | Beta | 1.27 |
176 + | `PodSchedulingReadiness` | `true` | Beta | 1.27 |

@@ -664,6 +665,7 @@ Each feature gate is designed for enabling/disabling a specific feature:
664 - `PodHostIPs`: Enable the `status.hostIPs` field
665 - `PodHostIPs`: Enable the `status.hostIPs` field
```

***Also ... look for
post-release blogs for
user-friendly docs!***

To Wrap Up 📺

Some things we learned ...

- The KEP Template README, when viewed in its raw/code form is very helpful
 - <https://git.k8s.io/enhancements/keps/NNNN-kep-template/README.md>

- API compatibility, both upward and downward is taken really seriously
 - https://git.k8s.io/community/contributors/devel/sig-architecture/api_changes.md
- New API added behind feature gate
- Feature gate is off by default at alpha
- If someone uses the feature when gate is off -> default / fall back to safe settings / or drop fields automatically etc

How Feature Gates work in code

- Alpha gated code, off by default
- Beta gated code, on by default
- Stable - gate removed, logic no longer wrapped in feature gate check

Things that came in handy

- The ability to work with Kind tooling.
Lets us test changes quickly.

Like the ability to create Node Images etc ...)

- <https://kind.sigs.k8s.io/docs/user/configuration/>
- <https://kind.sigs.k8s.io/docs/design/node-image/>


- **Contact Info!**
 - <https://git.k8s.io/enhancements/keps/README.md#faqs>
 - If in doubt ask SIG Architecture and they can advise.
- **Thank you!**

To Wrap It All Up ...




DON'T FORGET

DRINK WATER.



GET SUNLIGHT.



YOU'RE
BASICALLY A
HOUSE PLANT
WITH MORE
COMPLICATED
EMOTIONS.



Contact & Slides

slack.k8s.io:

#sig-architecture

#enhancements

Slack handles:

@psaggu,

@jasonbraganza

Email:

priyankasaggu11929@gmail.com

jason@janusworx.com

<https://sched.co/1mVS3>